

The Apache Software Foundation: No Jerks Allowed!



Dr. Justin R. Erenkrantz
The Apache Software Foundation
OSBC 2010
<http://www.erenkrantz.com/>
justin@erenkrantz.com
Twitter: @jerenkrantz

Why should I pay attention?

- *Committer to Apache HTTP Server, APR, Subversion (just graduated!), and Serf*
- *President, The Apache Software Foundation*
- *Ph.D. from Univ. of California, Irvine*
 - *Computational REST (next-gen web arch.)*
 - *<http://www.erenkrantz.com/CREST/>*
- *CTO, Project WBS*

Open Source & Me

- *Started contributing to Apache HTTP Server in early 2001...*



<http://www.flickr.com/photos/x180/2704965763/>
(James Duncan Davidson)

<http://www.flickr.com/photos/bluxte/2103773754/>
(Sylvain Wallez)

Apache Top-level Projects

Abdera	CXF	James	POI	Tiles
ActiveMQ	DB (Derby)	Lenya	Portals	Tomcat
Ant	Directory	Logging	Qpid	Turbine
APR	Excalibur	Lucene	Roller	Tuscany
Archiva	Felix	Maven	Santuario	UIMA
Axis	Forrest	Mina	ServiceMix	Velocity
Buildr	Geronimo	MyFaces	Shindig	Wicket
Camel	Gump	ODE	Sling	Web Services
Cassandra	Hadoop	OFBiz	SpamAssassin	Xalan
Cayenne	Harmony	OpenEJB	STDCXX	Xerces
Click	HC	OpenJPA	Struts	XML
Cocoon	HTTP Server	OpenWebBeans	Subversion	XMLBeans
Commons	iBATIS	PDFBox	Synapse	XML Graphics
Continuum	Jackrabbit	Perl (mod_perl)	Tapestry	<i>...and Incubator</i>
CouchDB	Jakarta	Pivot	TCL	<i>podlings...</i>

Apache's tagline

We are more than a group of projects sharing a server, we are a community of developers and users.

Committer geographic diversity



<http://people.apache.org/map.html>

Challenges and advantages arise from this!

Project Diversity in Apache

- *Seventy top-level projects (TLPs)*
 - *Java, C, C++, Perl, etc. projects...*
- *It's okay for projects to be in "same" space*
 - *Ant/Maven, Pig/Hive, Axis/CXF...*
- *Over thirty projects currently in "pipeline"*

Apache aims to provides a pragmatic non-technical framework to its projects

Apache's mission

The Apache Software Foundation provides support for the Apache community of open-source software projects. The Apache projects are characterized by a collaborative, consensus based development process, an open and pragmatic software license, and a desire to create high quality software that leads the way in its field.

Apache's mission (redux)

- *Let developers focus on what they do best: code. Foundation exists to do rest.*
- *“The Apache Way”*
 - *Open development vs. open source*
 - *All technical decisions about a project are made in public (mailing lists)*

Indirect financial support

- *Apache does not pay for development*
- *Many (not all!) developers are paid by a third-party to work on the project*
 - *Participate in Google Summer of Code*
- *Foundation bears indirect support costs*
 - *Infrastructure, publicity, etc.*

Motivations & Business Models

- *Each contributor has own motivations*
- *Seeking help from others on a big project*
- *Commercial support: who to call at 2am?*
- *Common/shared platform: value add on top of Apache products*
- *...or it's fun/interesting/challenging!*

Leaders? We don't need leaders!

- *“Apache” and our feather logo refer to the Native American Apache Nation (Inde tribe)*
 - *Not really “A Patchy Web Server”*
- *Whoever has the best idea “leads”*
- *...until a better idea is presented to the group and then that person “leads”*

“Meritocracy in Action”

- *meritocracy: power distributed based upon demonstrated talent*
 - *Young’s Rise of the Meritocracies (1958)*
 - *Negative connotations...but...oh well!*
- *Duncan’s “Rules for Revolutionaries” dives into the operational aspects seen in Apache*

<http://incubator.apache.org/learn/rules-for-revolutionaries.html>

Founding of Apache

- *Started as “Apache Group” (8 members)*
- *Resumed work on NCSA httpd in Feb. 1995*
 - *UIUC put httpd in public domain, but essentially abandoned it*
 - *Chose permissive licensing (more later)*
- *Informal corporate structure until...*

