

University of Groningen

The Evolving Landscape of Software Performance Engineering

Kudrjavets, Gunnar; Jeff Thomas; Nachiappan Nagappan

Published in:
EASE 2022

DOI:
[10.1145/3530019](https://doi.org/10.1145/3530019)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2022

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Kudrjavets, G., Jeff Thomas, & Nachiappan Nagappan (2022). The Evolving Landscape of Software Performance Engineering. In *EASE 2022: The International Conference on Evaluation and Assessment in Software Engineering 2022* (pp. 260-261). ACM Press. <https://doi.org/10.1145/3530019>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

The Evolving Landscape of Software Performance Engineering

Gunnar Kudrjavets
University of Groningen
Groningen, Netherlands
g.kudrjavets@rug.nl

Jeff Thomas
Meta Platforms, Inc.
Menlo Park, CA, USA
jeffdthomas@fb.com

Nachiappan Nagappan
Meta Platforms, Inc.
Menlo Park, CA, USA
nnachi@fb.com

ABSTRACT

Satisfactory software performance is essential for the adoption and the success of a product. In organizations that follow traditional software development models (e.g., waterfall), Software Performance Engineering (SPE) involves time-consuming experimental modeling and performance testing outside the actual production environment. Such existing SPE methods, however, are not optimized for environments utilizing Continuous Integration (CI) and Continuous Delivery (CD) that result in high frequency and high volume of code changes. We present a summary of lessons learned and propose improvements to the SPE process in the context of CI/CD. Our findings are based on SPE work on two products conducted over 5 years at a major online services company. We find that (a) SPE has mainly become a post hoc activity based on data from the production environment, (b) successful application of SPE techniques require frequent re-evaluation of priorities, and (c) engineers working on SPE require a broader skill set than one traditionally possessed by engineers working on performance.

CCS CONCEPTS

• **Software and its engineering** → *Software post-development issues*; **Software performance**; *Programming teams*.

KEYWORDS

Performance analysis, software performance engineering, SPE

ACM Reference Format:

Gunnar Kudrjavets, Jeff Thomas, and Nachiappan Nagappan. 2022. The Evolving Landscape of Software Performance Engineering. In *The International Conference on Evaluation and Assessment in Software Engineering 2022 (EASE 2022)*, June 13–15, 2022, Gothenburg, Sweden. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3530019.3534977>

1 INTRODUCTION

Software systems often not meeting performance requirements is a known problem in the industry [6]. The accepted definition of SPE is “a proactive approach that uses quantitative techniques to predict the performance of software early in design to identify viable options and eliminate unsatisfactory ones before implementation begins” [10]. In this paper, we enumerate the lessons learned from our experience while working on tasks related to SPE at a major

online services company. Most of our experience comes from being responsible for the SPE of two products for 5 years.

Context. The organization in question is a large online services company. Software developed by this company includes closed-source mobile applications, open-source software (such as a memory allocator and a key-value store) and significant contributions to the Linux kernel. The first product is an open-source software database storage engine. Classic database metrics such as queries per second and transactions per second are this product’s most relevant performance metrics. The second product is a collection of mobile applications used on various devices running either Android or iOS. Several metrics quantify the performance of mobile applications [4]. The metrics for a second product include application start time [11], memory usage [5], and responsiveness [13].

Past approaches to SPE. In organizations that use a classical waterfall model, the performance work follows a specific model. Engineers design experiments to measure the performance of a component. The experiment measures a component’s performance either in isolation or in conjunction with the rest of the system. Based on the results from these experiments, engineers optimize either design or code until performance metrics meet specific criteria. A feature that meets the performance criteria is either integrated into the main development branch or released to a subset of users. The product teams make design and code changes based on the data provided by the performance team. The rationale for this approach is the assumption that software is modified infrequently once released, and patching is an involved process (e.g., updating Microsoft Windows).

