

Modern Java on FreeBSD

An interview with Harald Eilertsen

Tom Jones

In September 2025, Harald Eilertsen gave a talk at EuroBSDcon titled “Improving OpenJDK on FreeBSD – Where we are, and what I’ve learned”. Harald has been working with the upstream OpenJDK project on improving the support for FreeBSD. Sitting with Harald is Tom Jones. Tom is a committer, FreeBSD Journal editorial board member, and regular contributor to the journal.

Tom: How did you get started with FreeBSD?

Harald: I’ve been using FreeBSD for ages, more or less. I got started trying it out on the desktop system in the early 2000s. I don’t remember exactly but for me at least it didn’t work out too well. There were a lot of problems at the time running it on a desktop.

I’ve been using it on the server for a long time. I essentially have a jail on a friend’s server so I kind of have my own virtual little space where I host various web things and my own social network, my own fediverse instance and a bit... yeah a lot of different blogs, my bands, my different stuff, also for other people sometimes.

Tom: What was the process for you going from being a FreeBSD user to doing development?

Harald: I've been a developer for my whole life essentially, but the reason I got into doing this project with OpenJDK for FreeBSD was essentially that the FreeBSD foundation was looking for someone to do it.

Tom: Where you doing Java development before?

Harald: I have done some Java development, but that was decades ago. For me the JVM is more interesting than Java the language itself.

It's more C++ and a bit of Assembler than Java itself. It is the system that is powering Java. I see that there's been a lot of development on the language side of Java too. So it has become a much more modern language over the past decade than it was in the beginning.

Tom: Yeah, I've not touched Java for a very long time and I'm very happy about it.

Harald: Most people are.

Tom: Do you have a sense from doing this work the sorts of things people want to run in OpenJDK on FreeBSD?

Harald: Yeah, some of it and that's in part because I know some people who are using Java on FreeBSD in production.

But also because what I see coming in from bug reports and other things, it seems like it's tied to essentially big critical systems.

It's like Elasticsearch. It's like Tomcat and kind of those mostly server side things is what I get the impression of right now, at least.

A company I used to work for previously was running Java, powering 'card not present' payment authentication. Authenticating the cards, not the payment. That was running on Java. So it seems like it's still used in a lot of

critical systems and big important systems that's mostly hidden from mostly normal people in a way.

Tom: What was the state of things when you got started?

Harald: There was a port for JDK 23. This was just before JDK 24 was about to be released by Upstream.

I spent some time getting to know the code, getting to know how it works, getting to know how to run the unit tests, and stuff like that. One of the issues I had was that a lot of unit tests... not a lot, but there were several quite essential unit tests that didn't pass.

There were many different reasons for that. So a lot of the stuff I've been doing so far has been both fixing the unit tests or the test framework so that it works on FreeBSD, but also fixing the code that the unit test or the tests flagged as not working.

I think the state now is that JDK 24, and the coming 25 when I'm able to get it out, is that it's a much more solid Java or JDK implementation that works better and has better performance on FreeBSD than the previous ones have.

Tom: Do you have a sense of what the performance improvement is and where?

Harald: Part of it has been because I've been notified by users that, oh, we have a performance issue in this or that.

There was a pretty big Elasticsearch installation that had performance issues on systems with more than 128 gigs of RAM. Looking at the code, the part was Elasticsearch wants to know the amount of committed memory and how we account for that and the way we were doing that was less than optimal.

It was a fairly trivial change once you first find it.

FreeBSD essentially has an API that gives you the number right away so we don't have to wade through files and do lots of stuff to calculate it. So once

that was in place, it was confirmed by the user that, yeah, this works. So there's issues like that.

Tom: There's no, percentage improvement you just thought off the top of your head?

Harald: No, and I think that will depend heavily on the workload you're having too. One of the things I'm doing and also that I think [bapt@](#) has been helping with is to reduce our reliance on the `procfs` file system.

When the code as it is from Linux relies heavily on the `procfs` system and that seems to be a lot faster on Linux than on FreeBSD.

We usually have better and easier ways to get the information on FreeBSD. So getting rid of it entirely, it will be an improvement.

I haven't done any performance testing as such and especially and specifically this issue I was talking about is I don't have the machine to test that kind of stuff. I don't have a machine with that kind of memory. And also, I guess, setting up a heavily loaded Elasticsearch cluster or installation takes some effort and takes some time to kind of measure.

Tom: It's certainly my experience that one part is: access to big machines is hard. It'd be great to be able to connect to more people in the community who could say, I show the machine on our test workload, here's how you do everything.

Harald: It's the same with this Elasticsearch thing. I got pretty good contact with the people who were reporting it and they they were absolutely willing to help and in this case I essentially just said "this looks suboptimal let's try a different approach." And I sent them the patch and they essentially just confirmed that yes this solves the problem so that's the way you've got to test it. In this case they did the testing. I tried to figure out where in the code to do the change.

Tom: For me when I ported VPP to FreeBSD there was a lot of stuff in `procfs` where we're just reading strings from a file to do configuration and the changes for that were very straightforward most of the time it was either use a `sysctl` interface or we had a native interface. It feels like we should have a guide for porting software which is like, grep for this pattern and then just replace that.

