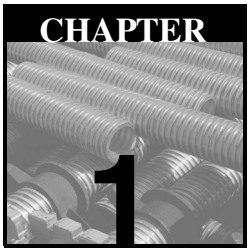

UNIX Evolution and Standardization

This chapter introduces UNIX from a historical perspective, showing how the various UNIX versions have evolved over the years since the very first implementation in 1969 to the present day. The chapter also traces the history of the different attempts at standardization that have produced widely adopted standards such as POSIX and the Single UNIX Specification.

The material presented here is not intended to document all of the UNIX variants, but rather describes the early UNIX implementations along with those companies and bodies that have had a major impact on the direction and evolution of UNIX.

A Brief Walk through Time

There are numerous events in the computer industry that have occurred since UNIX started life as a small project in Bell Labs in 1969. UNIX history has been largely influenced by Bell Labs' Research Editions of UNIX, AT&T's System V UNIX, Berkeley's Software Distribution (BSD), and Sun Microsystems' SunOS and Solaris operating systems.

The following list shows the major events that have happened throughout the history of UNIX. Later sections describe some of these events in more detail.

UNIX History Page 2 Tuesday, December 10, 2002 7:02 PM

1969. Development on UNIX starts in AT&T's Bell Labs.

1971. 1st Edition UNIX is released.

1973. 4th Edition UNIX is released. This is the first version of UNIX that had the kernel written in C.

1974. Ken Thompson and Dennis Ritchie publish their classic paper, "The UNIX Timesharing System" [RITC74].

1975. 6th Edition, also called V6 UNIX, becomes the first version of UNIX to be used outside Bell Labs. The University of California at Berkeley starts development on the *Berkeley Software Distribution* or more commonly called BSD.

1977. At this stage there were 500 sites running UNIX. Universities accounted for about 20 percent of those sites.

1979. 7th Edition UNIX was rewritten to make it more portable. Microsoft licenses 7th Edition and starts development of Xenix.

1980. Microsoft releases Xenix, a PC-based version of UNIX.

1982. AT&T's UNIX Systems Group releases System III UNIX. The Santa Cruz Operation (SCO) licenses Xenix from Microsoft.

1983. AT&T's UNIX System Development Labs release System V Release 1 UNIX.

1984. 4.2BSD is released including TCP/IP. System V Release 2 is released and the number of installations of UNIX worldwide exceeds 100,000. Digital Equipment Corporation's (DEC's) 4.2BSD-based Ultrix is released.

1986. 4.3BSD is released. 4.2BSD-based HP-UX first appears. IBM releases AIX 2 for the RT server.

1987. AT&T releases System V Release 3, which includes STREAMS, the Network File System (NFS), and the Transport Level Interface (TLI).

1989. As a joint venture between AT&T's Unix System Laboratories (USL) and Sun Microsystems, System V Release 4.0 is released.

1990. Based on SVR2 with enhancements from 4.2BSD and 4.3BSD, IBM releases AIX 3.1.

1991. Linus Torvalds announces Linux 0.0.1.

1992. USL releases System V Release 4.2 that includes the VERITAS filesystem VxFS and Volume Manager VxVM.

1993. 4.4BSD, the last release from Berkeley, is released. SVR4.2MP is released by Novell following their purchase of USL from AT&T.

1994. 4.4BSD Lite, which was free of copyrighted UNIX source code, is released.

1995. SCO buys Novell's UNIX business.

1996. The Linux 2.0 kernel is released.

1997. UnixWare 7, a merge of SVR4.2MP and SCO OpenServer, is released.

UNIX History Page 3 Tuesday, December 10, 2002 7:02 PM

2001. SCO's UNIX business is sold to Linux distributor Caldera. The Linux 2.4 kernel emerges after many delays.

How Many Versions of UNIX Are There?

Most versions of UNIX have stemmed from System V or BSD, with many taking enhancements from both. The 1980s saw a proliferation of versions of UNIX. Although it is difficult to give an exact figure on how many versions of UNIX have actually been produced, by the late 1980s it is safe to say that there were close to 100 different UNIX variants employed throughout the world. It is no wonder why UNIX has had its critics and not surprising that attempts at standardization have proceeded for much of UNIX's history.

