

How a Large-Scale Open Source Project Works



Robert N. M. Watson

FreeBSD Foundation

Computer Laboratory
University of Cambridge

Introduction

- What do we mean by Open Source Project?
- FreeBSD Project as a case study:
 - What is the FreeBSD Project?
 - How does the project work?
 - And does it all depend on who you ask?



What is an Open Source project?

- Open source software: software under open source license
- But what is an open source project?
 - source code and its history?
 - people who write, support, distribute, use, or advocate the software?
- What is a large-scale open source project?
 - It's all about a sustainable community model



FreeBSD

- Open source BSD UNIX OS (1978, 1992)
- ISP network service platform
 - Yahoo!, Verio, NY Internet, ISC, Demon, ...
- Appliance/product/embedded OS foundation
 - Mac OS X, VXWorks, NetApp, Secure Computing, Nokia, nCircle, Isilon, Symmetricon, NetScaler, Juniper, Thomson, Panasas, Cisco, Palisade, Avid, The Weather Channel, Sandvine, Blue Coat, ...
- Can't use the Internet without using FreeBSD



What is FreeBSD?

- Complete, integrated UNIX system
 - Multi-processing, multi-threaded, preemptive kernel
 - Intel/AMD 32/64-bit, ia64, sparc64, ARM, PPC, MIPS
 - UNIX, POSIX, BSD programming interfaces
 - Multi-protocol network stack
 - IPv4, IPv6, IPX/SPX, AppleTalk, IPSEC, ATM, Bluetooth, 802.11, SCTP, ...
 - Unified, coherent build system across components
 - Extensive documentation
- 17,700 third party software ports



The FreeBSD Project

- Online development community
 - Central source repository and revision control
 - Extensive online community
 - 340 CVS committers, thousands of contributors
 - Extensive user community
- Liberal Berkeley open source license
 - Designed to maximize commercial reuse
 - No requirement that derived works be open source
 - Extensive use in commercial, research systems



FreeBSD Foundation

- Non-profit organization based in Boulder, CO
 - Intentionally independent from FreeBSD Project
 - Sponsored development
 - Intellectual property, contracts, licensing, legal
 - Developer travel grants, event sponsorship
 - Hardware purchase
 - Collaborative R&D agreements
- Support the FreeBSD Project
 - Donate online today!



What the Project Produces

- Integrated FreeBSD kernel, user space
- Security officer, release engineering
- Ports collection, binary packages
- FreeBSD releases
- Manuals, handbook, web pages, marketing
- Technical support, debugging, etc.
- A variety of user/community events



Things We Consume

- Beer, soda, chocolate, and other vices
- Donated and sponsored hardware
 - Especially in racks, with hands
- Bandwidth in vast and untold quantities
- Travel grants, salaries, contracts, grants
- Thanks, user testimonials, good press



FreeBSD People and Processes

- FreeBSD committers
- Core Team
- Ports committers and maintainers
- Groups/projects
- Derived projects
- Mailing lists
- Web sites
- Events
- Development cycle and branches, releases
- CVS and Perforce
- Clusters
- Conflict resolution



FreeBSD Committers

- Committer is someone with CVS commit rights
- Selected based on key characteristics
 - Technical expertise
 - History of contribution to the FreeBSD Project
 - Ability to work well in the community
 - Having made these properties obvious!
- Key concept: mentor
 - Mentor proposes to core@ (portmgr@, doceng@)
 - Guide through first few months of committing



