

Performance Engineering and Database Development at MongoDB

David Daly
MongoDB Inc
david.daly@mongodb.com

ABSTRACT

Performance and the related properties of stability and resilience are essential to MongoDB. We have invested heavily in these areas: involving all development engineers in aspects of performance, building a team of specialized performance engineers to understand issues that do not fit neatly within the scope of individual development teams, and dedicating multiple teams to develop and support tools for performance testing and analysis.

We have built an automated infrastructure for performance testing [6] that is integrated with our continuous integration system [4]. Performance tests routinely run against our development branch in order to identify changes in performance as early as possible. We have invested heavily to ensure both that results are low noise and reproducible [5] and that we can detect when performance changes [3]. We continue to invest to make the system better [2] and to make it easier to add new workloads [1].

All development engineers are expected to interact with our performance system: investigating performance changes, fixing regressions, and adding new performance tests. We also expect performance to be considered at project design time. The project design should ensure that there is appropriate performance coverage for the project, which may require repurposing existing tests or adding new ones.

Not all performance issues are specific to a team or software module. Some issues emerge from the interaction of multiple modules or interaction with external systems or software. To attack these larger problems, we have a dedicated performance team. Our performance team is responsible for investigating these more complex issues, identifying high value areas for improvement, as well as helping guide the development engineers with their performance tests.

We have experience both hiring and training engineers for the performance engineering skills needed to ship a performant database system. In this talk we will cover the skills needed for our performance activities and which skills, if added to undergraduate curricula, would help us the most. We will address the skills we would like all development engineers to have, as well as those for our dedicated performance team.

CCS CONCEPTS

• **Software and its engineering** → **Software performance**; *Designing software*; • **Information systems** → **Database performance evaluation**; • **Social and professional topics** → **Computer science education**.

KEYWORDS

education, performance, testing, continuous integration, variability

ACM Reference Format:

David Daly. 2021. Performance Engineering and Database Development at MongoDB. In *Companion of the 2021 ACM/SPEC International Conference on Performance Engineering (ICPE '21 Companion)*, April 19–23, 2021, Virtual Event, France. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/3447545.3451199>

REFERENCES

- [1] 2020. Genny workload generator. <https://github.com/mongodb/genny> original-date: 2018-02-12T19:23:44Z.
- [2] David Daly. 2021. Creating a Virtuous Cycle in Performance Testing at MongoDB. In *Proceedings of the 2021 ACM/SPEC International Conference on Performance Engineering (ICPE '21)*, April 19–23, 2021, Virtual Event, France. ACM. <https://doi.org/10.1145/3427921.3450234>
- [3] David Daly, William Brown, Henrik Ingo, Jim O'Leary, and David Bradford. 2020. The Use of Change Point Detection to Identify Software Performance Regressions in a Continuous Integration System. In *Proceedings of the ACM/SPEC International Conference on Performance Engineering (ICPE '20)*. Association for Computing Machinery, Edmonton AB, Canada, 67–75. <https://doi.org/10.1145/3358960.3375791>
- [4] Kyle Erf. 2016. Evergreen Continuous Integration: Why We Reinvented The Wheel. <https://engineering.mongodb.com/post/evergreen-continuous-integration-why-we-reinvented-the-wheel>
- [5] Henrik Ingo and David Daly. 2019. Reducing variability in performance tests on EC2: Setup and Key Results. <https://engineering.mongodb.com/post/reducing-variability-in-performance-tests-on-ec2-setup-and-key-results>
- [6] Henrik Ingo and David Daly. 2020. Automated system performance testing at MongoDB. In *Proceedings of the workshop on Testing Database Systems (DBTest '20)*. Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3395032.3395323>

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 '21 Companion, April 19–23, 2021, Virtual Event, France

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8331-8/21/04.

<https://doi.org/10.1145/3447545.3451199>