Developing operating systems is a costly business, however, which has resulted in consolidation throughout the mid to late 1990s. On the low end, Linux and SCO variants tend to dominate while in the mid to high-end, Sun's Solaris, Hewlett Packard's HP-UX, and IBM's AIX account for most of the market share.

As time goes by there is likely to be further consolidation at the low to mid-range of PCs, workstations, and servers with Linux clearly leading the way. UNIX is still a clear leader at the high end of the market with Sun, HP, and IBM all having enterprise level capabilities in their UNIX offerings. While it is difficult to see further consolidation at this end of the market, only time will tell.

Why Is UNIX So Successful?

Although the large number of different versions of UNIX can be considered a weakness since it raised the possibility of incompatibilities, it also demonstrates one of the great strengths of UNIX: the ease by which UNIX has been ported to a wide number of different hardware architectures and platforms, a task that was addressed very early on during development at Bell Labs.

Even though the number of versions of UNIX increased dramatically over the years, porting applications between one version and the next was still considerably easier than porting between one proprietary OS and the next. This task has been made easier by the introduction of relevant standards that nearly all of the major UNIX vendors have adopted. No other operating system can claim this level of dominance across such a wide range of platforms. The proliferation of UNIX resulted in a huge amount of development pushing UNIX way ahead of its nearest proprietary competitors.

The Early Days of UNIX

The research arm of the Bell Telephone Company, Bell Labs, had seen the need for a new computer operating system in the late 1950s. This resulted in the BESYS

UNIX History Page 4 Tuesday, December 10, 2002 7:02 PM

operating system which, although used internally, had limited distribution outside of Bell Labs. By the mid 1960s, third-generation computer equipment was emerging and the people at Bell Labs had to decide whether to create a new operating system or to adopt a third party OS. The decision was eventually made to join General Electric and MIT to create a new timesharing system called the *MULTIplexed Information and Computing Service* (MULTICS). This collaborative venture aimed to show that general purpose, multiuser operating systems were a viable solution. Based on a research operating system from MIT called the *Compatible Time Sharing System* (CTSS), the MULTICS project resulted in a wide range of new approaches. Amongst those working on MULTICS were Bell Lab researchers Ken Thomson and Dennis Ritchie, who would take many of the ideas from MULTICS and evolve them into what would become UNIX. As an example, many of the UNIX commands were based on commands of the same name in MULTICS, namely `ls`, `cd`, and `pwd`. Due to the high cost of development and the amount of time that it was believed MULTICS would take to complete, AT&T withdrew from the MULTICS project in 1969.

On an internal GE-645 computer at Bell Labs, the GECOS operating system was installed, which proved inadequate to many of the researchers. For many this was seen as being back at square one. This resulted in a proposal by Thompson and Ritchie to get Bell Labs to buy a new machine so they could start work on their own interactive time-sharing system. The proposal was rejected and Thompson started work on an old PDP-7. Developing initially on a GE-635, Thompson produced a primitive kernel, assembler, shell, and a few utilities (`rm`, `cat`, `cp`) that were then transported to the PDP-7 by paper tape. Even at this stage the new primitive OS was sufficient for all further development to proceed on the PDP-7 itself. As a pun on MULTICS, the new environment was named UNIX.

In 1970 the UNIX pioneers acquired a DEC PDP-11 that was initially diskless and had 24KB of memory. They used 12KB for the operating system, allowed a small amount of space for user programs, and the rest was used as a RAM disk. It was around this time that the first 1st Edition UNIX appeared.

