Twitter Position Statement
JCP Elections 2021

Why Twitter?

Twitter relies on the JVM as the language runtime for its services, with almost all of our services being written in Scala or Java. We are solving uniquely challenging problems when it comes to designing software at our scale and are developing libraries for scalable I/O, distributed scatter-gather, and system monitoring and management. We want to ensure that our systems will continue to cope with current traffic levels (for example, hundreds of millions of tweets sent per day) and will be able to scale to higher traffic levels in the future. We also want to ensure that our services run efficiently at data center scale, maintaining high quality of service, while minimizing data center footprint.

We hope to bring our insight in high traffic, low latency, large deployment, multi-language systems when evaluating and giving feedback on new JSRs in order to expand and enhance the Java platform in the most constructive and useful way.

Twitter and OpenSource

Since its inception, Twitter has had a very strong record of community involvement, with most of our internally developed software (Finagle, our asynchronous RPC framework, Heron, Scalding, etc.) being open sourced, and with additional contributions to many other open-source projects.

Twitter gets enormous value from OpenJDK being an open-source platform. We have a dedicated team, the #TwitterVMTeam, that is responsible for building, deploying, maintaining, and extending OpenJDK with Twitter-specific changes to optimize it for our workloads and use-cases. We have been an early Graal adopter and are actively working to improve its performance on our Scala code. Many of these enhancements and improvements are generally applicable, and we have contributed back such changes to the relevant upstream projects (OpenJDK, Graal, etc.).

The main goal of Twitter's JCP involvement is to ensure that OpenJDK remains a viable, high performance, and fully-featured open source platform.

Twitter’s JCP Representatives

Tony Printezis is a Staff Software Engineer at Twitter and a member of the #TwitterVMTeam in the Infrastructure Organization. He has over 20 years of (mainly Java) virtual machine and language runtime implementation experience, with special focus on memory management. Most of his projects have involved improving the performance, scalability, responsiveness, parallelism, concurrency, monitoring, and visualization of garbage collectors. He was one of the designers (and tech lead) of the Garbage-First GC and the original implementer of the widely-used CMS GC. Before Twitter, Tony worked at Adobe, Oracle, and Sun Microsystems. He holds a PhD and a BSc(Hons) in Computing Science, both from the University of Glasgow in Scotland.

Flavio Brasil is a Staff Software Engineer at Twitter and a member of the Strato team in the Core Services organization. He has worked with several technologies during his 14 years of experience in software engineering. His diverse background, which includes multiple programming languages and runtimes, distributed systems, language and library design, big data, and others, brings a broad perspective to the projects he works on. Flavio is an open source contributor and has created groundbreaking projects like Quill (http://getquill.io), one of the most advanced LINQ solutions available for a programming language. Over the past 9+ years, he has specialized in performance on the JVM and is has worked on Graal JIT optimizations tailored to the specific needs of Scala codebases. He holds a BSc in Computing Science from the University of Sao Carlos in Brazil.