

# MariaDB Improvements

in the Latest Versions

September 14-16, 2020

Sveta Smirnova



PERCONA

# Sveta Smirnova

---



- MySQL Support Engineer
- Author of
  - [MySQL Troubleshooting](#)
  - JSON UDF functions
  - FILTER clause for MySQL
- Speaker
  - Percona Live, OOW, Fosdem, DevConf, HighLoad...



# Table of Contents

---

- Storage Improvements
- Optimizer and SQL Improvements

# MySQL

---

1995 Initial Release by MySQL AB

# MySQL

---

1995 Initial Release by MySQL AB  
2006 Percona Server for MySQL

# MySQL

---

1995 Initial Release by MySQL AB

2006 Percona Server for MySQL

2008 Sun owns MySQL

# MySQL

---

1995 Initial Release by MySQL AB

2006 Percona Server for MySQL

2008 Sun owns MySQL

2010 Oracle owns Sun

- and MySQL

# MySQL

---

1995 Initial Release by MySQL AB

2006 Percona Server for MySQL

2008 Sun owns MySQL

2010 Oracle owns Sun

■ and MySQL

2010 MariaDB



# MariaDB

---

- Started as fork of MySQL

# MariaDB

---

- Started as fork of MySQL
- Independent product

# MariaDB

---

- Started as fork of MySQL
- Independent product
- Hundreds of unique features

# MariaDB

---

- Started as fork of MySQL
- Independent product
- Hundreds of unique features
- Important improvements are backported
  - from MySQL

# This Talk

---

- Features, ported from MySQL 8.0

# This Talk

---

- ~~Features, ported from MySQL 8.0~~
  - Another talk



MySQL 8.0 and Percona Improvements

# This Talk

---

- ~~Features, ported from MySQL 8.0~~

- Another talk



MySQL 8.0 and Percona Improvements

- MariaDB Unique Features

- Storage Improvements

- Alternative storage engines
- Others

- Optimizer Improvements

- Advanced SQL
- Performance and diagnostics



# Storage Improvements



# S3

---

- Storage Engine

# S3

---

- Storage Engine
- Data stored on Amazon S3

# S3

---

- Storage Engine
- Data stored on Amazon S3
- Any original engine

```
ALTER TABLE my_table ENGINE=S3;  
ALTER TABLE my_table ENGINE=INNODB;  
ALTER TABLE my_table ENGINE=S3;
```

# S3

---

- Storage Engine
- Data stored on Amazon S3
- Any original engine
- Underlying engine is Aria

# S3

---

- Storage Engine
- Data stored on Amazon S3
- Any original engine
- Underlying engine is Aria
- Read-only tables

# S3

---

- Storage Engine
- Data stored on Amazon S3
- Any original engine
- Underlying engine is Aria
- Read-only tables
- Shared and separated storage for replication

# S3

---

- Storage Engine
- Data stored on Amazon S3
- Any original engine
- Underlying engine is Aria
- Read-only tables
- Shared and separated storage for replication
- Own layer to connect to S3: `libmarias3`

# ColumnStore

---

- Storage Engine
- Available since 2017 as a plugin
- Part of 10.5 Community Server release



# ColumnStore

---

- Columnar Storage Engine
- Based on InfiniDB

# ColumnStore

---

- Designed for parallel processing
- SQL interface
- JOINS with SQL engines (InnoDB etc.)
- Easy integration with existent SQL setup
- All advantages of the column storage

# Many More Storage Engines

---



MyRocks



Connect



Mroonga



OQGRAPH



SphinxSE



Spider

# System-Versioned Tables

---

- Change history

# System-Versioned Tables

---

- PITR on the fly

# System-Versioned Tables

---

- PITR on the fly
- Create a versioned table

```
MariaDB [test]> alter table employees ADD SYSTEM VERSIONING;  
Query OK, 300024 rows affected (1.060 sec)  
Records: 300024 Duplicates: 0 Warnings: 0
```

# System-Versioned Tables

---

- PITR on the fly
- Check table options

```
MariaDB [test]> show create table employees\G
***** 1. row *****
Table: employees
Create Table: CREATE TABLE 'employees' (
  'emp_no' int(11) NOT NULL,
  'birth_date' date NOT NULL,
  'first_name' varchar(14) COLLATE utf8mb4_unicode_ci NOT NULL,
  'last_name' varchar(16) COLLATE utf8mb4_unicode_ci NOT NULL,
  'gender' enum('M','F') COLLATE utf8mb4_unicode_ci NOT NULL,
  'hire_date' date NOT NULL,
  PRIMARY KEY ('emp_no')
) ENGINE=InnoDB DEFAULT CHARSET=utf8mb4 COLLATE=utf8mb4_unicode_ci WITH SYSTEM VERSIONING
1 row in set (0.001 sec)
```



