The Debut of a New UNIX

Upon examining the Open Software Foundation's new operating system, OSF/1, you'll find it's very much like one of the many UNIX variants based on AT&T's System V, with Berkeley enhancements.

That's good news for programmers, because software may port more easily. OSF/1's interface conforms to two application programmer interfaces: System V Interface Definition (SVID Release 3) and the latest X/Open Portability Guide (XPG3). Porting will be an important point, because even System V-based software, such as The Santa Cruz Operation's 3.2-based applications, will require porting because of differences in system call semantics.

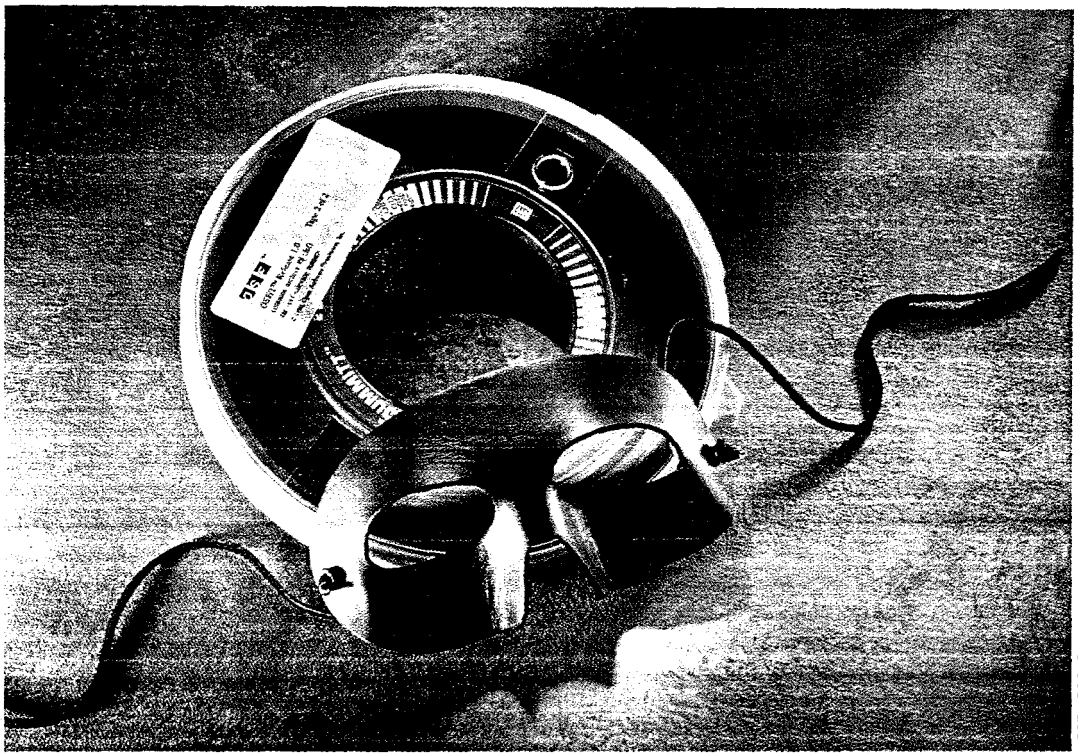

ROSE HODGES

Then there are more subtle distinctions. For example, OSF/1 will be the first to provide an application program-level multiprocessing interface. Also, OSF/1 uses the Berkeley sockets networking interface, but SVR4 uses Streams libraries to provide a networking interface, with sockets as an option.

Still, on the issues of kernel-level multiprocessing, security, supported file systems, graphical user interfaces, logical volumes, internationalization, and memory-mapped files, both versions have strong similarities.

When the first snapshot of OSF/1 was released last October, the tape contained three reference ports made by OSF: Digital Equipment Corp.'s 3100 workstation; the Intel 386 microprocessor; and the Encore Multimax, using the National Semiconductor 32000 CPU. The snapshot is intended to show how other OSF members–notably Hewlett-Packard and IBM–can port OSF/1 to their own hardware.

You have to look hard to differentiate OSF/1 from its rival, AT&T's System V Release 4

By Rik Farrow

Despite the initial hoopla surrounding the founding of OSF, only DEC seems interested in providing OSF/1 any time soon (see sidebar). HP has promised to provide a port on one of its four workstation lines late this year, while IBM is planning to port to the entire AIX family, though it has set no delivery date.

