Brian Behlendorf
Apache

Scott Bradner
Internet Engineering Task Force

Jim Hamerly
Netscape

Kirk McKusick
Berkeley Unix

OPEN SOURCES

Voices from the Open Source Revolution

Tim O'Reilly
O'Reilly & Associates, Inc.

Tom Paquin
mozilla.org

Bruce Perens
Open Source Initiative

Eric Raymond
Open Source Initiative

Richard Stallman
Free Software Foundation

Michael Tiemann
Cygnus Solutions

Linus Torvalds
Linux

Paul Vixie
Bind

Larry Wall

Open Sources: Voices fro
Open Source Revolution

1st Edition January 1999
1-56592-582-3, Order Number: 5823
280 pages, $24.95

Twenty Years of Berkeley Unix

From AT&T-Owned to Freely Redistri

Marshall Kirk McKusick

Early History

Ken Thompson and Dennis Ritchie presented the first Unix paper a
Symposium on Operating Systems Principles at Purdue University
1973. Professor Bob Fabry, of the University of California at Berke
attendance and immediately became interested in obtaining a copy
to experiment with at Berkeley.

At the time, Berkeley had only large mainframe computer systems
processing, so the first order of business was to get a PDP-11/45 su

Although Ken Thompson at Purdue was not involved in the installa
Berkeley as he had been for most systems up to that time, his exper
needed to determine the cause of several strange system crashes. Be
Berkeley had only a 300-baud acoustic-coupled modem without aut
capability, Thompson would call Standiford in the machine room ar
insert the phone into the modem; in this way Thompson was able to
debug crash dumps from New Jersey.

Many of the crashes were caused by the disk controller's inability to
overlapped seeks, contrary to the documentation. Berkeley's 11/45 w
the first systems that Thompson had encountered that had two disks
controller! Thompson's remote debugging was the first example of
cooperation that sprang up between Berkeley and Bell Labs. The w
the researchers at the Labs to share their work with Berkeley was in
the rapid improvement of the software available at Berkeley.

Though Unix was soon reliably up and running, the coalition of Co
Science, Mathematics, and Statistics began to run into problems; M
Statistics wanted to run DEC's RSTS system. After much debate, a
was reached in which each department would get an eight-hour shif
would run for eight hours followed by sixteen hours of RSTS. To p
fairness, the time slices were rotated each day. Thus, Unix ran 8 a.m
one day, 4 p.m. to midnight the next day, and midnight to 8 a.m. the
Despite the bizarre schedule, students taking the Operating Systems
preferred to do their projects on Unix rather than on the batch mach

Professors Eugene Wong and Michael Stonebraker were both stymi
confinements of the batch environment, so their INGRES database
among the first groups to move from the batch machines to the inte
environment provided by Unix. They quickly found the shortage of
time and the odd hours on the 11/45 intolerable, so in the spring of
purchased an 11/40 running the newly available Version 5. With the
distribution of INGRES in the fall of 1974, the INGRES project bec
group in the Computer Science department to distribute their softwa

instructional 11/45's for the Computer Science department's own us
1975, the money was obtained. At nearly the same time, DEC anno
11/70, a machine that appeared to be much superior to the 11/45. M
two 11/45s was pooled to buy a single 11/70 that arrived in the fall
Coincident with the arrival of the 11/70, Ken Thompson decided to
year sabbatical as a visiting professor at the University of Californi
his alma mater. Thompson, together with Jeff Schriebman and Bob
brought up the latest Unix, Version 6, on the newly installed 11/70.

Also arriving in the fall of 1975 were two unnoticed graduate stude
and Chuck Haley; they both took an immediate interest in the new s
Initially they began working on a Pascal system that Thompson had
together while hanging around the 11/70 machine room. They expa
improved the Pascal interpreter to the point that it became the progr
system of choice for students because of its excellent error recovery
fast compile and execute time.

With the replacement of Model 33 teletypes by ADM-3 screen term
and Haley began to feel stymied by the constraints of the *ed* editor.
from an editor named *em* that they had obtained from Professor Geo
Coulouris at Queen Mary's College in London, they worked to prod
at-a-time editor *ex*.

With Ken Thompson's departure at the end of the summer of 1976,
Haley begin to take an interest in exploring the internals of the Unix
Under Schriebman's watchful eye, they first installed the fixes and
improvements provided on the "fifty changes" tape from Bell Labs.
learned to maneuver through the source code, they suggested severa
enhancements to streamline certain kernel bottlenecks.

Early Distributions

Meanwhile, interest in the error recovery work in the Pascal compil
requests for copies of the system. Early in 1977, Joy put together th

Joy was finally able to write *vi*, bringing screen-based editing to Be soon found himself in a quandary. As is frequently the case in unive strapped for money, old equipment is never replaced all at once. Ra support code for optimizing the updating of several different termin decided to consolidate the screen management by using a small inte redraw the screen. This interpreter was driven by a description of th characteristics, an effort that eventually became *termcap*.