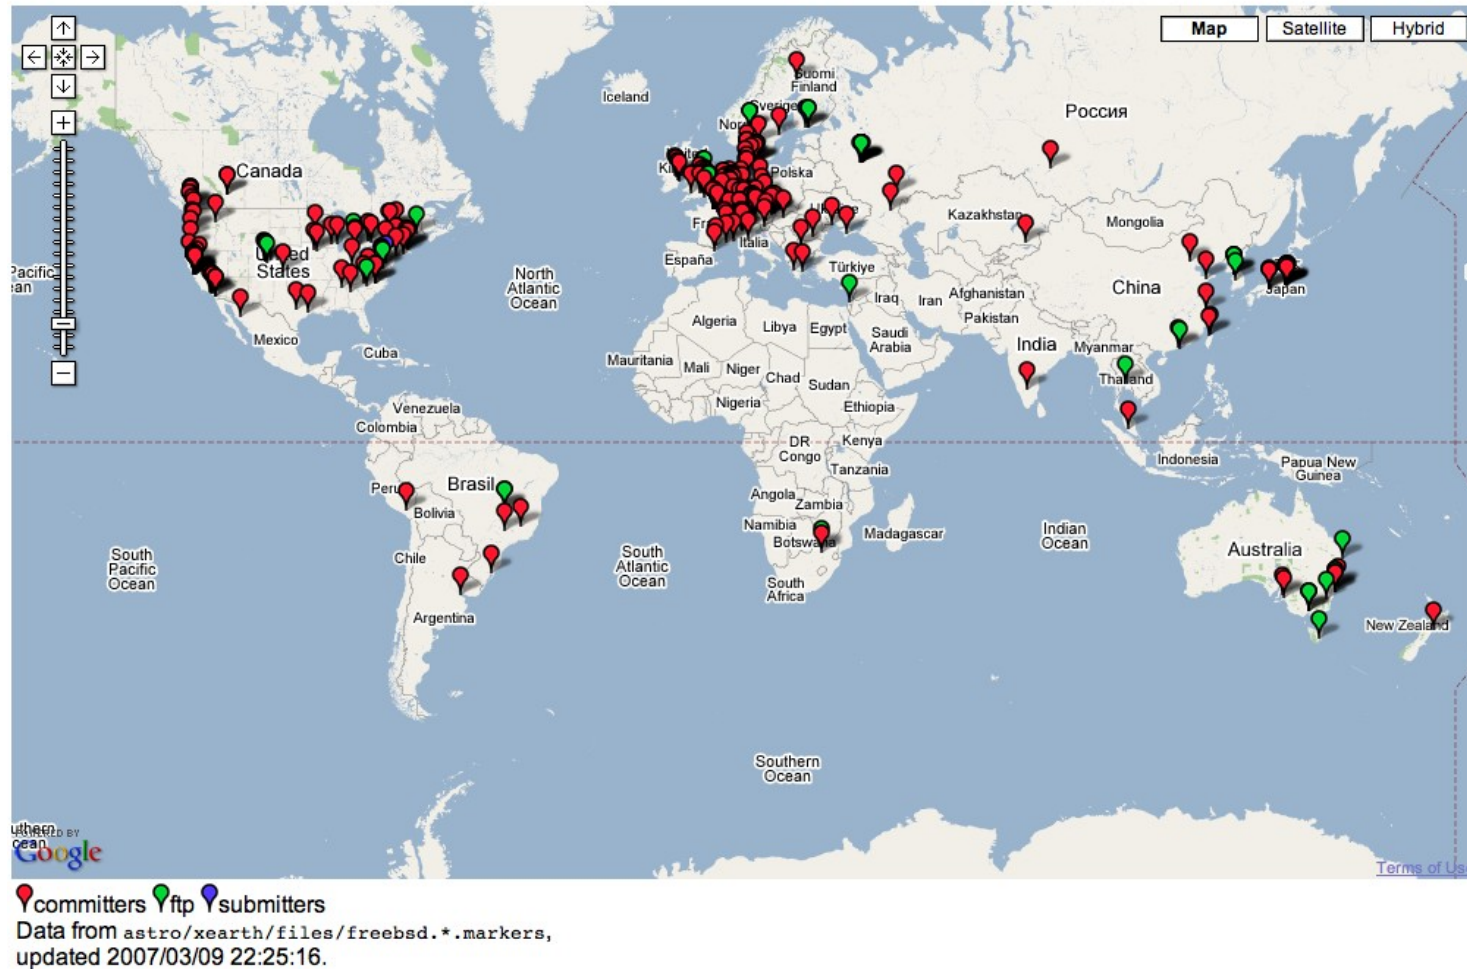
Who are the Committers? (2007)

- Locations
 - 34 countries
 - 6 continents
- Ages
 - Oldest (disclosed) committer born 1948
 - Youngest committer born 1989
 - Mean age 32.5, median age 31, stddev 7.3
- Professional programmers, hobbyists, consultants, university professors, students ...




Locations of FreeBSD Committers (March 2007)


FreeBSD locations



FreeBSD Developer Age Distribution (March 2007)



Number of Commit Bits by Type (March 2007)



FreeBSD Core Team

- Historically “key” developers
- Now 9-member elected management body
 - Votes and candidates from FreeBSD committers
 - Core secretary
- Responsibilities
 - Administrative (commit bits, hats, team charters)
 - Strategic (project direction, coordination, cajoling)
 - Rules, conflict resolution, enforcement

