The Apache Software Foundation: An Overall Look at the Foundation's Structure and Projects

By Vikram Gaur

Note: *More than a few readers have requested an explanation of the Apache Software Foundation and its many projects to eliminate the confusion that exists among IT managers and staff. The primary reason for the confusion seems to be the use of abbreviated terms and nicknames in the IT industry when referring to the Foundation, projects, or subprojects. Eliminating this confusion is the goal of this and future articles. This first article provides a base of information upon which numerous future articles will provide explanation and insight into each of the most popular projects.*

In the early '90s, a portion of the Software Development Group of The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign created an HTTP/1.0-compatible server for making hypertext and other documents available to Web browsers. The name of this server was NCSA HTTPd. The NCSA also is the origin of the NCSA Mosaic Web browser, which was a driving force behind the creation of the NCSA HTTPd server.

Several early Web visionaries (e.g., Marc Andreessen, Rob McCool, Eric Bina, Carlos Varela, Brandon Long, and several others) all happened to be at the NCSA, working together to bring new, improved Web browsing tools to the Internet. Many people needing a Web server quickly adopted the NCSA HTTPd server, since it provided a free, good quality, easy to use, and smaller solution than the CERN Web server developed by Tim Berners-Lee and colleagues. Source code also was available, enabling users to contribute fixes to the NCSA to be incorporated into future releases.

In 1994, the principle developer of the NCSA HTTPd server, Rob McCool, left the NCSA. While the NCSA HTTPd server was enjoying widespread use, new development and the application of contributed fixes soon declined. In 1995, Brian Behlendorf and a group of additional developers joined forces to create an alternative server solution based on the codebase of the NCSA HTTPd 1.3 server, since it looked as if the NCSA wouldn't be continuing development and support. This group of individuals called themselves The Apache Group.

The Apache Group began collecting and applying its own patches to the NCSA HTTPd codebase and created enhancements to deliver a new alternative server product—the Apache HTTP Server. Within a year of being introduced, the Apache HTTP Server overtook the use of the NCSA HTTPd server, and remains the most widely used Web server in the world.

The Apache Group also used a merit system that provided rules and rewards for individuals who contributed to the Apache HTTP Server effort. Since all the work to create the Apache HTTP Server was performed on a voluntary basis, the initial members of The Apache Group needed a method to determine which new contributors would be permitted to join the community of developers. The system they used and still use today is referred to as a meritocracy, whereby those who contributed the most earned the merit to become a member of the development community and were granted direct access to the code repository. This system resulted in a scalable method of increasing the developer base in an equitable manner, and has been a driving factor for years for many developers to contribute. It's the community management philosophy of the organization that continues to attract developers to volunteer their efforts, and makes the name Apache more than simply a name.

For many years, the only product provided by The Apache Group was the Apache HTTP Server. Looking back, it's easy to see why people commonly referred to this popular Web server and the organization with the nickname "Apache," since there was a one-to-one relationship between the organization and the product. While you still needed to listen to the context of the conversation to know if a person was talking about the organization or the product, it was still easy to understand. The legacy continues today, and those of us who have either been involved with Apache since the early days, have contributed, or are contributing, continue to use the nickname "Apache" when referring to the Web server or the organization. Today, there's no longer only one product being created or supported by this large non-profit organization now known as The Apache Software Foundation (ASF).

Today, the ASF encompasses more than 70 projects, subprojects, and sister projects. Some of the ASF projects include Ant, APR,

Figure 1: ASF Projects

JANUARY/FEBRUARY 2006 | ENTERPRISE OPEN SOURCE JOURNAL | 41

Cocoon, Commons, DB, HTTP Server, and others noted in Figure 1. ASF projects are elevated to the status of an "official" project when the community can manage itself independently and the board of directors approves the project. Sister projects are those that have a close relationship with the ASF but aren't officially part of the ASF. The current sister projects include Module Registry, PHP, mod-ssl, and Apache-SSL. The ASF also has other entities that aren't focused on creating software, but are instrumental in the operation (Apache Infrastructure), community building (Apache Incubator), ApacheCon conference planning (Apache Conference Planning), and community interests (Java Community Process) of the ASF. Currently, there are more than 1,100 active committers contributing code to one or more of the projects or subprojects.

