

Contents

Articles

MariaDB	1
Introduction	2
Storage Engines	4
Features	6

References

Article Sources and Contributors	7
----------------------------------	---

Article Licenses

License	8
---------	---

MariaDB

MariaDB is a fork of MySQL. When MySQL was bought by Oracle, the original author of MySQL decided to start a new branch, which will always be Free Software and open to the contributions of the community.

The purpose of this book is to provide a user-friendly community-contributed manual for MariaDB. Because MariaDB is a fork, this manual won't explain the features of the "parent" software, which are described in the book MySQL.

Unlike the official MySQL documentation, this book is a Free project. It is available under the Creative Commons Attribution/Share-Alike License 3.0 (Unported) ^[1] and the GNU Free Documentation License ^[2].

Table of contents

Introduction

1. What is MariaDB?
2. Differences between MariaDB and MySQL
3. History
4. Resources

Storage Engines

1. Aria
2. PBXT
3. XtraDB
4. FederatedX

Features

1. Pool of threads
2. Table elimination
3. Microseconds granularity
4. Performance

Miscellaneous

Related wikibooks

- Structured Query Language: about general SQL programming, with lots of exercises.
- Converting MySQL to PostgreSQL: also valid for MariaDB
- (IT) MariaDB

References

[1] <http://creativecommons.org/licenses/by-sa/3.0/>

[2] <http://www.gnu.org/copyleft/fdl.html>

Introduction

What is MariaDB?

MariaDB is a fork of MySQL. MySQL is the world's most popular RDBMS; for information about its history, see the official documentation or the MySQL Wikibook.

In 2008 Sun Microsystems bought MySQL. After the acquisition, the development process has changed. The team has started to release new MySQL versions less frequently, so the new code is less tested. There were also less contributions from the community.

In 2009 Monty Widenius, the founder of MySQL, left the company and created a new one, called The Monty Program ^[1]. He started a new fork called MariaDB.

When Oracle announced the acquisition of Sun Microsystems (and thus MySQL software), most of the MySQL developers left Sun to join its forks: MariaDB and Drizzle.

The scopes of MariaDB are:

- import all the new code that will be added to the main MySQL branch, but enhancing it to make it more stable;
- clean the MySQL code;
- add contributions from the community (new plugins, new features);
- develop the Maria storage engine;
- adding new features to the server.

Many of the improvements to the code have been written by third parties. The Monty Programs wants to keep open the development process. Its resources are not competitive with other big companies, but it benefits from the community's work. Most of the improvements are imported from Percona's patches, which are included in OurDelta MySQL builds.

Percona also provides packages for some GNU/Linux distributions.

Differences between MariaDB and MySQL

The patches developed by the MariaDB team are focused on bug fixing and performance. Many improvements have been made for the test suite.

The following Storage Engines are included by default in MariaDB and not in MySQL:

- Maria (also used for internal tables)
- PBXT
- XtraDB
- FederatedX

Others may be added in the future.

Why MariaDB?

- Free as in Freedom - Released with GPL version 2 license (inherited from MySQL)
- Cost - Free!
- Support - MySQL has online tutorials, forums, mailing list (lists.mysql.com), paid support contracts; more MariaDB-specific sites / services are expected to come.
- Open - contributions and suggestions from the community are welcome.
- Speed - One of the fastest databases available; probably faster than MySQL.
- Functionality - support for all the MySQL features, plus others developed by the community.
- Ease of use - the syntax is flexible and managing a database is pretty simple.
- Portability - easily import / export from CSV and XML.
- Scalable - Useful for both small as well as large databases containing billions of records and terabytes of data in hundreds of thousands of tables.
- Plugins - Some Storage Engines and other plugins are pre-installed in MariaDB (and not in MySQL); you can get support for them.

Resources

Docs

Official sites and documentation.

- MariaDB official site ^[1]
- PBXT Docs ^[2]
- XtraDB Docs ^[3]
- OurDelta ^[4]

Blogs

Blogs of some MariaDB developers and blogs about MariaDB.

- Monty Says ^[5]
- Sergey Petrunia ^[6]
- Patrick Galbraith ^[7]

People

User groups and other people.

- MariaDB @ Meetup ^[8]
-

Misc

- GNU GPL v2 ^[9]
- Open Life ^[10]
- Zak Greant's blog ^[11]

References

- [1] <http://www.askmonty.org/>
- [2] <http://www.primebase.org/documentation/>
- [3] <http://www.percona.com/docs/wiki/percona-xtradb:start>
- [4] <http://ourdelta.org/>
- [5] <http://monty-says.blogspot.com/>
- [6] <http://s.petrunia.net/blog/>
- [7] <http://capttofu.livejournal.com/>
- [8] <http://mariadb.meetup.com/it/>
- [9] <http://www.gnu.org/licenses/old-licenses/gpl-2.0.html>
- [10] <http://openlife.cc/blog>
- [11] <http://zak.greant.com/>

Storage Engines

MariaDB includes all the Storage Engines which are included by MySQL 5.1, except for ClusterDB, IBMDB2I and Example.

In addition, it includes some more SEs.

Aria

The Aria storage engine is compiled in by default in MariaDB 5.1 and it is required to be 'in use' when mysqld is started.