Although OSF claims its operating system will be compatible with existing software, only one application, Informix Wingz, has actually been demonstrated. Application developers will be the first to work with OSF/1 when they decide to port software to computers running OSF/1. Unlike SVR4, OSF/1 provides no binary standards, so developers will have to support more ports than for SVR4-based computers with application binary interfaces.

continued

Right now, the acceptance of OSF/1 rests largely on the existence of applications running under it. Few vendors are willing to make the port without the support of hardware vendors.

Multiprocessing for the Masses

Some significant technologies will appear in OSF/1. The most heralded addition is a multiprocessing capability built into every OSF/1 kernel. Multiprocessing has become a critical issue as users and vendors have learned that adding processors to a system is an easy way to provide scalability. A computer that was designed to handle 20 users can be scaled up to handle many more users by adding processors, memory, and serial ports. Adding a processor is also seen as a way to improve performance without having to abandon an entire system.

Vendors receive OSF/1 with symmetric multiprocessing capabilities. Symmetric multiprocessing (SMP) means that the system load will be balanced equally among processors, unlike asymmetric versions, which force one processor to do more work. OSF/1 has parallelized system calls, libraries, file systems, the virtual memory subsystem, TCP/IP, Streams, and BSD sockets.

Missing in this list are parallelized device drivers, which are the software that actually talks to the serial ports and disks drives. The device drivers in OSF/1 are not multiprocessing–multiprocessing device drivers provide the next level of multiprocessing performance improvements. OSF/1's kernel-level multiprocessing is based upon its Mach kernel. UNIX International, the AT&T vendor advisory group, has promised to deliver an SMP version of SVR4 in the second half of 1991.

SMP capabilities have been developed previously for both System V-based versions and BSD versions of UNIX. OSF/1 has new user-level multiprocessing capabilities, based on threads. Threads give application programmers the ability to create applications with multiple, simultaneous, execution paths. In an ordinary, unthreaded program, there is one path of execution where, essentially, one list of instructions is followed. In threaded programs, each thread represents a path of execution that may be followed.

As an example, a database server application could be written to use threads to respond to each user's query. When the server receives the query, it dispatches the thread, and if there are multiple processors, the thread may be executing simultaneously with other threads. Dispatching a thread requires that the operating system allocate some memory for temporary use by the thread. Executing a new program, by comparison, involves reading a disk image and allocating a much larger amount of memory. Thus, threads are called "lightweight processes" because they are simpler and quicker to start up and manage.

OSF/1 uses the draft Posix standard for multiprocessing support, the 1003.4a Pthreads programming interface. UI is helping to determine the final draft of the Posix standard. The AT&T version of user-level threads won't be ready until 1992.

Although threads are an exciting new toy, they will also be a tricky one to master. Threaded software is more difficult to write and debug than non-threaded software. Each thread must cooperate because it shares the same global variables. Special routines are defined in the draft standard to prevent threads from corrupting shared data. Writing threaded software is similar to writing device drivers and is an art to be mastered by tomorrow's programmers.

B1 Security Issues

With the U.S. government providing much of the motivation behind the rush into open systems, security becomes an important issue for operating systems. OSF/1 uses SecureWare Inc.'s SMP+ (Secure Module Package Plus) product to provide Orange Book B1-level security. Security is designed into the OSF/1 kernel to be selectable at compilation time to either standard UNIX (C1), C2, or B1. In addition to B1, OSF/1 offers some features from higher security levels, such as least privilege (B2) and access control lists (B3). Auditing, which is activated at the C2 level, can be adjusted on a per-user granularity.

Security certifications are essential for selling UNIX systems to the federal government as it will require B1 security on many of its purchases by 1992. That may spell trouble for OSF members, however, because certification at the B1 level has taken more than two years in the past. Recently, the federal government's own National Institute of Standards and Technology has taken responsibility for certifying software from the National Computer Security Center, and the certifying process may be accomplished much faster.

AT&T SVR4, with its year-long lead time, will have a B2-certified secure product, SVR4ES (Enhanced Security), in the first half of 1991. The SVR4ES can be configured to provide lesser degrees of security, so it can be sold to any place that requires C2, B1, or B2 Orange Book security. Here AT&T has a slight edge. Also, vendors interested in SecureWare's SMP+ can license it directly from SecureWare for System V ports before SVR4.

