

Metadata locking in *MySQL 5.5*

Konstantin Osipov
Engineer @ Mail.Ru

kostja.osipov@gmail.com

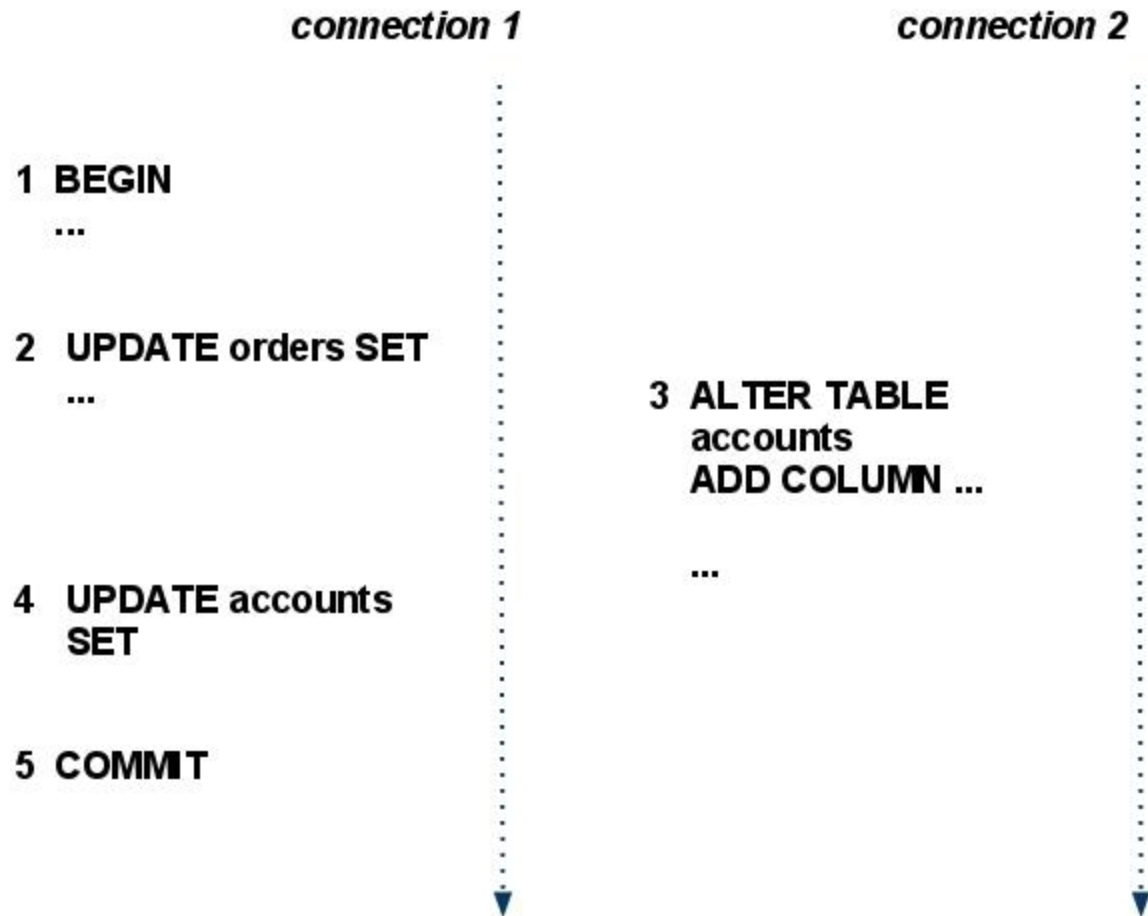


*These slides released under the Creative Commons
Attribution-Noncommercial-Share Alike License*

