



# Was ist neu in MariaDB 11.4?

FrOSCon 2024, St. Augustin

**Oli Sennhauser**

CTO, FromDual GmbH

<https://www.fromdual.com/presentations>

# Über FromDual GmbH



www.fromdual.com

## Support



## Beratung



## remote-DBA



## Schulung





# Content

## **MariaDB 11.4**

- › **MariaDB Support Matrix**
- › **Features in MariaDB 11.x (up to 11.4)**
- › **Upgrade to MariaDB 11.4**
- › **Migrate MySQL to MariaDB 11.4**
- › **Future (MariaDB 11.5, 11.6, 11.7)**
- › **Past (MariaDB 10.7 to 10.11)**



# MariaDB Support Matrix

- Up to **10.6 (LTS)** release cycles were clear and simple!
  - For about 12 years (2009 – 2021)
  - Every 1 – 1.5 years a new LTS release
  - Approx. 4 releases per year, 5 years support
  - No new features, only bug fixes
- **10.7 – 10.10 (STS)**: Innovation release model: Short term maintenance releases (1 year maintenance, quarterly releases)
- **10.11 (LTS)**: Long term maintenance release (5 years, Feb 2028)
- **11.0 – 11.2 (STS)**: Short term maintenance releases (1y + 4/y)
- **11.3 (RR)**: Rolling release model (no more releases after GA!)
- **11.4 (LTS)**: Long term maintenance release (5 years, May 2029)
  - GA: 29 May 2024
- **11.5 – 11.6 (RR)**



# Lifetime support matrix

MAJOR VERSION	GA DATE	LTS	SUPPORT END
MariaDB Database 10.6	July 2021	x	July 2026
MariaDB Database 10.7 <sup>1</sup>	February 2022		February 2023
MariaDB Database 10.8	May 2022		May 2023
MariaDB Database 10.9	August 2022		August 2023
MariaDB Database 10.10	November 2022		November 2023
MariaDB Database 10.11 <sup>2</sup>	February 2023	x	February 2028
MariaDB Database 11.0 <sup>4</sup>	June 2023		June 2024
MariaDB Database 11.1	August 2023		August 2024
MariaDB Database 11.2	November 2023		November 2024
MariaDB Database 11.3 <sup>5</sup>	February 2024		none
MariaDB Database 11.4	May 2024	x	May 2029
MariaDB Database 11.5	expected Q3 2024		none
MariaDB Database 11.6	expected Q4 2024		none
MariaDB Database 11.7	expected Q1 2025		none



# Features up to 11.4

- **At what release are you at the moment?**
- **Focus: 11.0 - 11.4**



# Optimizer

- **Biggest / most important change**
- **Configurable Cost Based Optimizer (CBO)**
  - **Before a pure? rule based optimizer**
- **Semi-join (Subqueries) optimization for single table UPDATE/DELETE**
- **DATE or YEAR functions comparing against constant can make use of an index**
  - **WHERE YEAR(dt) = 2024**
  - **WHERE DATE(dt) <= '2024-08-01'**
- **UCASE(varchar\_column) can use index if CI collation**
- **MIN() and MAX() are now optimized also with descending indexes**



# Practical example

- MariaDB 10.11

```
SQL> EXPLAIN SELECT * FROM test WHERE DATE(ts) <= '2024-08-01';
+-----+-----+-----+-----+-----+-----+-----+-----+
| table | type | possible_keys | key | key_len | rows | Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+
| test  | ALL  | NULL          | NULL | NULL    | 8148448 | Using where |
+-----+-----+-----+-----+-----+-----+-----+-----+

8 rows in set (3.433 sec)
```

- MariaDB 11.4

```
SQL> EXPLAIN SELECT * FROM test WHERE DATE(ts) < '2024-08-01';
+-----+-----+-----+-----+-----+-----+-----+-----+
| table | type | possible_keys | key | key_len | rows | Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+
| test  | range | ts           | ts  | 4       | 8    | Using index condition |
+-----+-----+-----+-----+-----+-----+-----+-----+

8 rows in set (0.001 sec)
```



