Meet MariaDB 10.1

Colin Charles, Team MariaDB, MariaDB Corporation
colin@mariadb.com / byte@bytebot.net
http://bytebot.net/blog/ | @bytebot on Twitter
FOSSASIA 2016, Singapore
19 March 2016

whoami

- Work on MariaDB at MariaDB Corporation (~~SkySQL Ab~~)
  - Merged with Monty Program Ab, makers of MariaDB
- Formerly MySQL AB (exit: Sun Microsystems)
- Past lives include Fedora Project (FESCO), OpenOffice.org
- MySQL Community Contributor of the Year Award winner 2014

Global Top 20 Sites

1. Google
2. Facebook
3. YouTube
4. Baidu
5. Yahoo!
6. Amazon
7. Wikipedia
8. QQ
9. Google.co.in
10. Twitter
11. live.com
12. Taobao
13. Msn.com
14. yahoo.co.jp
15. Sina
16. Linkedin.com
17. google.co.jp
18. Weibo
19. Bing.com
20. yandex.ru

What is MariaDB Server?

- GPLv2 branch of MySQL with a lot of added **community** development
- Feature enhanced
- Application compatible & feature complete with MySQL
- A “drop-in replacement” (upgrade standpoint)

MariaDB thru the years

- MariaDB 5.1
- MariaDB 5.2
- MariaDB 5.3
- MariaDB 5.5
- MariaDB 10.0
- MariaDB 10.1
- MariaDB Galera Cluster 5.5
- MariaDB Galera Cluster 10.0
- C Connector
- Java Connector
- ODBC Connector

Why MariaDB?

- MySQL has a single owner; MariaDB has the MariaDB Foundation (not just Corporation)
- MySQL ecosystem development is at its most vibrant now than it has ever been
- Community can get features inside a shipping server with ease
- Storage engine vendors get shipping & wide distribution including testing
- Roadmaps are public on Jira

MariaDB, the ecosystem

- Besides the Server, we do develop LGPL Connectors
- Focus on making “enterprise” features opensource
  - threadpool is a great example of working in the open
- Open content (& friendly licensed) Knowledge Base

MariaDB MaxScale

- Level 7 proxy router, that understands the MySQL protocol, with a pluggable architecture
- Possibilities are endless - use it for logging, writing to other databases (besides MySQL), preventing SQL injections via regex filtering, route via hints, query rewriting, have a binlog relay, etc.
- Load balance your Galera clusters

Where is MariaDB found?

- http://mariadb.org/
- Your Linux/BSD **distribution** - it is the default in RHEL 7, SUSE Enterprise 12, openSUSE, CentOS, Fedora, OpenBSD, etc.
  - and a choice in Debian & Ubuntu
- Pivotal.io CloudFoundry, RackSpace Cloud, Azure, etc.

MariaDB on AWS RDS!

What's New from Amazon Web Services

TA)

d monitor
s using

OCT 7 | AMAZON RDS FOR MARIADB

Introducing highly available, scalable and secure MariaDB deployment on AWS

OCT 7 | DATABASE MIGRATION SERVICE

Minimal downtime database migration to AWS using AWS Database Migration Service

The “community release”

- 30 Jun 2014 - MariaDB 10.1.0
- 17 Oct 2014 - MariaDB 10.1.1
- 7 Dec 2014 - MariaDB 10.1.2
- 2 March 2015 - MariaDB 10.1.3
- 13 April 2015 - MariaDB 10.1.4
- 4 June 2015 - MariaDB 10.1.5
- 27 July 2015 - MariaDB 10.1.6
- 9 September 2015 - MariaDB 10.1.7 RC
- 17 October 2015 - MariaDB 10.1.8 GA
- 23 Nov 2015, 24 Dec 2015, 29 Jan 2016, 25 Feb 2016

Google Summer of Code

- SQL Roles
- Kerberos authentication
- PCRE regular expressions
- InnoDB memcached interface
- InnoDB Redis interface
- Improvements in replication auto-discovery
- (per-query variable settings)

Themes

- Security
- High Availability
- High Performance
- Operational Ease
- Better for developers and DBAs

Security

Encryption

- Encryption: tablespace and table level encryption with support for rolling keys using the AES algorithm
  - table encryption — `PAGE_ENCRYPTION=1`
  - tablespace encryption — encrypts everything including log files
- New `file_key_management_filename`, `file_key_management_filekey`, `file_key_management_encryption_algorithm`
- Well documented — https://mariadb.com/kb/en/mariadb/data-at-rest-encryption/