The Early History of the C Language

Following the early assembler versions of UNIX, Thompson worked on a Fortran compiler that evolved to support the language B, a cut-down version of BCPL. The B compiler was provided as part of 1st Edition UNIX in 1971, and some of the first utilities were even written in B. It was Ritchie's work over the next two years that resulted in the first C compiler, which became part of 3rd Edition UNIX in 1973. Note that it would still be another 5 years before the appearance of Kernighan and Ritchie's book, *The C Programming Language* [KERN78]. Following an abortive attempt by Thompson to write part of the kernel in an early version of C which did not support structures, by 4th Edition UNIX that appeared in 1973, Thompson and Ritchie had rewritten the kernel in C.

UNIX History Page 5 Tuesday, December 10, 2002 7:02 PM

Research Editions of UNIX

There were a total of ten research editions of UNIX from Bell Labs. Perhaps the most famous was 6th Edition UNIX which was described in John Lions' book *Lions' Commentary on UNIX 6th Edition* [LION96], which until it was published in 1996 remained an underground work. One thing that distinguished each research edition was the introduction of a new *Programmer's Reference Manual*. Following is a brief summary of the different research editions and which UNIX features they introduced:

1st Edition—1971. As well as the B compiler, 1st Edition UNIX introduced a number of well-known UNIX commands including `cat`, `chdir`, `chmod`, `chown`, `cp`, `ed`, `find`, `mkdir`, `mkfs`, `mount`, `mv`, `rm`, `rmdir`, `wc`, and `who`.

2nd Edition—1972. One amusing note on 2nd Edition was a comment in the *Programmer's Reference Manual* that the number of UNIX installations had now grown to 10!

3th Edition—1973. The UNIX C compiler (`cc`) first appeared. The kernel was still written in assembler and the number of installations had grown to 16. Pipes were also introduced.

4th Edition—1973. The kernel was rewritten in C.

5th Edition—1974. This edition appeared around the time of Thompson and Ritchie's paper "The UNIX Time Sharing System" [RITC74]. The source code was made freely available to universities for educational purposes.

6th Edition—1975. This edition, also known as V6, was the first edition widely available outside of Bell Labs. Most of the operating system by this time had been written in C.

7th Edition—1979. The first K&R (Kernighan and Ritchie) compliant C compiler made its appearance with 7th edition together with Steve Bourne's shell (`sh`). The kernel was rewritten to make it more portable to other architectures. At this time the UNIX Systems Group was created and started working on enhancing 7th Edition (on which System V UNIX would be based). Microsoft also licensed 7th Edition, which it used to develop the Xenix operating system. Note that the size of the 7th Edition kernel was only 40KB, a tiny program by today's standards.

8th Edition—1985. 8th Edition UNIX picked up some enhancements developed from 4.1BSD. This edition was used as the basis for System V Release 3.

9th Edition—1988. This edition picked up enhancements made for 4.3BSD.

10th Edition—1989. This was the last edition.

AT&T's Commercial Side of UNIX

In the late 1960s, while Bell Labs was looking for a new timesharing system, the Bell Telephone company was looking for a way to automate their telephone

UNIX History Page 6 Tuesday, December 10, 2002 7:02 PM

operations using minicomputers to switch over from their existing system of people and paper.

It was Berkley Tague, the head of the computer planning department, who, having seen the capabilities of UNIX, realized its potential and saw how it could ease their job. By 1971 Tague gained approval for the adoption of UNIX to support Bell Telephone operations. By 1973 he formed the *UNIX Support Group* (USG) which worked closely with the UNIX team from Bell Labs. During the same year, the first UNIX applications started to appear, initially involved in updating customer directory information and intercepting calls to phone numbers that had been changed. 1973 also saw the first C version of UNIX released internally together with the first *Programmer's Work Bench*, which included `sccs` and other tools.

Around the time of 7th Edition UNIX, USG took responsibility for UNIX and after a number of internal-only releases, System III UNIX became the first version of UNIX that was available for use outside Bell Labs.

USG later became the *UNIX System Development Laboratory* (USDL). In 1984, this group released System V Release 2 (SVR2) which was the first version of UNIX to support paging, copy-on-write semantics, shared memory, and file locking. SVR2 UNIX is described in Bach's classic book *The Design of the UNIX Operating System* [BACH86]. At this time there were about 100,000 installations of UNIX worldwide running on a whole host of different platforms and architectures. The fact that no other operating system had achieved this goal was perhaps the single greatest reason why UNIX became so popular.

