


Behind the Scenes of The Apache Software Foundation

Lars Eilebrecht

<lars@apache.org>

Stefano Mazzocchi

<stefano@apache.org>




What is the ASF?

- ASF stands for Apache Software Foundation
- The ASF is a registered non-profit organization
- Incorporated in the state of Delaware, USA
- The ASF is a volunteer organization



What is the ASF?

- all work is done by volunteers
- nobody gets paid by the foundation
- except a contractor working full-time as system administrator (since April 2006)



Why the ASF?

- to provide a foundation for open, collaborative software development projects by supplying hardware, communication and business infrastructure
- to provide a means for individual volunteers to be sheltered from lawsuits directed at the foundation's projects



Why the ASF?

- to create an independent legal entity to which companies and individuals can donate resources and be assured that those resources will be used for the public benefit
- to protect the »Apache« brand, as applied to its software products, from being abused by other organizations



Bylaws of the ASF

- bylaws can only be changed by the members or board of directory, but not just by the president
- annual vote of a new board of directors
- board members do not need to be residents of the USA
- becoming a new member requires a majority vote of all existing members




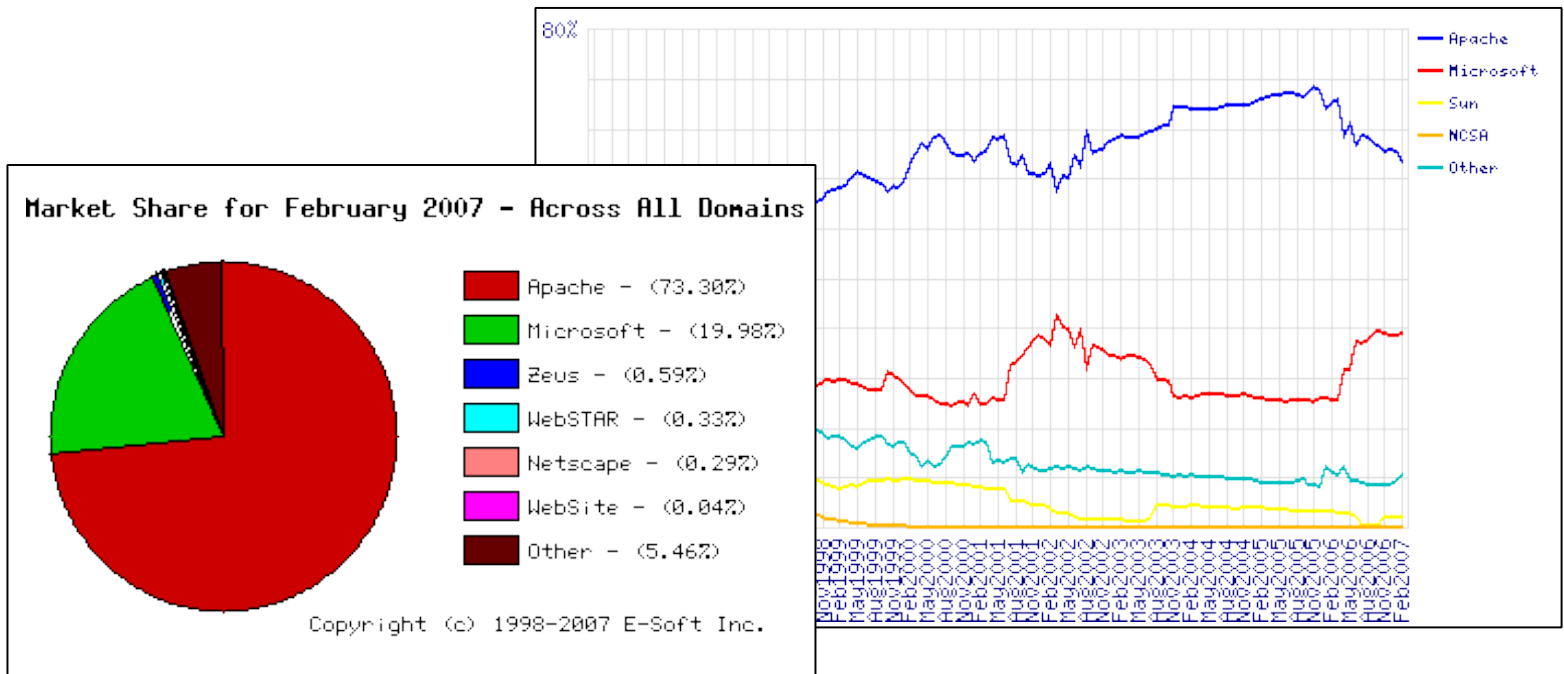
Before the ASF

- in 1995 a group of people got together to exchange software patches for the NCSA web server
- they started to call themselves the Apache Group
- later that year they released their own web server: the Apache HTTP Server




Before the ASF

- in less than 6 months (April 1996) they surpassed the NCSA web server adoption
- market share February 2007: 61-73%



Before the ASF

