


# The Apache Software Foundation: Brian Behlendorf

Charles Severance, *University of Michigan*



From “A Patchy Server” to the Apache Software Foundation that is an important part of the open source movement today, Brian Behlendorf describes the ASF’s history.

**T**he Apache Software Foundation (ASF) is seen by many as the leading example of how an open source software foundation should be run.

The ASF comprises more than 160 open source projects, each with a team of volunteers responsible for the design, planning, and development of the software tools in wide use in millions of technology projects worldwide. These tools include software like the Tomcat Java-based Web server, httpd Web server, the Lucene search engine, the Hadoop distributed computing engine, the Apache Shindig OpenSocial implementation, and many others.

Apache provides a meeting point where engineers from large companies like IBM, Google, Yahoo, Sun, and Oracle work as volunteers alongside talented individual contributors to build open source software infrastructure.

Brian Behlendorf, one of the ASF’s cofounders, describes the project’s history; visit [www.computer.org/computingconversations](http://www.computer.org/computingconversations) for the full interview.

## HUMBLE BEGINNINGS

The group that would become the ASF got its start in 1995 when the National Center for Supercomputing Applications (NCSA) dropped support for its open source Web server (httpd). As Behlendorf describes it,

We got our start in the early days as a group of webmasters who were using a piece of freely available Web software but had difficulty with it. We were fixing bugs and sharing these bug fixes—patches—with each other like baseball trading cards.

One day, we discovered that the group that had put out the webserver that we were using basically folded when all its developers left to join a new company called Netscape. We thought, “Hey, we’re dependent on this software, but we don’t want to become full-time webserver developers. We want to be able to use this thing that we’ve had for free and improve it.”

So we formed a mailing list of webmasters and people working at some early Internet service providers, website design companies, or places like Amazon and the Internet Movie Data-

base, and we combined patches and decided to call it “A Patchy Server.” The model of how we worked was based upon us as a group of peers proposing ideas, vetting each other’s ideas and patches, and fixing bugs as a group as a team.

The Apache Web server’s popularity grew along with the Web, and it was able to keep up with as well as fuel the rapid online innovation occurring during the late 1990s. Because the people who made up the Apache community were working at organizations pushing the leading edge of Web application technologies, it’s natural that the Apache webserver had the latest and greatest features.

## BECOMING A LEGAL ENTITY

In 1998, the community decided that it was important to create a more formal legal structure around its software efforts, so it formed the nonprofit ASF on 1 June 1999. It’s impressive that the group was able to coordinate its efforts without a formal legal structure or centralized governance for nearly five years:

It turns out to be not that hard to work together when people have the same common goal, which is, “let’s build a product that does all this great stuff.” One thing we did that made it easier to make decisions was to have a very modular API, which made it easy for us to say, “If you want that special cool feature, do it as a separate thing, and we’ll decide whether to bring it into the product once it has become successful or not.”

The ASF’s creators made sure to build a set of operating principles into the organization that were based on their experience of working together successfully. Core to these founding principles was that people had to work together as volunteers and peers. Because no member of the group could hold anything over another member of the group, there was no need to layer any kind of formal structure on top of a project’s activities or to assign team members official positions or job titles:

The style of leadership isn’t so much command and control and plotting moves ahead of time but instead being able to get people on your side to convince them that you’re going to value the contributions they make. That’s really the story of successful open source projects writ large: people working together on common technologies to solve common problems so they can go off and make money other places or so they can have fun and try new ideas. That’s really the same story of Apache, Linux, and other open source projects.

Beyond using a flat organizational structure within projects, the ASF also has a liberal license on its software. The Apache license encourages anyone to make copies or to alter and redistribute their own version of the software, even if the new version isn’t open source. This approach to intellectual property allows companies like Google, IBM, and Oracle to let

their best engineers work on a proprietary version of a webserver and still contribute freely to a parallel open source version of that same software. They can maintain their unique proprietary features and integrations and also leverage the talent available across an open source team.

### OPEN SOURCE REVOLUTION

Increasingly, companies view Apache as a natural place to contribute their proprietary code when they

**The Apache license encourages anyone to make copies or to alter and redistribute their own version of the software, even if the new version isn’t open source.**

see no further reason to “hide” their product’s source code. The Apache Tomcat Java-based webserver was originally developed at Sun Microsystems as a proprietary product and then later contributed to Apache. The Apache Shindig Project is an open source reference implementation of the OpenSocial API pioneered by Google.

Apache makes it quite natural for a company to interact with the open source community in a way that brings significant value to both the company and the community.

A core value of Apache is that everyone participates because they want to be there, not because they were coerced. Another core value is that Apache does not feel the need to be the only open source organization. Anyone has the right to make a copy of the entire code base and move it to another open source organization. Making a copy and starting a new thread of development is called a *fork* of the software. According to Behlendorf,

An open source license like we had on our product carries with it the right to fork, which means that if I were to say, “We’re going to go over here!” and no one else wanted to follow, they could decide to pick up the code and start a different project somewhere else. This right to fork means that you don’t have to have any tolerance for dictators, you don’t have to deal with people who make bad technical decisions—you can put the future into your own hands, and if you find a group of other people who agree with you, you can create a new project around it. I think the right to fork limits the excesses we see when we talk about how groups make decisions and conflict arises and how you deal with that conflict.

Many organizations have worked to adopt open source approaches to management and leadership for internal projects. In 1999, Behlendorf founded CollabNet along with Tim O’Reilly of O’Reilly publishing to help companies apply open source principles within their own organizations.

In a way, the ASF has grown up alongside the Web—but it eventually emerged to become one of the predominant sources of infrastructure for building large-scale Web applications. Because this infrastructure has been built as open source, everyone can use best-of-breed tools without worrying about cost. Having these powerful tools available has greatly contributed to the culture of innovation that keeps the Web evolving. ■

*Charles Severance, Computing Conversations column editor and Computer’s multimedia editor, is a clinical associate professor and teaches in the School of Information at the University of Michigan. You can follow him on Twitter @drchuck or contact him at csev@umich.edu.*

**cn** Selected CS articles and columns are available for free at <http://ComputingNow.computer.org>.