

MariaDB

**drop in replacement of
MySQL**

Agenda

- MySQL branch
- GPL licence
- Maria storage engine
- Virtual columns
- FederatedX storage engine
- PBXT storage engine
- XtraDB storage engine
- Additional storage engines
- Slow query log extended statistics
- Microsecond precision in processlist
- Table elimination optimization
- Thread pool support
- utf8_croatian_ci, ucs2_croatian_ci collations
- Compatibility

MySQL branch

MySQL branch, developed by Monty Program.

Versions:

- Last 5.1 version: 5.1.50 (09.09.2010)
- Last 5.2 version: 5.2.3 (10.11.2010)

OS:

- Debian
- Ubuntu
- CentOS / Redhat
- Windows
- Solaris i386
- Gentoo (5.1.47)
- Any (source)

GPL licence

- Will be always on GPL license

Maria storage engine

- Rename to Aria.
- Based on MyISAM code.
- The goal is “transactional MyISAM”.
- Version 1.X is “atomic” and “crash-safe”.
- Version 2.X will be transactional.
- Used for SQL runtime's temporary tables.
- Tables can be copy/move as files. (*.frm, *.mai, *.mad)
- Full text indexes are supported.
- “Merge”-like functionality still not implemented.

Maria storage engine

```
CREATE TABLE ppl(  
  id int primary key,  
  name varchar(100),  
  country char(2)  
)  
ENGINE maria  
ROW_FORMAT=page  
TRANSACTIONAL=1;
```

Maria storage engine

```
CREATE TABLE t1 (a int)  
ROW_FORMAT=FIXED;
```

```
CREATE TABLE t2 (a int)  
ROW_FORMAT=DYNAMIC;
```

Virtual columns

- Available in MariaDB 5.2
- Available for MyISAM and MariaDB
- Persistent (stored) and virtual (not stored)

```
CREATE TABLE table1 (  
  a int not null primary key,  
  b varchar(32),  
  c int as (a mod 5) virtual,  
  d int as (a mod 7) persistent  
);
```

FederatedX storage engine

- Oracle stopped support for Federated storage engine [can not be confirmed].
- Developed by original author, Patrick Galbraith.
- An improved version of MySQL's Federated storage engine.
- Transaction support.

PBXT storage engine

- Developed by PrimeBase Technologies (<http://www.primebase.org/>).
- Transactional, ACID-compliant, Multi Version.
- “Write once” log architecture.
- Engine level replication.
- Competitor to InnoDB.
- Current version – 1.0.11 - Pre GA.

XtraDB storage engine

- Developed by Percona (<http://www.percona.com/>).
- Same people behind <http://mysqlperformanceblog.com/>
- Percona Server – another MySQL branch.
- InnoDB plugin + patches.
- Improvements for multi-cpu, fast start, more diagnostic and statistics.

Additional storage engines

- OQGRAPH

<http://openquery.com/graph/doc>

- SphinxSE

<http://sphinxsearch.com/>

- NDB is disabled.

Slow query log extended statistics (by Percona)

```
# Time: ...
# User@Host: root[root] @ localhost []
# Thread_id: 1 Schema: test QC_hit: No
# Query_time: 4.605642 Lock_time: 0.000964 Rows_sent: 1
  Rows_examined: 10
# Full_scan: Yes Full_join: No Tmp_table: No Tmp_table_
  on_disk: No
# Filesort: No Filesort_on_disk: No Merge_passes: 0
SET timestamp=...;
select count(*) from ppl;
```

Microsecond precision in processlist (by Percona)

```
# MySQL [(test)]> select * from information_schema.processlist;
```

ID	USER	HOST	DB	COMMAND	TIME	STATE	INFO
2	root	localhost	test	Query	2	Sending data	select

```
MariaDB [(test)]> select * from information_schema.processlist;
```

ID	USER	HOST	DB	COMMAND	TIME	STATE	INFO	TIME_MS
2	root	localhost	test	Query	0	executing	select	1.363