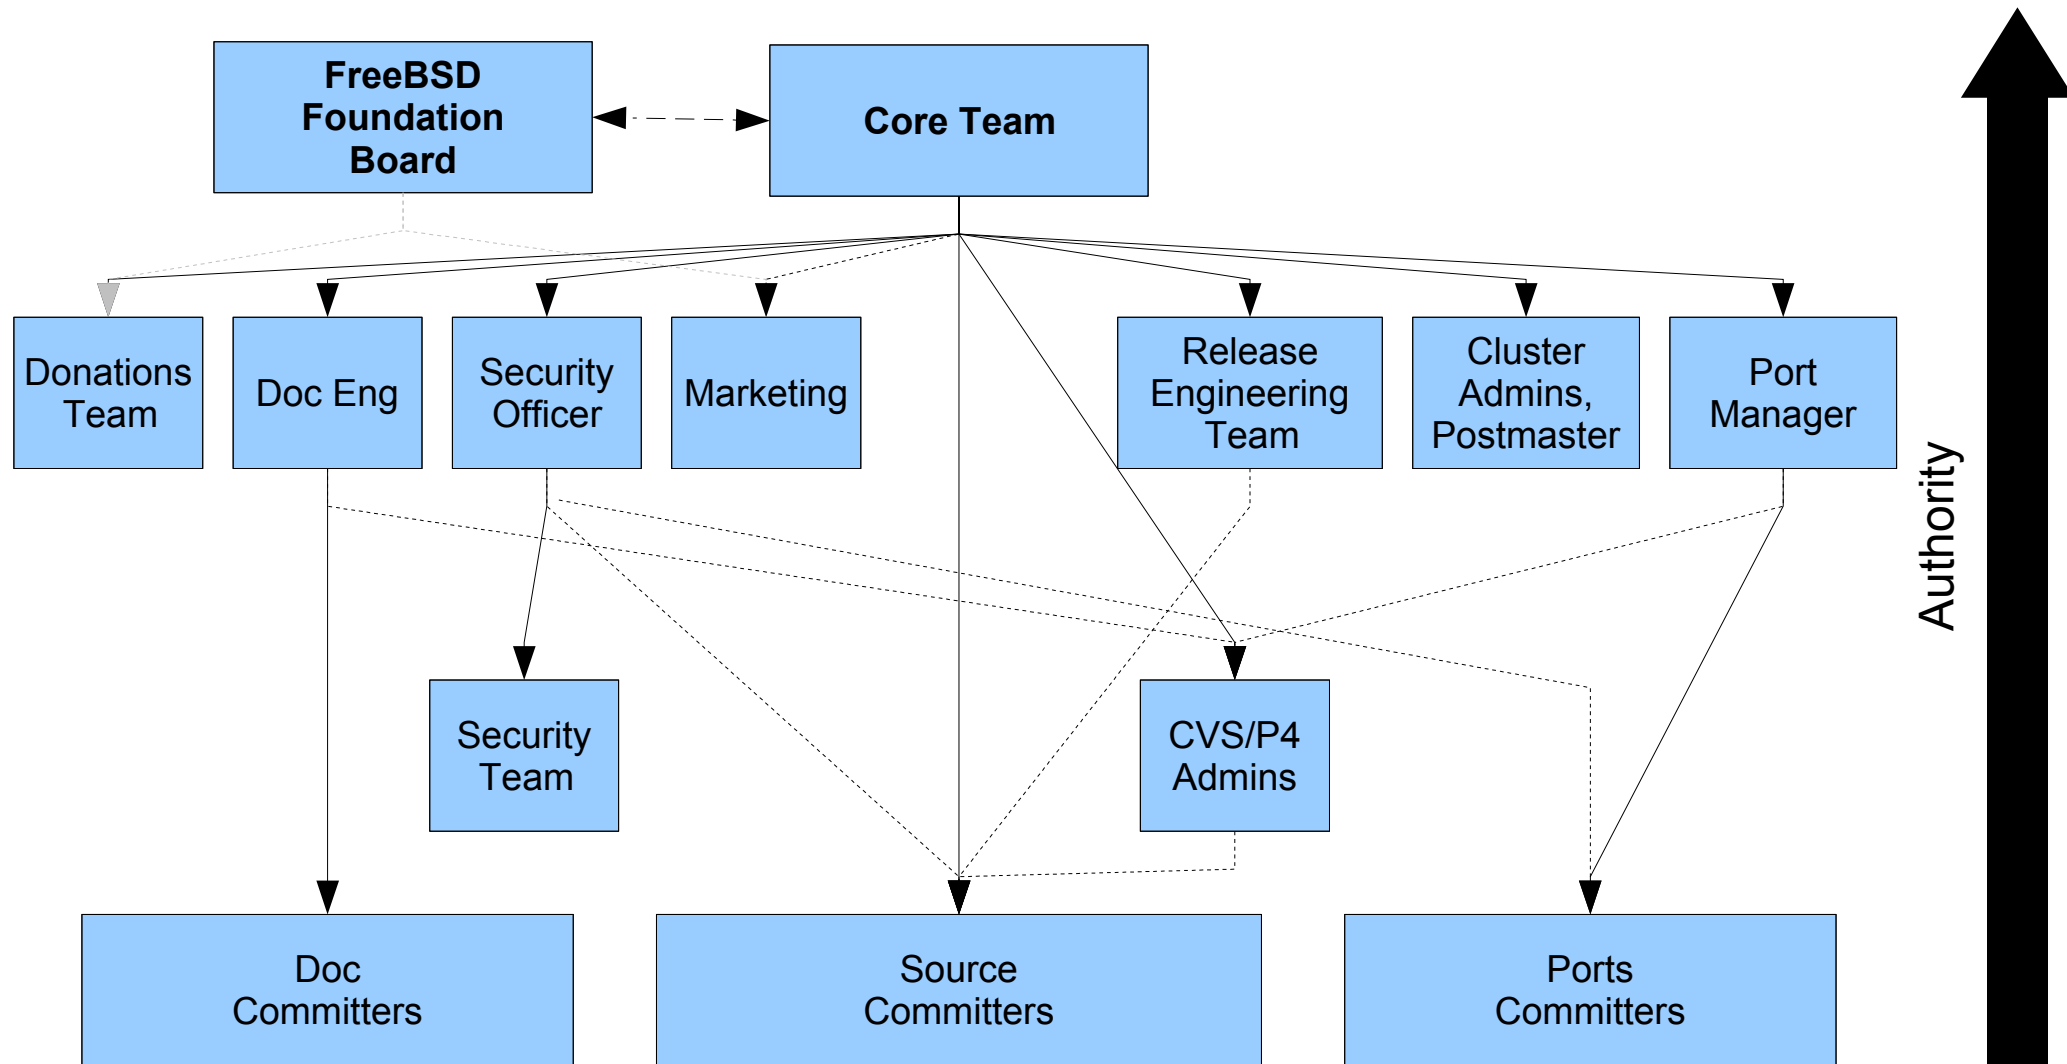


Ports Committers, Maintainers

- Framework for adapting and building third-party software for FreeBSD
- Snapshot (~2007)
 - 158 ports committers
 - Over 1,400 ports maintainers
 - Over 17,700 ports
 - 112 ports/commmitter
 - 12 ports/maintainer
 - 8 maintainers/commmitter



FreeBSD Project Org Chart



Groups and Projects

Development

- Source Developers
- Core Team
- Core Team Secretary
- Release Engineering Team
- Release Engineering Build Teams
- Security Officer
- Security Team
- Ports Team
- Port Managers
- Doceng Team
- Documentation Team
- Vendor Relations Team

Administrative

- Foundation Board of Directors
- Foundation Operations Manager
- FreeBSD.org admins@
- FreeBSD.org webmaster
- Sentex cluster admins
- ISC cluster admins
- Mirrors Team
- Donations Team

Administrative (cont)

- Marketing Team
- Perforce Admins
- CVS Admins
- Postmaster
- CVSUP Mirrors Team

Other Contributors

- Perforce Contributors
- Questions Subscribers
- Bugbusters

Software Adaptation Projects

- FreeBSD GNOME Project
- FreeBSD KDE Project
- Mono on FreeBSD
- OpenOffice.org on FreeBSD
- Java on FreeBSD
- X.org on FreeBSD

Special Projects

- Stress Testing
- FreeBSD Tinderbox

Special Projects (cont)

- Coverity Team
- FreeBSD Standards
- SoC Mentors
- Monthly Status Reports


External Projects

- KAME Project
- TrustedBSD Project
- PC-BSD
- DesktopBSD
- DragonflyBSD
- FreeNAS
- pfSense
- allbsd.org cluster

...



FreeBSD-centric View of the Open Source OS Ecosystem (simplified)