# Optimizer - Impact

- **Monty: Everything will be better/faster!!!**
  - **From Oracle CBO we know, this is nearly true...**
- **More complicated / possibilities to tune**
- **Should help in various cases**
  - **Complex queries (Frameworks! e.g. Hybernate, etc.)**
  - **Low cardinality indexes**
  - **Huge ranges (> 10%)**
  - **Flipping QEP**
  - **Subselects**
  - **ORDER BY / GROUP BY resolved via index**
- **We wait for field test results...**



# InnoDB

- Temporary tablespace (`ibtmp1`) can be shrunk online
  - `innodb_truncate_temporary_tablespace_now`
- System tablespace (`ibdata1`) can now be shrunk at startup
  - UNDO space
  - Tables in system tablespace (`innodb_file_per_table = on`)
  - Temporary tables earlier blew up `ibdata1`
- InnoDB Change Buffer was removed
  - Some corruptions should be eliminated now
- `innodb_file_per_table = off` is deprecated
  - No more tables in the system tablespace!
- InnoDB flushing (old: `innodb_flush_method`)
  - `innodb_log_file_buffering` / `innodb_data_file_buffering`
  - `innodb_log_file_write_through` / `innodb_data_file_write_through`



# Practical example

```
# ls -lah ibtmp1
-rw-rw---- 1 mysql mysql 1.2G Aug  2 15:27 ibtmp1

SQL> SET GLOBAL innodb_truncate_temporary_tablespace_now = 1;

# ls -lah ibtmp1
-rw-rw---- 1 mysql mysql 12M Aug  2 15:28 ibtmp1

[Note] InnoDB: Truncating temporary tablespace from 74496 to 768 pages
[Note] InnoDB: Temporary tablespace truncated successfully
```



# Practical example

```
# ll -h ibdata1
-rw-rw---- 1 mysql mysql 1.4G Jun 12 14:01 ibdata1

SQL> SET GLOBAL innodb_fast_shutdown=0;
SQL> SHUTDOWN;

[Note] bin/mariadb (initiated by: root[root] @ localhost []): Normal shutdown
...
[Note] InnoDB: Truncating system tablespace from 90880 to 768 pages
[Note] InnoDB: System tablespace truncated successfully
[Note] InnoDB: Starting shutdown...
...
[Note] InnoDB: Removed temporary tablespace data file: "./ibtmp1"
...
[Note] bin/mariadb: Shutdown complete

# ll ibdata1 -h
-rw-rw---- 1 mysql mysql 12M Jun 12 14:11 ibdata1
```



# Replication / Binary Log

- Binlog writing speed improved
- Binary log purging: `max_binlog_total_size`
- Disables automatic binary log purging until # connected slaves
  - `slave_connections_needed_for_purge`
- Reduce binlog size: After image only has columns that changed:
  - `binlog_row_image = full_nodup`
  - `mariadb-binlog --flashback`
- GTID index on binary log. Master finds position faster where to start
- Delayed replication with GTID:
  - `START SLAVE UNTIL {SQL_BEFORE_GTIDS | SQL_AFTER_GTIDS}`
- `binlog_*` variables are visible as system variables



# Practical example

```
# 11
-rw-rw---- 1 mysql mysql 141570270 Jun 12 chef_mariadb-114_binlog.000010
-rw-rw---- 1 mysql mysql      4096 Jun 12 chef_mariadb-114_binlog.000010.idx
-rw-rw---- 1 mysql mysql 519087231 Jun 12 chef_mariadb-114_binlog.000011
-rw-rw---- 1 mysql mysql      4096 Jun 12 chef_mariadb-114_binlog.000011.idx
-rw-rw---- 1 mysql mysql      1381 Jun 12 chef_mariadb-114_binlog.000012
-rw-rw---- 1 mysql mysql          0 Jun 12 chef_mariadb-114_binlog.000012.idx
-rw-rw---- 1 mysql mysql      1775 Aug  2 chef_mariadb-114_binlog.index
```

```
SQL> SHOW GLOBAL VARIABLES LIKE 'binlog%db';
```