- thanks to its modular design, sister projects for web server modules came to gravitate around the Apache Group
- being leader of the market, the Apache Group started to fear abuse and lack of legal protection
- all this lead to the creation of the Apache Software Foundation in 1999




Why »Apache« as a Name?

- no, it isn't an attack helicopter ...
- »a patchy« server? No, that's just a pun
- reference to the ASF's development philosophy:
 - »Characteristic of both Eastern and Western Apache, [...] was the lack of a centralized tribal organization. The band, an autonomous collection of small local groups within a given locality, was the primary political unit as well as the primary warring and raiding unit. The strongest headman of the local groups was recognized as an informal chief, and several bands might be united under one leader. Chieftainship was thus not generally hereditary.«
(*Encyclopaedia Britannica*)



Meritocracy

- govern of merit
- the more merit you have, the more power you get
- no old-men club because of openness



The Chain of Merit

- User
- Committer
- Project Member
- ASF Member



User

- anybody who uses the software
 - passive: does not contribute directly to the project (aka lurker)
 - active: contributes directly to the project (aka contributor), either with code patches, feature requests, bug reports, answering questions on mailing lists, etc ...



Committer

- an active user that was elected for merit and, for this, receives:
 - write access to the project's source code repositories (SVN)
 - an »apache.org« email address and access to certain ASF server infrastructure



Project Member

- an active committer that was elected for merit in the evolution of the project, for this, receives:
 - the ability to vote
 - the ability to propose other people for committership


