INTERNATIONAL STANDARD ISO/IEC 9945-3:2003
TECHNICAL CORRIGENDUM 1

Published 2004-09-15

Information technology — Portable Operating System Interface (POSIX®) —

Part 3:
Shell and Utilities

TECHNICAL CORRIGENDUM 1

Technologies de l'information — Interface pour la portabilité des systèmes (POSIX®) —

Partie 3: Enveloppe et services

RECTIFICATIF TECHNIQUE 1

Technical Corrigendum 1 to ISO/IEC 9945-3:2003 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 22, *Programming languages, their environments and system software interfaces.*

ICS 35.060

Ref. No. ISO/IEC 9945-3:2003/Cor.1:2004(E)

Published in Switzerland

1 Scope

This technical corrigendum addresses issues raised in defect reports and interpretation requests submitted up to 14 August 2003, and that meet all of the following criteria:

a. They are in the scope of the approved International Standard.

b. They contain no new APIs (functions/utilities), however, they may add enumeration symbols, non-function # defines, and reserve additional namespaces.

c. They address contradictions between different parts of the International Standard, or add consistency between it and overriding International Standards, or address security-related problems.

2 Changes to ISO/IEC 9945-3

Change Number: XCU/TC1/1 [XBD ERN 20]
On Page: xxxii Line: "ISO/IEC 8859" Section: Referenced Documents
Add after line starting "Part 10":
"Part 11: Latin/Thai Alphabet"
Add after line starting "Part 15":
"Part 16: Latin Alphabet No. 10"

Change Number: XCU/TC1/2 [XCU ERN 29]
On Page: 5 Line: 139-150,170 Section: 1.7.1.4
In Table 1-1, change the dash (-) entries to letter "U" on lines 140, 141 and 150 (rows for types B, C and T respectively).
Change all the dash (-) entries in the T column to letter "U".
Add after line 170
"U   The effect is unspecified unless specified by the utility description."

Change Number: XCU/TC1/3 [XRAT ERN 2]
On Page: 41 Line: 1704 Section: 2.6.4
Add a new paragraph before the existing paragraph that commences "As an extension, the shell may recognize arithmetic expressions beyond those listed...":
"All changes to variables in an arithmetic expression shall be in effect after the arithmetic expansion, as in the parameter expansion ${x=value}.
If the shell variable x contains a value that forms a valid integer constant then the arithmetic expansions $((x)) and $(($x)) shall return the same value."

Change Number: XCU/TC1/4 [XCU ERN 17]
On Page: 48 Line: 1986-1990 Section: 2.9.1.1
In Section 2.9.1.1 Command Search and Execution
Change From:
"If the execve() function fails due to an error equivalent to the [ENOEXEC] error defined in the System Interfaces volume of IEEE Std 1003.1-2001, the shell shall execute a command equivalent to having a shell invoked with the command name as its first operand, with any remaining arguments passed to the new shell."
To:
"If the execve() function fails due to an error equivalent to the [ENOEXEC] error defined in the System Interfaces volume of IEEE Std 1003.1-2001, the shell shall execute a command equivalent to having a shell invoked with the pathname resulting from the search as its first operand, with any remaining arguments passed to the new shell, except that the value of $0 in the new shell may be set to the command name."

Change Number: XCU/TC1/5 [XCU ERN 30]
On Page: 212 Line: 8323-8324 Section: c99
In the OPTIONS section
Change From:
"If a directory specified by the -L option contains files named libc.a, libm.a, libl.a, or liby.a, the results are unspecified."
To:
"If a directory specified by the -L option contains files with names starting with any of the strings "libc.", "libl.", "libpthread.", "libm.", "librt.", "libtrace.", "libxnet.", or "liby.", the results are unspecified."

Change Number: XCU/TC1/6 [XCU ERN 19]
On Page: 215,218 Line: 8466-8467,8583-8588 Section: c99
In the EXTENDED DESCRIPTION section, Programming Environments
Change From:
"shall be output by a getconf command using the _POSIX_V6_WIDTH_RESTRICTED_ENVS argument."
To:
"shall be output by a getconf command using the POSIX_V6_WIDTH_RESTRICTED_ENVS argument, as a <newline>-separated list of names suitable for use with the getconf -v option."
(Note the removed leading underscore on POSIX_V6_WIDTH_RESTRICTED_ENVS.)
In the EXAMPLES section
Change From:

```
"# ... if there are no additional constraints, the first one will do:

CENV=$(getconf _POSIX_V6_WIDTH_RESTRICTED_ENVS | head -n l)

# ... or, if an environment that supports large files is preferred,

# look for names that contain "OFF64" or "OFFBIG". (This chooses

# the last one in the list if none match.)

for CENV in $(getconf _POSIX_V6_WIDTH_RESTRICTED_ENVS)"
```

To:

```
"# ... if there are no additional constraints, the first one will do:

CENV=$(getconf POSIX_V6_WIDTH_RESTRICTED_ENVS | head -n l)

# ... or, if an environment that supports large files is preferred,

# look for names that contain "OFF64" or "OFFBIG". (This chooses

# the last one in the list if none match.)

for CENV in $(getconf POSIX_V6_WIDTH_RESTRICTED_ENVS)            "
```

Rationale:
This change corrects the name of the variable not to include the underscore.

