

Understanding and Improving Software Performance at MongoDB

David Daly
MongoDB Inc
USA
david.daly@mongodb.com

ABSTRACT

It is important for developers to understand the performance of a software project as they develop new features, fix bugs, and try to generally improve the product. At MongoDB we have invested in building a performance infrastructure to support our developers. The infrastructure automates the provisioning of systems under test, the running of performance tests against those systems, collecting many metrics from the tests and system under test, and making sense of all the results.

Our performance infrastructure and processes are continually changing. As the system has become more powerful we have used it more and more: adding new tests, new configurations, and new measurements. Tools and processes that work on one scale of use start to break down at higher scales and we must adapt and update. If we do a good job, we keep pace with the rising constraints. If we do a great job, we make the system fundamentally better even as we scale the system.

In this talk we describe our performance testing environment at MongoDB and its evolution over time. The core of our environment is a focus on automating everything, integrating into our continuous integration (CI) system (Evergreen), controlling as many factors as possible, and making everything as repeatable and consistent as possible. After describing that core, we will discuss the scaling challenges we have faced, before relating what we have done to address those scaling challenges and improve the system overall.

CCS CONCEPTS

- **Information systems** → **Database performance evaluation**;
- **Software and its engineering** → **Software performance**; Software development methods.

KEYWORDS

performance testing; database performance; continuous integration

ACM Reference Format:

David Daly. 2023. Understanding and Improving Software Performance at MongoDB. In *Companion of the 2023 ACM/SPEC International Conference on Performance Engineering (ICPE '23 Companion)*, April 15–19, 2023, Coimbra, Portugal. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/3578245.3584855>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ICPE '23 Companion, April 15–19, 2023, Coimbra, Portugal

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0072-9/23/04.

<https://doi.org/10.1145/3578245.3584855>

1 BIOGRAPHY

David is a staff engineer at MongoDB focused on server performance. He focuses on increasing the understanding of how MongoDB's software performs for its customers. In practice this includes: Asking hard questions about MongoDB performance and then trying to answer them (or having someone else try to answer them); Challenging assumptions and commonly accepted wisdom around MongoDB performance; Encouraging everyone at MongoDB to think about performance, including adding new performance tests relevant to their ongoing work (e.g., adding new performance tests for new features or refactors); And explaining the current state of performance to others. He helped build and design MongoDB's performance testing infrastructure from the bottom up. At various times this required focusing on complete end-to-end automation, control of test noise and variability, working around test noise, and building processes to make sure that issues identified by the infrastructure were properly recognized and addressed.