By mid-1978, the software distribution clearly needed to be updated system had been made markedly more robust through feedback fro expanding user community, and had been split into two passes so th run on PDP-11/34s. The result of the update was the "Second Berke Distribution," a name that was quickly shortened to 2BSD. Along w enhanced Pascal system, *vi* and *termcap* for several terminals was i Once again Bill Joy single-handedly put together distributions, ansv phone, and incorporated user feedback into the system. Over the ne seventy-five tapes were shipped. Though Joy moved on to other pro following year, the 2BSD distribution continued to expand. The fin this distribution, 2.11BSD, is a complete system used on hundreds still running in various corners of the world.

VAX Unix

Early in 1978, Professor Richard Fateman began looking for a mach larger address space on which he could continue his work on Macsy (originally started on a PDP-10). The newly announced VAX-11/78 the requirements and was available within budget. Fateman and thir faculty members put together an NSF proposal that they combined departmental funds to purchase a VAX.

Initially the VAX ran DEC's operating system VMS, but the depart gotten used to the Unix environment and wanted to continue using after the arrival of the VAX, Fateman obtained a copy of the 32/V p

megabyte on the new VAX.

To alleviate this problem, Fateman approached Professor Domenico
member of the systems faculty at Berkeley, to investigate the possib
having his group write a virtual memory system for Unix. Ozalp Ba
of Ferrari's students, set about to find some way of implementing a
paging system on the VAX; his task was complicated because the V
reference bits.

As Babaoglu neared the completion of his first cut at an implementa
approached Bill Joy for some help in understanding the intricacies o
kernel. Intrigued by Babaoglu's approach, Joy joined in helping to i
code into 32/V and then with the ensuing debugging.

Unfortunately, Berkeley had only a single VAX for both system de
and general production use. Thus, for several weeks over the Christ
the tolerant user community alternately found themselves logging in
"Virtual VAX/Unix." Often their work on the latter system would c
abrupt halt, followed several minutes later by a 32/V login prompt.
1979, most of the bugs had been worked out, and 32/V had been rel
history.

Joy saw that the 32-bit VAX would soon make the 16-bit PDP-11 o
began to port the 2BSD software to the VAX. While Peter Kessler a
the Pascal system, Joy ported the editors *ex* and *vi*, the C shell, and
other smaller programs from the 2BSD distribution. By the end of 1
complete distribution had been put together. This distribution inclu
virtual memory kernel, the standard 32/V utilities, and the additions
In December, 1979, Joy shipped the first of nearly a hundred copies
the first VAX distribution from Berkeley.

The final release from Bell Laboratories was 32/V; thereafter all Un
from AT&T, initially System III and later System V, were managed
different group that emphasized stable commercial releases. With th
commercialization of Unix, the researchers at Bell Laboratories wer

DARPA Support

Meanwhile, in the offices of the planners for the Defense Advanced
Projects Agency (DARPA), discussions were being held that would
influence on the work at Berkeley. One of DARPA's early successe
set up a nationwide computer network to link together all their majo
centers. At that time, they were finding that many of the computers
centers were reaching the end of their useful lifetime and had to be
heaviest cost of replacement was the porting of the research softwa
machines. In addition, many sites were unable to share their softwa
the diversity of hardware and operating systems.

Choosing a single hardware vendor was impractical because of the
varying computing needs of the research groups and the undesirabil
depending on a single manufacturer. Thus, the planners at DARPA
the best solution was to unify at the operating systems level. After
discussion, Unix was chosen as a standard because of its proven po

In the fall of 1979, Bob Fabry responded to DARPA's interest in m
towards Unix by writing a proposal suggesting that Berkeley develo
enhanced version of 3BSD for the use of the DARPA community. I
copy of his proposal to a meeting of DARPA image processing and
contractors, plus representatives from Bolt, Beranek, and Newman,
developers of the ARPAnet. There was some reservation whether B
could produce a working system; however, the release of 3BSD in
1979 assuaged most of the doubts.

With the increasingly good reputation of the 3BSD release to valida
Bob Fabry was able to get an 18-month contract with DARPA begi
April 1980. This contract was to add features needed by the DARP
Under the auspices of this contract, Bob Fabry sets up an organizati
was christened the Computer Systems Research Group, or CSRG fo
immediately hired Laura Tong to handle the project administration.
his attention to finding a project leader to manage the software deve

The project started promptly. Tong set up a distribution system that a higher volume of orders than Joy's previous distributions. Fabry n coordinate with Bob Guffy at AT&T and lawyers at the University to formally release Unix under terms agreeable to all. Joy incorpora Kulp's job control, and added auto reboot, a 1K block file system, a for the latest VAX machine, the VAX-11/750. By October 1980, a distribution that also included the Pascal compiler, the Franz Lisp s enhanced mail handling system was released as 4BSD. During its n lifetime, nearly 150 copies were shipped. The license arrangement institution basis rather than a per machine basis; thus the distributio about 500 machines.

