

Valerie Aurora on File Systems and the Ada Initiative

An Interview

RIKKI ENDSLEY AND VALERIE AURORA



Rikki Endsley is the Managing Editor of *login:* and the USENIX Association's Community Manager. rikki@usenix.org



Valerie Aurora (formerly Val Henson) is the Executive Director and co-founder of the Ada Initiative, a nonprofit dedicated to promoting women in open tech/culture. She is an experienced software engineer and was a leading file systems developer, researcher, and consultant for more than a decade. She invented several new file systems concepts, including a widely used power-saving feature in file systems called relative atime, and co-founded the Linux Storage and File Systems Summit. valerie@adainitiative.org

In July, Valerie Aurora received a 2013 O'Reilly Open Source Award for her long-time contributions to the Linux community and for advocating new developments in Linux file systems. Valerie's expertise in the area of file systems dates back to 2002, when she worked as a software engineer on ZFS at Sun Microsystems. She went on to design and implement chunkfs, union mounts, fsck parallelization, and relative access time (relatime). In January 2011, Valerie announced the launch of the Ada Initiative [1], which helps women get into and stay involved in free and open source projects. Valerie also has a long history with the USENIX Association and our conferences. She has been on the committee for USENIX events, including FAST '09 and the 2007 Linux Storage & Filesystem Workshop, which was co-located with the 5th USENIX Conference on File and Storage Technologies (FAST '07). At HotDep '06, Valerie and her co-authors presented a paper on chunkfs [2].

I've followed Valerie's career and her efforts to help open source communities become more inviting to women since I ran across her "HOWTO Encourage Women in Linux" document during my thesis research in 2006. The document played a huge part in my research into how I, as a tech journalist, could help women in tech by covering their projects and careers. Today the howto is still a thorough, practical resource for encouraging women in IT and reads as if it were written recently. On one hand, "HOWTO Encourage Women in Linux" shows Valerie's forethought on the topic; on the other hand, it shows how much more work needs to be done to make IT more inviting to women.

In this interview, Valerie discusses her work with Linux file systems, conferences, and the Ada Initiative.

Rikki: In 2011, you left your Linux kernel developer position at Red Hat and helped launch the Ada Initiative, but you still do consulting work. What kind of consulting work are you doing?

Valerie: Well, if you look at my consulting Web site, I'm obviously not doing Web design. I do short-term contracts mainly in the area of Linux storage and file systems, usually things like debugging silent data corruption, analyzing performance, and prototyping new features. My favorites are finding the root cause of data corruption and debugging race conditions.

Rikki: Are you doing any Linux/UNIX development right now?

Valerie: Not at the moment. My most recent contracts were all performance tuning or fixing data corruption. My last mainstream kernel patch was to fix a kernel configuration error, and before that a bug fix for a file system freeze locking problem.

Rikki: You were a Linux kernel developer. The kernel team doesn't have a reputation for being particularly inviting. What was your experience like working on the Linux kernel and working with other kernel developers?

Valerie Aurora on File Systems and the Ada Initiative

Valerie: The vast majority of Linux kernel developers I worked with were incredibly kind, thoughtful, intelligent people who just wanted to make Linux better. Unfortunately, it only takes a few jerks to make a working environment terrible, especially when they are in leadership positions. In my experience, at least 95% of kernel developers I know wish they didn't have to be humiliated and mocked in order to contribute to Linux. But like them, I felt helpless to change a system that condones and rewards nasty behavior, starting from the top. I hope that the work I do in the Ada Initiative will someday help change the Linux kernel development culture, because there were a lot of things I enjoyed about working in the kernel community and I would love to return.

Rikki: Linux Software Engineer at Intel and Linux kernel contributor Sarah Sharp recently called out Linus Torvalds (and others) for...um...less than professional communication on the kernel list. I assume you followed that whole exchange [3]? What are your thoughts on the incident? Do you agree with Sarah's request for everyone to "Keep it professional on the mailing lists"? And what do you think about how Linus handled the situation?