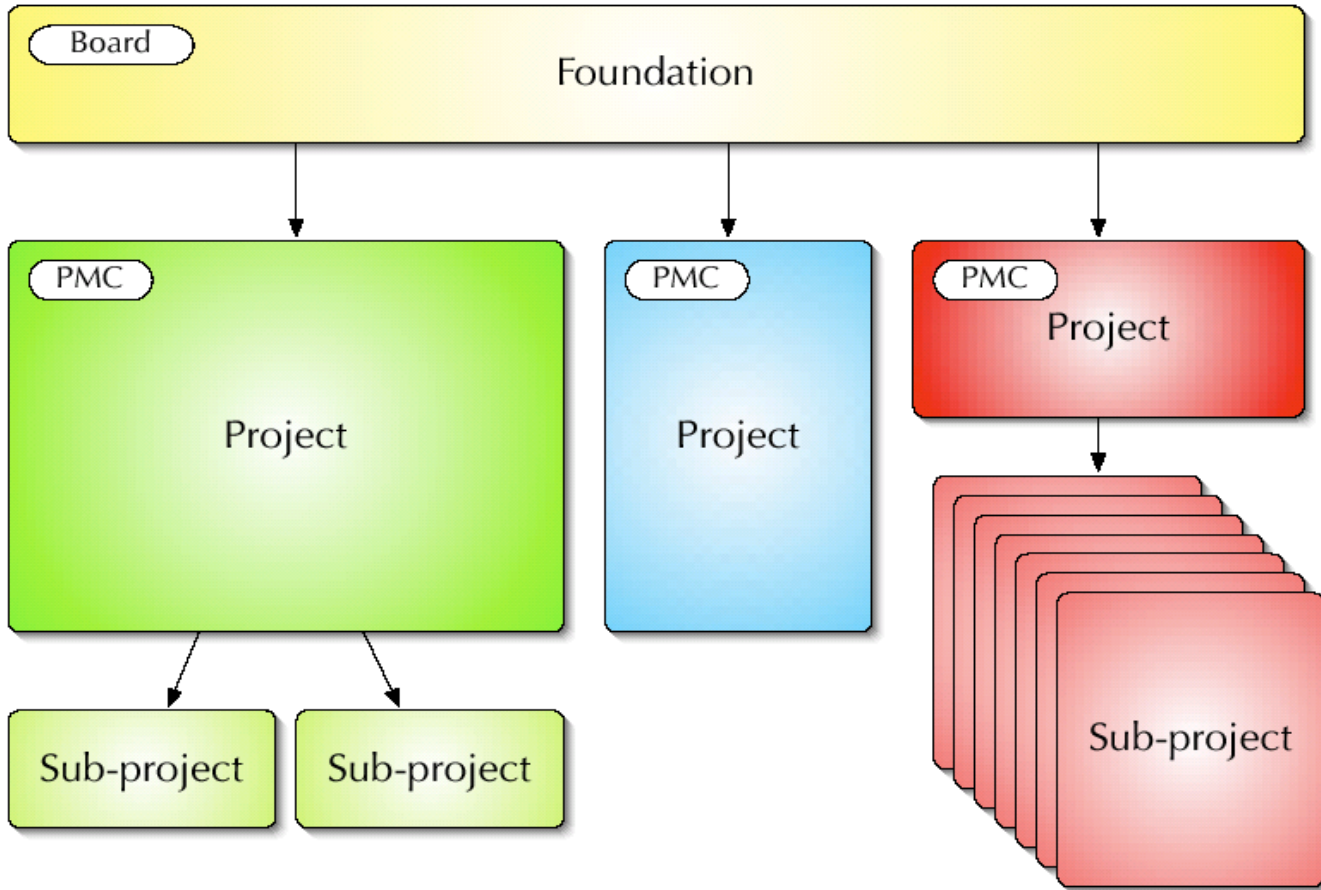


ASF Member

- a project member or committer that was elected for merit in the evolution of the foundation
 - shareholder of the foundation
 - can work on foundation projects/teams
 - can propose committers for membership
 - can elect the board
 - can be a candidate for the board election
 - can propose a new project for incubation



The Foundation Structure




Some Foundation Numbers


- 1239 Committers
- 202 Members
- 42 Officers and 9 Directors



(June 2006)



ASF Member Demographics



The Board of Directors

- manages corporate assets (funds, intellectual property, trademarks, and support equipment)
- allocates corporate resources to the various projects
- 9 members elected every year
- makes no technical decisions!



Project Management Committee

- the PMC manages a project (and its sub-projects) subject to the direction of the board
- each PMC has the faculty of establishing its own set of rules and procedures for day-to-day project management
... but many share the same



Communication

- communication is done mainly over asynchronous media (mostly mailing lists)
- all communication is kept archived, so that others can join or have a reference
- this allows geographic distribution



Decision Making

- lazy consensus: making a decision does not require the whole group to be involved
- voting rules:

yes	abstain	no (veto)
+1	0	-1
- three +1 and no veto required for approval
- a veto requires proposal of an alternative solution or explanation of reason



Philosophy

- collaborative software development
- high-quality software
- respectful, honest, technical-based interaction
- faithful implementation of standards
- security as a mandatory part of software development
- Open Source: making the software available to the general public, and usable for any purpose they see fit