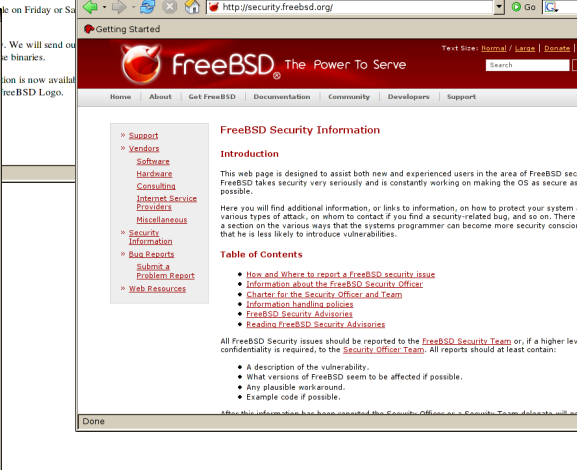




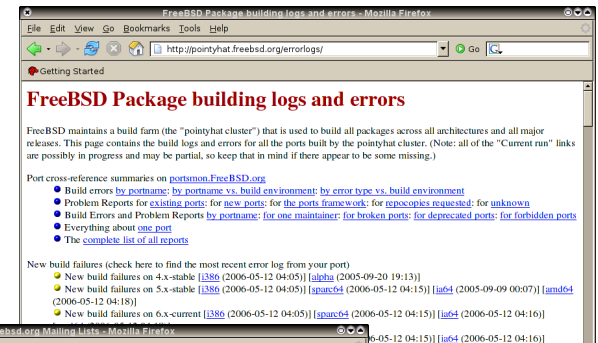


Mailing Lists

- Where the business of the project takes place
 - Developer and user lists
- Over 100 active topic mailing lists
 - -announce, -current, -arch, cvs-all, -security, ...
 - -chat, -hackers, -questions...
- Mostly public; a few private
 - E.g., security-officer



FreeBSD Project Web Pages (Just a few)



Events

- Conferences
 - USENIX ATC
 - BSDCan
 - BSDCon
 - EuroBSDCon
 - AsiaBSDCon
 - NYCBSDCon
 - MeetBSD
 - BSDConTR
- Developer Summits
 - Two day events
 - September 2007, EuroBSDCon, Copenhagen
 - March, 2008: AsiaBSDCon 2008, Tokyo, Japan
 - May: BSDCan 2008, Ottawa, Canada



FreeBSD Developer Summits

BSDCan 2006-2007




Recent Development Projects

- DTrace
- Network virtualization
- Xen
- Sun4v
- SCTP
- 32-core scalability
- Multi-threaded, multi-processor network stack
- 802.11n + Virtual AP
- ARM, MIPS, PPC
- Security event audit
- MAC Framework
- ZFS, GJournal
- gcc 4.2
- Coverity
- 10gbps optimization



Branched Development Model



- Concurrent development
- Divergence based on feature maturity
- MFC's
- Errata/security branches
- 18-24 month “dot zero” cycle



FreeBSD 7.0 due ~~2007Q4~~ really soon

- MP Scalability
 - Locking primitives
 - ULE2 scheduler
 - New threading library
 - jemalloc
 - Network stack/IPC
 - File systems
 - Sun's ZFS file system
 - GJournal for UFS
 - Sun4v
 - Networking
 - Direct dispatch
 - 10gbps optimizations + many drivers
 - SCTP
 - And much more ...
- Attend Kris Kenaway's talk on FreeBSD 7.0 in the BSD + PostgreSQL track to learn more!**



CVS

- Primary revision control system
 - Authoritative project activity is in CVS
 - Actually four repositories
 - /home/ncvs – FreeBSD src cvs
 - /home/pcvs – FreeBSD ports cvs
 - /home/projcv – FreeBSD project cvs
 - /home/dcv – FreeBSD documentation cvs
 - 10+ year revision history
 - One commit every 11.8 minutes for last three years
- Technical limitations becoming more apparent