Harald: I think yeah that's pretty much the same solution I went for. It was a `sysctl` or this had something that just gave me the result directly, so that's much much faster. I think there's one of the problems I have at least with the OpenJDK is that traditionally, especially in the hotspot code in OpenJDK, the BSD part of the source tree includes Mac OS because they traditionally see Mac OS as BSD.

Although the code that we need is usually a lot closer to, but not necessarily identical to the Linux code which is in a different subdirectory of the tree so it gets a bit confusing sometimes because sometimes you have this BSD-specific code that's essentially split into one or two or sometimes more because of the different BSD's also have different ways of doing things but then usually the Mac OS part is like just completely different from all of it.

Tom: How's upstream been for accepting patches?

Harald: They are very nice. They are very strict. So there's a really well-defined workflow and process for getting patches upstream. One of the things that they don't want to have is an incomplete patch set for a platform that's almost supported.

They are very supportive in getting the BSD's into the upstream but there is a defined process for doing that.

You have to follow their process of doing it so what we're doing now is we got our own repository under the OpenJDK organization on github where we will create some branches that we will develop and keep up to date with upstream. Once we are in a position where this works, it passes the tests, it has everything and we also need to write a Java Enhancement Proposal to

say that we want to make FreeBSD or the BSD's in general or however we go about it as supported platforms in upstream.

That has to be accepted and then essentially all of it will be merged at once, once we get there but I have had some patches that are not strictly BSD related but that kind of helps us in the process being accepted upstream and they are like, "yeah". They've been very nice about it and they're very supportive.

It's this old Java slogan: "Write once, run everywhere." It seems like it actually means something to them.

Tom: Do you see the end on the horizon? Are we getting close to doing the big merge where we add FreeBSD support?

Harald: I think we are starting on that work now. I don't know how long, but anything from six months to another year I think before we get it. It may go faster because the port works and we pretty much have all the changes.

There's still a few tests I'm not able to get to pass yet, but I think that once we have that in place, and also have working CI for FreeBSD that fits into their system. That's what has been worked on lately and it looks like it's happening so I don't know really.

I'm a programmer, I'm not able to plan time.

Tom: Once we're a supportive platform does that mean that the burden is on other people to not break support?

Harald: Yeah because then we will be a part of the CI process and if the CI process fails for a patch it will not be merged. That doesn't mean that they will necessarily have to fix it. It may be that yeah, they contact me or somebody else on the BSD porters group cool.

Tom: Thank you very much for your time today is there anything else you want to talk about before we go.

Harald: I've just come back from traveling and EuroBSDCon and meeting lots of people everywhere so my head is a bit blank. It was my first EuroBSDCon this year and I think it was a good experience. It was nice meeting people in real life and there was lots of interesting talks.

I think it also, for me at least, coming kind of from the outside, I'm not part of the BSD community. It's very nice for me to see all the stuff going on and what people are working on and yeah all the how to say it all the all the craft going into building and building a system

Tom: Did you get lots of feature requests from people when they learned you worked on OpenJDK?

Harald: Not so much it seems like it's more I got a lot of good support that's like: "it's really nice that somebody's working on this and it's really nice that we are working on getting OpenJDK to a better state so that it works better."

I think it's a really good thing and I'm also, of course, getting it upstream so that it's supported.

I think that's more the most important part. At least from my impression from what people are saying.

Tom: Did you tell people to thank the FreeBSD Foundation?

Harald: Yeah, I did say that. And it's kind of the point. It's like: "Why are you doing this work?"

It's like, "yeah, I get paid to do it."

But it's also a project that I find interesting and I find it challenging. I've been working as a software developer in various parts. Different systems for almost 30 years now. And I have a lot of experience.

It's not that easy to come across work that is both interesting and aligns more or less with my beliefs in free software and open source. And so to be able to combine that and get a challenging job that I feel is something that — Not everybody can do.

And it's not just like — there's a lot of stuff. At least when you're looking at what's available work-wise. It's a lot of web development and JavaScript and whatever is the latest framework. And it's like, I have no interest for that.

Tom: It's a lot of — it seems like a lot of churn for churn's sake. Make something new and exciting.

Harald: So that's what I feel.

This is a gig that kind of allows me to use some of my knowledge but also challenges me. I have a lot of things I need to learn, both about FreeBSD and about OpenJDK and how everything works. I notice on both of those sides, there are really knowledgeable people. And I can learn from them.

And that's — I find that very valuable.

Harald Eilertsen is software developer and information security professional. He has been programming computers since the early 1980's, and has worked for companies like Tandberg, Cisco, Modirum and Automattic. He is currently self employed and working as an independent consultant.

Tom Jones is a FreeBSD committer interested in keeping the network stack fast.

Copyright © 2026 Held by Owner/Author. Publication Rights Licensed to The FreeBSD Journal. This work is licensed under a [**Creative Commons Attribution International 4.0 License**](#).

3980 Broadway Street

STE #103-107

Boulder CO 80304

1-720-207-5142

info@freebsd.foundation.org