Following yet another name change to *AT&T Information Systems* (ATTIS), the group released System V Release 3 (SVR3) in 1987. This included a number of enhancements:

- The File System Switch (FSS) provided an architecture under which multiple filesystems could coexist in the kernel at the same time. The FSS provided a layer by which the rest of the kernel could perform file and filesystem related operations in a filesystem independent manner through a well defined interface.
- The RFS (Remote File Sharing) filesystem provided a fully distributed, cache-coherent file system.
- The STREAMS subsystem for building networking stacks. The initial implementation of STREAMS was first introduced in 8th Edition UNIX.
- The Transport Layer Interface (TLI) for network programming.
- Shared libraries which can reduce the amount of memory used.

System V Release 3.2 was released in 1987 which involved a merge of SVR3 and Xenix, produced by Microsoft and the Santa Cruz Operation (SCO).

One of the major releases of UNIX from AT&T was System V Release 4 in conjunction with Sun Microsystems. This is described in more detail in the section *System V Release 4 and Variants* later in the chapter.

UNIX History Page 7 Tuesday, December 10, 2002 7:02 PM

The Evolution of BSD UNIX

Following Thompson and Ritchie's paper on UNIX at the Symposium on Operating System Principles in 1974, Bob Fabry, a professor at the University of California at Berkeley wanted to get a copy to experiment with.

After buying a PDP 11/45, he received the tapes for 4th Edition UNIX which was installed in 1974. Due to disk related problems, Ken Thompson spent time dialed in over a modem debugging the system.

Following the purchase of a Vax 11/70 in 1975, Ken Thompson started a sabbatical at Berkeley during which time he brought up 6th Edition. Around this time, graduate students Bill Joy and Chuck Haley arrived and started working on the newly installed system, initially enhancing a Pascal system that Thompson had put together. The same year, they produced the `ex` editor and started working on the kernel following the departure of Thompson back to Bell Labs.

Following requests for the Pascal environment, Bill Joy put together the *Berkeley Software Distribution* (consider this as 1BSD) in 1977 and distributed thirty copies. Soon after, Joy wrote the `vi` editor, still hugely popular 25 years later.

In 1978 Joy released the second Berkeley Software Distribution which became known as 2BSD. This included an updated Pascal system, the `vi` editor, and `termcap` which could be used for driving multiple different terminal types, a must for supporting `vi`.

Needing more power, a Vax 11/780 was purchased and the 32/V port of UNIX, initiated at Bell Labs, was installed. Following a number of enhancements to make use of the new virtual memory capabilities of the machine, Joy started porting 2BSD to produce the third Berkeley distribution, 3BSD, which was released in 1979.

Around this time, DARPA (*Defense Advanced Research Projects Agency*) decided to standardize on UNIX in order to provide a network to link their major research centers. Based on Fabry's proposal to DARPA and the ensuing success of 3BSD, an 18 month contract was awarded to Berkeley. Fabry set up the *Computer Systems Research Group* (CSRG) to handle the contract and research. Bill Joy came on board and set to work on what would become 4BSD. Released in 1980, the new system included the Pascal compiler, job control, auto reboot, and a 1KB size filesystem. Joy then released 4.1BSD which contained numerous performance improvements to the kernel.

Following renewal of the contract by DARPA, the new project would produce what would become the *Berkeley Fast File System*, support for large virtual address spaces and better IPC mechanisms. The TCP/IP stack was integrated into BSD and a number of temporary tools were introduced on top of the networking stack. These temporary tools, namely `rcp`, `rsh`, `rlogin`, and `rwho` are a little more permanent than their original authors anticipated, still being used today.

Following Bill Joy's departure in 1982 to co-found Sun Microsystems, 4.2BSD was released in 1983. Due to the introduction of TCP/IP and the Fast File System, the number of 4.2BSD installations far exceeded System V from AT&T.