Encryption II

- The key file contains encryption keys identifiers (32-bit numbers) and hex-encoded encryption keys (128-256 bit keys), separated by a semicolon.
- don’t forget to create keys!
  - eg. `openssl enc -aes-256-cbc -md sha1 -k secret -in keys.txt -out keys.enc`

Encryption III

```
CREATE TABLE customer (
   customer_id bigint not null primary key,
   customer_name varchar(80),
   customer_creditcard varchar(20))
ENGINE=InnoDB

page_encryption=1

page_encryption_key=1;
```

Encryption IV

- Tablespace encryption (Google)
  - again, you need to pick an encryption algorithm
  - specify what to encrypt: `innodb-encrypt-tables, aria, aria-encrypt-tables, encrypt-tmp-disk-tables, innodb-encrypt-log`
- don’t forget key rotation:
  - `innodb-encryption-threads=4`
  - `innodb-encryption-rotate-key-age=1800`

Encryption V

- `/etc/my.cnf.d/enable_encryption.preset`
- Consider using Eperi Gateway for Databases
- MariaDB Enterprise will have a plugin for Amazon Key Management Server (KMS)
- mysqlbinlog has no way to read (i.e. decrypt) an encrypted binlog
- This does not work with MariaDB Galera Cluster yet (gcache is not encrypted yet), and also xtrabackup needs additional work (i.e. if you encrypt the redo log)

Password validation

- simple_password_check password validation plugin
  - can enforce a minimum password length and guarantee that a password contains at least a specified number of uppercase and lowercase letters, digits, and punctuation characters.
- cracklib_password_check password validation plugin
  - Allows passwords that are strong enough to pass CrackLib test. This is the same test that pam_cracklib.so does

SQL Error Logging Plugin

- Log errors sent to clients in a log file that can be analysed later. Log file can be rotated (recommended)
- a MYSQL_AUDIT_PLUGIN

```
install plugin SQL_ERROR_LOG soname
'sql_errlog.so';
```

Audit Plugin

- Log server activity - who connects to the server, what queries run, what tables touched - rotating log file or syslogd
- MariaDB has extended the audit API, so user filtering is possible
- a MYSQL_AUDIT_PLUGIN

```
INSTALL PLUGIN server_audit SONAME
‘server_audit.so’;
```

Authentication Plugins

- Shipped a PAM authentication plugin for the longest time
- Now you also have a Kerberos/GSSAPI based authentication plugin
  - works with Microsoft Active Directory too!

High Availability

MariaDB 10 replication: provisioning a new slave

```
SET GLOBAL GTID_SLAVE_POS =
BINLOG_GTID_POS("masterbin.00045",
600);

CHANGE MASTER TO
master_host="192.168.2.4",
master_use_gtid=slave_pos;

START SLAVE;
```

Replication: turning on GTID for slaves

```
STOP SLAVE;

CHANGE MASTER TO
master_use_gtid=current_pos;

START SLAVE;
```

Replication: change masters

```
STOP SLAVE;

CHANGE MASTER TO
master_host="10.2.3.5";

START SLAVE;
```

Multi-source replication

- Work from Taobao
- Many users partition data across many masters... now you can replicate many masters to a single slave
- Great for analytical queries, complete backups, etc.
- All master/slave commands take a connection name now (like CHANGE MASTER “connection_name”, SHOW SLAVE “connection_name” STATUS, etc.)

Galera Cluster integrated

- Full integration of Galera Cluster into MariaDB 10.1 — not a separate download
    - no lost transactions, optimisations for WAN replication, non-blocking DDL, no limits on transaction size
- `Server version: 10.1.3-MariaDB-wsrep MariaDB Server, wsrep_25.10.r4144`
- Granular monitoring in INFORMATION_SCHEMA — WSREP_MEMBERSHIP, WSREP_STATUS

Optimistic parallel replication

- Before, transactions committed in parallel on the master could be run in parallel
- Now, more than one transaction will be considered to be run in parallel giving another performance boost in master-to-slave replication
- Need a 10.1 master to work
- https://mariadb.atlassian.net/browse/MDEV-6676

Replication: START TRANSACTION WITH CONSISTENT SNAPSHOT

- Works with the binlog, possible to obtain the binlog position corresponding to a transactional snapshot of the database without blocking any other queries.
    - by-product of group commit in the binlog to view commit ordering
- Used by the command `mysqldump--single-transaction --master-data` to do a fully non-blocking backup which can be used to provision a new slave
- Works consistently between transactions involving more than one storage engine
- https://kb.askmonty.org/en/enhancements-for-start-transaction-with-consistent/