Valerie: I'm one of hundreds of Linux kernel developers, past and present, who agree with Sarah Sharp's request—she's just the person brave enough to directly call for change from Linus Torvalds and other community leadership. I was a little horrified to see how many top-notch kernel developers spoke up to say that this is one reason why they dropped out of kernel development. So I'm thrilled to hear this will be a topic of discussion at the next Linux Kernel Summit. I hope that other kernel developers will join her in standing up for a working environment without abuse.

I think Linus responded based on the information he has. For example, he's probably not aware of research showing that people's intuition that performance improves after severely criticizing someone is wrong: any improvement in performance is due to random chance, what many people are familiar with as "regression to the mean." It turns out that when you evaluate the effect of criticism vs. praise on performance scientifically, praise is the clear winner [5]. We as computer programmers should use the same scientific logical approach to community management as we do for software development.

Rikki: Let's talk file systems. Are you still working on union mounts?

Valerie: No, I already have one full-time job at the Ada Initiative! David Howells is continuing work on union mounts and doing a great job. Sometimes I wonder if the main use of the union mounts project is to find and fix lurking bugs in the VFS code. Certainly working on union mounts is a great way to understand the design and rules of the VFS.

Rikki: Are you working on any other file system-related projects? What's "next" for file systems?

Valerie: At this point, I just applaud from afar whenever Linux file systems hit another milestone. Getting btrfs stable and ready for production is in my mind the top priority for now. I sympathize with their main showstopper bug right now because when I left ZFS development, we were working on the same problem. With a copy-on-write file system, you have to be sure that there is enough space on disk to write out all the changes in the current transaction, before you free the blocks with the original data. It is hard to predict how much space you'll need to do this, so the default is to overestimate the space. (If you underestimate the space, the file system gets wedged and you'll probably have to reboot.) The problem with overestimation, of course, is that writes fail with ENOSPC (out-of-space error) when there is really plenty of disk space left. That's better than crashing but not good. I don't know what the ZFS solution was, but perhaps it could be shared with the btrfs developers.

Rikki: You've been on several USENIX conference committees, including the 2007 Linux Storage & Filesystem Workshop. More recently, you've helped launch the AdaCamps and Allies workshops under the Ada Initiative umbrella. Tell us about those events. What inspired you to start them, and do you think they've been successful? Is an AdaCon in the works?

Valerie: I organized the first Linux File Systems Workshop in 2006 because it was clear that Linux file systems development was stalled, and I wanted it to get moving again. It worked—the formal ext4 development branch was announced two weeks after the 2006 workshop, and the first btrfs announcement came three months after the 2007 summit. Chris Mason explicitly credited that meeting with inspiring btrfs. This meeting continues being useful today, under the ungainly but accurate name of Linux Storage, File Systems, and Memory Management Summit.

AdaCamp is a medium-sized—about 200 people—unconference. My co-founder Mary Gardiner and I started AdaCamp with the goal of increasing women's commitment to open technology and culture. Women are excited by and want to be part of open source or Wikipedia or what have you, but then get more or less subtle "Get out, you don't belong here" messages from their communities. AdaCamp brings women together to support each other in their enthusiasm and commitment to open tech/culture. I thought it would be hard to tell whether AdaCamp worked, but it's been pretty easy. AdaCampers email us regularly to tell us they landed an open source internship, or started women's edit-a-thons in India, or learned how to solder. Our post-AdaCamp surveys show that 85% of attendees thought that AdaCamp increased their commitment to open tech/culture; we're not sure if the other 15% were just already highly committed.

Valerie Aurora on File Systems and the Ada Initiative

The Allies Workshop was inspired by a workshop taught by Caroline Simard at Grace Hopper Celebration. Her workshop taught men how to support women in tech by role-playing through common scenarios, like having a woman's idea credited to a man in a meeting. We changed it to be short discussions about scenarios in which men could intervene to make women feel more welcome: how to introduce yourself to a woman at a conference, how to respond to harassment on IRC, how to respond to a sexist argument on a Wikipedia talk page. The most interesting success of the Allies Workshop was unexpected: it gave co-workers an opportunity to talk candidly and safely about their experiences with sexism. People who had worked together for years would discover for the first time that their female colleagues were regularly harassed online or propositioned in the lunch room. It gives people a chance to learn and ask questions in a non-judgmental setting.