UNIX History Page 8 Tuesday, December 10, 2002 7:02 PM

Following criticism of 4.1BSD performance, a two year period of tuning and refining produced 4.3BSD which was released in 1986. Two years later, completing the work started by Joy to divide the BSD kernel into machine dependent and machine independent layers, CSRG released the finished work under 4.3BSD-Tahoe. Further development which resulted in a rewrite of the virtual memory subsystem, based on the Mach microkernel, together with NFS, produced 4.3BSD-Reno in 1990.

BSD Networking Releases

To avoid BSD recipients having to obtain an AT&T source license while wanting to have source access to the networking components of BSD, the *Networking Release* of BSD was released in 1989. An expanded version, which involved rewriting all except six kernel files, was distributed as the Networking Release 2 in 1991. This involved a huge effort by many people.

Bill Jolitz continued the work by rewriting the remaining six kernel files to avoid AT&T copyrighted source code and porting the system to the Intel 386, resulting in 386/BSD which was distributed over the internet.

UNIX Goes to Court

Following the Net/2 release of BSD, the *Berkeley Software Design, Incorporated* (BSDI) company was formed to develop a fully supported, commercial version. The BSDI version, released in 1992, included replacements for the six kernel files, was considerably cheaper than System V UNIX from USL, and used *UNIX* as part of the telephone number in their advertisements to call for questions. This was followed by a lawsuit from AT&T, initially aiming to prevent BSDI from promoting their product as UNIX. This was then followed by an additional lawsuit that claimed that the BSDI version contained proprietary USL source code and secrets.

While the lawsuit continued, USL was purchased by Novell in 1993. Novell founder and CEO, Ray Noorda, wanted to drop the lawsuit and in 1994 an agreement was finally reached. As part of the agreement, 5 of the 18,000 files that made up the distribution were removed. With some minor changes to other files and the addition of copyright notices in an additional 70 files, the new, 4.4BSD-Lite version was released.

The NetBSD Operating System

386/BSD was extremely successful. Unfortunately Jolitz was unable to work full time and keep up with his work on 386/BSD. Frustrated with the way that development of 386/BSD was progressing, others started working on a parallel development path, taking a combination of 386BSD and Net/2 and porting it to large array of other platforms and architectures.

UNIX History Page 9 Tuesday, December 10, 2002 7:02 PM

The FreeBSD Operating System

Following work on Jolitz's 386/BSD system, Jordan Hubbard, Rod Grimes, and Nate Williams released the *Unofficial 386BSD Patchkit* which contained a number of changes. Jolitz denounced approval of the project in 1993, which was followed by discussions between Hubbard and Walnut Creek to produce a new operating system, which they called FreeBSD. The first CDROM version of FreeBSD, version 1.0, was released in December of 1993.

Following the USL lawsuit, the base operating system was upgraded from Net/2 to 4.4BSD-Lite, which resulted in the release of FreeBSD 2.0 in November of 1994. Enhancements continue to be added with the latest stable release being FreeBSD 4.2.

FreeBSD has been relatively successful on its own ground. It was also used as the basis for Apple's Mac OS X operating system.

The OpenBSD Operating System

Following a disagreement between Theo de Raadt, who had been responsible for the SPARC port of NetBSD, and the NetBSD core team, de Raadt founded OpenBSD. The new OS started to diverge from NetBSD 1.1 in 1995 and this was followed by the first release, OpenBSD 3.0 in October of 1996. The core focus of OpenBSD was security.

Although not as portable as NetBSD, OpenBSD still runs on a wide range of machines and architectures and continues to lead the way as the most secure BSD release available.

Sun Microsystems and SunOS

Sun Microsystems was founded in 1982 by four people including current CEO Scott McNeally and BSD developer Bill Joy. In their first year they released their first workstation based on hardware developed at Stanford University and on the BSD operating system.

Sun has continued from day one to innovate and enhance UNIX. In order to provide remote file access they introduced the *Network File System* (NFS) and the VFS/vnode architecture to support it.