Change Number: XCU/TC1/7 [XCU ERN 33]
On Page: 343 Line: 13238-13240,13451 Section: ed
In the EXTENDED DESCRIPTION section (List command)
Delete the sentence:
"If the size of a byte on the system is greater than nine bits, the format used for non-printable characters is implementation-defined."
In the RATIONALE section:
Change From:
"Systems with bytes too large to fit into three octal digits must devise other means of displaying non-printable characters. Consideration was given to requiring that the number of octal digits be large enough to hold a byte, but this seemed to be too confusing for applications on the vast majority of systems where three digits are adequate. It would be theoretically possible for the application to use the getconf utility to find out the CHAR_BIT value and deal with such an algorithm; however, there is really no portable way that an application can use the octal values of the bytes across various coded character sets, so the additional specification was not worthwhile."
To:
"Prior versions of this standard allowed for implementations with bytes other than eight bits, but this has been modified in this version."

Change Number: XCU/TC1/8 [XCU ERN 31]
On Page: 370 Line: 14333-14334 Section: ex
In the EXTENDED DESCRIPTION section
Change From:
"Implementations may restrict the set of characters accepted in lhs or rh, except that printable characters and <blank>s shall not be restricted. Additional restrictions shall be implementation-defined."
To:
"Implementations may restrict the set of characters accepted in lhs or rhs, except that printable characters and <blank>s shall not be restricted. Additional restrictions shall be implementation-defined."
Note to editor: The change is from "rh" to "rhs".

Change Number: XCU/TC1/9 [XCU ERN 33]
On Page: 380 Line: 14705-14709,16154 Section: ex
In the EXTENDED DESCRIPTION section (Print command)
Delete the sentence:
"If the size of a byte on the system is greater than 9 bits, the format used for non-printable characters is implementation-defined."
Add to the RATIONALE section:
"Prior versions of this standard allowed for implementations with bytes other than eight bits, but this has been modified in this version."

Change Number: XCU/TC1/10 [XCU ERN 26,28]
On Page: 447-448 Line: 17290-17292,17297-17300,17317-17322 Section: file
On lines 17290-17292
Change From:
"The type shall consist of the type specification characters c, d, f, s, and u, specifying character, signed decimal, floating point, string, and unsigned decimal, respectively."
To:
"The type shall consist of the type specification characters d, s, and u, specifying signed decimal, string, and unsigned decimal, respectively."
On lines 17297-17300
Change From:
"The type specification characters d, f, and u can be followed by an optional unsigned decimal integer that specifies the number of bytes represented by the type. The type specification character f can be followed by an optional F, D, or L, indicating that the value is of type float, double, or long double, respectively."
To:
"The type specification characters d and u can be followed by an optional unsigned decimal integer that specifies the number of bytes represented by the type."
On lines 17317-17322
Delete the paragraph:
"For the type specifier f, the default number of bytes shall correspond to the number of bytes in the basic double precision floating-point data type of the underlying implementation. The implementation shall support values of the optional number of bytes to be converted corresponding to the number of bytes in the C-language types float, double, and long double. These numbers can also be specified by an application as the characters F , D , and L , respectively."

Change Number: XCU/TC1/11 [XCU ERN 27]
On Page: 447-448 Line: 17280 Section: file
In the EXTENDED DESCRIPTION section
On line 17280
Change From:
"Each line shall be composed of the following four <blank>-separated fields:"
To:
"Each line shall be composed of the following four <tab>-separated* fields:
 Footnote:* Implementations may allow any combination of one or more white space characters other than <newline> to act as field separators."
On lines 17336-17341
Change From:
"\character

The backslash-escape sequences as specified in the Base Definitions volume of IEEE Std 1003.1-2001, Table 5-1, Escape Sequences and Associated Actions ('\\' , '\a' , '\b' , '\f' , '\n' , '\r' , '\t' , '\v'). The results of using any other character, other than an octal digit, following the backslash are unspecified."

To:

"\character

The backslash-escape sequences as specified in the Base Definitions volume of IEEE Std 1003.1-2001, Table 5-1, Escape Sequences and Associated Actions ('\\' , '\a' , '\b' , '\f' , '\n' , '\r' , '\t' , '\v'). In addition, the escape sequence '\ ' (the <backslash> character followed by the <space> character) shall be recognized to represent a <space> character. The results of using any other character, other than an octal digit, following the backslash are unspecified."