Creation of Foundation

- *Incorporated with 21 members in 1999*
 - *~2,300 committers, 274 members, and 52 emeritus members today*
- *Membership-based organization*
- *IRS 501(c)3 public charity status*
 - *Donations by individuals tax-deductible*

Organization of ASF

- *Each Apache project is independent*
- *Grouped as 'top-level' PMCs (TLP)*
 - *Board: Social - not technical - guidance*
- *Some TLPs have 'sub-projects'; discouraged*
- *Karma in one PMC doesn't grant rights in another PMC - earn karma independently!*

Relationships between projects

- *HTTP Server and Tomcat connect via AJP*
- *Subversion and HTTP Server use APR*
- *Commons logging API supports log4j*
- *Tika uses POI for extracting OOXML/ODF*
- *Sling uses Jackrabbit for content storage*
- *Synapse uses Axis for SOAP bindings*

Board of Directors

- *Elected yearly by membership*
- *Shane Curcuru (IBM), Doug Cutting (Cloudera), Justin Erenkrantz (Project WBS), Roy T. Fielding (Day), Jim Jagielski (VMWare), Geir Magnusson, Jr. (Gilt), Brian McCallister (Ning), Brett Porter (G2iX), Greg Stein (Independent)*

Executive Officers and Staff

- *Exec officers appointed by Board (unpaid!)*
 - *Chairman, President, Secretary, Treasurer, Executive Vice President*
- *Chairman and President typically split external visibility roles; Chairman responsible for Board; President oversees day-to-day operations*

Hardware Infrastructure

- *Bulk of work done by volunteers*
 - *If a project wants it, they must volunteer!*
- *2 paid system administrators fill “holes”*
- *Main data centers: OSU OSL, SURFnet (.nl)*
- *Build farms at Traci.net and Yahoo!*
- *Off-site backups at UC Irvine*

Other Committees

- *ConCom: ApacheCon and smaller events*
- *Legal: work w/EFF, SFLLC, & Larry Rosen*
- *Publicity: Deal with press/outside world*
- *Brand Management: Trademarks/licensing*
- *Fundraising: Sponsorship program*

Budget

- *2009-2010FY Budget (May 1 - April 30)*
- *Projected income: \$540,000*
 - *Rely upon public and corporate sponsors*
- *Projected expenses: \$404,000*
 - *Biggest expense is infrastructure: \$150k*

<http://www.apache.org/foundation/records/>

ASF Sponsorship Program

- *Platinum (\$100,000 per year)*
 - *Yahoo!, Microsoft, Google*
- *Gold (\$40k/yr): HP, Facebook*
- *Silver (\$20k/yr): VMWare, Progress (IONA)*
- *Bronze (\$5k/yr): AirPlus International, BlueNog, Intuit, Joost, Matt Mullenweg, Two Sigma Investments*

Contributions

- *Apache wants voluntary contributions*
 - *Not copyleft!*
- *Many forms of contributions*
 - *evangelism, bug reports, testing, documentation, code, design feedback*
- *Contributing is easier than you think!*

Where decisions happen

- *Can't kibitz in the morning over coffee*
- *First time meet face-to-face is at events*
- *Mailing lists are the pulse of the project*
 - *IRC, AIM, Jabber, etc. not for decisions*
- *Roy's mantra: "If it doesn't happen on-list, it didn't happen."*

Finding that list!

- *Should be listed on every project's site*
- *dev@ - dev-to-dev discussions (**primary list**)*
- *commits@ - automated source changes*
- *users@ - user-to-user help (optional)*
- *http://mail-archives.apache.org/mod_mbox/*

Scoping contributions

- *Do not do “powerplant” contributions*
 - *Large changes are hard to review*
 - *Break patches up into small chunks*
- *Before starting large changes, email plan to developer list - explain how and why*
- *Try to get feedback early and often*

Voting

- *Vote on releases, code, and ideas...*
- *+1? -1? +0? -0?*
- *“Binding” vote given to committers*
 - *Everyone should feel free to vote*
- *Cast your vote on-list*
- *...what about vetos?*

The dreaded -1 (veto)

- *Power given to all committers to absolutely stop an action*
- *Code can be vetoed - releases can not*
- *Some projects have an 'override'*
- *Vetos should only be cast as a measure of last resort; use judiciously!!*

