



# MariaDB: Community Driven SQL Server

Kristian Nielsen  
MariaDB developer  
Monty Program AB



Background

# Why



MySQL is a high-profile Free Software package and very important to the Free Software world in general.

Ensure the continued maintenance of the MySQL codebase (under whatever name) at a high level of quality.

Ensure the continued availability of the LAMP stack.



**Monty Program**



# Why



## Lack of confidence in MySQL/Sun/Oracle:

- ♦ MySQL management toying with cripplewaring MySQL.
- ♦ Sun takeover causing core developers to leave the project.
- ♦ Poor community involvement.
- ♦ Uncertain about Oracles intentions and motives.

"We need a backup plan B"



**Monty Program**



# What



- ♦ MariaDB - a branch of the MySQL codebase.
  - ♦ GPL v2, but "MySQL" is trademarked.
  - ♦ 100% compatible with MySQL.
  - ♦ MariaDB 5.1.39 contains all of MySQL 5.1.39.
  - ♦ MariaDB is a viable alternative for everyone currently using MySQL.
- 
- ♦ Merge all MySQL code.
  - ♦ Merge and maintain community contributions.
  - ♦ Bug fixing.
  - ♦ Own development.



**Monty Program**



# Who



## Monty Program AB

Founded by Michael "Monty" Widenius.

~15 people in many countries, including a number of old-time core MySQL engineers.

**MariaDB development:** NRE/bugfixing, partnerships, ODBA

Partners:

- ◆ OurDelta (Open Query): packaging
- ◆ Percona: support



Monty Program



# Where



- ♦ apt and yum repositories: <http://ourdelta.org/>
- ♦ Debian 4/5; Ubuntu 8.04/8.10/9.04/9.10; Centos 5; tarball binaries. amd64 and i386.
- ♦ Source tarballs.
- ♦ Launchpad bzd trees:  
<https://code.launchpad.net/maria>
  
- ♦ `sudo apt-get install mariadb-server`
- ♦ Upgrade existing MySQL installation to MariaDB.



**Monty Program**



# Community



## Open development:

- ♦ All trees on Launchpad.
- ♦ Open Buildbot testing of all pushes.
- ♦ Open mailing list  
`maria-developers@lists.launchpad.net`
- ♦ Commit rights to external community members (lp:~maria-captains)
- ♦ Day-to-day work on public IRC channel #maria on FreeNode
- ♦ Less strict copyright assignment for contributions (but still require MCA or BSD for core server modifications).



**Monty Program**





# Monty Program – Our partners



redhat.

READY  
BUSINESS PARTNER

Business Partner



IBM

Business  
Partner



PartnerNet.  
TECHNOLOGY  
PARTNER



Percona

JASPER SOFT



Monty Program



# The Open Database Alliance



- ◆ Database agnostic
- ◆ Vendor neutral
- ◆ Provides a single point of contact for those that need support, services, NRE, etc for open databases
- ◆ Democratic policies reflect the open design of the consortium. Reflected in the ODBA's referral incentive program.
- ◆ More flexible solutions than a single large company could provide




**Monty Program**





MariaDB: Community Driven SQL Server

# State of MariaDB 5.1




- ♦ Recommended beta currently 5.1.39
- ♦ MariaDB 5.1 is kept up to date with MySQL 5.1 (pulls are done monthly or when there is a new MySQL build)
- ♦ Binaries with up to date versions of Maria, PBXT and XtraDB
- ♦ Expect final release late 2009 or early 2010



**Monty Program**



# New features in MariaDB 5.1



## XtraDB


- ◆ Percona-enhanced InnoDB plugin
- ◆ Performance
- ◆ Scalability
- ◆ Maintenance
- ◆ DBA/monitoring
- ◆ Backup
- ◆ <http://www.percona.com/docs/wiki/percona-xtradb:start>



**Monty Program**



# New features in MariaDB 5.1



## PBXT

- ◆ New storage engine for MySQL
- ◆ Log-based storage
- ◆ Data written only once
- ◆ Transactional, MVCC, good benchmarks
- ◆ Good example of the power of pluggable storage engines
- ◆ <https://launchpad.net/pbxt>



**Monty Program**



# New features in MariaDB 5.1



## Maria 1.5


- ◆ New storage engine for MySQL
- ◆ MyISAM replacement
- ◆ Logging/journaling (recover from crash)
- ◆ Roadmap: transactional, MVCC, Data Warehouse



**Monty Program**



# New features in MariaDB 5.1



## Table elimination

```
CREATE TABLE B (id int primary key);
```

```
Select  A.colA from  tableA A  
left outer join tableB B  
on B.id = A.id;
```

In this case we can remove table B and the join from the query.