More in replication

- Enhanced semi-sync replication (like in 5.7 - thanks FB/Google)
- `domain_id` based replication filters
- Slaves can execute triggers now (in RBR)
- Dump thread enhancements (remove binlog lock LOCK_log) from 5.7 included (Google)

High Performance

An opensource threadpool

- Modified from 5.1 (libevent based), great for CPU bound loads and short running queries
- No minimization of concurrent transactions with dynamic pool size
- `thread_handling=pool-of-threads`
- https://mariadb.com/kb/en/mariadb/thread-pool-in-mariadb/
- now you can also have a priority mode for tickets

Threadpool

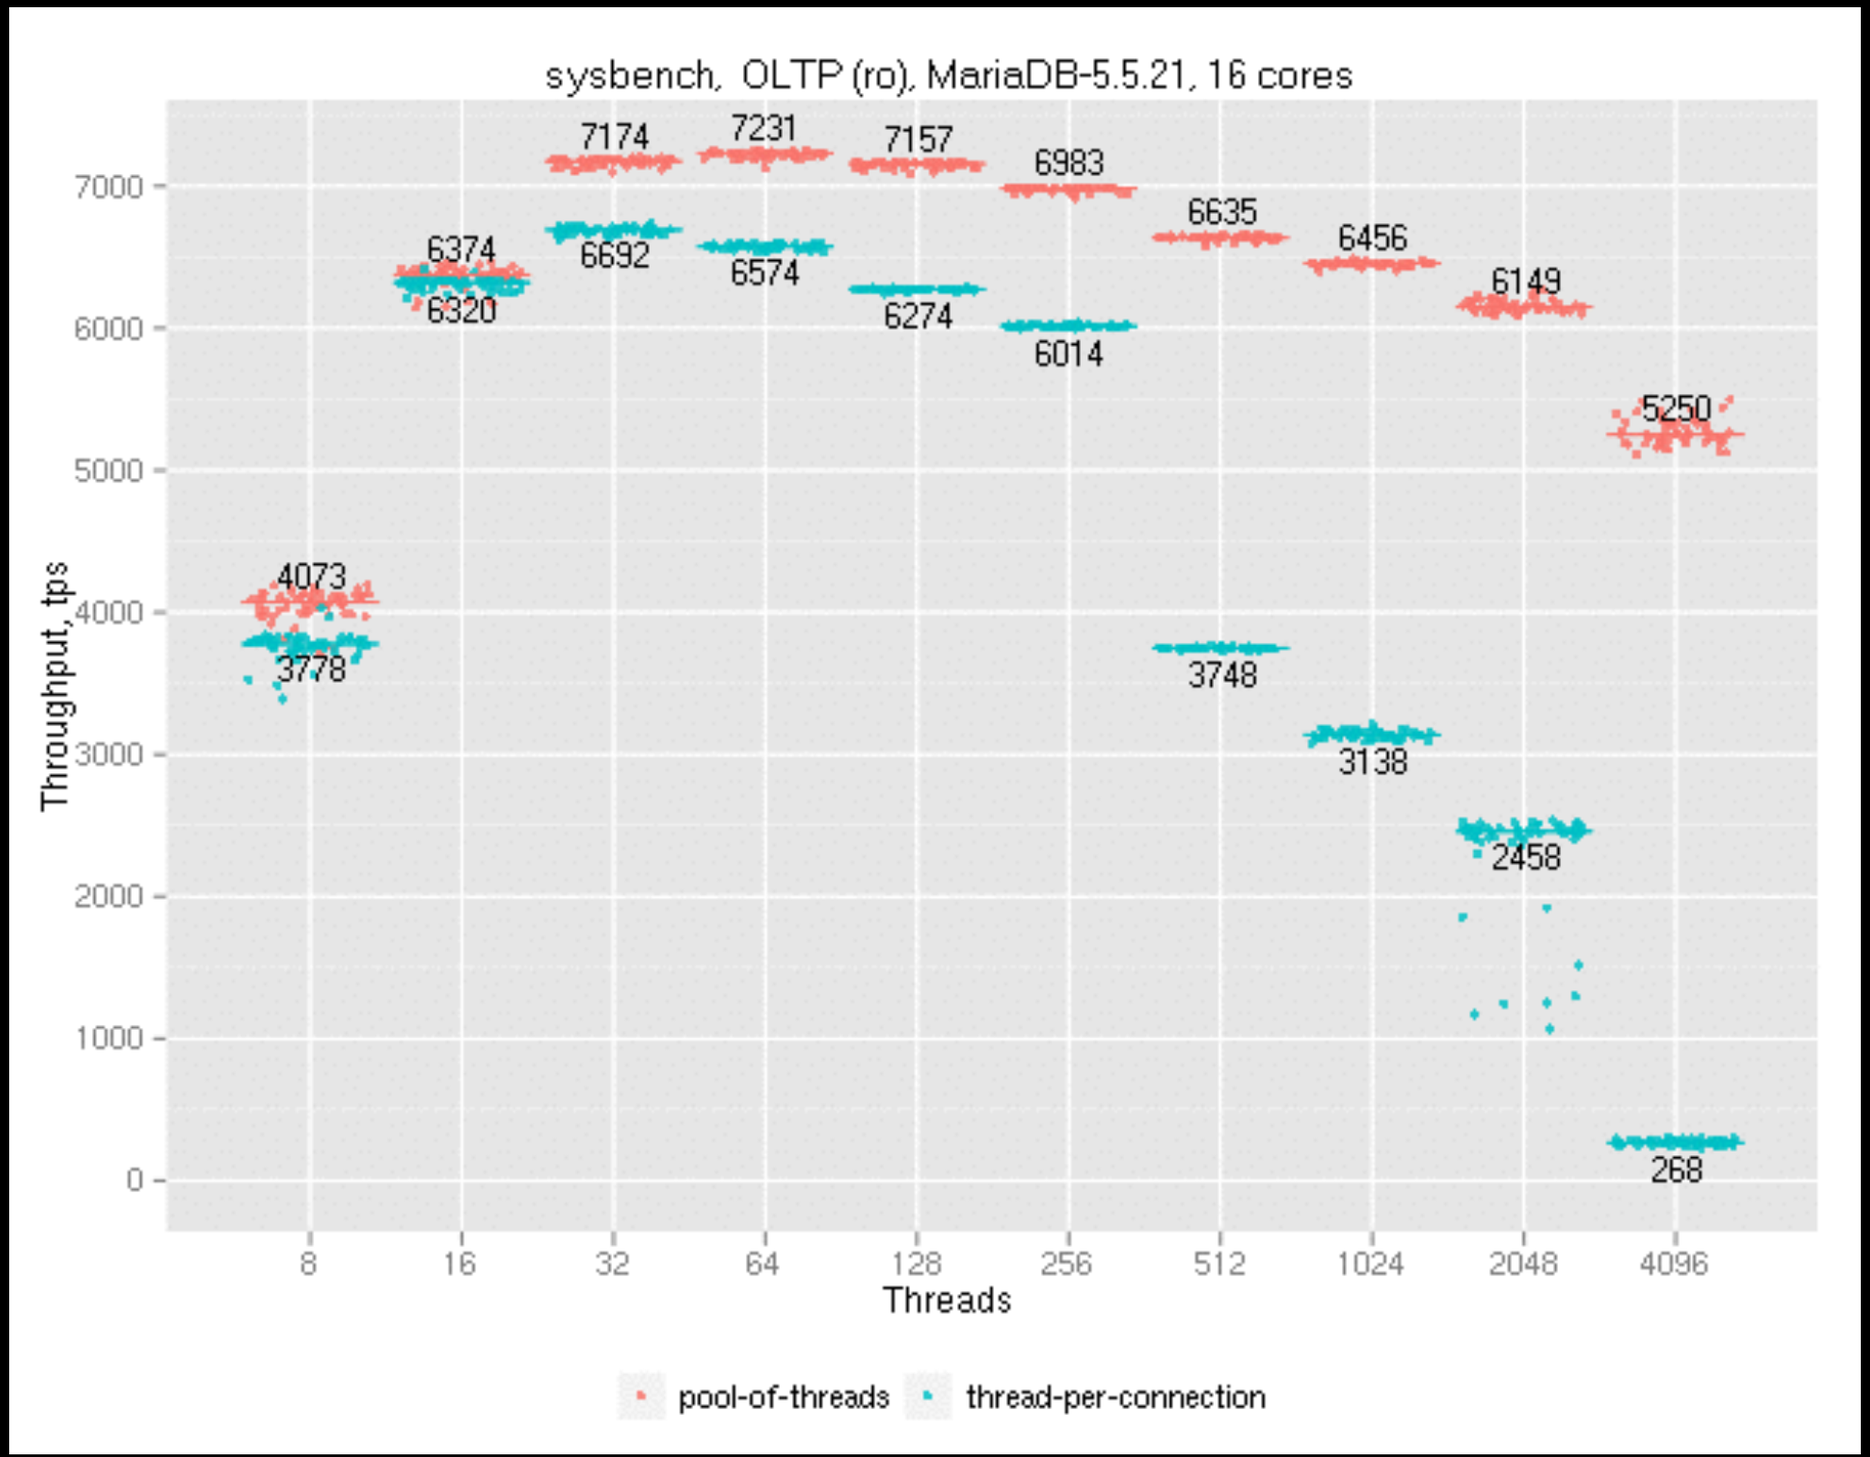
sysbench, OLTP (ro), MariaDB-5.5.21, 16 cores
Throughput, tps
7000
6000
5000
4000
3000
2000
1000
0
7174
7231
7157
6983
6635
6456
6149
5250
6374
6320
6692
6574
6274
6014
4073
3778
3748
3138
2458
268
8
16
32
64
128
256
512
1024
2048
4096
Threads
pool-of-threads
thread-per-connection

