



AIX 4.1 Now Implements XPG4

By Mark S. Brown

One major change in AIX Version 4.1 is the move to conform to X/Open's Portability Guide, Issue 4, popularly known as XPG4. This standard and its upcoming revision, currently known as SPEC 1170, consolidate the UNIX variations offered today. XPG4 and SPEC 1170 are the standards used in AIX to ensure application portability in the open systems market.

X/Open is a consortium composed of most leading UNIX systems suppliers, software vendors, and user groups. Unlike other open standards organizations such as the Institute of Electrical and Electronics Engineers (IEEE) or American National Standards Institute (ANSI®) that only produce standards guidelines, X/Open also provides certification and testing. X/Open test suites are used to certify conformance to these standards. Representatives from the X/Open members form working groups to produce open specifications for operating system interfaces. The X/Open body then votes on whether to accept these specifications.

The *Portability Guides* are collections of X/Open-approved specifications for a given product. Issue 3 (also known as XPG3) contains standards for a set of basic system interfaces and commands known as the *Base*—commonly used programming languages, curses, terminal interfaces, and international language-support features.

X/Open prefers to build specifications based on industry standards already in place, and then expand on those standards (without contradicting them) to meet the needs of its members and the market.

Although XPG3 was based primarily on the System V Interface Definition (SVID)—the only publicly distributed UNIX operating system specification available at the time the *X/Open Portability Guides* were being developed—XPG3 was still considered too proprietary by some. At the time,

IEEE POSIX 1003.1 was being developed, and ANSI was finalizing the C language standard. As these standards became accepted, X/Open realized the need to update XPG. The membership also believed that although POSIX and ANSI C were valuable, the two standards were not useful to application developers because they left too much unspecified or undefined, including some functions that UNIX developers relied on daily.

XPG4 and SPEC 1170

XPG4 was developed to address these issues. It is based on POSIX 1003.1 and the new 1003.2 (shell and utilities) standards, with additions from SVID Issue 3 and some invention by the working group to complete gaps in the other standards. It is a superset standard meant to provide a useful set of interfaces available on all certified operating systems. If vendors adhere to XPG4, application developers can develop their products, including previously used interfaces, without fearing the "slight difference" in code portability that always seems to occur between two vendor platforms. For this reason alone, IBM considers XPG4 an important standard for AIX 4.1 to meet.

Users and developers, however, believe that even XPG4 is incomplete, leaving too many uncertainties in UNIX system portability. This belief has led five of the leading UNIX system vendors, including IBM, to sponsor an updated and expanded specification, currently known as SPEC 1170, that will be even more comprehensive. It is hoped that X/Open will adopt SPEC 1170 as the next version of XPG and make it a requirement for using the UNIX trademark. As one of the sponsors of SPEC 1170, IBM intends to develop AIX to meet this specification for easier application portability to and from AIX.



Mark S. Brown

XPG4 Components

Although XPG4 is composed of several different standards, the most important is the Base Specification, which consists of three documents: *System Interfaces and Headers*, *Shell and Utilities*, and *Base System Definitions*. A companion volume, *XPG3-XPG4 Base Migration Guide*, lists the differences between XPG3 and XPG4 for developers who are moving applications. For vendors who wish to certify their systems as conforming, a system conformance profile book describes conforming system configurations.

Currently, there are two different levels of certification for XPG4. Since the commands specification was based on the new 1003.2 standard, X/Open has not yet offered a certification suite for it. Therefore, systems branded "XPG4" are only required to have XPG3 commands. Once the test suite is available, systems will be required to re-brand to XPG4 command level. AIX 4.1 commands and interfaces, already enhanced to meet the XPG4 level, will be certified for commands when the test suite is available.

These standards are modeled on the 1003.1 and 1003.2 standards, even though they are a proper superset. The third standard avoids duplicating certain material that would otherwise have been described in both standards.

System Interfaces and Headers

Because IBM has already implemented changes in Version 3 of AIX to meet XPG4, AIX 4.1 interfaces covered by the XPG4 will look similar to Version 3.2.4 or 3.2.5 interfaces today. In fact, none of the changes required to meet certification will break binary portability. Application sources will require only minor changes to meet the new function prototyping or minor header file modifications. One exception is the curses library discussed later in this article.

Most XPG4 work was "under the covers" or in the development of new interfaces, such as the new regular expression interfaces required by 1003.2. There are new interfaces in AIX 4.1 for shell globbing and parsing, allowing applications to parse input in a similar way to shells. Developers whose applications accept regular or extended regular expressions are encouraged to evaluate these interfaces for their reliability and portability. The portion of the work in international locale support was greatly revised and expanded by 1003.2 and XPG4. Locale source files can now provide much cultural variation, such as multilevel sorting ability.

XPG4 will make locale source files more portable between locales and between platforms. However, locale description file binaries must be remade from the locale source file when moved from platform to platform. A few new interfaces in this area cover date and time functions and character I/O in a more portable manner. The `iconv` converter utility has also been enhanced to be more useful.

