

DBMS
Version 5.15.0

Michael Fruth
michael.fruth@uni-passau.de
University of Passau
Passau, Bavaria, Germany

DBMS
Version 5.15.1



MariaDB MDEV-20753

Patch - if (tables_list && !procedure)

issue: + if (tables_list && top_join_tab_count && !procedure)

Release new version



Critical (security)
issue found

How to update the DBMS as cloud provider? Problems:


① A sense of urgency as security leaks might be exploited.

② For self-service products offered to the customer:
When are downtimes acceptable to the customer?

③ Rolling updates or hot standby is very expensive.

Solution: Dynamic Software Updating (DSU), a.k.a. **live patching**, allows to patch a running process without restart.

Global Quiescence




• **Each** thread

needs to reach its barrier before a patch can be applied.

Can cause unbounded wait times:

- (1) Thread waits for user input that may never be provided.
- (2) Thread performs a long running task.
- (3) Deadlock can occur with inter-thread dependencies.

Local Quiescence



• Each thread is “patched” **individually**, regardless of the status of the other threads.

• WFPATCH [1], a novel approach proposed by Rommel et al., enables local quiescence using a modified linux kernel and quiescence points implemented into the source code of the application to patch.

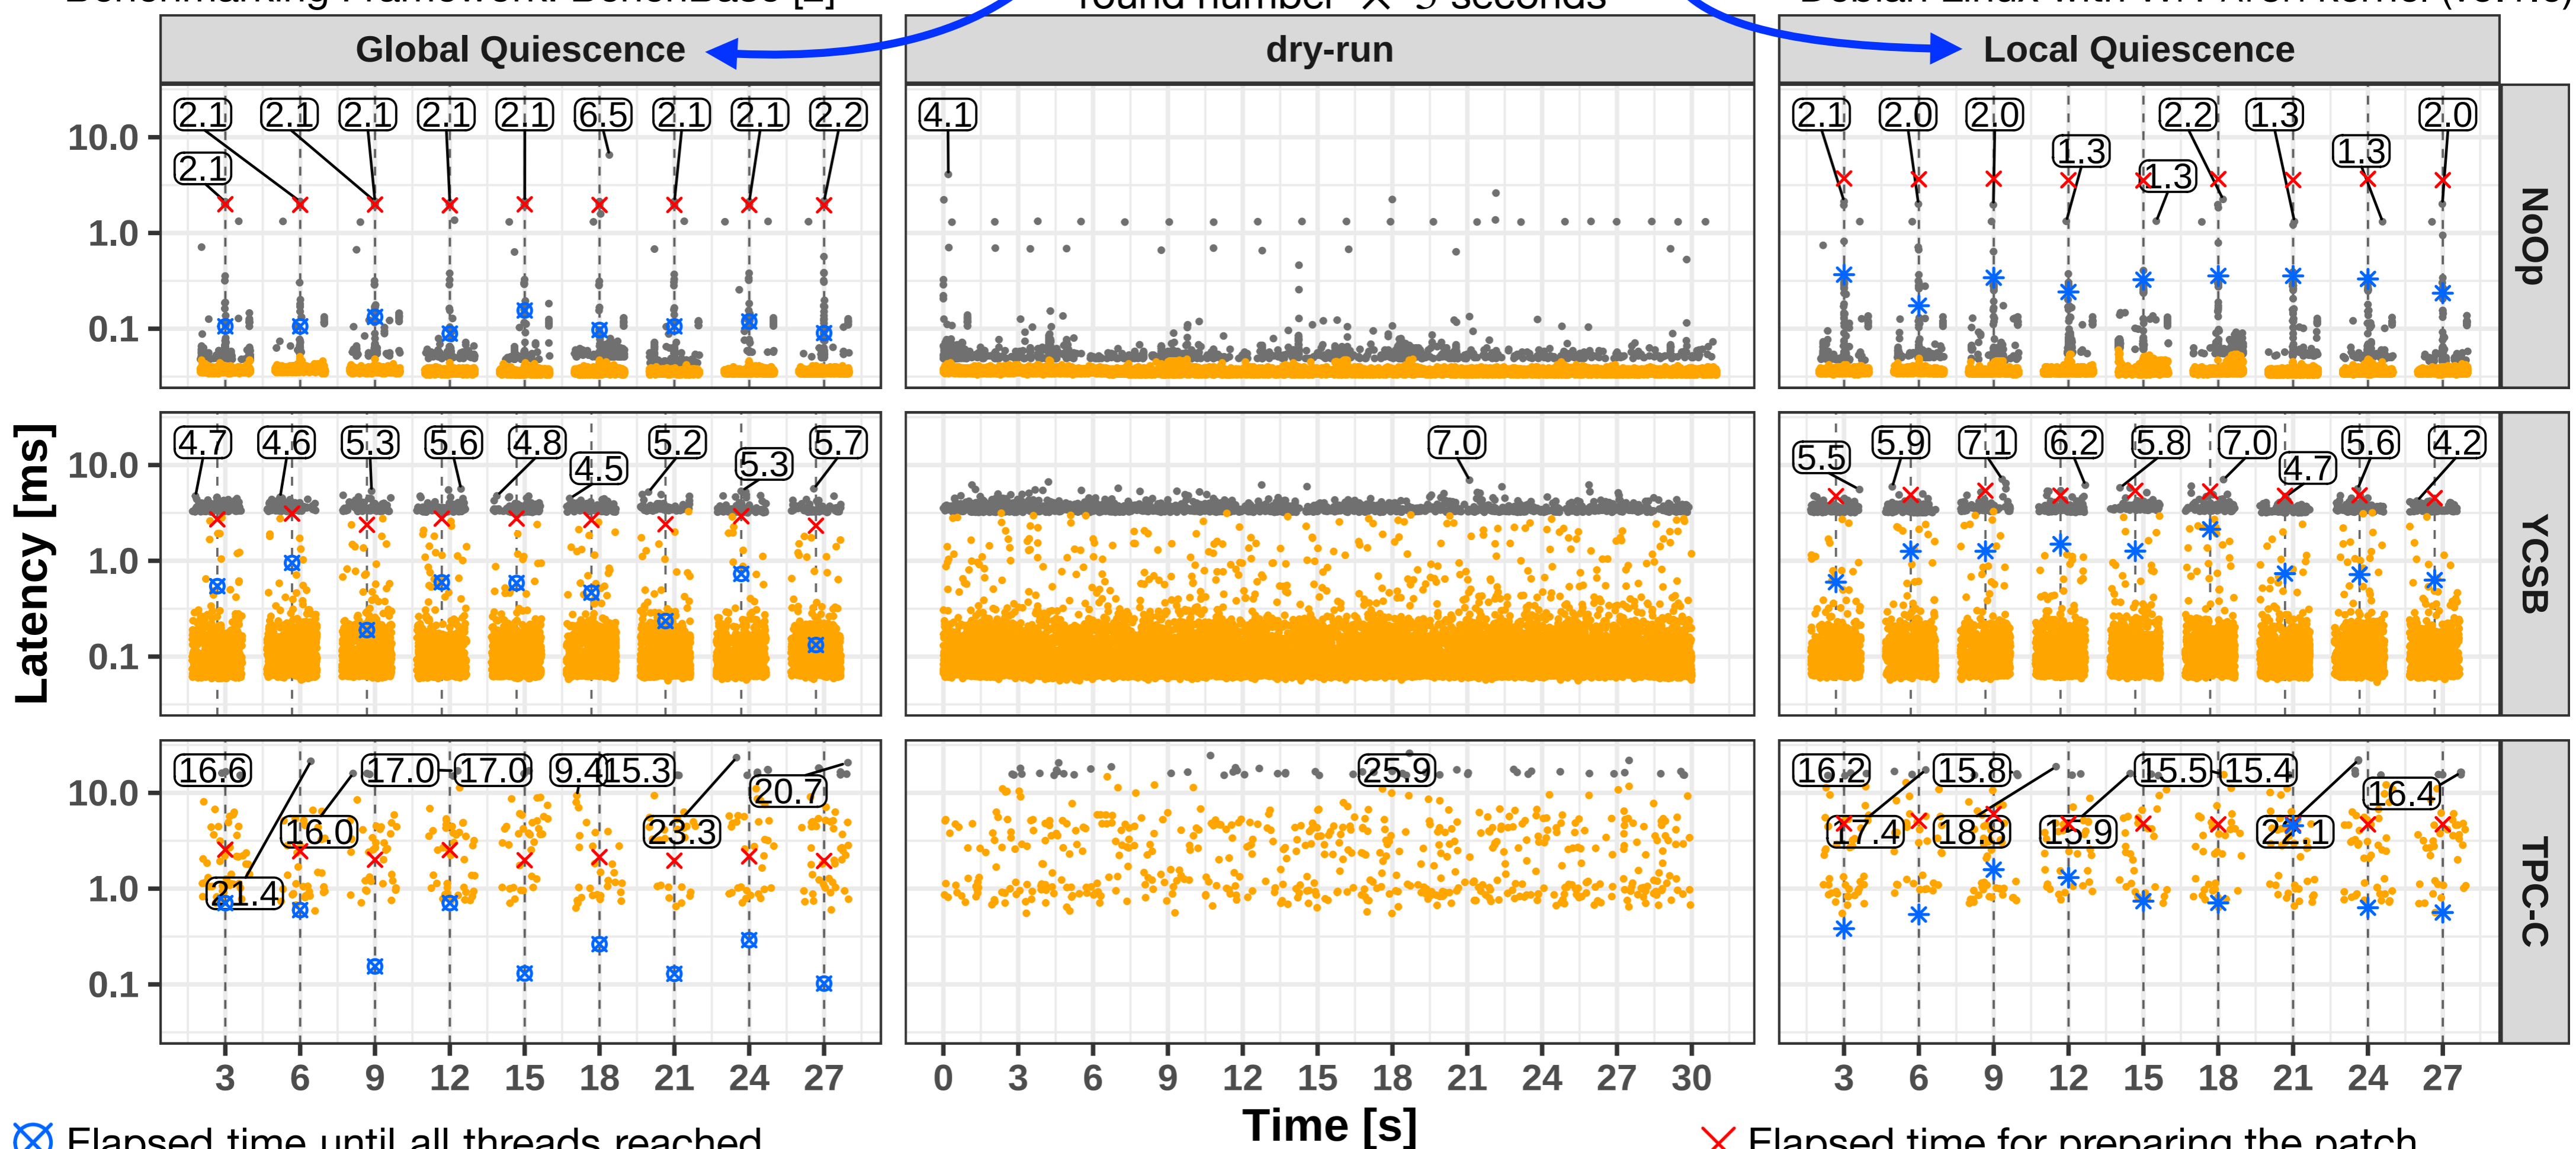
Experiments:

