Twitter Position Statement
JCP Elections 2016

Why Twitter?

Twitter relies on the JVM as the base runtime for its services, with almost all of our new software being written in Scala or Java. We are solving some 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 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

Twitter has a very strong record of community involvement since its inception, with most of our internally developed software being open sourced and additional contributions to many other open source projects. We get enormous value from OpenJDK being an open-source platform. We are able to build and deploy OpenJDK with Twitter-specific changes to optimize it for our workloads and use cases. Many of these enhancements and improvements are generally applicable to most workloads and we contribute back such changes to OpenJDK. One of the goals of our 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 VM Team in the Infrastructure Organization. He has over 15 years of (mainly Java) virtual machine 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 is a committer and reviewer for the HotSpot group at the OpenJDK project. He was one of the designers of the Garbage-First GC and the original implementer of the CMS GC. Before Twitter, Tony worked at Adobe, Oracle, Sun Microsystems, and Sun Microsystems Laboratories. He holds a PhD and a BSc(Hons) in Computing Science, both from the University of Glasgow in Scotland.

Ramki Ramakrishna is a Staff Software Engineer in the Platform Infrastructure Group at Twitter. He's also currently serving on the Twitter Architecture Group. Ramki has worked with several generations of JVMs, including at Sun and Oracle, before Twitter. He is a committer and reviewer for the HotSpot group at the OpenJDK project. His principal contributions have been in the areas of performance analysis, tuning and adaptive optimization, parallel and concurrent garbage collection, and synchronization infrastructure within the JVM. Before coming to industry, Ramki worked at SUNY Stony Brook, the Tata Institute of Fundamental Research in India, and Aalborg University in Denmark, dividing time between teaching and research into formal verification of concurrent systems, based primarily on process algebras, temporal logics and automatic theorem-proving. Ramki obtained a Ph.D. in Electrical and Computer Engineering from the University of California at Santa Barbara and a B.Tech. in Electrical Engineering from IIT Kanpur in India.