We're not sure how to do an AdaCon yet but we'd like to have the chance. AdaCon would be a much larger (400?) person conference with more structure and a lower bar to entry. Some of the features of AdaCamp that people love are hard to scale. For example, AdaCamp's open application/invitation-only selection process is an overwhelming amount of work for 200 attendees, but it also creates a safe, welcoming, productive environment that attendees love.

Rikki: The Ada Initiative anti-harassment work seems to be gaining traction, with more than 100 conferences adopting anti-harassment policies in the past few years. Obviously, adopting a policy doesn't "fix" everything. What's next for conferences? What else should conference organizers be doing to make their events inviting and safe for a diverse group of attendees?

Valerie: PyCon US 2013, Open Source Bridge, and AdaCamp are three conferences that show what the next steps are in concrete form. PyCon US did "all the things": anti-harassment policy; outspokenly pro-women community leaders; travel scholarships; inviting women to speak personally; organizing women's events; free booths for women's groups; and explicitly women-friendly spaces, like the Women's Office Hours room and the Ada Initiative Feminist Hacker Lounge. They had 20% women speakers and attendees at a 3000+ person conference.

Open Source Bridge is an explicitly social justice oriented open source conference in Portland, Oregon. The organizers are themselves fairly diverse, and that is reflected in the breadth of topics they cover and the speakers they attract. Like many open source conference organizers, they view their conference as an opportunity to promote social justice and diversity in open source, and spend time and money to accomplish that. This year they improved the accessibility of their conference by adding "travel lanes" with blue tape on the floor, which let people who use mobility devices or have vision impairment move more easily

around the conference and around everyone else, too. I hear estimates that attendees are around 30% women and speakers are around 40%.

AdaCamp has a pretty good record of being diverse in a number of dimensions, relative to most open tech/culture events: age, race, place of origin, first language, etc. Being mostly women, we were actually not that diverse in gender per se, though more so in gender identification and expression.

From my experience, the general principles of attracting a diverse audience are:

1. Have a diverse organizing committee.
2. Ask for, listen to, and implement suggestions ASAP.
3. Communicate early and often how much you appreciate diversity of attendees.

For example, a surprisingly effective way to improve diversity is by having a variety of food that caters to people's dietary requirements. At the last AdaCamp, we had people who were vegan, gluten-free, fructose-intolerant, celiac, and allergic to mushrooms, soy, and lettuce, to name a few. We had to fight with the caterers for weeks to get food that had something edible for everyone, tasted great, and was cheap enough for a nonprofit budget. It wasn't easy to accomplish, but that's kind of the point: the fact that we cared enough to go to the effort to make sure everyone could eat lunch together was a signal to our attendees that we cared about their needs. Think about it—if you have to leave the conference to get lunch or snacks, what kind of slap in the face does that feel like? And once you get used to putting in the time and effort to meet people's food needs, it becomes a habit to do so in other ways.

Rikki: Your work with the Ada Initiative has allowed you to meet a variety of interesting people. Has anyone stood out as being particularly effective or innovative when it comes to encouraging women in technology? Which organizations or events do you think are standing out when it comes to encouraging women in tech?

Valerie: Wow, hard question. How do I mention just a few? I apologize to everyone I left out in this answer—I assume *:login:* can't publish a novella!

The Outreach Program for Women [4] (formerly GNOME Outreach Program for Women) has been a stellar success for training and recruiting women developers in open source. OPW is the product of many people's hard work, with Marina Zhurakhinskaya currently leading the project. Increasingly, whenever I meet a new woman open source developer, it turns out she got her start through an OPW internship. Thanks to the Linux Foundation and Sarah Sharp, this year the OPW awarded seven Linux kernel internships [5].

The Geek Feminism Wiki and blog have been around for years, but are starting to reach their full potential as tools to support women in tech. In the past year, the Geek Feminism Wiki has been cited by several mainstream media outlets, such as *The New Yorker*, as supporting evidence for stories on women in tech. In particular, the “Timeline of Incidents” has been crucial supporting evidence in many discussions of whether or not women in tech face systemic discrimination. Geek Feminism is the work of many people, but two of the most prolific contributors and founders are Alex “Skud” Bayley and Mary Gardiner. Without the Geek Feminism Wiki, the Ada Initiative and many other groups could not have made the progress we have over the past few years.