2 INDUSTRY CHALLENGES

Disruption by CI/CD. Environments that use CI/CD present different constraints for performance engineering than those utilizing the traditional development models. With the advent of “trunk based development” [12, p. 339], code is no longer developed in hierarchical child branches. Hundreds to thousands of engineers can commit code to the main branch every day. *Each commit can potentially cause or fix a performance regression.* Though we cannot publish a specific number about the ratio of engineers contributing to the product and performance engineers at a company that we study, we can say that the number of product engineers greatly exceeds the number of performance engineers. Functionality in products is typically enabled or disabled by using *feature flags* (*feature toggles*) [3] depending on variables such as deployment goals, target audience, or state of the feature (e.g., GateKeeper [8]). Given N feature flags, there is a possibility of 2^N combinations of features being active simultaneously. Testing all these combinations before the release is not practical (or even possible, given that N may be in hundreds).

Limitations of dogfooding. Microsoft popularized the concept of “dogfooding” in the software industry [2]. Dogfooding is the

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).
EASE 2022, June 13–15, 2022, Gothenburg, Sweden
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9613-4/22/06.
<https://doi.org/10.1145/3530019.3534977>

practice of engineers using and testing the software they develop prior to the release to users. While this technique has been, in our experience, a valuable tool, we observe that engineers form a biased test sample to detect performance regressions. Engineers (a) tend to have better hardware than an average user, (b) have different usage patterns, and (c) use the software using outlier configurations settings (e.g., specific feature flags, internal tracing mode, debug builds). Dogfooding enables the early identification of obvious performance issues but does not serve as a good predictor for how an average user will behave.

Increased required skill set. Modern software systems are complicated and have many dependencies. Detecting, debugging, and fixing performance-related issues in a product requires an ever-expanding skill set and mastery of different tools [1]. In our experience, a wide array of knowledge spanning from the internals of operating systems development, code generation by compilers, internals of a particular programming language to product code itself is required. We observe that debugging and profiling kernel mode and user mode code is necessary to solve complex performance issues. Assuming that each product team has engineers with that skill set is unrealistic. Therefore, more requirements are placed on engineers working on SPE.

3 METHOD

Given the relatively small number of engineers working on SPE for products we research, we elect not to perform formal personal opinion surveys. However, we conducted annual “performance summits” (informal focus groups with a mediator) with engineers involved in various SPE-related tasks. These events, follow-up discussions between engineers, and internal technical white-papers documenting practical SPE challenges serve as a basis of this paper.

4 FINDINGS AND RESULTS

An existing study summarizing 9 years of SPE experience at Fidelity states that their SPE team was forced to evolve “into a rag tag team of system integrators” who were, amongst other things, “capable of profiling, instrumenting and changing application code” [9]. Our experiences with SPE are similar to the following findings:

- (1) A large part of SPE is conducted based on data collected from the production environment (e.g., various counters, logs, snapshots). Simulating the production environment in-house is a challenging research problem. For example, to correctly evaluate mobile software performance and limit the presence of confounding variables, a Faraday cage may be necessary.
- (2) A software performance engineer in a CI/CD environment works in a constantly evolving environment given the large volume of code changes due to commits in application code base and related dependencies.
- (3) The skill set required to be a successful software performance engineer has changed. Software performance engineers design experiments, measure their results and present the data to the product teams. They are also generalists who can work with a variety of teams and have the ability to debug and profile code at any abstraction level.
- (4) In our experience, it is valuable to have a small number of dedicated performance engineers working alongside a larger team of product engineers.
- (5) Classic performance metrics include battery (power consumption), CPU, I/O, and memory [1]. Products may demand a different focus and can value one metric over another. For example, a long-running process executing on a server in a data center is not concerned as much about the initial process start time as an application executing on a mobile device.

5 CONCLUSION

The role of engineers working in the field of SPE for software used daily by millions of people has changed. The focus has shifted from traditional, primarily theoretical scientific methods (e.g., experimental design, modeling, projections) to a more practical hands-on approach (e.g., debugging, profiling, fixing performance regressions).