Additionally, internal on-disk tables are in the Aria table format instead of the MyISAM table format. This should speed up some GROUP BY and DISTINCT queries because Aria has better caching than MyISAM. The inclusion of Aria is one of the differences between MariaDB 5.1 and MySQL 5.1.

PBXT

The PBXT storage engine is included in the MariaDB source and binaries by default.

PBXT versions in various releases are:

version 1.0.11 in MariaDB 5.1.47 version 1.0.08d in MariaDB 5.1.44b PBXT is a general purpose transactional storage engine. PBXT is fully "ACID" compliant, which means it can be used as an alternative to other MariaDB transactional engines (such as XtraDB or InnoDB).

PBXT features include the following:

- MVCC Support: MVCC stands for Multi-version Concurrency Control. MVCC allows reading the database without locking.
 - Fully ACID compliant: This means that all transactions are: atomic, consistent, isolated and durable.
 - Row-level locking: When updating, PBXT uses row-level locking. Row-level locking is also used during SELECT FOR UPDATE.
 - Fast Rollback and Recovery: PBXT uses a specialized method to identify garbage which makes "undo" unnecessary. This make both rollback of transactions and recovery after restart very fast.
-

- **Deadlock Detection:** PBXT identifies all kinds of deadlocks immediately.
- **Write-once:** PBXT uses a log-based storage which makes it possible to write transactional data directly to the database, without first being written to the transaction log.
- **Referential Integrity:** PBXT supports foreign key definitions, including cascaded updates and deletes.
- **BLOB streaming:** In combination with the BLOB Streaming engine PBXT can stream binary and media directly in and out of the database.
- PBXT will not take any resources (disk space or CPU processing) until you create your first PBXT table.

XtraDB

XtraDB is a fork of InnoDB created and maintained by Percona.

XtraDB is completely compatible with InnoDB. The InnoDB codebase is constantly updated when new InnoDB versions are released. Percona applies to it several patch, which fix known bugs and add new features. This SE is completely compatible with InnoDB.

XtraDB's documentation: <http://www.percona.com/docs/wiki/percona-xtradb:start>

In most builds InnoDB is disabled and XtraDB is instead installed under the name InnoDB. This means that if you specify InnoDB you actually get XtraDB.

FederatedX

FederatedX is based off of the Federated Storage Engine, which is no longer maintained. It is an attempt to move the Federated Storage Engine forward to fix bugs, add new features and develop new concepts that are easier to achieve as a pluggable storage engine.

FederatedX at Launchpad: <https://code.launchpad.net/federatedx>

Features

The features that are discussed in this page have been developed for MariaDB and probably are not present in MySQL.

Microseconds precision

SHOW PROCESSLIST, INFORMATION_SCHEMA.PROCESSLIST and the Slow Log now have a microsecond precision.

A microsecond is a unit of time equal to one millionth (10^{-6}) of a second.

SHOW PROCESSLIST

In MariaDB, an extra column `TIME_MS` has been added to the INFORMATION_SCHEMA.PROCESSLIST table, as well as to the output of SHOW FULL PROCESSLIST. This column shows the same information as the column `TIME`, but in units of milliseconds with microsecond precision (the unit and precision of the `TIME` column is one second).

The value displayed in the `TIME` and `TIME_MS` columns is the period of time that the given thread has been in its current state. Thus it can be used to check for example how long a thread has been executing the current query, or for how long it has been idle.

In MariaDB the `TIME` column and the `TIME_MS` column are not affected by any setting of `@TIMESTAMP`. This means that it can be reliably used also for threads that change `@TIMESTAMP` (such as the replication SQL thread).

The `TIME` column of SHOW FULL PROCESSLIST and INFORMATION_SCHEMA.PROCESSLIST can not be used to determine if a slave is lagging behind. For this, use instead the `Seconds_Behind_Master` column in the output of SHOW SLAVE STATUS.

The addition of the `TIME_MS` column is based on the `microsec_process` patch, developed by Percona.

Slow Log

The Slow Log includes details of execution plan and microsecond-precision resolution.

This addition is based on the `microslow` patch from Percona.

Performance

MariaDB's code has several optimizations which have been made after the fork has been launched.

Some speed enhancements are:

- CHECKSUM TABLE is faster;
- performance improvements for common cases of character set conversion;
- test suite speedups.

The list is not complete.

Article Sources and Contributors

MariaDB *Source:* <http://en.wikibooks.org/w/index.php?oldid=2178421> *Contributors:* Adrignola, F razzoli, Sante Caserio, 2 anonymous edits

Introduction *Source:* <http://en.wikibooks.org/w/index.php?oldid=2604086> *Contributors:* BestBackwards, QuiteUnusual, Sante Caserio, 1 anonymous edits

Storage Engines *Source:* <http://en.wikibooks.org/w/index.php?oldid=2521183> *Contributors:* F razzoli, QuiteUnusual, Roeme, Sante Caserio, 8 anonymous edits

Features *Source:* <http://en.wikibooks.org/w/index.php?oldid=2083432> *Contributors:* QuiteUnusual, Sante Caserio, 4 anonymous edits

License

Creative Commons Attribution-Share Alike 3.0
[//creativecommons.org/licenses/by-sa/3.0/](https://creativecommons.org/licenses/by-sa/3.0/)