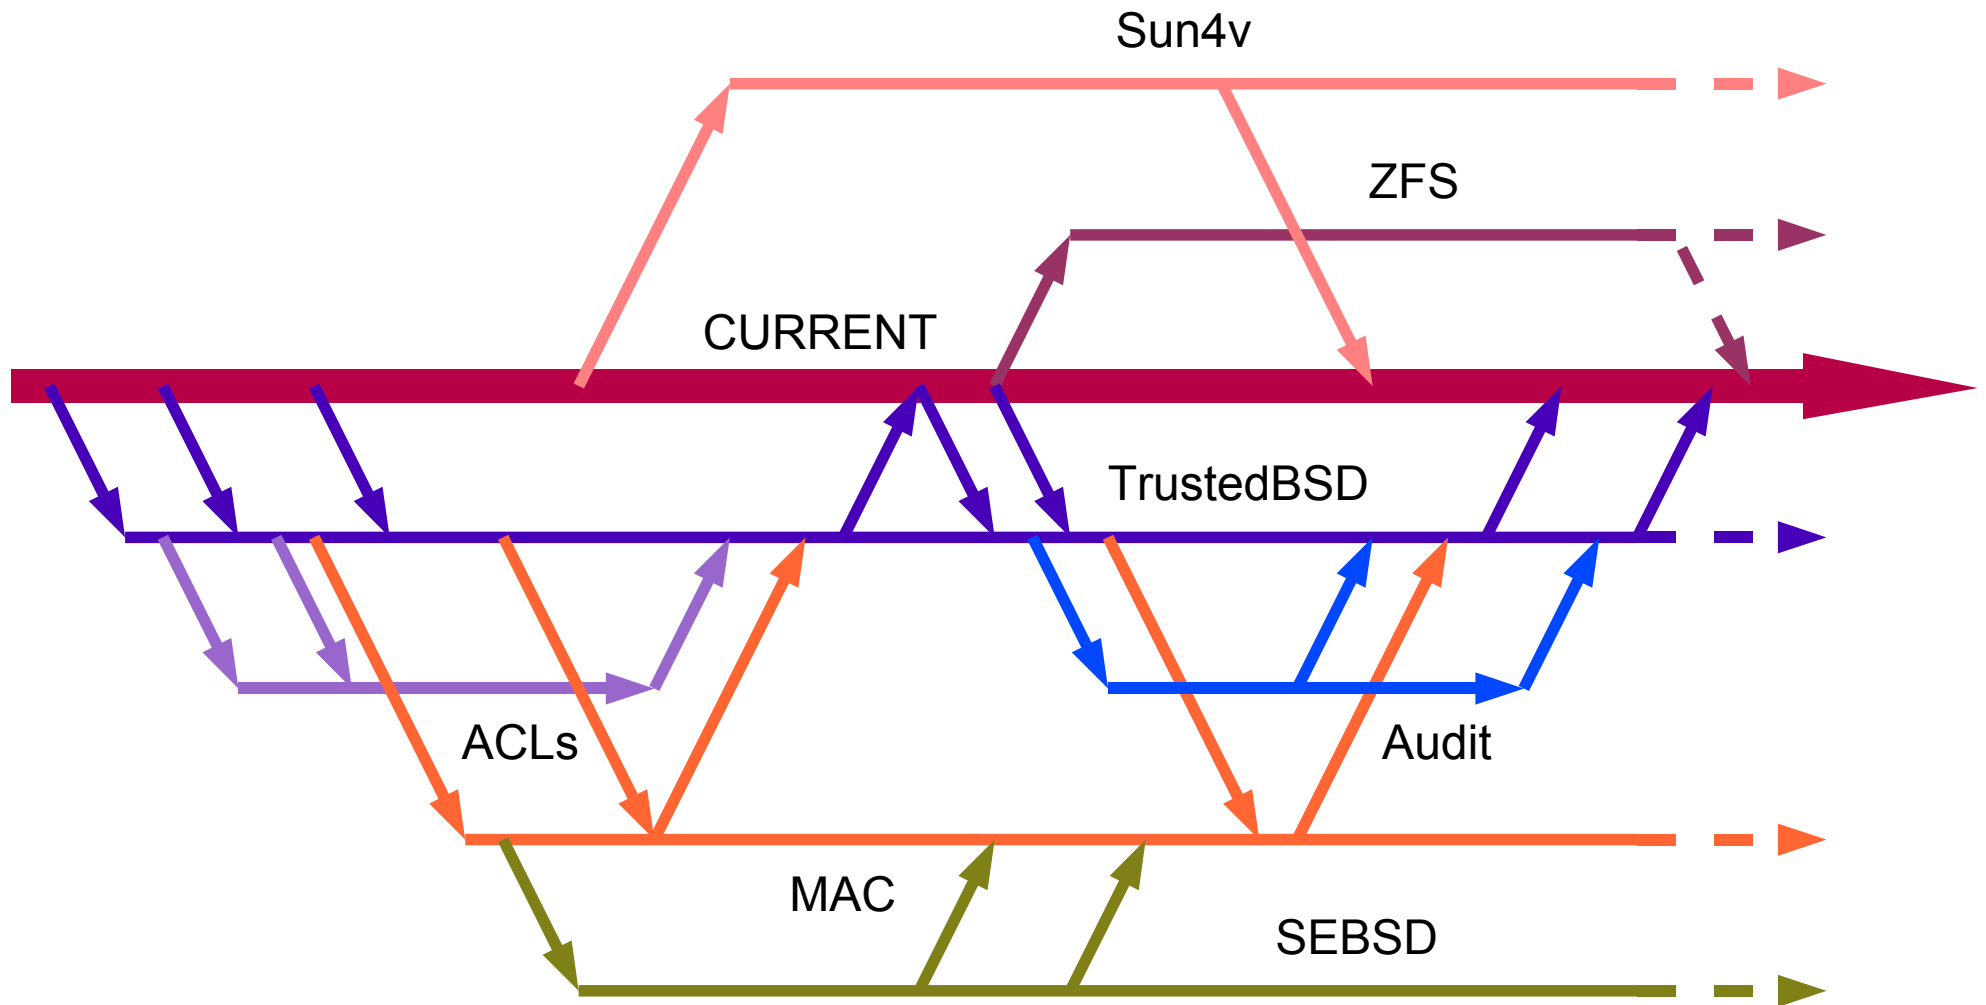


Perforce

- Secondary revision control system
 - Supports heavily branched development
 - FreeBSD developers
 - Guest accounts and accounts
- Dozens of active projects, including:
 - SMPng, TrustedBSD Audit+MAC+SEBSD, Superpages, uart, ARM, MIPS, Summer of Code, DTrace, Xen, Sun4v, GEOM, GJournal, AFS, ZFS, highly parallel network stack, ...