The ASF is a large organization with a wide variety of projects. You now have a good basic understanding of the ASF organization and this article will help others understand the breadth and depth of the ASF. For more information, visit www.apache.org. You'll find a wealth of information here.

Future issues of *Enterprise Open Source Journal* will provide an explanation about each of the major ASF projects and subprojects. We can't cover every project, but we can bring you a recurring stream of articles that will more fully explain the major projects and subprojects. The series will give readers interested in open source solutions and the ASF the knowledge needed to also understand projects and subprojects not covered.

Vikram Gaur is an electronics graduate from Kurukshetra University, India. In 1999, he and Neeraj Gaur formed IIRA Technologies, which later became IIRA Compuserv Technologies (P) Ltd. As CTO, he has headed numerous projects based on Linux and Windows. Currently, he and his team are working on Centralized Network Management. He is also heading projects based on PHP, Java, and SAP.
e-Mail: vikramgaur@iiratechnologies.com

1 figure
(Denny retyped and gave to Martin)

Callout: The community management philosophy of the organization continues to attract developers to volunteer their efforts and makes the name Apache more than simply a name.

The Apache Foundation

Is the Apache Software Foundation (ASF) a non-profit organization?

Yes, the Apache Software Foundation is a non-profit, 501(c)(3) corporation. It's membership-based and registered in Delaware. As a non-profit charity, donations to the ASF are tax deductible.

Is it true that some companies are part of Apache?

No, membership of the ASF is composed of individuals only; companies have no role as far as the composition of ASF is concerned. Moreover, all the members have a legal stake in the ASF. However, individuals who work for companies are encouraged to contribute to Apache. In fact, there's a specific extra Contributor License Agreement (CLA) to ensure individuals can contribute to the needs of ASF during "work hours." It's also expected that the members are committed and acting solely on behalf of the ASF when wearing their ASF hats, regardless of their employer.

Does ASF mean Apache Web Server?

No, ASF doesn't mean Apache Web Server. Rather, ASF is the foundation that maintains several projects, including the Web server.

Does Apache 2.0 mean ASF 2.0?

No, ASF is a foundation, and foundations can't have versions. When someone talks about Apache 2.0, it doesn't mean ASF 2.0. Apache 2.0 is the version of Web server (i.e., Apache HTTP Server 2.0).

Should the term "Apache 2.0" be used to refer to Apache HTTP Server 2.0?

No, it's misleading. Apache isn't only a Web server, but also a group of projects managed by the ASF. So, instead of Apache 2.0, one should say Apache HTTP Server 2.0.

Who owns the Apache code?

The authors contributing to an ASF project own their own contribution, but license it to the ASF and its recipients. No author ever assigns copyright to the ASF. However, the ASF owns each project's official distribution, which is a collection of various contributions.

When talking about the direction of the code, it's more appropriate to talk about the Project Management Committee (PMC) members, not the ASF members. The ASF members play more of a role in guiding the direction of the foundation as a whole, not any one project.

What is the ASF all about?

The ASF places a large emphasis on the community around a project's code base. All of our projects operate through public consensus-driven decision making with the goal of making the best decisions for the projects. For example, an individual committer can't issue a release by themselves; rather, a minimum of three representatives must agree to issue the release. Individual rights are bestowed based on demonstrated merit through constructive and sustained involvement in a project. The ASF believes that a living, growing project with a healthy community is more important than having a great piece of software with no one to maintain it.

-V.G .-

JANUARY/FEBRUARY 2006 | ENTERPRISE OPEN SOURCE JOURNAL | 42

Guar article ends

JANUARY/FEBRUARY 2006 | ENTERPRISE OPEN SOURCE JOURNAL | 43