Variable_name	Value
binlog_do_db	
binlog_ignore_db	



# Security

- **SSL/TLS**
  - Enabled by default
  - Clients can validate self-signed server certificates
  - Clients require SSL and server side certification validation enabled by default
  - Same for replication clients
  - `--disable-ssl` or `--disable-ssl-verify-server-cert`
  - New client options: `--tls-fp=<fp>` `--tls-fplist=file`
  - Think about old clients!!!
- New privileges:
  - To see NOT our own routines but executable routines as well:
  - `SHOW CREATE ROUTINE`
- Various Security fixes... (CVE-...)



# Practical example

```
SQL> status
-----
mariadb Ver 15.1 Distrib 10.11.8-MariaDB, for linux-systemd (x86_64)

...
SSL:                Not in use
```

```
SQL> status
-----
mariadb from 11.4.2-MariaDB, client 15.2 for linux-systemd (x86_64)

...
SSL:                Cipher in use is ECDHE-RSA-AES128-GCM-SHA256, cert is OK
```





# Interesting for Dev

- **Online Schema Change**
  - `ALTER TABLE ... ALGORITHM=COPY, LOCK=NONE` → simultaneous DML allowed in most cases, Caution: COPY
- **New/changed JSON functions**
  - `JSON_TABLE` allows retrieval of the key
  - `JSON_OBJECT_FILTER_KEYS`, `JSON_OBJECT_TO_ARRAY` and `JSON_ARRAY_INTERSECT`
  - `JSON_SCHEMA_VALID`
- **Data types**
  - `GEOMETRY` data types can now be partitioned
  - `INET4` can be casted into `INET6`, server can automatically convert
- **Application-Time Periods / System versioned tables**
  - New `I_S` views: `PERIODS`, `KEY_PERIOD_USAGE`
  - New columns in `COLUMNS` view: `IS_SYSTEM_TIME_PERIOD_START` and `IS_SYSTEM_TIME_PERIOD_END`
- **Miscellaneous Functions**
  - `CONV()` up to base 62 (before base 36)
  - `KDF` slow function to protect against passwords hashing generators
  - `FORMAT_PICO_TIME` function (`P_S???`)
  - `DATE_FORMAT` time zone and offset from UTC with `%Z %z`
  - Functions `AES_ENCRYPT` and `AES_DECRYPT` extended
- **Miscellaneous**
  - `CREATE PACKAGE` in `sql_mode = default` for `SQL/PSM`
  - `ALTER TABLE ... IMPORT TABLESPACE` enhancement
  - Stored Programs validation improved



# Practical example

```
SQL> SELECT FORMAT_PICO_TIME(0000000000043) AS ps  
        , FORMAT_PICO_TIME(00000000004321) AS ns  
        , FORMAT_PICO_TIME(00000043211234) AS us  
        , FORMAT_PICO_TIME(00432112344321) AS ms  
        , FORMAT_PICO_TIME(43211234432123) AS s;
```

ps	ns	us	ms	s
43 ps	4.32 ns	43.21 us	432.11 ms	43.21 s

```
SQL> SET SESSION time_zone = 'Europe/Moscow';
```

```
SQL> SELECT DATE_FORMAT(CAST('2024-05-05 14:56:00' AS DATETIME), '%Z') AS time_zone;
```

time_zone
MSK

```
SQL> SELECT DATE_FORMAT(CURRENT_TIMESTAMP(), '%Z') AS time_zone;
```

time_zone
+0300



# Practical example

```
SET SESSION sql_mode = ''; -- NOT oracle!

DELIMITER //

CREATE OR REPLACE PACKAGE TestPackage
  PROCEDURE TestProcedure(INOUT pText VARCHAR(64));
  FUNCTION TestFunction(IN pText VARCHAR(64)) RETURNS VARCHAR(128);
END;

CREATE OR REPLACE PACKAGE BODY TestPackage

  PROCEDURE TestProcedure(INOUT pText VARCHAR(64))
  BEGIN
    SELECT CONCAT('Hello ', pText) INTO pText;
  END;

  FUNCTION TestFunction(IN pText VARCHAR(64)) RETURNS VARCHAR(128)
  BEGIN
    RETURN CONCAT('Hello ', pText);
  END;

END;
//

DELIMITER ;
```