InnoDB improvements

- Multi-threaded flush (also in 5.7, different implementation + we’re first)
- 64KB pages in InnoDB (old limit = 16KB).
- Defragementation (FB, backported by DaumKakao)
- `I_S.INNODB_SEMAPHORE_WAITS, I_S.INNODB_MUTEXES`
- Forced primary key
  - If option is true, create table without primary key or unique key where all keyparts are NOT NULL is not accepted. Instead an error message is printed.

InnoDB WebScaleSQL

- MDEV-6936: Buffer pool list scan optimisation
- MDEV-6929: Port Facebook Prefix Index Queries Optimization
- MDEV-6932: Enable Lazy Flushing
- MDEV-6931: Page cleaner should do LRU flushing regardless of server activity
- fixes mysql#71988, mysql#70500
- DB-746 merge clustering key is covering key for mariadb 10 (TokuDB)
- MDEV-6933: Spurious lock_wait_timeout_thread wakeup in lock_wait_suspend_thread()
  - fixes mysql#72123

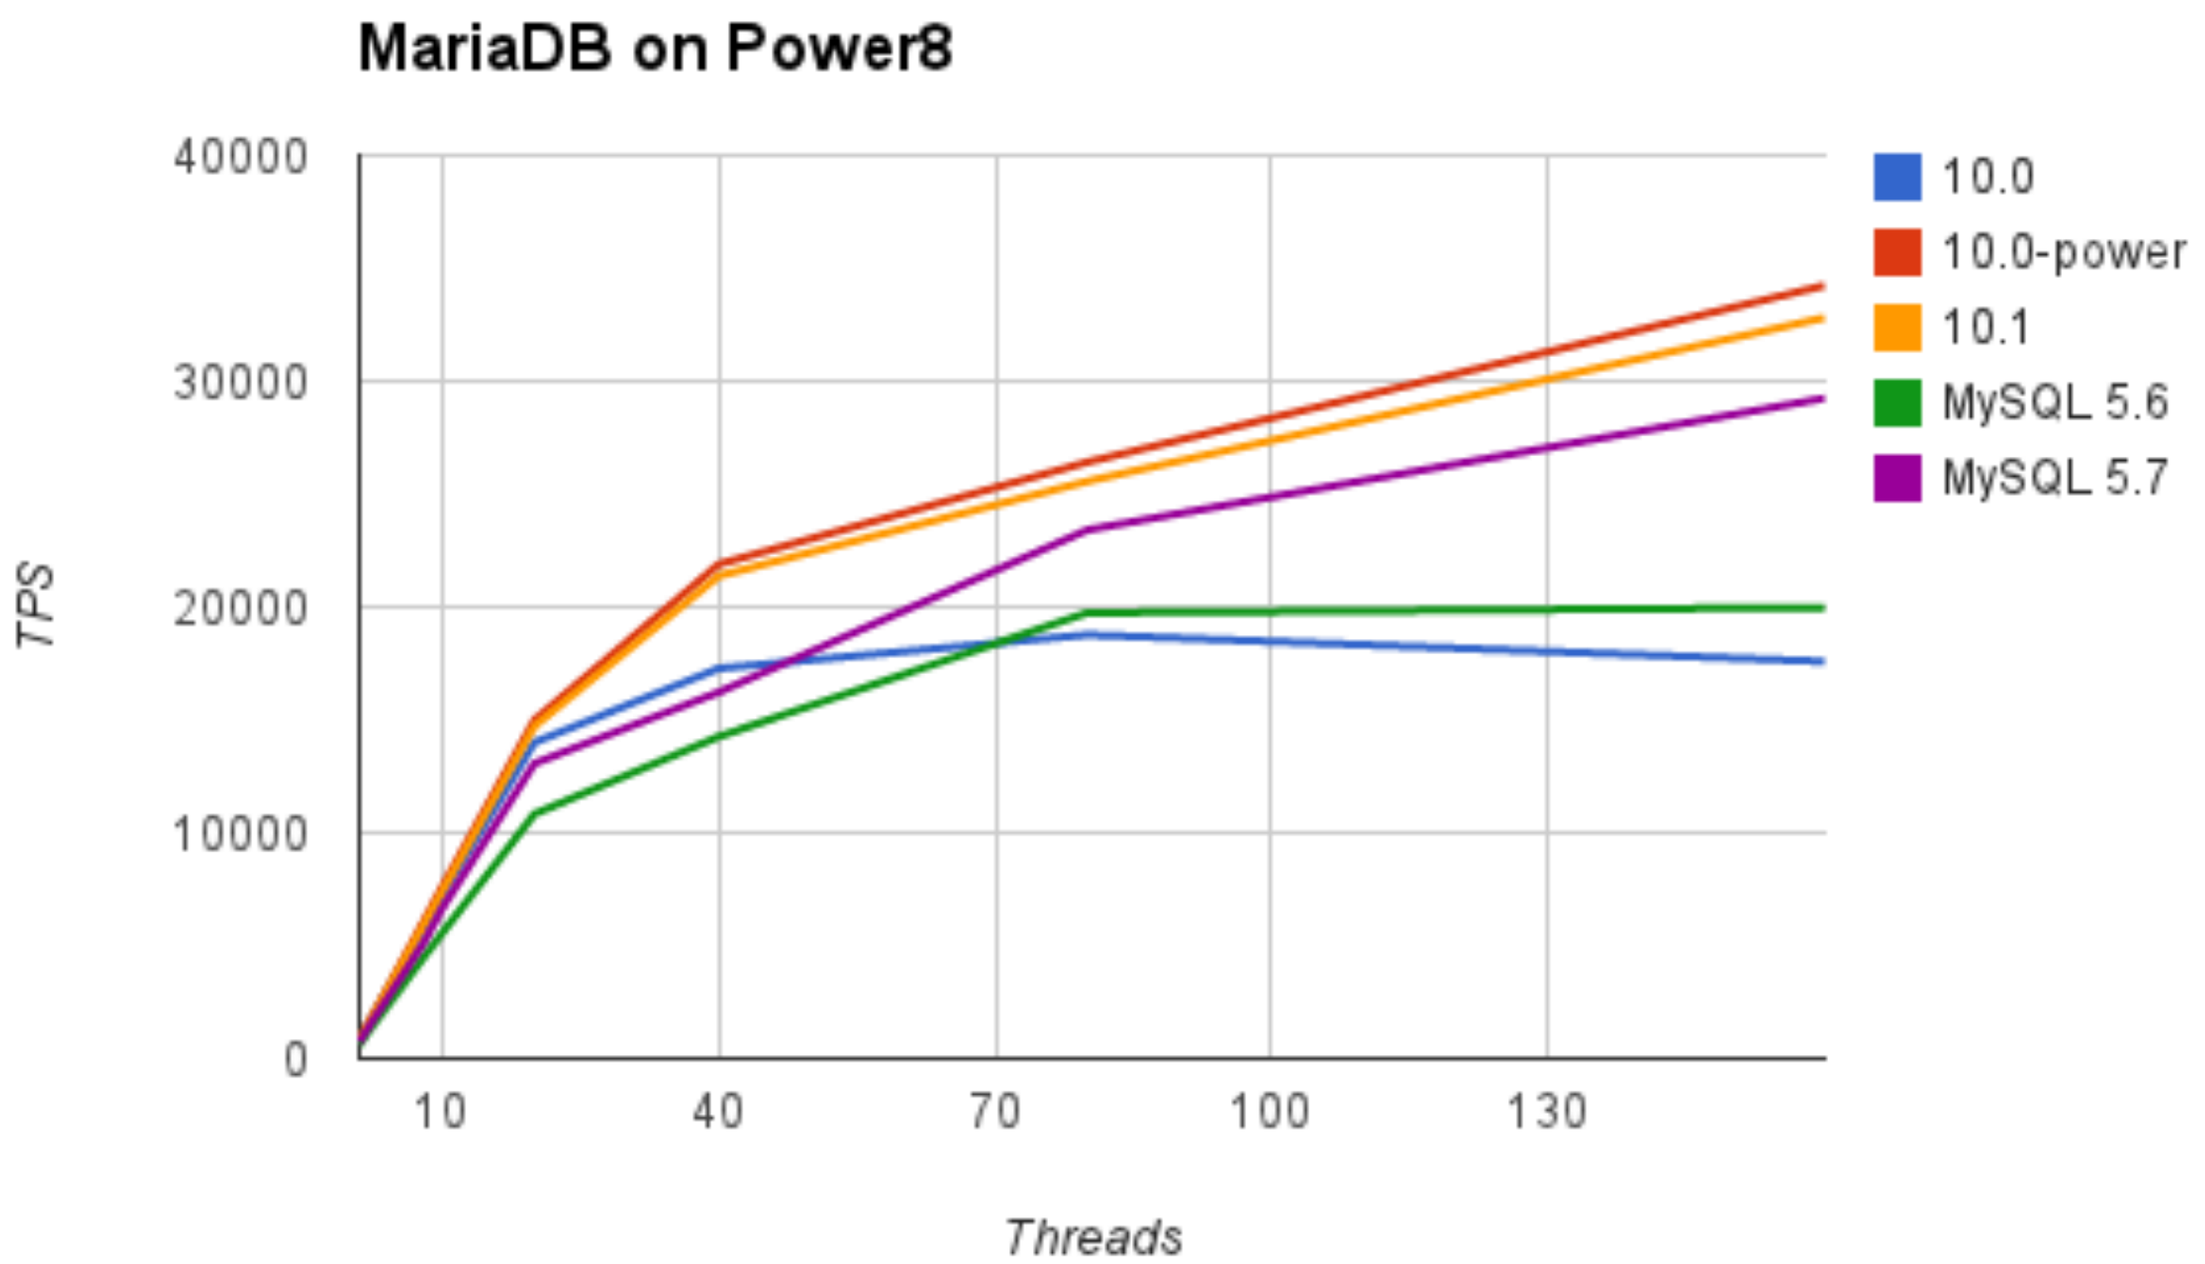

