

The Open Source Way

Episode 22: Open Source at VMware – Key Driver of Innovation



Transcript

Karsten: Welcome to the Open Source Way. This is our podcast series, SAP's podcast series about the difference that open source can be. And in each episode, we'll talk with experts about open source and why they do it the open source way. I'm your host, Karsten Hohage, and in this episode I'm going to talk to Dawn Foster of VMware about many things and possibly some more. Hi, Dawn. Nice to have you here.

Dawn: Thank you so much for having me on the podcast.

Karsten: Great to have you here. Dawn Foster is the head of the Open Source Community Strategy at VMware. She started there, actually, she got started there by an acquisition VMware made in 2020. But as Dawn has had quite an interesting path through the IT industry over the years, we'll just let her tell that story herself in a minute. Until then, Dawn, I can see you as opposed to our listeners, and you have this slightly or definitely purplish hair. Would you happen to know my SAP colleague, Josh, who's also like in the same range of community and things, who's wearing a blue beard?

Dawn: I don't think I know Josh, but I will definitely look for him at upcoming conferences.

Karsten: I think you should somewhere be in a video conference together, which would be greatly colorful in that case. Now, you can't know them all, but as I said, you do have quite a history in the industry. Want to sketch out your path so far for us?

Dawn: Yeah, sure. I have been in the industry for well over 20 years. I came into the path kind of from a traditional background. I came out of university with a computer science degree back in the mid-nineties and I managed to luck into a Unix system administration job right out of college at a manufacturing company in the US Midwest, where I used a lot of open source tools, primarily because manufacturing companies don't like to pay for IT. So, that was really my only choice. And I learned more about open source over the years, but I also got an MBA a few years later, just before I joined Intel in 2000, and since previously I'd used open source, and Linux is a little like Unix, they asked me to start looking at open source and Linux developer tools that were likely to be strategic for Intel

over the next few years as we enabled various different types of processors. I later left Intel in about 2006 and worked out a few startups in various director of community and developer relations roles focused again on open source software. And I did a bit of consulting around building communities for a few years after that. In 2000, I actually went back to Intel where I led the community team within the Open Source Technology Center, where I spent a few years working mostly on open source mobile operating systems like MeeGo and Tizen, if anyone remembers those, and after I left Intel in 2012, I spent a couple of years as director of Community for Puppet, which took me back to my system administration roots, which was a lot of fun. And then in 2015, for my mid-life crisis, I decided to quit my job at Puppet, sell my house, sell my car in Portland, and I moved to London to get a PhD and I spent the next three and a half years researching how people collaborate within the Linux kernel, which did end indeed in a PhD at the end of that. And then I decided to stay in the UK. I really like it here and I got a job at Pivotal driving our upstream Kubernetes contributor strategy and when we were acquired by VMware, they already have people working on their Kubernetes strategy. Some of the people like Craig McLuckie and Joe Beda, who actually started Kubernetes. So, I knew I wasn't going to get to do that, but I ended up in a really nice role in the Open Source Program Office running the community strategy team.

Karsten: Wow. That's quite an all IT industry CV, I should say. Let me pick out two points out of that. The one is you like it in the UK? Actually, this summer we plan to check out how the post-Brexit UK is doing. So, we're planning to go to southern England on vacation. Haven't been there in a while, so now I found it interesting that you so pointedly say I like it in the UK. Great. Makes me hopeful. The other one is you mentioned the open source mobile operating systems, the back in the days operating systems for mobile phones there. My office mate, Sebastian Wolf, who has been on this podcast a couple of times, some people may remember, he still fiddles around with contributions to Jolla, and that is an open source mobile operating system left over from back in these days; has roots in for all I know, Nokia and MeeGo, and there must have been some historic and game-changing announcement in like 2010 or 2011 in that entire history. Can you tell us about that and how it influenced your work or how the day of that announcement was and everything? I just have a nebulous idea of it.

Dawn: Yeah. Those were interesting times. MeeGo was an incredibly fun project to work on and I worked with loads of folks from Nokia, so it was a Linux Foundation project, but

it was mostly people from Intel and Nokia who were working on the project. And then Nokia got a new CEO and they told us that they were going to make a big announcement that had something to do with MeeGo, but they wouldn't tell us what it was. I think maybe they told our CEO, but nobody else was allowed to know what this announcement was going to be. And they did it at a really convenient time for those of us in Europe, which was something like two or 3:00 in the morning for me in Portland. And previous to this announcement, our PR folks had put together this huge like flowchart document, which was if Nokia announces this, this is what Intel is going to do next; If Nokia announces this, this is what Intel is going to do next. So, I got up in the middle of the night, made some coffee or tea, and sat there with a webinar that Nokia was doing in one window, which was how we were going to find out what they were going to do. It was kind of like an earnings announcement or something, I think. And then in the other window I had the IRC channel open for the MeeGo chats and so I just sat there and watched the Nokia announcement where they basically said that they were not going to do any open source mobile operating systems and that they were going to work entirely with Microsoft on Microsoft based mobile open source operating systems, which effectively killed the MeeGo project, which was frankly pretty sad. But yeah, I sat there on IRC for a few hours in the middle of the night and just talked to people from the community as we processed what was honestly sort of fairly traumatic information for those of us within the community.