SecureWare's SMP+ includes much of what you will find in UI's secure product. In particular, the password management in the vanilla SVR4 is very similar to SecureWare's Password Management System. This is not surprising, because both are based upon the U.S. Defense Department *Password Management Guideline*. The overlap in security features is expected because all products must conform to the Orange Book to succeed in certification.

One area of difference is SecureWare's

Selected OSF/1 and AT&T SVR4 Features

	OSF/1	*System V Release 4*
Multiprocessing	Encore Multimax SMP user level threads	SMP in summer 1991 user threads in 1992
Security	SecureWare level B1	SVR4ES level B2 in 1st quarter 1991
File Systems	System V, FFS, NFS, and customer defined	System V, FFS, NFS, RFS, /proc, and customer defined
Logical Volume	available	announced for 1st quarter 1991
Disk Mirroring	available	announced for 1st quarter 1991
Networking Interface	BSD Sockets and Streams from Mentat	Streams and Sockets on top of TLI Core OSI stack 1st quarter 1991
Graphical User Interface	OSF/Motif and X11R4	Open Look, OSF/Motif, NEWS X11R4
Internationalization	European and Asian environment	European and full Asian environment; multibyte support; supports multiple languages
Portability	ANDF under discussion	ABIs for Intel 80x86, i860, Motorola 68000 and 88000, WE32000, MIPS, and SPARC

Trusted Application Programming Interface (TAPI). The TAPI is used to provide application programmers with a method for using least privilege. The principle of least privilege, a B2-level feature, specifies that each process has, and uses, the least powerful level of privilege to perform each action. In standard UNIX systems, the set-user-id principle provides a dangerous form of least privilege. A set-user-id program, such as the **passwd** command, permits users to modify files that would otherwise be protected.

TAPI permits programmers to adjust the privilege level during program execution, so that only those sections of code that require special privileges have those privileges. Although TAPI will provide a finer degree of control over privilege, the set-user-id principle in SVR4 has been enhanced so that it also has a means to turn privilege on or off during portions of a secure program.

File System Comparisons

The Berkeley Fast File System (FFS) is the most popular enhancement made to UNIX file systems, and both OSF/1 and SVR4 will support it. In addition, SVR4 and OSF/1 support a variety of other file system types: System V, the Network File System (NFS), and customer-defined file systems. The SVR4 version is based on Sun Microsystems' "vnode" architecture, while OSF/1 uses the BSD vnodes, an adaptation of Sun's vnodes. Hardware vendors have been adding the FFS and NFS to System V products for years–FFS for its increased performance and NFS for its almost universal acceptance as a networking standard. SVR4 also offers AT&T Remote File Sharing, a proprietary networked file system architecture, and the **/proc** file system, a new development used for process tracing.

For system administrators, it looks like SVR4 has the edge in file system architecture, but OSF/1 has a slight lead here. OSF/1 includes the concept of logical volume management today, which permits a file system to span physical disks, or to be increased *while in use*. Logical volume management is a UNIX feature whose time has come, providing to UNIX systems the type of capability that has been available in MIS environments for years. SVR4 was expected to have a logical volume manager early in 1991.

An important feature of IBM's AIX 3.1, the journaled file system, will not initially appear on either operating system, although both UI and OSF have plans for it. IBM's use of the journaled file system provides both performance and reliability benefits. UI plans to incorporate a journaled file system by the end of 1991, and OSF will leave implementation up to individual vendors.

A new addition to the UNIX system is memory-mapped files. These files permit programmers to treat files as if they were part of a program's data. With older UNIX versions, files are first copied into the kernel's private buffers, then to program data spaces as requested. Memory-mapped files allow a programmer to include the image of a file inside the program's data area. The file is still opened as before, but then is memory mapped and addressed as if it was memory without seeks or reading. The operating system uses its paging system to move portions of the memory-mapped file between the disk and memory as needed. Memory-mapped files offer performance improvements to any disk I/O limited application, and are offered by both OSF/1 and SVR4.