Rikki: What else is the Ada Initiative focusing on right now?

Valerie: We’d like to expand the use of the Allies Workshop as corporate training for technology companies. Many corporations want to hire more women in tech, but aren’t aware of the ways their internal culture are off-putting or frankly hostile to women. If companies succeed in hiring women despite their culture, those women often end up fighting an uphill battle to change the internal culture, and often end up leaving out of exhaustion. With the Allies Workshop, we teach men how to fight these battles and change their culture to be more welcoming to women and many different kinds of people. It’s fun, too; after one Allies Workshop, an attendee asked HR if they could get “more training like that.”

We’d like to contribute to the trend of community-wide codes of conduct going on in open source, Wikipedia, and similar online communities. This is not as easy as banging out an example anti-harassment policy like we did for conferences—at heart, conferences have much more in common with each other than open tech/culture communities. For example, conferences usually have clear-cut leadership who can kick people out of a clearly defined physical space. Online communities are much more varied in governance and structure, so there’s no one-size-fits-all way of effectively implementing a code of conduct.

Rikki: Did you know that I referenced your “HOWTO Encourage Women in Linux” [6] article in my Master of Science in Journalism thesis? Even though you wrote the document back in 2002, it passed the test of time and you covered topics that are still relevant today. If you were to update the document now, what (if anything) would you add or change?

Valerie: No, I had no idea—thanks for letting me know! Dozens of people helped me write that HOWTO, mostly other LinuxChix members, and I’m glad so many people found it useful. That HOWTO showed me how powerful the written word could be, and I’ve never forgotten that lesson.

For years, I had “Update HOWTO Encourage Women in Linux” on my to-do list, but the thought of reading something I’d written that long ago made me cringe. I did finally bring myself to reread it a couple of years ago and was pleasantly surprised with how much of it I still agreed with, enough that I stopped planning to update it. I’d improve some language, I’d replace the example of sexism in the introduction with a link to the Geek Feminism Wiki, and maybe add a few more items. But overall, I think my time is better spent on new projects than in bikeshedding one that is successful enough.

Rikki: Can you tell us about your interest in labyrinths? I know that a lot of developers are also runners. When I interviewed Nick Lang and Jacob Kaplan-Moss about the PyCon 2012 5k [7], Nick told me that running helps him figure out programming problems he’s stuck on. Do you feel that way about labyrinths?

Valerie: As a programmer, I’m deeply interested in ways to encourage that unconscious intuitive leap that shows you the bug fix or the solution to the design problem. A common theme in people’s stories about “Aha!” moments is that they were doing something else at the time that didn’t take up all of their concentration. In my experience, that something else is often walking—but also showering, driving, falling asleep, etc. For me, any kind of walking helps me come up with creative solutions or new insights. Labyrinths are neat both because they let you walk without needing a destination, but also because they have so much history tied up in them.

References

- [1] Ada Initiative: <http://adainitiative.org/>.
- [2] Val Henson et al., “Chunkfs: Using Divide-and-Conquer to Improve File System Reliability and Repair”: <https://www.usenix.org/conference/hotdep-06/chunkfs-using-divide-and-conquer-improve-file-system-reliability-and-repair>.
- [3] Linux-kernel list email exchange: <http://marc.info/?l=linux-kernel&m=137390362508794&w=2>.
- [4] Outreach Program for Women: <https://wiki.gnome.org/OutreachProgramForWomen>.
- [5] Linux Kernel Internships 2013: <http://sarah.thesharps.us/2013/05/23/%EF%BB%BF%EF%BB%BFopw-update/>.
- [6] Val Henson, “HOWTO Encourage Women in Linux”: <http://tldp.org/HOWTO/Encourage-Women-Linux-HOWTO/>.
- [7] Rikki Endsley, “How to Pound the Pavement with Programmers at PyCon” (PyCon 2012 5K), *Network World*, Mar. 6, 2012: <http://www.networkworld.com/community/blogs/pycon5k>.