# Practical example

```
SQL> SET @text := 'world';
SQL> CALL TestPackage.TestProcedure(@text);
SQL> SELECT @text AS greetings;
+-----+
| greetings |
+-----+
| Hello world |
+-----+

SQL> SELECT TestPackage.TestFunction('world') AS greetings;
+-----+
| greetings |
+-----+
| Hello world |
+-----+
```



# Interesting for Ops

- **Backup and Restore**
  - `mariadb-dump --tab --parallel=#` → `mariadb-import`
  - **MariaDB Backup:** `xtrabackup_*` files renamed to `mariadb_backup_*`
- **Partitioning**
  - `EXCHANGE PARTITION / CONVERT TABLE: {WITHOUT | WITH} VALIDATION`
- **sys Schema**
  - `sys.privileges_by_table_by_level`
- **Connection redirection (atm client-side support is missing)**
  - `SET GLOBAL redirect_url="mysql://192.168.0.2:3306";`
  - **Bypass proxy, load balance without proxy, MaxScale?**
- **Spider**
  - **Spider option table (not COMMENT) any more**
- **Processlist**
  - `I_S.PROCESSLIST SENT_ROWS` column (currently sending rows)
- **Variables added, removed, deprecated**



# Practical example

```
SQL> SELECT distinct grantee FROM sys.privileges_by_table_by_level
      WHERE table_schema = 'test' AND table_name = 'test';
```

```
+-----+
| grantee          |
+-----+
| 'brman'@'127.0.0.1' |
| 'mysql'@'localhost' |
| 'root'@'localhost' |
| 'PUBLIC'@''       |
+-----+
```

```
SQL> SELECT id, user, host, db, command, sent_rows
      FROM information_schema.processlist;
```

```
+----+-----+-----+-----+-----+-----+
| id | user  | host          | db   | command | sent_rows |
+----+-----+-----+-----+-----+-----+
| 79 | dummy | localhost    | test | Query   | 8343395  |
| 77 | root  | localhost    | sys  | Query   | 0        |
+----+-----+-----+-----+-----+-----+
```



# Upgrade to MariaDB 11.4

- **Minor / major (~~11.4.2~~ → ~~11.4.3~~ / 10.11 → 11.4)**
- ***You should be able to trivially upgrade from ANY earlier MariaDB version to the latest one, usually in a few seconds.***
- **Logical or physical?**
- **Use Master (old)/Slave (new) replication?**
  - **For testing**
  - **For fall-back**
- **Sources / Literature:**
  - **<https://mariadb.com/kb/en/upgrading-between-major-mariadb-versions/>**
  - **<https://mariadb.com/kb/en/upgrading-from-mariadb-10-11-to-mariadb-11-4/>**



# How to upgrade?

- Ideally do all this on a testing system first!
- Read the upgrade notes and fix issues
- Backup the database
- `SQL> SET GLOBAL innodb_fast_shutdown = 0;`
- Shutdown database server
- Uninstall old MariaDB version
- Install new MariaDB version (Repo!?!)
- Change `my.cnf`
- Start database server again
- `# mariadb-upgrade --force`
- Check the error log!
- Backup database again under new version