# System-Versioned Tables

---

- PITR on the fly
- Perform activities

```
MariaDB [test]> select count(*) from employees where hire_date < '1986-01-01';
+-----+
| count(*) |
+-----+
|      35316 |
+-----+
1 row in set (0.188 sec)
```

```
MariaDB [test]> delete from employees where hire_date < '1986-01-01';
Query OK, 35316 rows affected (0.837 sec)
```



# System-Versioned Tables

---

- PITR on the fly
- Lost rows

```
MariaDB [test]> select count(*) from employees where hire_date < '1986-01-01';
```

```
+-----+  
| count(*) |  
+-----+  
|          0 |  
+-----+
```

```
1 row in set (0.103 sec)
```

# System-Versioned Tables

---

- PITR on the fly
- Restore deleted rows

```
MariaDB [test]> insert into employees select * FROM employees
                -> FOR SYSTEM_TIME between (now() - interval 1 hour) and now()
                -> where hire_date < '1986-01-01';
```

```
Query OK, 35316 rows affected (0.721 sec)
Records: 35316 Duplicates: 0 Warnings: 0
```

```
MariaDB [test]> select count(*) from employees where hire_date < '1986-01-01';
```

```
+-----+
| count(*) |
+-----+
|    35316 |
+-----+
1 row in set (0.216 sec)
```



# System-Versioned Tables

---

- Change history
- PITR on the fly
- Tunable
  - Time
  - Transaction ID
  - With History, stored separately
  - Partitioned
  - With columns, excluded from versioning
  - With time periods: application history
  - With both system and application period levels

# Invisible columns

---

- Columns, not visible for SELECT \*

```
MariaDB [test]> alter table employees add column address JSON invisible;  
Query OK, 0 rows affected (0.005 sec)  
Records: 0 Duplicates: 0 Warnings: 0
```

- Update address for some employees

```
MariaDB [test]> select * from employees limit 3;
```

emp_no	birth_date	first_name	last_name	gender	hire_date
10001	1953-09-02	Georgi	Facello	M	1986-06-26
10002	1964-06-02	Bezalel	Simmel	F	1985-11-21
10003	1959-12-03	Parto	Bamford	M	1986-08-28

```
3 rows in set (0.001 sec)
```

# Invisible columns

---

- Columns, not visible for SELECT \*
- Accessible for the direct query

```
MariaDB [test]> select first_name, last_name, json_extract(address, "$.City")  
-> from employees where address is not null;
```

```
+-----+-----+-----+  
| first_name | last_name | json_extract(address, "$.City") |  
+-----+-----+-----+  
| Tianruo    | Jenevein  | "Espoo"                          |  
| Dulce      | Kolinko   | "Espoo"                          |  
| Masasuke   | Gill      | "Espoo"                          |  
| Toshimi    | Karner    | "Espoo"                          |  
| Danco      | Yetto     | "Espoo"                          |  
+-----+-----+-----+  
5 rows in set (0.173 sec)
```

# Sequences

---

```
MariaDB [test]> create sequence odds start with 1 increment by 2;  
Query OK, 0 rows affected (0.001 sec)
```

```
MariaDB [test]> create sequence evens start with 2 increment by 2;  
Query OK, 0 rows affected (0.001 sec)
```

```
MariaDB [test]> create table numbers(  
    -> id int not null auto_increment primary key,  
    -> odd int default next value for odds,  
    -> even int default next value for evens  
    -> );  
Query OK, 0 rows affected (0.004 sec)
```

# Sequences

---

```
MariaDB [test]> insert into numbers values (), (), (), (), ();  
Query OK, 5 rows affected (0.002 sec)  
Records: 5 Duplicates: 0 Warnings: 0
```

```
MariaDB [test]> select * from numbers;
```

```
+----+-----+-----+  
| id | odd  | even |  
+----+-----+-----+  
|  1 |    1 |    2 |  
|  2 |    3 |    4 |  
|  3 |    5 |    6 |  
|  4 |    7 |    8 |  
|  5 |    9 |   10 |  
+----+-----+-----+
```

```
5 rows in set (0.001 sec)
```