- http://svoj-db.blogspot.ru/2014/12/mariadb-on-power8-2014-wrap-up.html

Operational Ease

Per query variables

- Long history (http://www.bytebot.net/blog/archives/2014/05/04/per-query-variable-settings-in-mysqlpercona-serverwebscalesql)
- `SET STATEMENT max_statement_time=1000 FOR SELECT name FROM name ORDER BY name;`

Statement timeouts

- `MAX_STATEMENT_TIME` to abort long running queries
  - We call it “query timeouts” + have a compatible syntax
  - https://mariadb.atlassian.net/browse/MDEV-4427

New KILL syntax

- HARD | SOFT & USER USERNAME are MariaDB-specific (5.3.2)
- KILL QUERY ID query_id (10.0.5) - kill by query id, rather than thread id
- SOFT ensures things that may leave a table in an inconsistent state aren't interrupted (like REPAIR or INDEX creation for MyISAM or Aria)

```
KILL [HARD | SOFT] [CONNECTION | QUERY]
[thread_id | USER user_name]
```

Progress reporting

- ALTER TABLE & LOAD DATA INFILE

```
MariaDB [mail]> alter table mail engine = maria;

Stage: 1 of 2 'copy to tmp table' 17.55% of stage done

MariaDB [mail]> select id, user, db, command, state,
-> time_ms, progress from information_schema.processlist;

+--------+-------------------+----------+----------+
| command | state             | time_ms   | progress |
+--------+-------------------+----------+----------+
| Query   | copy to tmp table | 23407.131 |   17.551 |
+--------+-------------------+----------+----------+
1 row in set (0.47 sec)
```

Better for Developers and DBAs

Optimiser enhancements

- UNION ALL without temporary tables (5.7)
- Improve ORDER BY … LIMIT in optimiser
- `ANALYZE <statement>`
- EXPLAIN JSON (like 5.6)*
- `EXPLAIN ANALYZE with FORMAT=JSON`
  - includes data from the query execution itself — this is MariaDB only
  - https://mariadb.com/kb/en/mariadb/analyze-formatjson-examples/