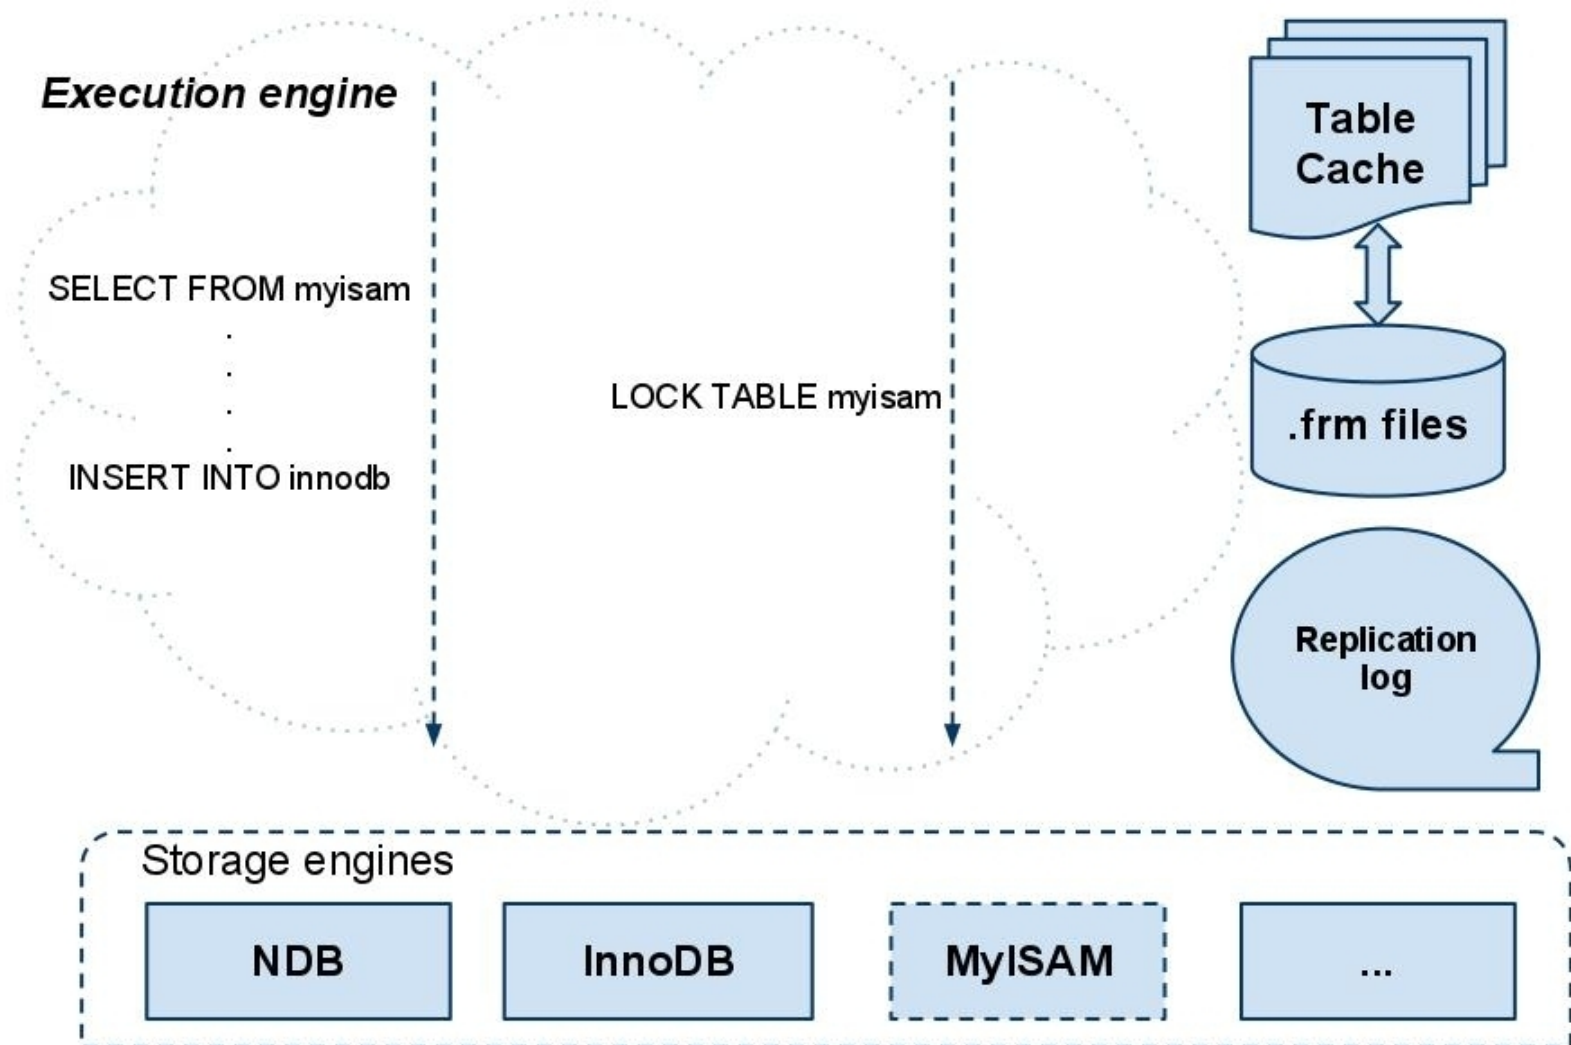
Mastering a 25-foot saw



DDL vs. DML: an illustration



Locking agents



Define the problem (2)

Four groups of bugs:

- ALTER, DROP violate ACID
- ... or ... break replication
- deadlocks in LOCK TABLES, FLUSH, SET GLOBAL read_only and ALTER
- DML doesn't scale or THANK YOU
LOCK_open

Define the problem (3)

[Bug#989 "If DROP TABLE while there's an active transaction, wrong binlog order"](#)

[Bug#25144 "replication / binlog with view breaks"](#)

[Bug#39675 rename tables on innodb tables with pending transactions causes slave data issue"](#)

[Bug#30977 "Concurrent statement using stored function and DROP FUNCTION breaks SBR"](#)

[Bug#37346 "innodb does not detect deadlock between update and alter table"](#)

[Bug#45225 "Locking: hang if drop table with no timeout"](#)

[Bug#46224 "HANDLER statements within a transaction might lead to deadlocks"](#)

[Bug#41804 "purge stored procedure cache causes mysterious hang for many minutes"](#)

[Bug#26141 "mixing table types in trigger causes full table lock on innodb table"](#)