Monty Program





# New features in MariaDB 5.1



## Pool of threads

- ◆ Allows hundreds of thousands of connections to MariaDB with a few worker threads doing the task.
- ◆ Backport from 6.0 with new option `-extra-port=#`
  - ◆ This allows an admin to login and resolve lock issues with threads in the pool and also allows different schedulers for different connections



Monty Program



# New features in MariaDB 5.1



## Virtual columns

```
CREATE TABLE t1 (  
  <name> <type> [GENERATED ALWAYS] AS  
<expression> [MATERIALIZED|VIRTUAL]  
)
```


Patch originally by Andrey Zhakov



Monty Program



# New features in MariaDB 5.1



## Percona patches


- ◆ `microslow_innodb.patch`
- ◆ `microsec_process.patch`
- ◆ Microsecond resolution in slow log and `SHOW FULL PROCESSLIST`
- ◆ Improved DBA tools for understanding DB server behaviour



Monty Program



# New features in MariaDB 5.1



## Speed improvements

- ♦ Index merge optimizations (faster queries when using multiple possible index)
- ♦ Faster complex queries (that uses disk based temporary tables)
- ♦ Speed improvements; sql\_bench benchmarks shows 10 % speedup for some common queries that returns lots of simple character data



**Monty Program**



# New features in MariaDB 5.1



## Code cleanups

- ◆ We have fixed a lot of compiler warnings and errors found by valgrind
- ◆ Most posted patches on [internal@lists.mysql.com](mailto:internal@lists.mysql.com)
- ◆ Detector if mutex are wrongly used
- ◆ Misc. bug fixes



**Monty Program**



# New features to be added to MariaDB 5.1



## Need to be added before 'first release'

- ♦ The new Federated storage engine.
- ♦ Proven and stable community patches we feel offer important performance or feature enhancements. Percona and Ourdelta patches exemplify this work.
- ♦ Memory tables with efficient VARCHAR/BLOB support



Monty Program



# Features planned for MariaDB 5.4



## Features to be nominated during all-company on-site planning meeting in August

- ◆ Stable features from MySQL 5.4
- ◆ Group commit for Maria storage engine (done)
- ◆ Extra arguments to storage engines for columns, index and tables in CREATE TABLE
- ◆ Google replication patches
- ◆ Facebook patches
- ◆ Stable patches from Ourdelta
- ◆ Recursive queries
- ◆ Optimizer extensions
- ◆ Dynamic columns
- ◆ Phone home (code shared with Drizzle)



**Monty Program**





# Buildbot Testing



# Buildbot testing



We use Buildbot to test every code change on different platforms and test configurations.

- ♦ <http://askmonty.org/buildbot/>
- ♦ Detect new problems as early as possible
- ♦ Facilitate developers integrating big changes in a smooth fashion.
- ♦ Contributions are required to pass in Buildbot in a branch before going into trunk.
- ♦ Buildbot is not the simplest to set up, but very strong in providing needed information to debug and fix problems.



**Monty Program**



# Buildbot Cross Reference Search



[http://askmonty.org/buildbot/reports/cross\\_reference](http://askmonty.org/buildbot/reports/cross_reference)

Very useful for

- Tests that fail only on specific platforms
- Tests that fail only with specific compile options
- Tests that fail only occasionally (races)

Example:

- `test=rpl_timezone`
- `output=session.time_zone`



Monty Program



# Testing MariaDB packages with KVM



- ♦ Build and test packages automatically on lots of platforms (currently 13)
- ♦ Use virtual machines with no "pollution" of extra packages
- ♦ Test installing on a machine with no build tools
- ♦ Test upgrading from MySQL (ToDo: upgrading from earlier MariaDB).
- ♦ Fully automatic; runs after every change made to MariaDB source.



**Monty Program**



- Based on QEMU userspace, but with kernel module for hardware acceleration (requires CPU support).
- Very easily available (in mainline kernel, `sudo apt-get install kvm`).
- Good support for command-line use and scripting (`man qemu`)



# Using KVM



- ◆ Each build steps boots up a VM in KVM, uses ssh to run commands inside it, then shuts it down.
- ◆ Simple, low hardware needs, robust.

```
kvm -hda vm-tmp-2226.qcow2 -redir 'tcp:2226::22' -cpu qemu32,-nx -nographic -net user
```

- ◆ Clone-and-discard approach. No pollution from one build to the next

```
qemu-img create -b /kvm/vms/vm-debian5-i386-install.qcow2 -f qcow2 vm-tmp-2226.qcow2
```



**Monty Program**



# Maria/MariaDB - resources



- ◆ Up to date Information about MariaDB/Maria at <http://askmonty.org>
- ◆ MariaDB source code at LaunchPad
- ◆ Maria roadmap, algorithms and internals are described in detail at <http://forge.mysql.com/worklog/task.php?id=3871>
- ◆ End user documentation about Maria can be found at <http://dev.mysql.com/doc/refman/5.1/en/storage-engines.html>
- ◆ Monty's blog at: <http://monty-says.blogspot.com/> will have all important updates to the MariaDB/Maria project.
- ◆ You can follow and participate in all MariaDB development discussions by joining us on Launchpad and #maria on freenode.



**Monty Program**





QUESTIONS?