Internationalization Support

Internationalization deals with the problems encountered when moving UNIX

Vendors Slow to Support OSF/1

In a great show of unity, IBM Chairman John Akers, Digital Equipment President Kenneth Olsen, and Hewlett-Packard President John Young stood together with four other corporation heads to announce the formation of the Open Software Foundation in May of 1988. OSF would create a new open systems standard that would seize control of the open systems operating systems market from AT&T and its chosen favorite, Sun Microsystems Inc.

OSF founders feared that Sun, already a successful upstart, would gain an unfair advantage if permitted to participate in the design of the next release, the fourth, of UNIX System V.

Alas, this unity was not to prevail. When the first distribution of the new operating system, OSF/1, was announced, only two former supporters promised to use it to replace their own versions of the UNIX operating system. Only Digital Equipment and France's Groupe Bull actually plan to provide OSF/1 as a replacement for UNIX on their computer systems this year. Other vendors are slower to embrace OSF/1 because it means letting go of their own versions of UNIX and of the applications ported to it.

Tom Chmielewski, Ultrix product manager of DEC, says that the next release of Ultrix, Version 4.1, will be the steppingstone to Version 5.0, the first OSF-based version of Ultrix. DEC plans to be shipping Ultrix 5.0 before the end of 1991. The version shipped on the DECstation 3100 is a developer's kit, an aid in porting software to Ultrix 5.0. Coincidentally, Sun Microsystems is also providing a 4.1 version of its own operating system as a steppingstone to SunOS version 5.0, the SVR4-based product.

Hewlett-Packard, which bought out Apollo, another OSF founding member, says it plans to provide OSF/1 on only one of its four workstation lines–those based on its own RISC architecture. The remaining workstations and minicomputers will continue to use HP-UX, a System V-based variant of UNIX. Jan Silverman, marketing manager for System Software and Communications at HP's Apollo Systems Division, says that HP plans to integrate other OSF modules, such as the Distributed Computing Environment (DCE), into HP's base product line before incorporating the operating system component, OSF/1. It's not surprising, because HP-Apollo gave OSF the DCE technology used as the basis for remote procedure calls.

IBM has been very cagey about letting its OSF-related plans be known. At this point, IBM will say only that it is planning to experiment with OSF/1 on the PS/2, the RS/6000, the IBM 370, and System 390. IBM plans to continue using its System V-based AIX 3.1 on its own line of workstations, the RS/6000, for the time being.

This lukewarm reception of the long-heralded first release is not surprising. OSF/1 represents new, untested technology. Although OSF/1 draws on many other technologies, ask systems integrators what they think about getting many products to work together, and you'll find that the more products involved in the integration, the more troublesome it will be.

OSF vendors must also garner applications that will run on their OSF-supported computers. This is a chicken-and-the-egg problem, because computers cannot be sold without software, and applications won't be ported until the software vendors believe there's a market for their software. What has happened in the past with new hardware must occur again. Hardware vendors must supply OSF/1 on their computers and offer incentives to application vendors interested in porting to their combination.

—R.F.

systems and their applications to environments where the native language is not English. Previously, lack of support for other character sets has been a bone of contention between European UNIX users and U.S. suppliers. To solve this problem, UNIX commands and library routines needed to be revised to support non-English character sets. Both SVR4 and OSF/1 provide commands that recognize a 256-character set, support for different sorting sequences, and message libraries for different languages. The message libraries permit application developers to produce one version of an application that uses different message libraries for presenting the end user's interface.

SVR4 also supports multibyte character sets. OSF/1 works with European character sets, but cannot support Asian languages that use pictographs for characters. A Japanese typesetting system might provide for 20,000 different pictographs, requiring two bytes to represent each character. OSF/1 currently supports only one-byte character sets.

Compatibility Standards: ANDF vs. ABI

Until recently, you could not buy UNIX software for an Intel-based system and expect it to work under SCO UNIX and Interactive Systems' UNIX. Subtle differences made the software incompatible. To solve this problem, some architecture vendors, including Intel, Motorola, Sun, and MIPS, worked with AT&T, SCO, and Interactive Systems to produce application binary interfaces (ABIs) for System V Release 3.2. Software that conforms to an ABI will run without porting on different computers that support the ABI for that processor. This means that an application only needs to be ported once for each CPU architecture, instead of once for each CPU and operating system combination.