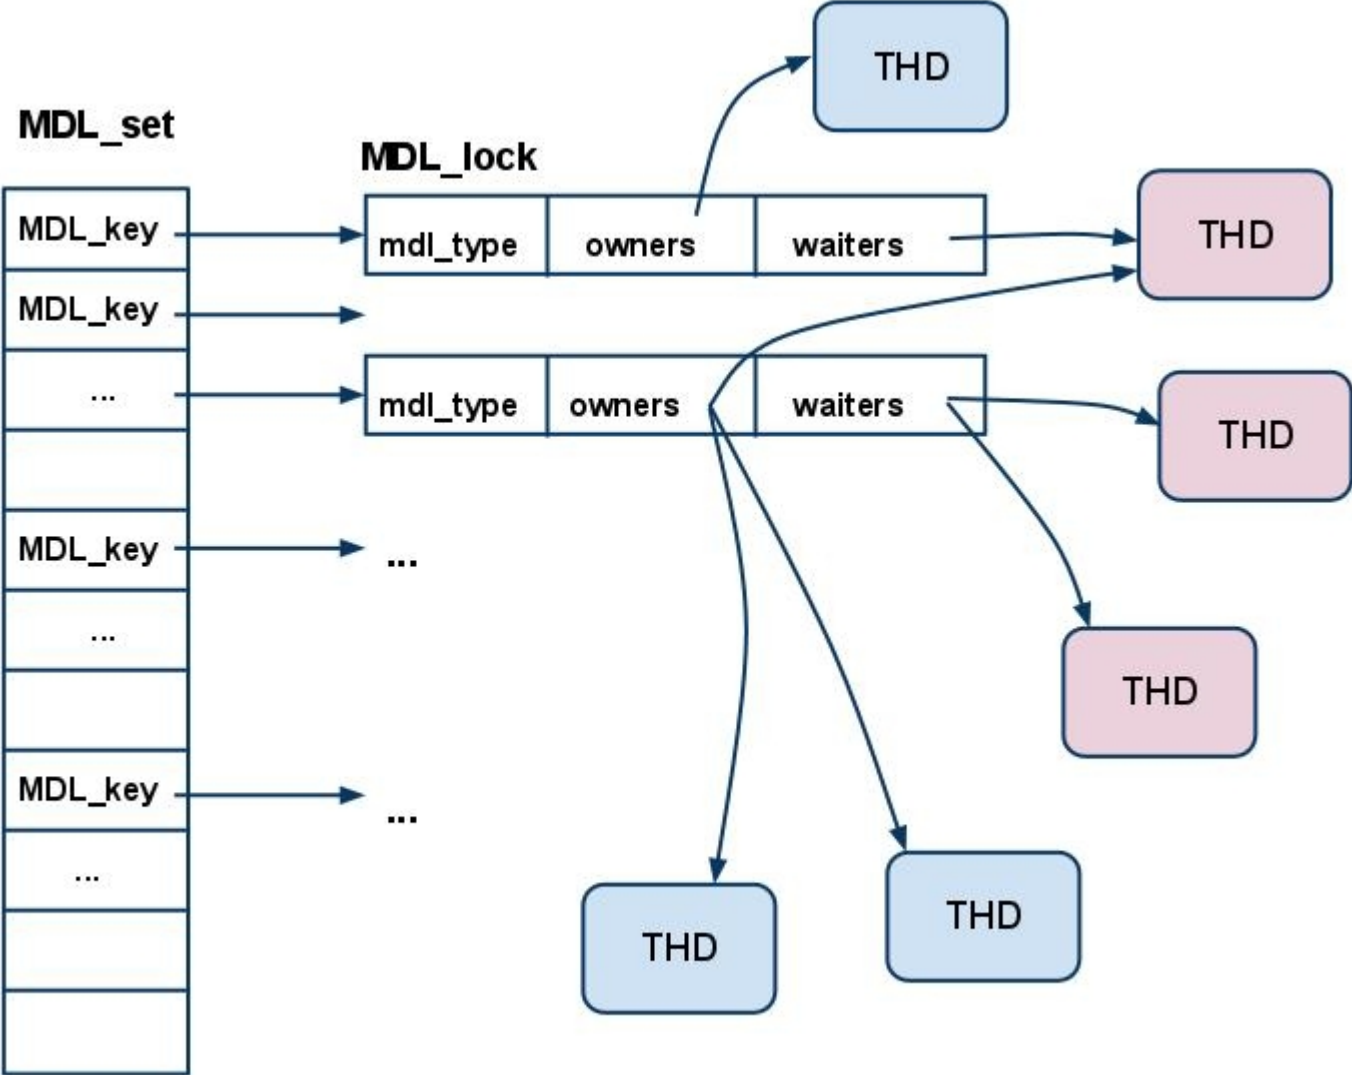
[Bug#43685 "Lock table affects other non-related tables"](#)

[Bug#22876 "Four-way deadlock"](#)

[Bug#44613 "SELECT statement inside