# Downgrade from MariaDB 11.4

- ***Downgrading MariaDB is not officially supported between major versions.***
- **Why?**
  - **Changes in `mysql` schema**
  - **Disk storage format changes**
  - **Substantial feature changes (e.g. REDO log)**
- **Downgrade to 10.11 should work, to 10.6 probably not**
- **And now?**
  - **Have a good backup!**
  - **Test your upgrade well!**
  - **Think about Master/Slave replication for fall-back**



# Migration from MySQL

- *You should be able to trivially upgrade from ANY earlier MariaDB version to the latest one, usually in a few seconds. **This is also mainly true for any MySQL version  $\leq 5.7$  to MariaDB 10.4 and newer.***
  - Logical or physical?
  - Logical: Got an error! Worked around...
  - Physical: Found crashing bug: MDEV-34706 → wip
- Use Master (old)/Slave (new) replication?
  - For testing
  - For fall-back
  - Bug in 10.6 and 10.11: MDEV-34412 → no action :-(
    - Not sure about 11.4 :-(



# Practical example

```
# mysqldump --user=root --all-databases > full_dump_57.sql

# cat /tmp/full_dump_57.sql | mariadb -uroot
ERROR 1050 (42S01) at line 833: Table 'user' already exists

# mysqldump --user=root --databases a b c > full_dump_57.sql
Works!
```

```
# systemctl start mariadb
# mariadb-upgrade --user=root --force
# mariadb-upgrade --user=root --force
mariadb-check: Got error: 2026: TLS/SSL error: The TLS connection was non-
properly terminated. when executing 'CHECK TABLE ... FOR UPGRADE'
FATAL ERROR: Upgrade failed

--> Crash https://jira.mariadb.org/browse/MDEV-34706
```



# Literature / Sources

- **MariaDB Lifetime Support Matrix**  
<https://fromdual.com/enterprise-support-for-mariadb-and-mysql>
- **Shrinking the InnoDB system tablespace**  
<https://fromdual.com/shrinking-the-innodb-system-tablespace>
- **Changes and Improvements in MariaDB 11.4**  
<https://mariadb.com/kb/en/changes-improvements-in-mariadb-11-4/>
- **MariaDB 11.4 LTS review: cost-based optimiser, autoconfigured SSL and more**  
<https://vettabase.com/mariadb-11-4-lts-review/>
- **The Optimizer Cost Model from MariaDB 11.0**  
<https://mariadb.com/kb/en/the-optimizer-cost-model-from-mariadb-11-0/>
- **Release Notes - MariaDB 11.4 Series**  
<https://mariadb.com/kb/en/release-notes-mariadb-11-4-series/>
- **Packages for Store Routines in MariaDB 11.4**  
<https://programmingbrain.com/2024/04/packages-for-store-routines-in-mariadb-11-4>
- **Upgrading Between Major MariaDB Versions**  
<https://mariadb.com/kb/en/upgrading-between-major-mariadb-versions/>
- **Downgrading between Major Versions of MariaDB**  
<https://mariadb.com/kb/en/downgrading-between-major-versions-of-mariadb/>



# Future 11.5

- Limit of created temporary files and tables
- `TIMESTAMP` year problem 2038 is solved (moved to 2106)
- Optimizer improvements (partitions, selectivity of pushed index conditions)
- MariaDB protocol improvements
- `INFORMATION_SCHEMA`
  - `USERS` table: `password_expiration_time`, `password_errors`
  - New columns in `TABLE_STATISTICS`, `CLIENT_STATISTICS`, `USER_STATISTICS` and `INDEX_STATISTICS`: `KEY_READ_{HITS|MISSES}`, `QUERIES`, `ROWS_{I|U|D}`
- `FLUSH GLOBAL STATUS`
- InnoDB async IO statistics: `Innodb_async_{reads|writes}_*` → **Monitoring!**
- Show variable deprecation: `mariadb --verbose --help | grep -B5 Depre`
- Extend Query Response Time plug-in (compatible with Percona)
- Sequences improved
- Default collation: `uca1400_ai_ci`
- `mariadb-dump` parallel dump with `--dir`
- Source: <https://mariadb.com/kb/en/changes-improvements-in-mariadb-11-5/>