Karsten: But that was at the same time then the beginning of the fully independent open source mobile systems like Jolla for instance, or...?

Dawn: Yeah. So those weren't things that Intel worked on. But what happened was a bunch of people who were pretty enthusiastic about MIMO, which was the predecessor to MeeGo, that was a Nokia project. And so, there are people that had been in this MIMO community and then the MeeGo community, and then they effectively took a lot of this open source technology and ran with it, which is eventually what became Jolla.

Karsten: Okay. So now that was our Sebastian reference, who we kind of thought could have been on this podcast. But then we decided, no, he always appears. So, let's just talk to Dawn, our guest from VMware here. That sounds like exciting times back then, but let's maybe get a little more to the nearer past or even the present or even the future. Since back in those Intel days, you have been all about open source and mostly about community, right?

Dawn: Yeah, exactly. Mostly open source. I, of course, dabble in various things because that's what I do.

Karsten: Okay. Then in all these companies that you have been at, what would you say is similar about how all of them deal with Open source?

Dawn: Yeah. Most of the companies that I've worked for took a really strategic approach to open source and had a deep understanding and appreciation for open source despite being very different types of companies, right. At VMware, we really see open source as a key driver of our innovation strategy. And by investing in and then contributing to open source projects, we get better productivity, we get faster time to market, and we end up at the end with products that are innovative, interoperable, scalable, and mature, and we see open source as strategic for our future, VMware. When I worked at Intel, it was important that key open source technologies like the Linux kernel, for example, ran really well on Intel hardware, right? The hardware isn't any good if it doesn't have software that runs really well on top of it. So, we needed to provide support and optimizations beyond just providing drivers, which is, I think what people tend to think of so that our customers would have the best experience possible when running open source technologies on top of Intel hardware. And again, these contributions to the kernel and a loads of other open source projects at the same time were really strategic for Intel.

Karsten: This is historic in your life, basically having been at Intel and so on. And you said it is or it has always been strategic. Has it really always been that way? Because with most others - at SAP, Microsoft - who we talk to, there is always the acknowledgement, yeah, it changed in this century over the past two decades, the perception of open source. Has it been like better from the start at Intel and VMware?

Dawn: No, I would say that certainly when I first started at Intel in 2000, those of us that were working on open source, were kind of an anomaly and there was a lot of distrust because frankly, at the time, you know, people talked about Wintel, right? So, Intel had this incredibly strategic partnership with Microsoft who wasn't particularly happy about the work that we were doing in Linux and open source software. So, at that point and a lot of companies back in those days, right, the early 2000's open source really wasn't strategic for a lot of companies. It became strategic over time, but certainly there was, yeah, open

source was not particularly strategic within Intel. And I know that VMware went through kind of a similar thing. They hired Dirk Hohndel to be our chief open source officer. He's since left VMware. But when they hired him, they were doing very little in open source and really wanted someone to come in and kind of kick start the whole program. And so that was, I don't know, not that many years ago, maybe six or seven years ago. So, I think companies do go through kind of a transition. But I would say most of the companies that I've worked in, at least in the last 10 to 15 years, have considered open source to be something that was strategic. But you're right, it was not always like that for sure.

Karsten: Just making sure, it kind of a little bit sounded like it has always been like that before. And I was just wondering if we'd now hit the two or three different examples while everyone else confesses to, okay, we know we were doing open source these days, but it hasn't always been like that. But great. So, same as everybody else really.

Dawn: Yeah.

Karsten: Now, you said most of the companies you've encountered treat it as a strategic topic. Is it also sometimes - that's always a picture in my mind, but I don't know because it's not my space - is open source sometimes used to basically create these neutral zones where one rather cooperates and competes? Sort of like, I don't know, a pacifistic defensive space between the different realms, maybe.

Dawn: Yeah, absolutely. And I think that really highlights the role of neutral foundations in open source projects. And by putting projects into these neutral foundations. So, you can think about this like the CNCF, the Linux Foundation, Eclipse, the Apache Software Foundation, and a bunch of different companies can work within these foundations together as equals to collaborate and innovate on projects that benefit the entire software ecosystem. One of the best examples of this is Kubernetes. By putting Kubernetes into the Cloud Native Computing Foundation, we can all work together, right, to produce a critical bit of infrastructure that everyone can use as we build our products, applications, and services on top of Kubernetes. And as companies who contribute, we don't need to worry about one company exerting too much control over the project. And our customers can use Kubernetes-based technologies knowing that they aren't going to be locked into a specific solution from a single company.