In 1987 Sun and AT&T joined forces to develop UNIX System V Release 4, which combined the best of SunOS and System V Release 3.2. SVR4 encompassed many of the ideas that Sun had implemented including VFS/vnodes, NFS, and their virtual memory architecture, which cleanly divides memory management into machine dependent and machine independent layers. Sun, together with IBM and HP, continues to take UNIX to the enterprise, continually enhancing their UNIX offerings while retaining compatibility at the standards level.

UNIX History Page 10 Tuesday, December 10, 2002 7:02 PM

System V Release 4 and Variants

System V Release 4 set the standard for everyone else to follow producing an extremely feature-rich operating system that combined the best of the historical versions of UNIX with many new ideas from Sun. The following list shows some of the major enhancements that came with SVR4:

- The VFS/vnode architecture that replaced the FSS from SVR3. The VFS/vnode architecture was originally developed as part of SunOS.
- Symbolic links.
- The C and Korn Shells along with job control.
- Memory mapped files.
- The UFS filesystem derived from the BSD Fast File System. UFS became the defacto standard on most versions of UNIX. It is still the default filesystem on Solaris and is still undergoing major development. SVR4 also included the NFS filesystem. At this stage, the largely unsuccessful RFS was starting to fade.
- STREAMS-based console and TTY (teletype) management.
- Real-time scheduling and a partial implementation of kernel preemption.

Enhancements continued thereafter. SVR4.1 included Asynchronous I/O. SVR4.2 included Access Control Lists (ACLs), the VERITAS Filesystem (VxFS), and VERITAS Volume Manager (VxVM). Following this, with a major rewrite, SVR4.2MP introduced Symmetric Multiprocessing (SMP) capabilities and kernel threads.

Novell's Entry into the UNIX Market

The UnixWare 1.0 release of UNIX was released in 1992 as a joint venture between Novell and USL under the name Univel. Novell completed the acquisition of USL in 1993, and both USL and Univel were merged to form the Novell UNIX Systems Group.

UnixWare 1.0 was based on SVR4.0. This was followed by UnixWare 1.1, which was based on SVR4.2. With the introduction of UnixWare 2.0, the kernel (SVR4.2MP) had changed significantly, introducing SMP support and kernel threads.

In 1993 Novell transferred the rights to the UNIX trademark to the X/Open organization (now the Open Group). Two years later they sold their UNIX business to SCO who in turn sold a dwindling UNIX business to Caldera in 2001.

UNIX History Page 11 Tuesday, December 10, 2002 7:02 PM

Linux and the Open Source Movement

One could argue that if readers didn't have to purchase Andrew Tanenbaum's MINIX operating system that accompanied his book *Operating Systems: Design and Implementation* [TANE87], there would be no Linux.

However, the *Free Software Foundation*, founded by Richard Stallman, had already been working for a number of years on a free version of UNIX. The compiler, utilities, and just about everything except the kernel had been written under the auspices of the GNU license which allowed the source to be freely distributed.

Linus Torvalds, a research assistant at the University of Helsinki in Finland, released Linux 0.0.1 in August of 1991, and the rest, as they say, is history. Popularity of Linux continues to grow. Although it originally took many of its ideas from Minix, Linux has been influenced by all versions of UNIX and non-UNIX systems. Linux followed in the success of UNIX by being ported to just about every hardware architecture and platform available from IBM mainframes down to hand-held organizers.

Users of Linux will find a number of components from many different authors and organizations. A Linux OS is comprised of the Linux kernel, much of the Free Software Foundation's GNU software, and a number of other free applications and utilities. There are many distributors of Linux, with the top players being Red Hat, SuSe, TurboLinux, and Caldera.

UNIX Standardization

The section *A Brief Walk through Time* earlier in the chapter showed how the different versions of UNIX came into existence through the 1980s. Although most of these versions stemmed from either System V or BSD, each OS vendor added its own enhancements, whether to increase performance or add new interfaces in response to internal or customer demands. Because application portability was crucial to the success of application developers, it soon became clear that a level of standardization was needed to prevent this divergence from going too far.