With the increasingly wide distribution and visibility of Berkeley U critics began to emerge. David Kashtan at Stanford Research Institu paper describing the results of benchmarks he had run on both VMS Berkeley Unix. These benchmarks showed severe performance prol the Unix system for the VAX. Setting his future plans aside for sev Joy systematically began tuning up the kernel. Within weeks he hac paper written showing that Kashtan's benchmarks could be made to on Unix as they could on VMS.

Rather than continue shipping 4BSD, the tuned-up system, with the Robert Elz's auto configuration code, was released as 4.1BSD in Ju Over its two- year lifetime about 400 distributions were shipped. Th intent had been to call it the 5BSD release; however, there were obj AT&T that there would be customer confusion between their comm release, System V, and a Berkeley release named 5BSD. So, to reso Berkeley agreed to change the naming scheme for future releases to 4BSD and just increment the minor number.

4.2BSD

With the release of 4.1BSD, much of the furor over performance di

Based on the needs of the DARPA community, goals were set and w to define the modifications to the system. In particular, the new syst expected to include a faster file system that would raise throughput of available disk technology, support processes with multi-gigabyte space requirements, provide flexible interprocess communication fa allow researchers to do work in distributed systems, and would inte networking support so that machines running the new system could participate in the ARPAnet.

To assist in defining the new system, Duane Adams, Berkeley's con at DARPA, formed a group known as the "steering committee" to h design work and ensure that the research community's needs were a This committee met twice a year between April 1981 and June 198 Bob Fabry, Bill Joy, and Sam Leffler of the University of Californi Alan Nemeth and Rob Gurwitz of Bolt, Beranek, and Newman; De of Bell Laboratories; Keith Lantz of Stanford University; Rick Rash Carnegie-Mellon University; Bert Halstead of the Massachusetts In Technology; Dan Lynch of The Information Sciences Institute; Dua and Bob Baker of DARPA; and Jerry Popek of the University of Ca Los Angeles. Beginning in 1984, these meetings were supplanted b that were expanded to include many more people.

An initial document proposing facilities to be included in the new s circulated to the steering committee and other people outside Berke 1981, sparking many lengthy debates. In the summer of 1981, I bec with the CSRG and took on the implementation of the new file syst the summer, Joy concentrated on implementing a prototype version interprocess communication facilities. In the fall of 1981, Sam Leff CSRG as a full-time staff member to work with Bill Joy.

When Rob Gurwitz released an early implementation of the TCP/IP Berkeley, Joy integrated it into the system and tuned its performanc this work, it became clear to Joy and Leffler that the new system w provide support for more than just the DARPA standard network pr

and rwho were intended to be temporary tools that would eventually by more reasonable facilities (hence the use of the distinguishing "r This system, called 4.1a, was first distributed in April 1982 for loca never intended that it would have wide circulation, though bootleg system proliferated as sites grew impatient waiting for the 4.2 relea

The 4.1a system was obsolete long before it was complete. Howeve from its users provided valuable information that was used to create proposal for the new system called the "4.2BSD System Manual." T document was circulated in February 1982 and contained a concise of the proposed user interfaces to the system facilities that were to b implemented in 4.2BSD.

Concurrent with the 4.1a development, I completed the implementa new file system, and by June 1982, had fully integrated it into the 4 The resulting system was called 4.1b and ran on only a few select d machines at Berkeley. Joy felt that with significant impending chan system, it was best to avoid even a local distribution, particularly si required every machine's file systems to be dumped and restored to 4.1a to 4.1b. Once the file system proved to be stable, Leffler proce the new file system related system calls, while Joy worked on revis interprocess communication facilities.

In the late spring of 1982, Joy announced he was joining Sun Micro Over the summer, he split his time between Sun and Berkeley, spen his time polishing his revisions to the interprocess communication f reorganizing the Unix kernel sources to isolate machine dependenci Joy's departure, Leffler took over responsibility for completing the Certain deadlines had already been established and the release had b promised to the DARPA community for the spring of 1983. Given t constraints, the work remaining to complete the release was evaluat priorities were set. In particular, the virtual memory enhancements sophisticated parts of the interprocess communication design were low priority (and later shelved completely). Also, with the impleme

than a year old and the Unix community's expectations heightened,
decided an intermediate release should be put together to hold peop
final system could be completed. This system, called 4.1c, was distr
April 1983; many vendors used this release to prepare for ports of 4
hardware. Pauline Schwartz was hired to take over the distribution
with the 4.1c release.