CONNECT

- CONNECT having full JSON/BSON support
- Can read/write filename.json files with ease
- Writing — INSERT, UPDATE, DELETE is supported
- Naturally the other good use? ODBC connections to other databases

Other bits

- CREATE or REPLACE for most database objects minus indexes
- SET DEFAULT ROLE (there is a default role now for current user)
- FRM files are now not created for temporary tables
- `INFORMATION_SCHEMA.SYSTEM_VARIABLES` - information for system variables
- Microseconds in GET_LOCK()

Other bits

- Compiled with security hardening options (fortify source - https://mariadb.atlassian.net/browse/MDEV-5730)
- `@@sql_log_slow` can now be controlled on a session basis (not just globally)
- Sequence engine enabled by default
- `default_tmp_storage_engine` option
- ALGORITHM column in I_S.VIEWS

GIS

- Full compliance for the OGC standards around GIS.
  - MDEV-4045 Missing OGC Spatial functions
    - ST_Boundary, ST_ConvexHull, ST_IsRing, ST_PointOnSurface, ST_Relate
  - MDEV-60 Support for Spatial Reference systems for the GIS data. MDEV-12 OpenGIS: create required tables: GeometryColumns, related views.
- I_S tables: `GEOMETRY_COLUMN SPATIAL_REF_SYS`
- `REF_SYSTEM_ID per GEOMETRY` column

“Community Release” Thanks!

- Google - encryption, scrubbing, enhanced semisync, dump thread, thd_specifics plugin service
- Eperi - encryption
- DaumKakao - defragmentation, online alter progress monitoring
- Antony Curtis - compound statements
- Sriram Patil (GSoC) - CREATE or REPLACE/IF NOT EXISTS
- Daniel Black - finer grained status variables for replication monitoring
- FusionIO - atomic writes, page compression, TRIM, multi-threaded flushing
- Facebook — defragmentation, prefix index queries optimization, lazy flushing, buffer pool list scan optimization, configurable long, semaphore wait timeout
- Percona - SET STATEMENT, enforce_storage_engine

Welcoming 5.7 features MariaDB had

- Multi-source replication
- Dynamic replication filters
- SHOW EXPLAIN for connection_id
- GIS functionality
- Statement timeouts
- Change master without stopping SQL thread
- Online GTID implementation
- GTID no longer requires `log-slave-updates` to be enabled
- Virtual columns (generated columns)
- SHUTDOWN command
- FusionIO functionality

MariaDB still rocks!

- START TRANSACTION WITH CONSISTENT SNAPSHOT
- Integrated Galera Cluster
- Table/tablespace encryption
- Optimistic parallel replication
- Enhanced semi-sync replication
- InnoDB defragmentation
- ANALYZE <statement>
- Threadpool
- cracklib_password_check
- SQL error logging plugin
- Extended REGEXP
- Roles

Participate!

- Contribute code - github.com/mariadb/server
- Write KB articles - http://mariadb.com/kb/
- Report bugs: http://mariadb.org/jira
- Join us on #maria at irc.freenode.net
- Enable the feedback plugin ([enable-feedback] in my.cnf)
- Mailing lists: maria-discuss, maria-developers
- Tweet us @mariadb, Like us on FB, + on GPlus

Books!

1. MariaDB Crash Course, Ben Forta (September 2011)
2. Getting Started with MariaDB, Daniel Bartholomew (October 2013)
3. MariaDB Cookbook, Daniel Bartholomew (March 2014)
4. Real MariaDB, Matt Lee (April 2014)
5. Building a Web Application with PHP & MariaDB: A Reference Guide, Sai Srinivas Sriparasa (June 2014)
6. MariaDB: Beginners Guide, Rodrigo Ribeiro (August 2014)
7. Mastering MariaDB, Federico Razzioli (September 2014)
8. MariaDB High Performance, Pierre Mavro (September 2014)
9. Learning MySQL & MariaDB, Russell Dyer (April 2015)

In conclusion

- MariaDB is GPLv2 licensed, freedom guaranteed
- Its feature complete with MySQL + loaded with extras
- Enterprise features made open is great for users
- Its distributed everywhere

Thank you!

Colin Charles
colin@mariadb.com / byte@bytebot.net
http://bytebot.net/blog | @bytebot on twitter
slides: slideshare.net/bytebot