



Meet MariaDB 10.1

The community release

Sergei Golubchik &
Michael “Monty” Widenius
MariaDB

Santa Clara, CA. April 2015



PERCONA
LIVE

What is MariaDB and where to get it

What is MariaDB

- A free fork of MySQL with extra features
 - features, requested by the users
- Backward compatible
 - file formats, replication, configuration files
- aim at
- Community
 - 50% of the community
 - no hidden

Sergei Golubchik, MariaDB, Percona

Where to get it

- In distributions
 - openSUSE, Gentoo, Mageia, Slackware, ALTLinux
 - FreeBSD
 - MacOS X with MacPorts or Homebrew
- From MariaDB.org:
 - sources, binaries in **.tar.gz** (Linux)
 - Windows MSI installer
 - MariaDB **apt** and **yum** repositories



Sergei Golubchik, MariaDB — Now Features.

What is MariaDB

- A free fork of MySQL with extra features
 - features, requested by the users
- Backward compatible

Where to get it

- Distributions:
 - RedHat Enterprise Linux, Fedora, Mandriva, openSUSE, Gentoo, Slackware, Arch, ALTLinux, TurboLinux, Chakra Project
 - FreeBSD, OpenBSD
 - Mac OS X with MacPorts or Homebrew
- From MariaDB.org
 - sources, binaries in **.tar.gz** or **.zip** (Windows)
 - Windows MSI installer
 - MariaDB **apt** and **yum** repositories

Sergei Golubchik, MariaDB, Percona Live Santa Clara, April 2014



What is MariaDB

- A **branch** of MySQL with extra features
 - features, requested by our users
- Backward compatible
 - file formats, replication, configuration files

Where to get it

- Distributions:
 - RedHat Enterprise Linux, Fedora, Mageia, openSUSE, Gentoo, Slackware, Arch, ALTLinux, TurboLinux, Chakra Project, Kdu, ...
 - FreeBSD, OpenBSD
 - Mac OS X with MacPorts or Homebrew
- From MariaDB.org
 - sources, binaries in **.tar.gz** or **.zip** (Windows)
 - Windows MSI installer
 - MariaDB **apt** and **yum** repositories



ADC 2012

4



10.1 Overview

Optimizer

Performance & Concurrency

XtraDB & InnoDB

Administration & Maintenance

Security

Replication

Compiling & Packaging

Plugins & Plugin API

GIS



The community release

**Thank you,
MariaDB community!**




Optimizer enhancement

- ANALYZE <statement>
- ORDER BY ... LIMIT optimizer
- UNION ALL optimization
- EXPLAIN FORMAT=JSON
- engine condition pushdown always on
- affected rows in the slow log

Performance and Scalability

- Removing locks — using lock-free data structures
 - in MDL (meta-data locking)
 - in TDC (table definition cache)
 - in XA (XID hash)
- Numerous scalability improvements
 - up 80% speedup in sysbench on Power8
- No .frm files for temporary tables

sysbench OLTP R0 on Power8





Replication

- **Galera!**
- **Optimistic parallel replication!**
- RBR triggers
- Enhanced semisync replication
- Dump thread enhancements
- domain_id based replication filters



- Better OpenGIS standard compliance

- INFORMATION_SCHEMA tables

GEOMETRY_COLUMN SPATIAL_REF_SYS

- REF_SYSTEM_ID per GEOMETRY column
 - Missing OGC spatial functions

ST_Boundary ST_ConvexHull ST_IsRing ST_PointOnSurface ST_Relate

Miscellaneous

- SET STATEMENT
- Compound statements
- CREATE OR REPLACE, CREATE IF NOT EXISTS, DROP IF EXISTS
- INFORMATION_SCHEMA.SYSTEM_VARIABLES
- Sequence engine enabled by default
- Microseconds in GET_LOCK()
- Assisted discovery in OQGraph
- default_tmp_storage_engine

Admin/Maintainability

- Default role
- Statement timeouts
- PERFORMANCE_SCHEMA without .frm files
- enabling slow log per session
- ALGORITHM column in INFORMATION_SCHEMA.VIEWS
- better help for SET/ENUM command-line options
- --getopt-prefix-matching



Security

- **Encryption!**
- Scrubbing
- Password validation
- Security-enhanced binaries



Compilation and Packaging

- Link with dynamic system libjemalloc.so
- GUI-friendly cmake options to enable/disable plugins
- Special compilation options for security-enhanced binaries
- Support for pkg-config, a mariadb.pc file added



XtraDB/InnoDB

- Encryption
 - tablespaces and logs
 - background key rotation
 - per-table or global
 - key identifiers specified per table
- Scrubbing
 - tablespaces and logs

- FusionIO related
 - CREATE TABLE ... ATOMIC_WRITES=...
 - CREATE TABLE ... PAGE_COMPRESSION=...
 - zlib, lz4, lzo, lzma, snappy
 - TRIM
 - Multi-threaded flush
 - support for 4K sector size



XtraDB/InnoDB

- I_S.INNODB_SEMAPHORE_WAITS, I_S.INNODB_MUTEXES
- Support for up to 64K pages
- Defragmentation
- Monitor progress of online ALTER TABLE
- Enforce the PRIMARY KEY
- Prefix index queries optimization
- Lazy flushing
- ...



Plugin API

- Two new plugin types
 - Password verification
 - `simple_password_check`
 - `cracklib_password_check`
 - Encryption
 - `file_key_management`
 - `example_key_management`
 - `password_key_management`



Services for Plugins

- Five new services
 - md5
 - sha1
 - encryption and key management
 - thd_specifics
 - wsrep (write set replication)



Contributors

- **Google** — encryption, scrubbing, enhanced semisync, dump thread enhancements, thd_specifics plugin service
- **Eperi GmbH** — encryption
- **Kakao Inc** — defragmentation, online alter progress monitoring
- **Codership** — Galera (wsrep patches)
- **Antony Curtis** — compound statements
- **Sriram Patil** — CREATE OR REPLACE/IF NOT EXISTS
- **Daniel Black** — new status variables for replication



In cooperation with

- **IBM** — Power8 performance and scalability optimizations
- **FusionIO** — atomic writes, page compression, trim, multi-threaded flush for XtraDB/InnoDB



Also used code snippets by

- **Facebook** — defragmentation, prefix index queries optimization, lazy flushing, buffer pool list scan optimization, configurable long semaphore wait timeout
- **Percona** — SET STATEMENT, enforce_storage_engine
- **Oracle** — UNION ALL optimization, default_tmp_storage_engine



You can participate too!

- 10.2 planning — what do you want to see in 10.2?
- Contribute code — <https://github.com/mariadb/server>
- Write KnowledgeBase articles — <https://mariadb.com/kb/>
- Report bugs — <https://mariadb.org/jira>
- Help others on the IRC — #maria on Freenode
- And on mailing lists — maria-discuss@lists.launchpad.net
- Enable Feedback plugin — enable-feedback in my.cnf



Questions?



Thank you!