# INET6 data type

---

- Native IPv6 support

```
MariaDB [test]> create table inet6_test(ip inet6);
Query OK, 0 rows affected (0.001 sec)
MariaDB [test]> insert into inet6_test values('fe80::e43b:a1ff:fe32:cb0b'); - IPv6
Query OK, 1 row affected (0.009 sec)
MariaDB [test]> insert into inet6_test values('::192.168.60.100'); - IPv4 compatible
Query OK, 1 row affected (0.001 sec)
MariaDB [test]> insert into inet6_test values('::ffff:192.168.60.100'); - IPv4 mapped
Query OK, 1 row affected (0.001 sec)
```

```
MariaDB [test]> SELECT ip, if(IS_IPV4_COMPAT(ip), 'deprecated', 'OK') FROM inet6_test;
```

ip	if(IS_IPV4_COMPAT(ip), 'deprecated', 'OK')
fe80::e43b:a1ff:fe32:cb0b	OK
::192.168.60.100	deprecated
::ffff:192.168.60.100	OK

```
3 rows in set (0.002 sec)
```






# WAIT and NOWAIT

---


- Normal table scan



# WAIT and NOWAIT

---


- WAIT n



# WAIT and NOWAIT

---

- NOWAIT or WAIT 0



# Optimizer and SQL Improvements

# Aggregate stored functions

---

- Your own aggregate functions

# Aggregate stored functions

---

- Your own aggregate functions

```
MariaDB [test]> create aggregate function count_positive(val int) returns int
-> begin
->   declare result int unsigned default 0;
->   declare exit handler for not found return result;
->   main_loop:
->   loop
->     fetch group next row;
->     if sign(val) = 1 then
->       set result := result + 1;
->     end if;
->   end loop;
-> end
-> |
```

Query OK, 0 rows affected (0.001 sec)

# Aggregate stored functions

---

- Your own aggregate functions

```
MariaDB [test]> select color, count_positive(num) from numbers group by color;
```

```
+-----+-----+
| color | count_positive(num) |
+-----+-----+
| NULL  |                    1 |
| blue  |                    0 |
| green |                    2 |
| red   |                    2 |
+-----+-----+
```

```
4 rows in set (0.001 sec)
```

```
MariaDB [test]> select group_concat(num) from numbers;
```

```
+-----+
| group_concat(num) |
+-----+
| 1,-2,3,-4,5,-6,7,-8,9,0 |
+-----+
```

```
1 row in set (0.001 sec)
```



# MariaDB 10.0+ Features, Announced in MySQL 8.0


---

- Histogram-based statistics
  - And MariaDB-only independent table statistics
- ANALYZE
- CTEs and WITH statement
- Window functions
- CHECK constraint
- Roles




# Data Distribution

---



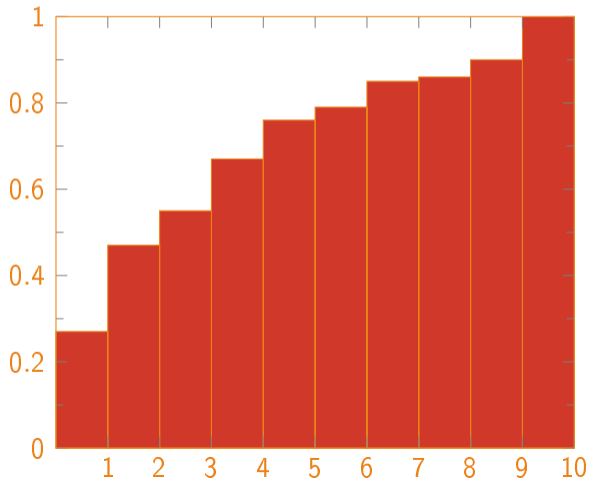
# Indexes: Cardinality

---



# Same Data in Histograms

---



# ANALYZE

---



## Based on MySQL bug #78651

- EXPLAIN is telling lies

```
MariaDB [test]> explain select * from ol
      -> where thread_id=10432 and site_id != 9939 order by id limit 3\G
***** 1. row *****
      id: 1
      select_type: SIMPLE
      table: ol
      type: index
      possible_keys: thread_id
      key: PRIMARY
      key_len: 4
      ref: NULL
      rows: 33
      Extra: Using where
```

# ANALYZE

---



## Based on MySQL bug #78651

- ANALYZE executes the statement

```
MariaDB [test]> analyze select * from ol
      -> where thread_id=10432 and site_id != 9939 order by id limit 3\G
***** 1. row *****
...
      type: index
possible_keys: thread_id
      key: PRIMARY
      key_len: 4
      ref: const
      rows: 100000
      r_rows: 100000.00
      filtered: 8.96
      r_filtered: 0.00
      Extra: Using where
```