Karsten: Yep. Okay. Fits the picture. And your role in all this has always mostly been governance, processes, community. Have you always been a process - I mean, you said you started as a UNIX administrator - so you probably haven't always been a people process person? Are you more of a code person originally?

Dawn: I would say that I am more of a community person than a process person. And I also still do some development. So, I'm also, I guess, a code person as well at the same time. But the reason that community work looks a lot like process is because that's how you scale communities and that's how you keep them healthy. Having processes defined early helps avoid misunderstandings later. So, things like documented governance, communication, security processes and loads of other processes help set expectations and help people understand how they can best participate in your community. And these documented governance structures show people what they need to do in order to eventually move into leadership positions if that's what they want to do.

Karsten: Okay. And one of the projects, as far as I know, itself open source, that tries to measure the things you just mentioned, like quality of governance in open source projects is the CHAOSS Project, right? You want to tell us about that a bit?

Dawn: Yeah, absolutely. So, I mentioned that processes help you have a healthy community, but you need to have metrics if you want to measure project health and help your open source projects improve over time. So, this is where I become a code person. So, I spend a quite a bit of time actually writing Python scripts that use the GitHub API and tools like Augur from the CHAOSS Project to learn more about how our open source projects at VMware, how healthy they are, and to help our maintainers improve the health of their projects as well.

Karsten: Okay. Now, I'm just noticing there's one part here that I don't quite get. You said you spend time writing Python scripts, and the CHAOSS project, is that does it consist of these Python scripts or is other Python scripts or other content of a larger framework that CHAOSS is representing?

Dawn: Yeah. So, I can explain this. So, the way this works is that. So, I'm using a tool called Augur right now, which is a CHAOSS project health tool. And just to be honest, the Augur tool there, their user interface is slightly immature. But I also one of the things that

I found is that with metrics dashboards, especially big metrics dashboards that have all sorts of things on them, people get confused, and they don't know what to focus on. So, you set them down in front of this big dashboard and tell them to measure their project health and they get overwhelmed and go away and never look at it again. So, what I wanted to do at VMware was I'm using Augur, which pulls all of this data out of the GitHub API and it stores it all into a really nicely formatted PostgreSQL database. So, my Python scripts query the post graphs database for a few things that I really care about, and I generate four charts that I give to people. And it's like, this is a snapshot of the very basics of your project health. So those are four charts that any maintainer across VMware can truly understand without any experience really in community and understanding metrics. And it's an easy way for them to get a snapshot. And then what I recommend for projects, some of our bigger projects that we care a lot about is that they have their own metrics dashboards and a community manager who can monitor those and look at those over time. I think that's important. But yeah, a lot of what I do is I write scripts that query the Augur database, but I also do a lot with the GitHub API because we ... there are things unrelated to Project Health that I don't necessarily want to put in Augur. One of them - and this is something that those of us that work in an open source program office see a lot of - is that it is super easy to create new orgs on GitHub. And so, what happens is we frequently find organizations that VMware employees have created within their VMware jobs on GitHub that are what I will just affectionately call rogue organizations that we need to figure out who owns them, what they do. And so, I also have some other scripts that hit the GitHub API that I just write natively against the GitHub API that don't use any of the CHAOSS tools that do things like figure out who might be responsible for this org and what the repos are like, when were things last updated. And it's just some basic data about this. And I've got this in a project on GitHub.

Karsten: On a side note, there, I find rogue one of the very interesting words of the English language, because it has so many connotations depending on where and how you use it. Anyway, CHAOSS, I've encountered some other projects that measure things like liveliness, trustworthiness, health of a community project, or at least so they claim. How does CHAOSS relate to all of them?

Dawn: Yeah. So CHAOSS, the CHAOSS project has quite a few different elements. So, I think quite a few people think of CHAOSS in terms of the tools which are primarily GrimoireLab and Augur, and people use those on a day to day basis to measure project

health. But a lot of people don't realize that a lot of the sort of daily work within the CHAOSS project focuses on metrics definitions to provide a common way to define metrics that can be implemented in a variety of different tools. So, we define metrics for diversity, equity, and inclusion, how projects evolve over time, measuring different types of risk, and measuring the value of a project. And there are certainly other metrics approaches outside of chaos that people can use. So, for example, I do a lot of work with the CNCF and all of those projects have dev stats dashboards. And as part of my work within the CNCF Contributor Strategy Technical Advisory Group, I wrote a guide to help CNCF projects, use dev stats to measure project health. And part of that was again because of this issue that people do get overwhelmed with these metrics. So, this project health guide has just a handful of things that I think people should look at to start and where to find those in dev stats because again, it can be super overwhelming because dev stats has hundreds of things that you can measure and it's hard for people to kind of figure out where to start.

