

Apache Subversion: 11 Years of Solidity

Greg Stein
Apache Software Foundation

My Background

- Industry for 27 years
 - Oracle, Microsoft, and Google (and others)
- Open Source for 17 years
 - Apache, Python, DAV, Subversion (and others)
- Apache Software Foundation
 - VP of Subversion
 - Member, Director, Chairman, Vice Chairman

Subversion Begins

- Concept begins at CollabNet
 - CVS hosting service was very troublesome
 - Karl Fogel (a CVS developer) was hired
 - “Replace CVS”
- Karl's friend Jim Blandy (another CVS dev) had a basic idea for “Inversion”
 - This became our repository “filesystem” design
- Renamed to Subversion in early 2000

February 2000

- Brian Behlendorf, the CollabNet CTO, invited me to join the project
 - Apache and WebDAV (DeltaV) experience
- Began updating IETF DeltaV specification for the CVS model of version control
 - DeltaV was used for 1.0 through 1.6

Early Life

- Initial planning in April and May, 2000
- Coding began in June, 2000
 - CVS on tigris.org
 - “Working copy” code came first
- Self-hosting in August, 2001
 - Missing features, but enough to move forward
- Nearly 60 releases before 1.0 in February, 2004

Pre-1.0 Principles

- Will a feature be included? “Is it in CVS?”
- “Ship when it is ready”
 - Larry McVoy, BitKeeper: “It takes longer than you think to build a version control system.”
- Trust: “\$100 million of IP checked-in”
- Design for large scale deployment

History

- Begins in June, 2000
- 1.0 in February, 2004
- 1.1 in September, 2004
- 1.2 in May, 2005
- 1.3 in December, 2005
- 1.4 in September, 2006
- 1.5 in June, 2008 (merge tracking)
- 1.6 in March, 2009

Subversion 1.7

- “Working copy” code has been rebuilt
 - Performance
 - Platform for future features
 - Refined features (eg. conflicts)
- Faster HTTP protocol (discarded DeltaV)
- Repository performance

- When? “When it is ready”

Enterprise Features

- Authorization
- Large, binary assets
- Massive repository sizes
- Sophisticated conflict handling
 - Tree conflicts from merge and update
- Multi-repository working copies (svn:externals)

Community Building

- From the start, we took Apache as a model
 - Open and inclusive
 - Involve people by giving commit privileges
 - Groomed into stakeholders
- Long-term health
 - Some people leave, some arrive

Testing

- Extensive testing
 - 1500 cases in the test suite
 - New cases added as issues are found/fixed
- Test suite was started very early
 - Part of the community's culture
- Continuous integration, build, and testing

Architecture

- Client/server design
 - Isolates server from problems on the client
 - Multiple access model to fit user needs
- C language, library-based model
 - Enables third-party ecosystem
 - Broad talent base

Review

- Strong principles
- Designed to scale
- Community as caretakers
- Testing culture
- Architect for robustness