The Base specifies several traditional interfaces not found in the POSIX or ANSI standards, such as shared memory and encrypt/decrypt. The intent is to codify what users and developers expect to see in a UNIX-based system.

Shell and Utilities

Users will see the largest difference resulting from the move to XPG4, which is a direct result of XPG4 and compliance with 1003.2. The Shell and Utilities volume codifies the behavior of the shell and standardizes the behavior of over 100 of the most commonly used system commands—from `at` to `yacc`. While much of the groundwork for these changes was laid in AIX 3.2, more was needed to truly meet the specifications.

The default AIX XPG4 shell is similar to the Korn Shell (`ksh`), with some modifications in script parsing, regular expression handling, and a few other areas. Dave Korn has developed a new version of `ksh` that conforms to XPG4; IBM intends to support this new version for AIX. (For AIX 4.1, IBM modified the old version of `ksh` to conform to XPG4.) This version specifies several shell utilities in more detail than users have seen in any system man page.

Utilities listed in XPG4 must follow a set of utility syntax guidelines that describes how options and option arguments are to be handled. A common set of commands means that users do not have to relearn commands when switching platforms. Certain command input and output specifications differ from AIX 3.2, causing a potential for user shell scripts to break. To avoid this problem, the old syntax was left intact whenever possible. In fact, many commands with changed syntax completely support the older syntax.

Grammars are provided for the shell, `lex`, `awk`, and `yacc`, making the output of these utilities more easily predicted. Exit codes and `stderr` reporting formats are also now more predictable. All XPG4 commands use `getopt()`, a change that was expanded in AIX to include several non-XPG commands. The new backup utility `pax` has also been extended.

A common set of commands means that users do not have to relearn commands when switching platforms.

Users of symbolic links will see standardization in how commands parse symbolic links and how interfaces expand pathnames containing symbolic links. The addition of the `lstat()` system call gives status such as the `stat()` interface, with added information for `symlink` files. There is now a consistent behavior based on the Berkeley Software Distribution (BSD) style of symlinks applicable to all commands that may need to handle a symbolic link reference in a special manner (such as `rm`, `find`, or `ls`).

Most changes required by the XPG commands specification are not actually changes from previous standards, but an attempt to codify and make portable behavior that has never been documented explicitly. Users who have argued that AIX behaved in a non-standard manner will be greatly pleased by the changes in Version 4.1, most of them based on this volume alone. In addition, since Sun, Novell/USL, DEC®, IBM, HP, Bull®, Siemens™, and many other major system vendors adhere to or plan to meet this specification, portable scripting between systems will be even more practical than ever.

SPEC 1170 and the Future

The next move for AIX is to meet the emerging SPEC 1170 specification, named for its 1,170 interfaces at the time of initial industry review. SPEC 1170 was intended to answer the market demand to make the UNIX environment more open and portable through further specification. SPEC 1170 supersedes XPG4 by adding support for Streams, sockets, XTI, an enhanced version of curses for the international community, and many additional SVID3 functions—including signals interfaces, memory mapping, `password/group/utmp` file handling, TCP/IP, and many more. Many of the specifications of SPEC 1170 have already been met in AIX or have been added in Version 4.1. For instance, AIX 4.1 currently supports Streams, XTI, sockets, and most C library functions, with curses on the way. However, since SPEC 1170 was still undergoing review and revision as IBM prepared to ship AIX 4.1, IBM expects to make more changes in future updates of AIX so that it will conform with the final SPEC 1170.

System V Curses

Another big standards change in AIX 4.1 is in the curses library `libcurses.a`. This change is IBM's

response to the many requests from users and developers to update its `libcurses/terminal` interfaces package. IBM could not simply follow SPEC 1170 requirements to meet this request, since SPEC 1170 was still undergoing major revisions in the curses area during the development period of AIX 4.1. Instead, IBM chose System V R3/R4 curses because this source contained most of the changes requested by customers. In addition, using the latest System V sources would make it easier to modify AIX 4.1 to meet the full SPEC 1170 requirements when they became final.

Because several major structural elements and data definitions in these new curses are different from the AIX Version 3 curses, application programs may need modification when recompiled on AIX Version 4. Only recompiles are mentioned, since AIX Version 4 maintains full binary compatibility with earlier releases using shared libraries and compatibility install options provided with the base system package.

When the version of curses that conforms to SPEC 1170 is available, it will be fully backward compatible with AIX 4.1. It will also contain new functions, such as integrated internationalization support and color. Other benefits will include more application portability, more standard curses, and `terminfo` tools to make portable configuration and installation procedures easier for users and administrators.

Summary

Moving to XPG4 is a major step for AIX, providing many benefits to users and developers seeking cross-platform portability and an industry-standard UNIX operating system on the RISC System/6000. Although AIX will soon conform to SPEC 1170, it will not leave current users behind.



Mark S. Brown, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Internet: mbrown@austin.ibm.com. Mr. Brown is a senior programmer with the AIX Base OS Development group, working in AIX BOS development since 1988. He holds a BS in Computer Science from Southwest Texas State University in San Marcos.

SPEC 1170 was intended to answer the market demand to make the UNIX environment more open and portable through further specification.