At the very least, OSF is to be commended for spurring AT&T into greater concern for user and vendor issues

ABIs lead to "shrink-wrapped" software–software that can run on any machine with the same processor. What's more important, most software that conforms to ABIs for System V Release 3.2 will run on SVR4 without change or recompilation. The majority of existing SCO 386 Release 3.2 software will run without modification on SVR4. Perhaps 10 percent requires modification of some X and networking code where the Release 3.2 ABI does not correspond to SVR4 and iBCS-2, the newer ABI for the Intel architecture.

OSF/1 stakes out new ground in compatibility standards. The Application Neutral Distribution Format (ANDF) is not an ABI. With ANDF, each application vendor ships only a single version of an application. The shipped version is in pseudo-code (the ANDF), and must be converted into native code during the installation process.

ANDF looks like a great idea for application vendors, because there is only one set of source code to support for all OSF/1 platforms. However, the conversion from pseudo-code to native code will introduce some coding inefficiencies, so the translated application will be slower to execute than one delivered in native code. Some small degradation in performance, say less than five percent, would be tolerable to vendors. But with any greater performance loss, the ANDF would be ignored and application programmers would still have to maintain ports to all OSF/1 platforms. OSF has not yet announced an acceptable ANDF technology.

Is OSF/1 All that It Can Be?

OSF was founded to create an operating system product that was not under AT&T's control. Fear of Sun Microsystems' involvement in SVR4 also spurred the creation of OSF. OSF/1 was supposed to be a completely new version of UNIX, free of any involvement of AT&T source code, with a smaller and more efficient kernel.

The OSF/1 kernel still uses code that requires a System V Release 2 or 3 source license. The command set, taken from AIX 3.1, also incorporates code taken from System V Release 2. The new kernel, still named `/vmunix`, will be no smaller than the kernels it replaces–OSF/1 will be about the same size as the Ultrix kernel on a DECstation 3100.

It's not really surprising, because even though the heart of the kernel has been replaced by Mach 2.5, the remainder and the vast majority of the kernel involve the UNIX system and related services, which have not shrunk, but grown even greater. Questions about increased efficiency have yet to be answered, as OSF/1 is still in an early release form and not ready for benchmarking.

Other questions about OSF/1 remain. Vendors who have decided to use OSF/1 must retrain their system programming staffs, support teams, and field engineers to handle a new operating system. While SVR4 vendors will also have some retraining, moving from an earlier version of System V to SVR4 is more of a migration. For companies like IBM and Hewlett-Packard, the issue of migrating their own support staffs has certainly affected their time frame for replacing their versions of UNIX with OSF/1.

OSF says its pricing structure for OSF/1 is lower than for System V. A license for a one-user copy of OSF/1 is only $65. However, that's not the end of the story. It is only what the hardware vendor providing OSF/1 on a computer pays to OSF. The vendor must also pay a System V Release 3 license fee of approximately $100 (which includes a volume discount) to AT&T for every copy of OSF/1. By comparison, a multiuser license for SVR4, using the same volume discount and computer list price, is $126–$39 cheaper than OSF/1. OSF/1 only becomes cheaper when the cost of the computer exceeds $280,000, because the SVR4 license fee structure is based upon the list price of the computer.

Finally, operating system success inevitably comes down to applications. If there are no applications for a computer, only a fool or a programmer would buy it. Application developers will get their first crack at porting to OSF/1, or testing to see if their applications will work without porting, when DEC is supposed to provide OSF/1 for the DECstation 3100 sometime this month. With other vendors holding back, and no ABIs for OSF/1, there are bound to be far fewer applications running on OSF/1 than SVR4 at year's end–thousands less.

At the very least, OSF is to be commended for spurring AT&T into greater concern with user and vendor issues. One of the complaints that helped get OSF started was that AT&T had its own hidden agenda for upgrading its operating system, and that end users had no say in priorities, or even getting needed features.

The founding of UNIX International has brought greater responsiveness from AT&T in making decisions about technologies to incorporate. The entire pace of development has heated up. The end result, for all users, is a much better UNIX operating system. Even though there is no single version of UNIX, we now have two that are closer than ever before. ■

Editor-at-Large Rik Farrow is a software engineer, and author of a book on UNIX system security. He teaches courses on system administration and system security.