Karsten: Then one other thing. There are some projects that are truly open source. Some have this kind of like veto right proprietary software vendor as an essential part in them. But CHAOSS is not like a VMware project, right. Or what's your particular involvement with it?

Dawn: Yeah, exactly. So CHAOSS is an open source project. It's under the Linux Foundation, so it's a Linux Foundation project. I'm on the governing board and I'm a maintainer for CHAOSS and I just sort of have a passion for data and metrics. So CHAOSS is something I've been involved in since before I started at Pivotal or VMware. So, for more than ten years, I've actually been using the tools that later evolved to become GrimoireLab, which is one of the CHAOSS tools. But we also, like I mentioned earlier, we use some of this work within VMware in particularly Augur to measure project health metrics. So, my role in CHAOSS is also then related back to my work at VMware. But you know, and I'll continue to contribute to CHAOSS regardless of where I work in the future.

Karsten: So, from the perspective of VMware, it's not like you're only doing this on the side, like in your spare time, but you do have an assignment to like be the liaison to the CHAOSS Project.

Dawn: Yeah. So, they consider my participation in CHAOSS as part of my role since my team is responsible for project health metrics. So, it's part of my day job, which is nice.

Karsten: Just making sure I hear because Sebastian, whom we've mentioned, on the one hand, he's with our open source program office and, of course, as that joining in many open source activities as part of his job. But then on the other hand his Jolla activities I think are not defined as in SAP's interest, really. So, CHAOSS is of interest for all of us, but not specifically VMware driven. Where could someone go to learn more about, in the first place, maybe the open source activities with any VMware participation.

Dawn: Yeah, absolutely. So, a lot of people don't realize that Spring, RabbitMQ, Salt and Greenplum are VMware open source projects among many, many others, and we're also one of the largest contributors to projects like Kubernetes and a lot of other CNCF projects. And we regularly contribute open source projects to foundations when they start to get traction, and we think that moving them to a foundation would help make the project stronger. So, for example, we recently put Harald into the Linux Foundation, the LF Public Health Foundation. And so, Harald is a contact tracing technology for things like pandemics. And we've contributed several projects to the CNCF over the years, Apache Software Foundation. We've contributed lots of projects to lots of different foundations over the years. And the VMware Open Source Blog and the VMW Open Source Twitter account are great places to learn more about our open source efforts at VMware.

Karsten: And I'm sure we're going to put these links in the metadata or the additional information under this podcast post there.

Dawn: Perfect.

Karsten: All right. Finally, the famous last question. If our listeners are supposed to remember three main points, key takeaways from this episode, what would they be?

Dawn: Yeah. So, it's hard to pick three, but I'll give it a go. First, having well documented processes like governance, for example, is how you scale communities and keep them healthy, which is super important. And second, it's important to think about how you measure project health to help you find areas where you can improve the health of your project and know whether you're improving or not. And third, tools like the ones from the

CHAOSS Project are a great way to start measuring the health of your open source project.

Karsten: All right. Thanks. You said it's hard to pick three. Did you want to add a fourth or a fifth? You still have time for that.

Dawn: Oh, I don't know. What would I pick? I mean, there are so many things I would say if I had to add a fourth, I would say take a strategic approach to your open source project participation. Think about how your participation in open source matches back to your company's strategy because it's a lot easier to justify what you're doing if you can say, we contribute to Kubernetes because it helps these products, or it helps this piece of corporate strategy. So, I think that's another really important element.

Karsten: Yeah. So basically, know what you're doing and why the hell you're doing it at all.

Dawn: Exactly. And measure it. Measurement is important.

Karsten: Yeah. Measure if you're doing what you wanted to be doing in the first place.

Dawn: Exactly.

Karsten: Okay. Thanks, Dawn. For all this. Wow. We could actually probably branch off into one of these subtopics and talk for another half hour, but that is currently not the point. Maybe we'll do another episode in a half a year or so. So, thanks very much for being our guest today, Dawn. It was nice to have you here.

Dawn: Yeah. Thank you so much. This was a lot of fun. I really enjoyed it.

Karsten: All right. Great. And thank you all for listening to the Open Source Way. If you enjoyed this episode, please share it. Don't miss our next episode. It's always published every last Wednesday of the month, you'll find us on openSAP and in most other places where you would usually find your podcasts, Apple Podcasts, Spotify, and the likes. Thanks again for listening and thanks again, Dawn, for being here and talk to you next time. Bye bye.

Dawn: Thanks.