Various bodies have been responsible for driving the standardization of UNIX interfaces, whether at a command level, library, or system call level; or newer initiatives such as the Large File Summit for 64-bit file access and the Data Management Interfaces Group (DMIG) for interfaces relating to Hierarchical Storage Management. This section describes the main standards bodies, their goals, and the standards that they have produced.

IEEE and POSIX

The /usr/group organization was formed by a group of individuals in 1980 with the intention of standardizing user-level interfaces with the goal of application portability.

UNIX History Page 12 Tuesday, December 10, 2002 7:02 PM

They reached consensus in 1984, and their work was used by the ANSI X3J11 committee, the same group who were working on standardization of the C language. As the number of versions of UNIX started to increase, divergence continued, and the /usr/group standard became less and less effective. This led to the formation of the *Portable Operating System Interface for Computing Environments* (POSIX) in 1995 which used the /usr/group standard as its base working document. As a point of interest, the name POSIX was suggested by Richard Stallman, founder of the *Free Software Foundation* (FSF).

The standard produced by this group, POSIX 1003.1-1998 became the most widely recognized standard throughout the UNIX industry and is available on many non-UNIX platforms. The initial standard was revised throughout the next three years and adopted by the Institute of Electrical and Electronics Engineers (IEEE) organization to become IEEE Std 1003.1-1990 although it is still more commonly known as POSIX.1 or simply the POSIX standard. In 1989 the /usr/group changed its name to Uniforum.

The POSIX working committees did not stop there and produced a number of other standards of which some are shown in Table 1.1.

The X/Open Group

With the same goals as the /usr/group, a number of European computer companies formed a non profit organization in 1984 called X/Open.

Although many of the players were not specifically UNIX based, application portability was still key. The first published standard from X/Open was the *X/Open Portability Guide* (XPG). The third draft of this standard, XPG3, included both POSIX 1003.1-1998 and a number of interfaces pertaining to the X Window System. The XPG3 test suite contained over 5,500 different tests that exercised system calls, library interfaces, and the C language.

The XPG4 standard was released in October of 1992. This encompassed not only POSIX.1, but also POSIX.2 and ISO C. A successful branding program was put in place so that companies could claim XPG4 compliance.

The System V Interface Definition

The UNIX System Group (USG) released the *System V Interface Definition* (SVID) version 1 with System V Release 2 in 1994. The SVID was a two-volume book that described all user accessible interfaces that were available with SVR2. SVID version 2 accompanied SVR3 in 1996.

With the introduction of SVR4 in 1989, version 3 of the SVID became available, this time a four-volume set. To accompany the SVID, USG produced SVVS, the *System V Verification Suite*, an exhaustive test suite that exercised all of the visible interfaces. Any vendors licensing System V were required to run and pass SVVS in order to use the name System V.

Since by this stage the SVID effectively encompassed the POSIX.1 standard, it was used as the main document in producing what would become the *Single UNIX Specification*.

UNIX History Page 13 Tuesday, December 10, 2002 7:02 PM

Table 1.1 POSIX Standards

STANDARD	DESCRIPTION
1003.1	System call and library routines
1003.2	The shell and UNIX utilities
1003.3	Test methods and conformance
1003.4	Real-time interfaces

Spec 11/70 and the Single UNIX Specification

In order to combine the existing UNIX standards such as POSIX.1 and XPG4, a group was formed by Sun Microsystems, HP, IBM, Novell/USL, and the Open Software Foundation (OSF) to provide a single unified standard based on existing standards and additional features provided by the different UNIX versions. Using XPG4 as a base which already encompassed POSIX.1 and ANSI/ISO C, a collection of 1,170 APIs were specified in total, and thus the name Spec 11/70 was given to the group and the specification.

The Spec 11/70 API was delivered to X/Open in 1983 resulting in the *Single UNIX Specification*, which was published in 1994. Various names have since followed this publication including UNIX 95 and the enhanced version renamed UNIX 98.

