

The FreeBSD Project: A Replication Case Study of Open Source Development

Peter C. Rigby

September 14, 2006



What were the processes used to develop Free BSD.

- OS derived from BSD Unix
- Two branches current and Stable
- Core member, committer, contributor
 - 2 year terms 7-9 member (Core)
 - Mentor process for new committer
- PR has standard fields
- Testing (Countdown period)
- Inspection
 - Active maintainer
 - Significant change

Release process and defect density

- 45 day prior, 15 days to merge
- 30 day prior, 15 day review of new code
- 15 days prior, code freeze and testing

TABLE 3

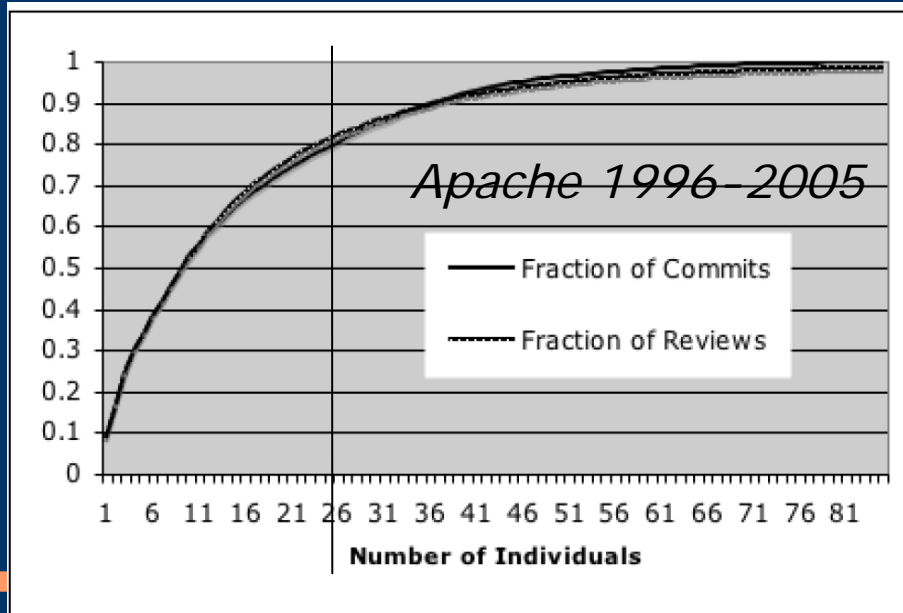
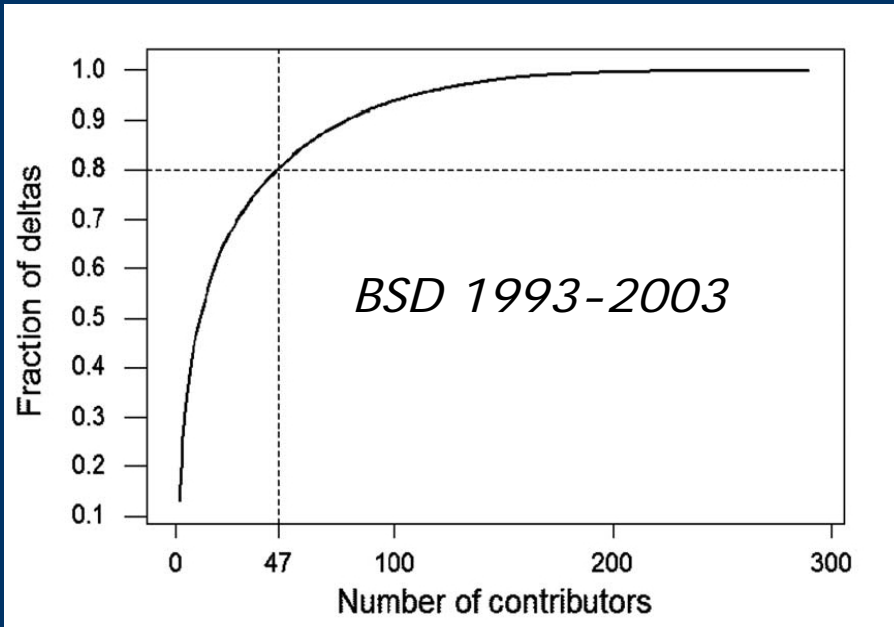
Defect Densities in FreeBSD, Apache, and Four Commercial Systems