- Benchmarks: NoOp (No Operation), YCSB, TPC-C
- Database: MariaDB (10.5.1)
- Benchmarking Framework: BenchBase [2]

- Patch: MDEV-21665
- 1 dry-run (without applying a patch)
- 9 runs while applying a patch after “round number × 3 seconds”

System Setup:

- Server: DELL R640 (2x Intel Gold 6248R; 384 GB RAM)
- Debian Linux with WFPATCH kernel (v5.1.0)



⊗ Elapsed time until all threads reached their barrier (maximum time a thread blocks and waits for other threads)

● Latency exceeding 99.95th latency percentile
● Down-sampled (0.5%) standard latency

⊗ Elapsed time for preparing the patch
⊗ Elapsed time until all threads migrated to the patched process version (corresponds to the quiescence time of global quiescence)

References

- [1] Florian Rommel, Christian Dietrich, Daniel Friesel, Marcel Köppen, Christoph Borchert, Michael Müller, Olaf Spinczyk, and Daniel Lohmann. 2020. From Global to Local Quiescence: Wait-Free Code Patching of Multi-Threaded Processes. In Proc. OSDI.
- [2] Djelle Eddine Difallah, Andrew Pavlo, Carlo Curino, and Philippe Cudré-Mauroux. 2013. OLTP-Bench: An Extensible Testbed for Benchmarking Relational Databases. Proc. VLDB Endow. 7, 4 (2013).

Acknowledgments

I thank Stefanie Scherzinger (University of Passau), Wolfgang Mauerer, Ralf Ramsauer (Technical University of Applied Sciences Regensburg), Daniel Lohmann, Florian Rommel (Leibniz University of Hannover) and Christian Dietrich (Hamburg University of Technology).