FUNCTION takes a shared lock"](#)

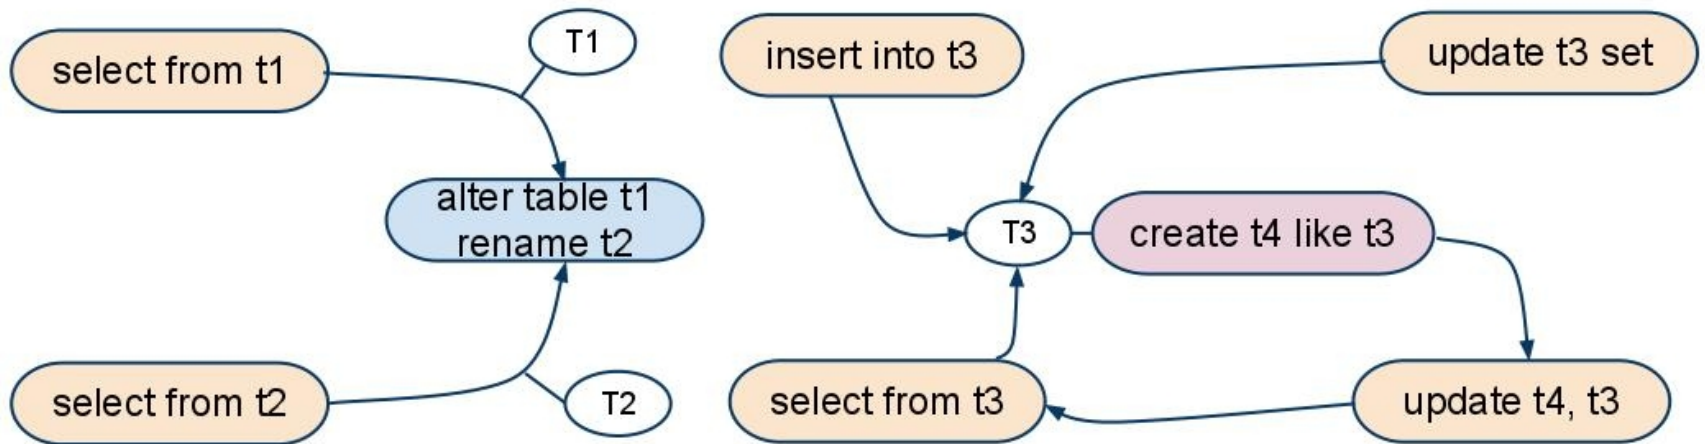
[Bug#33948 "performance issue when using sysbench benchmark on a multiple-core system."](#)

