

Apache

NCSA developed both a browser for viewing pages and server software for delivering web pages to people. The web server, HTTPd, was written by Rob McCool in 1993 and was based on the CERN server code. NCSA released the program and its source code for free (Moody, 2001). The source code is a human readable set of instructions for the computer. Consequently, the NCSA server quickly became the most popular web server for the Internet. Many sites chose the free NCSA HTTPd server over Netscape's web servers that cost several thousand dollars.

At first, the programmers at NCSA quickly patched any problems with HTTPd they received. But by 1995, the original team of programmers had left NCSA, and HTTPd was not updated in a timely manner. This led individuals outside of NCSA to begin to "patch" problems that they discovered. This was possible because the source code was in the public domain, and therefore, freely available. An example of a patch was the addition of password authentication by Brian Behlendorf for the Hotwired web site. Other patches improved the security and performance of HTTPd.

The NCSA's HTTPd server software was public domain through and including version 1.4. The last version released by NCSA, version 1.5, was not released as public domain and was instead copyrighted by the University of Illinois. According to Rob McCool, the creator of the NCSA HTTPd server, it was Marc Andreessen's decision to release the server as public domain, because of the problems Gopher with adoption because of its restrictive license.

Eventually, there were a number of patches for HTTPd circulating across the Internet. Most of the patches were posted to the mailing list [www-talk](#). However, if

someone wanted the benefit of these patches, they had to download the latest version of HTTPd, and then manually apply all the latest patches (Østerlie, 2002). This prompted users of HTTPd to consider updating NCSA's code. According to Østerlie, the individuals viewed themselves as disgruntled customers. They were simply filling the gap left by the departure of NCSA's original programmers to Netscape.

Behlendorf then began to contact other programmers. By February 1995, the group put together a mailing list called new-httpd and began circulating patches (Moody, 2001). The project's goal was to fix the existing problems and to add enhancements to the server. An example of an enhancement was the inclusion of Secure Sockets Layer. The first set of patches were applied to NCSA's HTTPd 1.3. The resulting code became the first official release of Apache in April 1995. The project was named Apache—after all—the joke name for the server was “A PatCHy server”.

The management structure for Apache is inspired by the IETF and its motto, “rough consensus and running code” (Moody, 2001). The procedural rules allow anyone to contribute code as they see fit. There is a voting system to decide what code will be released as the official Apache version. Only the core developers are allowed to vote. New voting members are added when a frequent contributor to the project is nominated and unanimously approved by the existing voting members (Fielding, 1999).

The core developers are located in the United States, Britain, Canada, Germany, and Italy and maintain contact through a public mailing list. The members are not teenage hackers, but consist of doctoral students, a Ph.D. in Computer Science, professional software developers, and a software business owner (Mockus, Fielding, & Herbsle, 2000). There are about fifteen core developers at any time. The core developers create

approximately 80% of the new functionality. However, over 400 individuals have contributed code, and over 3000 people have contributed problem reports.

During May and June of 1995, little work was done on Apache. The reason was described by Cliff Skolnick as follows, “[y]ou can add honey to shit, but you just get sweet tasting shit. No matter what you add to shit, you end up with some form of shit” (Østerlie, 2002). Apache had stagnated as developers didn’t see it as worthwhile to contribute their time and code. This would change after Robert Thau announced his “garage project”, new code named Shambhala, which was a rewrite of the server code. Within a few months, the Shambhala code became the basis of the Apache server. The new Shambhala code reignited discussion and work on the Apache server.

One important aspect of Shambhala was the separation of the functionality into a set of modules. The modules are mutually independent. People can work on individual modules and not affect ongoing work in other modules. This design feature supports a decentralized development process. This design change was extremely important, because it fostered the use of the open source distributed development model.

Shambhala also reinforced the community’s emphasis on open standards. Consider what happened during the Christmas of 1995 when AOL performed minor upgrades of their web proxies. Consequently, the web pages served by Apache returned an error to AOL users. This led to a debate in the Apache community about whether they should write a simple patch to fix the problem or dig in their heels and force AOL to fix their web proxies to comply with existing web standards. The community decided it was more important to stay with open standards, and in the end, AOL fixed its web proxies (Østerlie, 2002).

One of the reasons for Apache's success was the failure of servers from Netscape and Microsoft to meet the demands of the marketplace, specifically Internet Service Providers (ISPs). For example, ISPs widely embraced Apache, because it allowed them to offer web hosting for corporate web sites for less money than an in-house corporate web site would cost. Apache could host 10,000 web sites on a single web server. This functionality, virtual hosting, was included in Apache by the summer of 1995, but not found in servers from Netscape or Microsoft. Additionally, ISPs as well as other users also choose Apache because they could modify it for their own needs. They would simply have to modify the source code, which was freely available. With Netscape's and Microsoft's servers, a customer had to wait for them to add a new feature or fix a problem. Consequently, Apache's market share steadily grew from late 1995, and today, it is the most popular web server on the Internet.

Apache's success did not go unnoticed. IBM decided to adopt the Apache web server. In 1998, IBM announced it would ship the Apache web server with its Websphere product family as a commercial, enterprise-level package. IBM chose Apache over its own products, because Apache was the best server available. IBM understood that there was little money to be made from servers (Donahue, 1998). Instead, IBM would profit from service and support as well as from proprietary add-ons such as an online e-commerce system. In turn, IBM has contributed to the development of Apache. The relationship between IBM and Apache is still ongoing with both parties enjoying the benefits of the relationship.

References:

- Donahue, S. (1998, June 19). *IBM Picks Apache*. Retrieved June 22, 2004, from <http://www.wired.com/news/print/0,1294,13117,00.html>
- Fielding, R. T. (1999). Shared Leadership in the Apache Project. *Communications of the ACM*, 42(4), 42-43.
- Mockus, A., Fielding, R. T., & Herbsle, J. (2000). *A case study of open source software development: The Apache server*. Paper presented at the International Conference on Software Engineering, Limerick, Ireland.
- Moody, G. (2001). *Rebel Code*. New York: Penguin Press.
- Østerlie, T. (2002). *The User-Developer Convergence: Innovation and Software Systems Development in the Apache Project*. Norwegian University of Science and Technology, Trondheim, Norway.