The standard is still maintained by the *Open Group* which was formed by a merge of X/Open and OSF. The Single UNIX Specification can be viewed online at `www.opengroup.org`.

The main components of the Single UNIX Specification are:

System Interface Definitions (XBD). This document outlines common definitions used in the XSH and XCU documents.

System Interfaces and Headers (XSH). This document describes all programming interfaces and all header files. Most of the text provides UNIX manual style representations for each API.

Commands and Utilities (XCU). This document describes all of the commands and utilities in a UNIX manual page style format.

Networking Services. This document describes the X/Open Transport Interface (XTI), XPG4 sockets, and the IP address resolution interfaces.

X/Open Curses. This document describes X/Open version 3 *curses*.

UNIX International and OSF

The *Open Software Foundation* (OSF) was founded in 1988 by seven leading computer companies with the goal of producing an operating system together with an open and portable application environment.

As a reaction to OSF and with a consortium of over 200 vendors and users,

UNIX History Page 14 Tuesday, December 10, 2002 7:02 PM

UNIX International (UI) was founded in 1988 centered around AT&T's SVR4 version of UNIX. The goals of the organization were to drive the direction for SVR4 although in reality, UI turned out to be more of a marketing machine with little actual output. Within a few years, UI was dissolved, and the direction of SVR4 was left to Novell/USL and then SCO.

Both OSF and UI achieved some notable successes. The big battle predicted between the two never happened in reality. Through USL, UI pushed the SVID version 3, which became the basis for the Single UNIX Specification. OSF merged with X/Open to form the *Open Group* which still maintains the Single UNIX Specification today along with other UNIX related standards.

The Data Management Interfaces Group

A small number of independent software and hardware vendors were developing Hierarchical Storage Management (HSM) solutions, which involved modifications to the base UNIX kernel (see the section *Hierarchical Storage Management* in Chapter 12 for further details). Following publication of Neil Webber's USENIX paper "Operating System Support for Portable Filesystem Extensions" [WEBB93], a group of HSM, backup, OS, and filesystem vendors formed the Data Management Interfaces Group (DMIG) with the goal of producing an interface specification that the OS/filesystem vendors would implement to prevent the constant rewrite of HSM software with each iteration of the operating system.

X/Open adopted the Data Management API (DMAPI) and renamed it XDSM (*X/Open Data Storage Management)*.

The standard allows for applications to transparently migrate data from the filesystem (termed secondary storage) to tape or other offline storage devices (tertiary storage) bypassing the UNIX timestamping mechanisms and without knowledge of user-level applications. This allows HSM applications to achieve a virtual memory-like approach to storage.

The Large File Summit

32-bit operating systems imposed limits on the size of files that could be accessed due to limits imposed at various layers throughout the operating system, not least the fact that the value that could be held in a signed integer, the maximum value that could be held in a `size_t`, was limited to 2GB -1.

To provide an intermediate solution that could allow access to files greater than 2GB before the advent of 64-bit operating systems, the *Large File Summit*, a group of operating system and filesystem vendors, was formed to produce a specification that introduced a new set of data types and APIs that allowed for large file access.

Applications could access *large files*, files greater than 2GB, by either invoking 64-bit versions of the system calls or via compile time flags that switched the size

UNIX History Page 15 Tuesday, December 10, 2002 7:02 PM

of various data types. At the time of writing, much of this is now a moot point with 64-bit file access being the norm in UNIX.

Summary

This chapter highlighted the main events that show how the different versions of UNIX have evolved and where specific pieces of technology have come from. The history of UNIX could fill a book by itself. Indeed, Peter Salus' book *A Quarter Century of UNIX* [SALU96] describes UNIX history from 1969 to 1994.

Programmers wishing to follow UNIX standards should adhere to the Single UNIX Specification when striving for application compatibility across all the major versions of UNIX. Although Linux does not comply completely with the specification, most interfaces are supported. At a very minimum, the POSIX interfaces are supported by just about every operating system, UNIX and non-UNIX alike.