In come metadata locks



Deadlock detection

How does a typical wait-for graph look like?

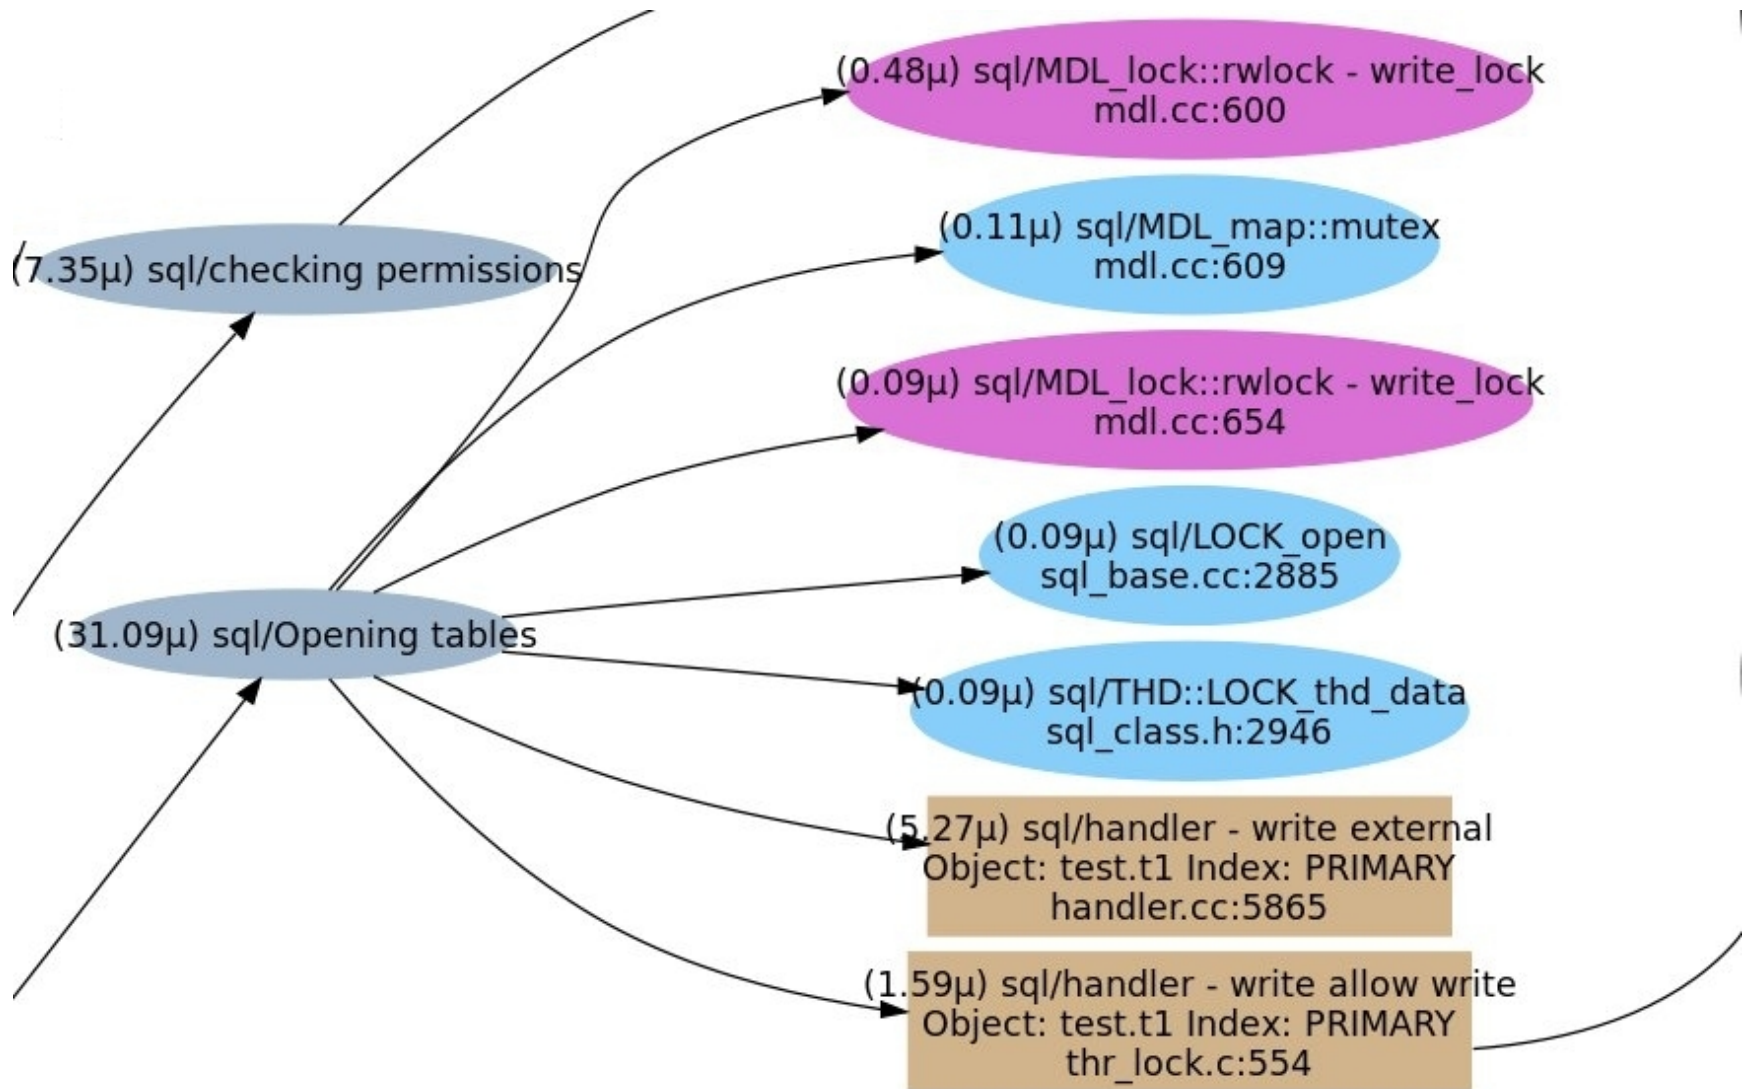


How to plug in

Engines need to plug-in to avoid cross-storage-engine deadlocks.

- it's safe to `#include "mdl.h"`, no extra deps
- **subclass** `MDL_waitfor_subgraph` to accept
- `MDL_waitfor_graph_visitor`
- let us know when you're waiting, by calling
- `MDL_context::will_wait_for()`


Performance impact



The roadmap

- help us add cross-storage-engine deadlock detection!
- thr_lock.c and handler::store_lock() need not be mandatory any more
- consistent DDL across all NDB cluster
- shard MDL_set (performance++)
- instrumentation

Performance impact



That's it.

Thank you for coming!

Further reading: WL#3873, WL#3726, WL#4284
Mdl.pdf is available in MySQL internals manual.