Rule of 3 and 72

- *3 voters is minimum acceptable quorum*
 - *Ensures diversity and true community*
- *72 hours minimum time frame for votes*
 - *Accommodate long weekends too*
 - *Generally err on side of caution*

Apache commit policies

- *RTC - Review then Commit*
 - *At least 3 +1s; no vetos*
- *CTR - Commit then Review*
 - *Lazy consensus - assumed okay*
- *Different policies per PMC or codebases*
 - *Stable: RTC, Trunk: CTR*

Invitation of Commit Access

- *Current PMC member nominates individual*
- *Discussions on private@ list*
 - *Key ? : Do we trust this individual?*
- *PMCs free to set own bar*
 - *HTTP Server: ~6 months of contributions*
 - *Subversion: 1-2 months of contributions*

No Jerks Allowed!

*“How Open Source Projects Survive
Poisonous People (And You Can Too)” by
Ben Collins-Sussman and Brian Fitzpatrick*

<http://video.google.com/videoplay?docid=-4216011961522818645>

- Most people are nice; there are dingbats, or may just be someone having a bad day*
- Trolls exist...don't feed them.*
- Don't become a poisonous person.*

Legal framework

- *ASF philosophy is that code can be used commercially without restrictions*
 - *Just don't call it Apache Foo!*
- *Apache License, version 2.0*
 - *Grants copyright and patent licenses*
- *Many non-ASF projects use ALv2 now*

Incorporating third-party code

- *ALv2 is GPLv3-compatible - can be used by GPL projects but no GPL within ASF*
- *<http://www.apache.org/legal/3party.html>*
- *<http://www.apache.org/legal/resolved.html>*
 - *List of licenses permitted / disallowed*
- *Questions to legal-discuss@apache.org*

What am I signing?

- *Contributor License Agreement (CLA)*
 - *Gives ASF 'license' - not ownership*
 - *Required before getting SVN account*
 - *Fax or email to secretary@apache.org*
- *Corporate CLA form - when needed?*

Growth of Foundation

- *Started with just HTTP Server in 1995*
- *Felt there was something repeatable*
- *Today, we have over 70 top-level projects*
- *It took over 15 years to get there...*
- *...but it wasn't smooth...*

Jakarta “Foundation”

- *Jakarta: “umbrella” for all Java efforts*
 - *Successful as a brand in its own right*
- *Tomcat, Ant, Struts, etc.: great innovation*
- *Started to copy foundation org structures*
 - *“Mini”-board...but problems arose...*
 - *Avalon: who was responsible?*

Importance of Oversight

- *Jakarta issues led to a lot of navel-gazing*
- *Ultimately agreed upon an extremely flat organizational structure: **umbrellas are bad!***
- *So, we killed Jakarta: spun-off projects*
- *Board requires **all** projects to submit reports each quarter: by far, most important thing that Board does in our monthly meetings.*

Let a thousand flowers bloom

- *Grassroots: interesting projects welcomed*
- *Board doesn't say "We want X", instead developers say "We think X is cool"*
 - *Helped keep us at forefront of innovation*
- *Community support is essential - we are not interested in "solo" projects, but how can we help create a viable community?*

Starting new Apache projects

- *Incubator - “podlings” can be nominated and eventually “graduate” to be a PMC*
 - *Needs foundation member to **mentor***
 - *Usually legal and/or community issues*
- *Labs - once you are a committer, you can have a sandbox (shared mailing list, no non-committers, no releases)*

Community departure?

- *What happens when the community leaves?*
- *Unable to muster 3 votes for a release...*
- *...no active committers...*
- *...or fail to report to Board*
- *Move project into “Attic”*
- *No one has returned from attic...yet.*

Beyond a committer...

- *Once you are committer, you can then become a PMC member, foundation member, Director...even President!*
- *You can nominate other foundation (or PMC) members; can serve as mentor for Incubating projects; vote for Board.*
- *Maybe you too can then give this talk!*



Thanks! Questions?

<http://www.erenkrantz.com/>

justin@erenkrantz.com

Twitter: @jerenkrantz

Other Recommended Resources

- *Producing OSS*
by Karl Fogel - <http://producingoss.com/>
- *Getting Started in Open Source: An Overview for Newbies*
by Leslie Hawthorn and Cat Allman (at SCaLE 7x and OSCON 2009)
 - http://scale7x.socallinuxexpo.org/sites/scale7x.socallinuxexpo.org/files/SCALE_Final_allman_lhawthorn.pdf