Software Projects

- main purpose is to create, maintain and distribute software
- Current top-level projects:
 - Active MQ, Ant, APR, Beehive, Cayenne, Cocoon, DB, Directory, Excalibur, Forrest, Geronimo, Gump, Harmony, HiveMind, HTTP Server, iBATIS, Jackrabbit, Jakarta, James, Lenya, Logging, Lucene, Maven, Mina, MyFaces, OFBiz, Apache-Perl, Portals, Santuario, SpamAssassin, Struts, Tapestry, Apache-TCL, Tomcat, Velocity, Web Services, Xerces, Apache-XML, XMLBeans, XML Graphics




Auxiliary Projects

- they do not distribute software, but help the foundation in specific domains
 - Incubator
 - Infrastructure
 - Security
 - Public Relations
 - Conference Planning
 - Legal
 - Java Community Process



Apache Incubator Project

- formed in October 2002
- entry point for all new software projects
- provides guidance for new committers
- focus on:
 - addressing legal issues with contributed IP
 - building healthy community around the project
- incubating projects are not officially endorsed



The Apache Security Team

- security@apache.org
- the central point of contact regarding potential security vulnerabilities in Apache Software
- manages vulnerability reports and coordinates efforts between projects
- reports are forwarded and handled by project-specific security teams (e.g., HTTP Server and Tomcat)



Apache Public Relations Team

- press@apache.org / fundraising@apache.org
- deals with public relations, trademark licensing, and fund-raising for the foundation
- common forms of donation:
 - money (check or via [paypal.com](https://www.paypal.com)), tax-deductable in the USA
 - buying ASF gear, e.g., T-Shirts
 - bandwidth and hardware



Apache Conference Committee

- concom@apache.org
feedback@apachecon.com
- organizes ApacheCon and similar events
- Upcoming Apache Conferences:
 - ApacheCon US 2007
Atlanta, Georgia, USA, 12-16 November



The Java Community Process

- represents the ASF in the Java Community Process, responsible for the maintenance and evolution of the Java Platform




ASF Infrastructure

- the ASF is a virtual organization: no offices or building
- only physical existence is the technical infrastructure that runs its services
 - web services
 - source code repositories
 - mail services
 - issue/bug tracking services
 - distribution mirroring



The Infrastructure Team

- manages the ASF infrastructure 24/7
- 16 servers located at 4 co-locations (mainly San Francisco and Amsterdam)
- people close to the location of the servers get notified if machines go down
- ASF employs one full-time system administrator (since April 2006)



Some Infrastructure Numbers


- Web servers:
 - about 50 requests per second
 - about 4 million hits per day
 - about 5 TByte per month
- (average numbers)

ASF Rack,
UnitedLayer Co-Location



Some Infrastructure Numbers

- Mail servers:
 - handle user accounts and more than 280 public lists with a total of about 100.000 subscribers
 - about 1 Mio connects per day (inbound)
 - about 300.000 incoming messages per day
 - Rejected (RBLs): about 700.000 per day
 - Filtered (spam): about 20.000 per day
 - Filtered (virus): about 7.000 per day



More Infrastructure Numbers

- Source Code Repository: 70 GB
- Binary Repository
 - Current 8.5 GB
 - Archive 30 GB
- Mirror servers:
 - currently 282 mirrors in 55 countries
 - about 450 GByte per month of rsync traffic
 - about 2 Mio file transfers per month



Web Sites and Mailing Lists

- Public- and Foundation-wide Mailing Lists:
 - announce@apache.org
 - community@apache.org
 - announce@apachecon.com
 - general@incubator.apache.org
 - legal-discuss@apache.org
(available to committers, others read-only)
- Web Site Resources:
 - <http://www.apache.org/foundation/>
 - <http://www.apache.org/licenses/>
 - <http://incubator.apache.org>




The Future of the ASF

- The future is what *you* make of it ...
 - all Apache work is based on someone needing something at some point and just doing it
 - all Apache technical decisions are discussed and implemented on public mailing lists
 - all you have to do to start is subscribe and participate

»The best way to predict the future is to invent it.«

(Dr. Alan Kay)



Any Questions?



Lars Eilebrecht <lars@apache.org>

ApacheCon