# Future 11.6 and 11.7

- **11.6: End of July 2024, GA Q3/Q4 2024**
- **Main feature: AI/Vectors → Next talk Vicențiu Ciorbaru**
- **Default character set `utf8mb4`**
- `mariadb-import --dir restore parallel backup`
- **Automatic SST user account management**
- **Optimizer: Improve selectivity computations for multi-part-keys**
- **Single table `DELETE` now supports table aliases**
- **New PARSEC authentication plug-in (default in 11.7?)**
- **Unix socket authentication supports `authentication_string`**
- **Replication:**
  - **New definition for `Seconds_behind_master`**
  - **New variables in `SHOW SLAVE STATUS`: `Master_last_event_time`, `Slave_last_event_time`, `Master_Slave_time_diff`**
- **`I_S.SLAVE_STATUS` view**
- **Add thread name to O/S thread**
- **Source: <https://mariadb.com/kb/en/changes-improvements-in-mariadb-11-6/>**
- **11.7: Catalog: Q4 2024 or Q1 2025?**



# Past features 10.7 to 10.11

- **Authentication**
  - GRANT TO PUBLIC
  - SUPER and READ ONLY ADMIN separated (`super_read_only` in MySQL)
  - Several addresses to bind server to: `bind_address`
  - New plug-in: `password_reuse_check`
- **Optimizer**
  - ANALYZE FORMAT=JSON time spent in optimizer
  - Descending indexes (`ORDER BY idx_fld DESC`)
  - Various other performance improvements
  - SHOW ANALYZE FORMAT=JSON
  - EXPLAIN FOR CONNECTION
- **Replication / Binary Log**
  - GTID base replication default
  - `slave_max_statement_time` (abort query!)
  - `MASTER_USE_GTID = Current_Pos` → `MASTER_DEMOTE_TO_SLAVE`
  - `--do-domain-ids --ignore-domain-ids --ignore-server-ids` for `mariadb-binlog`
  - ALTER TABLE is replicated at begin and not at the end
  - Multi source replication support MySQL style syntax
  - `mariadb-binlog` supports `--start-position` and `--stop-position` for GTID



# Past features 10.7 to 10.11

- Galera
  - IP whitelist for SST/IST
  - JSON file interface to wsrep node state
- InnoDB
  - InnoDB performance improvements
  - Redo log improvements (write)
  - `innodb_write_io_threads`, `innodb_read_io_thread` now dynamic
  - `innodb_log_file_size` is now dynamic
- General improvements
  - I\_S performance issues fixed
  - New and improved JSON functions
  - New UUID data type
  - System versioned tables (allow history modification, `mariadb-dump` dumps also history, dumps table as of a specified timestamp)
  - UCA-14.0.0 collation added (to become future default)
  - Windows password-less login, collation issues with `utf8mb4`
  - `CONVERT PARTITION TO TABLE` / `CONVERT TABLE TO PARTITION` instead of `CREATE/EXCHANGE/DROP`
  - New/improved functions (`RANDOM_BYTES`, `CRC32`, `CRC32C`, `NATURAL_SORT_KEY`, `SFORMAT`), data types (`INET4`), Stored Functions (`OUT`, `INOUT`), `my_print_defaults --default-*`, compression plug-ins
- Hashicorp Key Management plug-in
- Server variables, new, deprecated, removed, merged, dynamic
- Various security fixes
- Source: <https://mariadb.com/kb/en/changes-improvements-in-mariadb-1011/>



# Q & A



www.fromdual.com



**Fragen ?**

**Diskussion?**

**Wir haben Zeit für ein persönliches Gespräch...**

- **FromDual bietet neutral und unabhängig:**
    - **Beratung**
    - **remote-DBA**
    - **Support für MariaDB, MySQL und Galera Cluster**
    - **Schulung**
- [www.fromdual.com/presentations](http://www.fromdual.com/presentations)**