Measure	FreeBSD	Apache	A	C	D	E
Post-release defects/KLOC	0.55	2.64	0.11	0.1	0.7	0.1
Post-release defects/Kdelta	11.17	40.8	4.3	14	28	10
Post-feature defects/KLOC	1.89	2.64	*	5.7	6.0	6.9
Post-feature defects/Kdelta	38.58	40.8	*	164	196	256

Contributed 80% of Code

Project	Dev	Kdeltas	TopDev	% TopDev
FreeBSD (94-96)	212	180	28	13.2
FreeBSD (97-99)	161	172	27	16.7
FreeBSD (00-02)	265	174	42	15.8
Apache	388	15	18	4.6
Mozilla (layout)	174	42	35	20.1
Mozilla (js)	127	14	24	18.9
Mozilla (rdf)	123	12	26	21.1
Mozilla (network)	106	10	24	22.6
Mozilla (editor)	118	8	25	21.2
Mozilla (intl)	87	5	22	25.2
Mozilla (xpininstall)	102	5	22	21.5

Why is 3 years the magic number?



- There isn't a trend 27, 28, 42!
- Only justification is “Mockus did it that way”
- Who is responsible for the code?
- Developers turn-over. Name resolution.

Data and Tools

```
[1]=====
[2]RCS file: FreeBSD-CVS/src/bin/cp/cp.c,v
[3]-----
[4]revision 1.40
[5]date: 2002/07/23 00:42:56 author:aname;
    state: Exp: lines +9-5
[6]Add -n option, which answers "no"
    to the overwrite question.
[7]
[8]PR      7828
[9]Suggested by: Author Name <email address>
[10]Approved by: Mentor Name (mentor)
[11]MFC after: 2 weeks
[12]-----
```

netchild 2006-06-18 17:54:33
UTC

FreeBSD src repository

Modified files:

sys/dev/usb u modem.c

Log:

Add support for:

*- Curitel PC5740 Wireless Modem
(Verizon's PCMCIA card)*

*- Sierra MC5720 Wireless Modem
(Built in to Thinkpad X60s)*

PR: 98908

*Submitted by: Michael Collette
<metrol@metrol.net>*

Revision Changes Path

1.58 +2 -0

src/sys/dev/usb/u modem.c

Name resolution and PRs

Dean gaudet's identity: (there is at least one other dean on the project)
dean@arctic.org; dean gaudet

Ones that I feel should have been automatically added:

dgaudet@iacnet.com

dgaudet@arctic.org

dean-list-new-httpd@arctic.org

dgaudet@hotwired.com

dgaudet@wired.com

dgaudet@hyperreal.org

dgaudet-list-linux-kernel@arctic.org

See Bird's paper on how to do this properly.

You cannot talk about email without doing name resolution first!

The number of people to send PRs came from a mailing list.

Hypothesis 7

- “There will be rapid responses to customer problems.
- Unfortunately, we do not have enough data yet to evaluate H7.”
- What was the real reason?
- They didn't extract from GNATs just the incomplete bug mailing list
- This should be hard to do.
 - Look for the pr in a commit, and find out when it was submitted.

Validity

- Total deltas is useless.
Deltas per year, or
KLOC
- Time periods are not
comparable.

Comparisons between FreeBSD, Apache,
and Four Commercial Systems

Project	K Deltas	Years	Developers
FreeBSD	528	10	354
Apache	18	3	388
A	129	3	101
C	2.8	1.3	17
D	0.7	1.7	8
E	2.4	1.5	16