Modern software development (e.g., CI/CD) uses “release early, release often” type of methodology originating from the early years of the open-source software movement [7]. The nature of SPE has adapted to accommodate that approach. Contemporary SPE involves working closely with production data and debugging and profiling user and kernel model applications.

REFERENCES

- [1] Brendan Gregg. 2020. *Systems Performance: Enterprise and the Cloud* (2nd ed.). Addison-Wesley, Boston, MA, USA.
- [2] Warren Harrison. 2006. Eating Your Own Dog Food. *IEEE Software* 23, 3 (2006), 5–7. <https://doi.org/10.1109/ms.2006.72>
- [3] Pete Hodgson. 2017. *Feature Toggles (aka Feature Flags)*. Thoughtworks. Retrieved February 16, 2022 from <https://martinfowler.com/articles/feature-toggles.html>
- [4] Max Hort, Maria Kechagia, Federica Sarro, and Mark Harman. 2021. A Survey of Performance Optimization for Mobile Applications. *IEEE Transactions on Software Engineering* (2021), 1–1. <https://doi.org/10.1109/TSE.2021.3071193>
- [5] Jaehwan Lee and Sangoh Park. 2021. Mobile Memory Management System Based on User’s Application Usage Patterns. *Computers, Materials & Continua* 68, 3 (2021), 4031–4050. <https://doi.org/10.32604/cmc.2021.017872>
- [6] Daniel A. Menascé. 2002. Software, Performance, or Engineering?. In *Proceedings of the 3rd International Workshop on Software and Performance* (Rome, Italy) (WOSP ’02). Association for Computing Machinery, New York, NY, USA, 239–242. <https://doi.org/10.1145/584369.584407>
- [7] Eric S. Raymond. 1999. *The Cathedral & the Bazaar: Musings on Linux And Open Source by an Accidental Revolutionary* (1st ed.). O’Reilly, Sebastopol, CA, USA.
- [8] Chuck Rossi. 2017. *Rapid Release at Massive Scale*. Meta Platforms, Inc. Retrieved February 26, 2022 from <https://engineering.fb.com/2017/08/31/web/rapid-release-at-massive-scale/>
- [9] Jayshankar Sankarasetty, Kevin Mobley, Libby Foster, Tad Hammer, and Terri Calderone. 2007. Software Performance in the Real World: Personal Lessons From the Performance Trauma Team. In *Proceedings of the 6th International Workshop on Software and Performance* (Buenos Aires, Argentina) (WOSP ’07). Association for Computing Machinery, New York, NY, USA, 201–208. <https://doi.org/10.1145/1216993.1217027>
- [10] Connie U. Smith. 2015. Software Performance Engineering Then and Now: A Position Paper. In *Proceedings of the 2015 Workshop on Challenges in Performance Methods for Software Development* (Austin, Texas, USA) (WOSP ’15). Association for Computing Machinery, New York, NY, USA, 1–3. <https://doi.org/10.1145/2693561.2693567>
- [11] Natansh Verma. 2015. *Optimizing Facebook for iOS Start Time*. Meta Platforms, Inc. Retrieved January 25, 2022 from <https://engineering.fb.com/2015/11/20/ios/optimizing-facebook-for-ios-start-time/>
- [12] Titus Winters, Tom Manshreck, and Hyrum Wright. 2020. *Software Engineering at Google: Lessons Learned from Programming Over Time* (1st ed.). O’Reilly, Sebastopol, CA, USA.
- [13] Shengqian Yang, Dacong Yan, and Atanas Rountev. 2013. Testing for Poor Responsiveness in Android Applications. In *2013 1st International Workshop on the Engineering of Mobile-Enabled Systems (MOBS)*. IEEE, San Francisco, CA, USA, 1–6. <https://doi.org/10.1109/MOBS.2013.6614215>