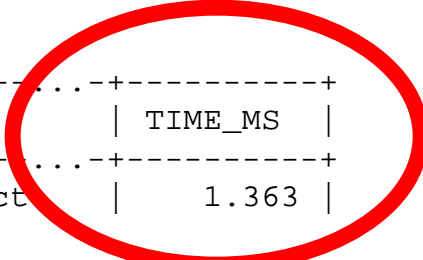


Table elimination optimization

- Developed by Monty Program

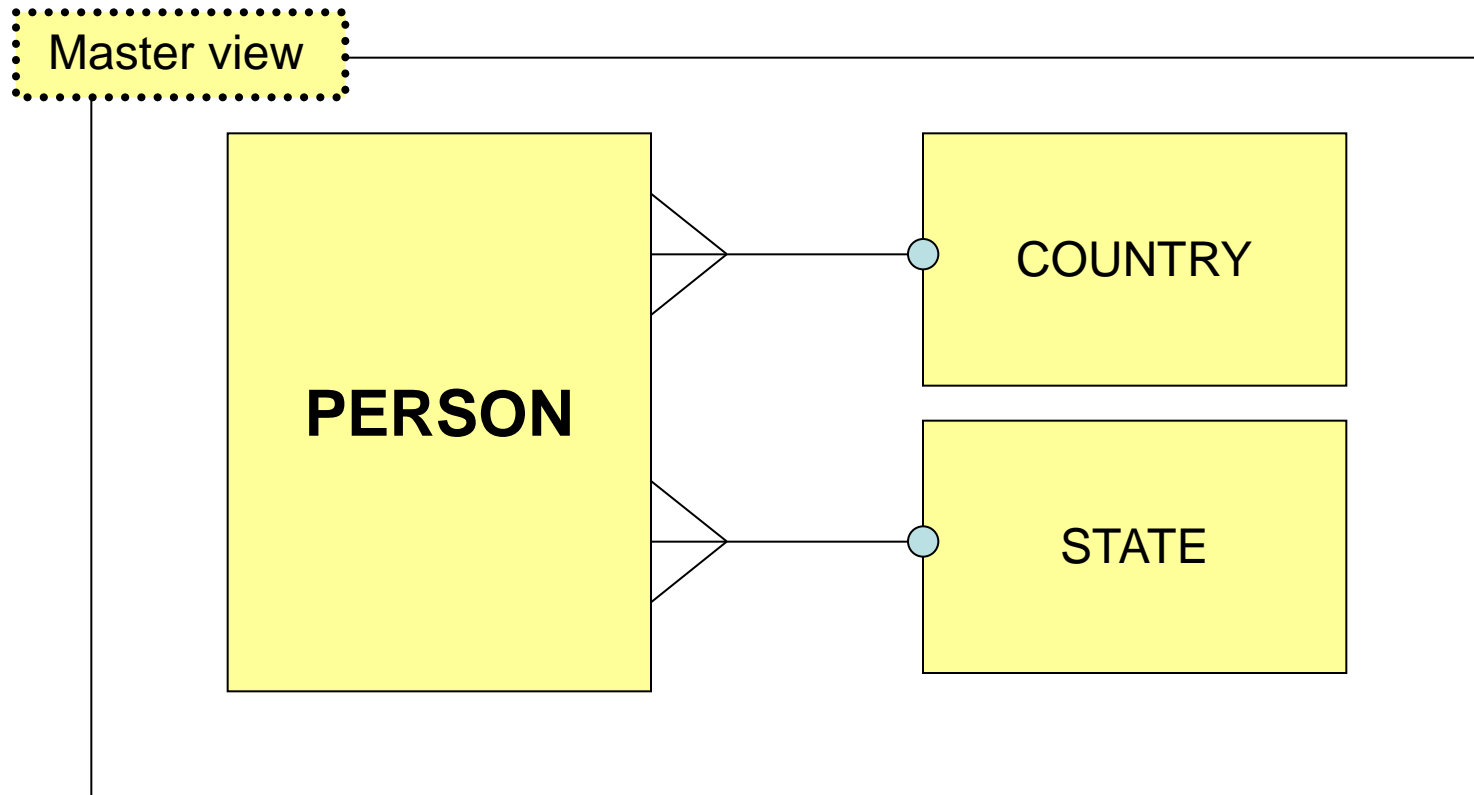


Table elimination optimization

```
mysql> create table a1 (id int primary key);  
Query OK, 0 rows affected (0.04 sec)
```

```
mysql> create table a2 (id int primary key);  
Query OK, 0 rows affected (0.01 sec)
```

```
mysql> insert into a1 values(1),(2),(3);  
Query OK, 3 rows affected (0.00 sec)  
Records: 3 Duplicates: 0 Warnings: 0
```

```
mysql> insert into a2 values(1);  
Query OK, 1 row affected (0.00 sec)
```




Table elimination optimization

```
mysql> explain
select
  a1.id
from
  a1 left outer join
  a2
on
  a1.id = a2.id;
```

id	select_type	table	type	possible_keys	key	key_len	ref	rows	Extra
1	SIMPLE	a1	index	NULL	PRIMARY	4	NULL	3	Using index
1	SIMPLE	a2	eq_ref	PRIMARY	PRIMARY	4	test.a1.id	1	Using index

```
2 rows in set (0.04 sec)
```

Hmmm...

Table elimination optimization

```
mysql> explain
select
  a1.id
from
  a1 left outer join
  a2
on
  a1.id = a2.id;
```

MySQL scans both tables.

id	select_type	table	type	possible_keys	key	key_len	ref	rows	Extra
1	SIMPLE	a1	index	NULL	PRIMARY	4	NULL	3	Using index
1	SIMPLE	a2	eq_ref	PRIMARY	PRIMARY	4	test.a1.id	1	Using index

2 rows in set (0.04 sec)

```
mysql> select version();
```

```
+-----+
| version() |
+-----+
| 5.0.44    |
+-----+
```

1 row in set (0.00 sec)

Table elimination optimization

```
MariaDB [test]> create view a3 as
select
  a1.id as a1id,
  a2.id as a2id
from
  a1 left outer join
  a2
on
  a1.id = a2.id;
Query OK, 0 rows affected (2.18 sec)
```

MariaDB scans single table.

```
MariaDB [test]> explain select a1id from a3;
```

id	select_type	table	type	possible_keys	key	key_len	ref	rows	Extra
1	SIMPLE	a1	index	NULL	PRIMARY	4	NULL	3	Using index

```
2 rows in set (0.04 sec)
```

Thread pool support

Conventional model:

- 1 user = 1 thread
- N users = N threads

Thread pool model:

- N users = M threads

utf8_croatian collations

- utf8_croatian_ci
- ucs2_croatian_ci

Compatibility

- Client libraries = YES
- Client-server protocol = YES
- Command line tool = YES [if using standard stuff]
- SQL “dialect” = YES [if using standard stuff]

- Server plugins = NO [actually not tested]
- Replication master-slave
 - MySQL master – MariaDB slave = YES
 - MariaDB master – MySQL slave = YES [if using standard stuff]

- Data directory = YES [if using standard stuff]

Links

- www.mysql.com - MySQL
- www.mariadb.com - MariaDB
- www.askmonty.org – Monty Program
- www.primebase.org – PBXT / Primebase
- www.percona.com – XtraDB / Percona
- www.innodb.com – InnoDB plugin

- www.tinyurl.com/maria2300 - this file

Thanks

Q&A