Perforce Development Branches

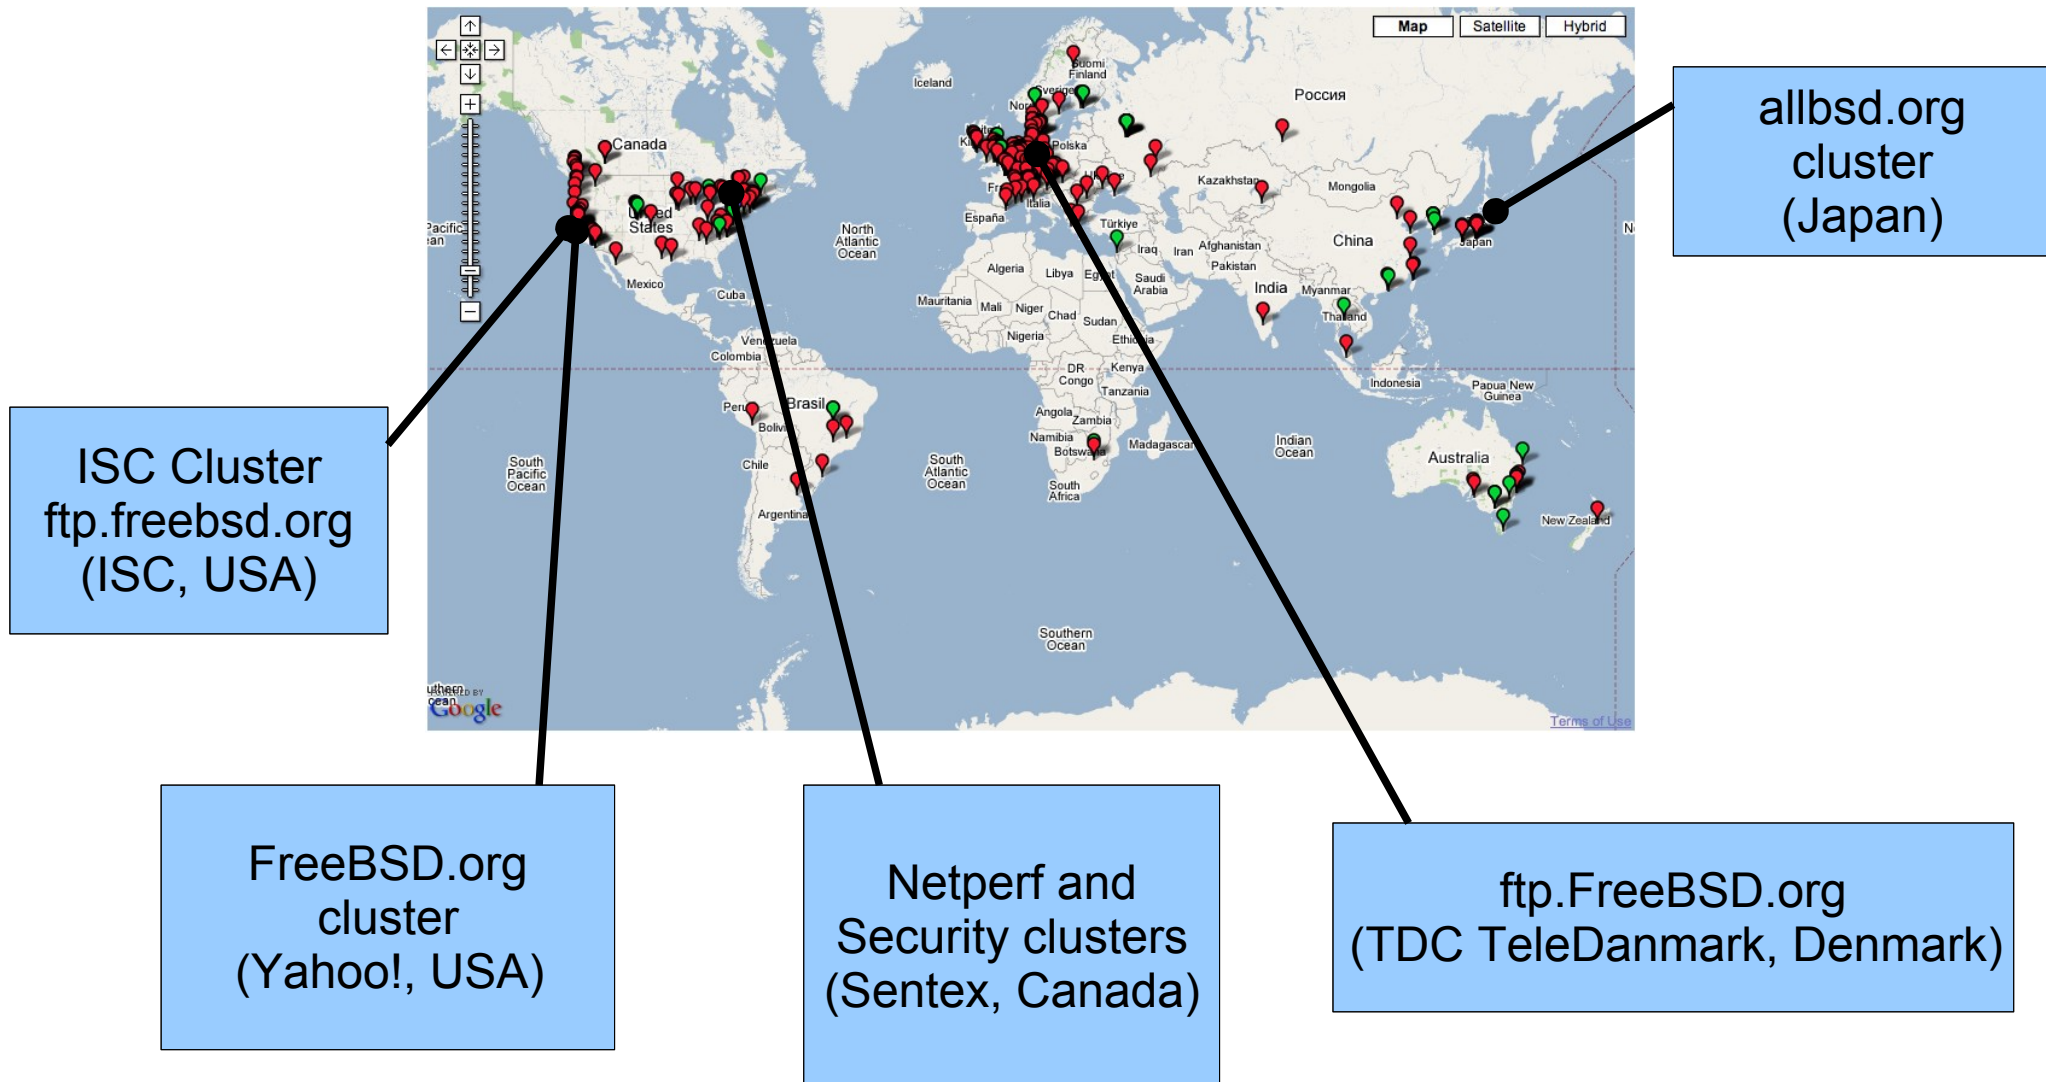


Revision Control: the Future

- Perforce use symptom of CVS weaknesses
 - Change sets, lightweight branching, history-aware merging, access control
- Every few years, reconsider options
 - Cost of migration very high
 - Interrupts development, retrain developers, etc.
- Currently evaluating several systems
- Sticking points for most scaling, obliteration



FreeBSD.org Clusters



Conflict Resolution

- Developers generally characterized by:
 - Independence
 - Cooperation
 - Common sense
- Facilitated by intentional avoidance of overlap
- Strong technical disagreements, personality conflicts, etc, can and will occur
- When they get out of hand, generally mediated by a member of core



Conclusion

- Just a glimpse into one of the largest, oldest, and most successful open source projects
 - Hundreds of committers, thousands of contributors
 - Millions of lines of code
 - Tens of millions of deployed systems
- Highly successful community model, not just good code, makes this possible
- <http://www.FreeBSD.org/>