# SQL DML

---

- CTEs
  - Not recursive

```
MariaDB [employees]> with
-> dept_data as
-> (select emp_no, dept_name from dept_emp join departments using (dept_no)
-> select first_name, last_name, dept_name
-> from employees join dept_data using(emp_no)
-> order by hire_date desc limit 3;
```

```
+-----+-----+-----+
| first_name | last_name | dept_name |
+-----+-----+-----+
| Bikash     | Covnot    | Quality Management |
| Yucai      | Gerlach   | Production          |
| Hideyuki   | Delgrande | Development         |
+-----+-----+-----+
```

```
3 rows in set (0.00 sec)
```



# SQL DML

---

- CTEs
  - Recursive

```
MariaDB [employees]> with recursive rand_generator(id, rand_value) as
->   (select 1, rand()
->   union all select id+1, rand()
->   from rand_generator where id < 5)
-> select * from rand_generator;
```

```
+-----+-----+
| id  | rand_value          |
+-----+-----+
|  1  | 0.5599308382346582 |
|  2  | 0.2151867702744778 |
|  3  | 0.39614136740205935 |
|  4  | 0.33514655692050843 |
|  5  | 0.4873087131300091 |
+-----+-----+
```

```
5 rows in set (0.00 sec)
```



# SQL DML

---

- Window functions

```
MariaDB [employees]> select  
    -> row_number() over win as id, dept_no, dept_name from departments  
    -> window win  
    -> as (order by dept_no);
```

id	dept_no	dept_name
1	d001	Marketing
2	d002	Finance
3	d003	Human Resources
4	d004	Production
5	d005	Development
6	d006	Quality Management
7	d007	Sales
8	d008	Research
9	d009	Customer Service



# CHECK Constraint

---

- Custom rules validation

# CHECK Constraint

---

- Custom rules validation

```
MariaDB [test]> create table even (even_value int check(even_value % 2 = 0)) engine=innodb;  
Query OK, 0 rows affected (0.004 sec)
```

```
MariaDB [test]> insert into even value(2);  
Query OK, 1 row affected (0.003 sec)
```

```
MariaDB [test]> insert into even value(1);  
ERROR 4025 (23000): CONSTRAINT 'even.even_value' failed for 'test'.'even'
```

# Roles

---

- No need to repeat GRANT

```
MariaDB [test]> create role read_only, admin;  
Query OK, 0 rows affected (0.001 sec)
```

```
MariaDB [test]> grant select on *.* to read_only;  
Query OK, 0 rows affected (0.001 sec)
```

```
MariaDB [test]> grant super on *.* to admin;  
Query OK, 0 rows affected (0.001 sec)
```

# Roles

---

- No need to repeat GRANT

```
MariaDB [test]> create user sveta;  
Query OK, 0 rows affected (0.001 sec)  
MariaDB [test]> create user kaj;  
Query OK, 0 rows affected (0.001 sec)  
MariaDB [test]> create user privileged;  
Query OK, 0 rows affected (0.001 sec)
```

```
MariaDB [test]> grant read_only to sveta;  
Query OK, 0 rows affected (0.001 sec)  
MariaDB [test]> grant read_only to kaj;  
Query OK, 0 rows affected (0.001 sec)  
MariaDB [test]> grant admin to privileged;  
Query OK, 0 rows affected (0.001 sec)
```

# Roles

---

- No need to repeat GRANT
- DEFAULT roles

# Conclusions

---

- New versions have new features
- MariaDB implements advanced features early
- Upgrade to new version
- Explore all MariaDB advantages!

# More Details

---



The S3 Storage Engine

MariaDB ColumnStore

Rewinding time with System Versioned Tables

Invisible Columns

Sequences

INET6 Data Type

WAIT and NOWAIT

# More Details

---



Stored Aggregate Functions

Histogram-Based Statistics

Engine Independent Statistics

ANALYZE Statement

Common Table Expressions

Window Functions



# More Details

---



CHECK Constraints  
Roles

# Thank you!

---



[www.slideshare.net/SvetaSmirnova](http://www.slideshare.net/SvetaSmirnova)



[twitter.com/svetsmirnova](https://twitter.com/svetsmirnova)



[github.com/svetasmirnova](https://github.com/svetasmirnova)



**DATABASE PERFORMANCE  
MATTERS**