Change Number: XCU/TC1/12 [XCU ERN 33]
On Page: 448 Line: 17345-17347,17462 Section: file
In the EXTENDED DESCRIPTION section
Delete the sentence:
"If the size of a byte on the system is greater than 9 bits, the valid escape sequence used to represent a byte implementation-defined."
In the RATIONALE section
Add the following text at the end:
"Prior versions of this standard allowed for implementations with bytes other than eight bits, but this has been modified in this version."

Change Number: XCU/TC1/13 [XCU ERN 4]
On Page: 458 Line: 17761-17762,17765-17766 Section: find
In the RATIONALE section
Change From:
"The descriptions of -atime, -ctime, and -mtime were changed from the SVID descrption of n "days" to "24-hour period"."
To:
"The descriptions of -atime, -ctime, and -mtime were changed from the SVID descrption of n "days" to n being the result of the integer division of the time difference in seconds by 86400".
Change From:
"For example, -atime 3 is true if the file was accessed any time in the period from 72 hours to 48 hours ago."
To:
"For example, -atime 2 is true if the file was accessed any time in the period from 72 hours to 48 hours ago."
Rationale:
The rationale for -atime/-mtime/-ctime was updated to match changes made in the normative text since POSIX.2-1992.

Change Number: XCU/TC1/14 [XCU ERN 33]
On Page: 540 Line: 20802-20806,21020-21023 Section: lex
In the EXTENDED DESCRIPTION section (Table 4-10)
Delete the sentence:
"If the size of a byte on the system is greater than nine bits, the escape sequence used to represent a byte is implementation-defined."
In the RATIONALE section
Change from:
"The description of octal and hexadecimal-digit escape sequences agrees with the ISO C standard use of escape sequences. See the RATIONALE for ed for a discussion of bytes larger than 9 bits being represented by octal values. Hexadecimal values can represent larger bytes and multi-byte characters directly, using as many digits as required."
To:
"The description of octal and hexadecimal-digit escape sequences agrees with the ISO C standard use of escape sequences.
Prior versions of this standard allowed for implementations with bytes other than eight bits, but this has been modified in this version."

Change Number: XCU/TC1/15 [XCU ERN 18]
On Page: 559 Line: 21535 Section: localedef
In the OPERANDS section
Change From:
"This capability may be restricted to users with appropriate privileges."
To:
"The ability to create public locales in this way may be restricted to users with appropriate privileges."

Change Number: XCU/TC1/16 [XCU ERN 13]
On Page: 560-561 Line: 21593-21594,21611 Section: localedef
In the EXTENDED DESCRIPTION section
Change From:
"If a non-printable character in the charmap has a width specified that is not -1, localedef shall generate a warning."
To:
"If a non-printable character in the charmap has a width specified that is not -1 the result will be undefined."
In the CONSEQUENCES OF ERRORS section
Delete the fourth bullet item:
"* If a non-printable character has a width specified other than -1."
Rationale:
The previous text was conflicting with the descriptions of WIDTH and WIDTH_DEFAULT in XBD.

Change Number: XCU/TC1/17 [XCU ERN 34]
On Page: 593 Line: 22946-22947 Section: mailx
In the EXTENDED DESCRIPTION section (Internal Variables in Mailx)
Change From:
"The defaults specified here may be changed by the implementation-defined system start-up file unless the user specifies the -n option."
To:
"The defaults specified here may be changed by the unspecified system start-up file unless the user specifies the -n option."
Rationale: The change is made for consistency with elsewhere in the EXTENDED DESCRIPTION.

Change Number: XCU/TC1/18 [XCU ERN 35]
On Page: 666 Line: 25769-25772 Section: nice
In the RATIONALE section
Delete the paragraph:
"Due to the text about the limits of the nice value being implementation-defined, nice is not actually required to change the nice value of the executed command; the limits could be zero differences from the system default, although the implementor is required to document this fact in the conformance document."

Change Number: XCU/TC1/19 [XCU ERN 33]
On Page: 682 Line: 26379-26380,26539 Section: od
In the EXTENDED DESCRIPTION section
Delete the sentence:
"If the size of a byte on the system is greater than nine bits, the format used for non-printable characters is implementation-defined."
Add to the RATIONALE section:
"Prior versions of this standard allowed for implementations with bytes other than eight bits, but this has been modified in this version."

Change Number: XCU/TC1/20 [XCU ERN 21]
On Page: 701 Line: 27113 Section: pax
In the SYNOPSIS section
Change From:
"pax -r -w[-diklntuvX][-H|-L][-p string]...[-s replstr]..."
To:
"pax -r -w[-diklntuvX][-H|-L][-o options]...[-p string]...[-s replstr]..."
Change From:
"pax [-cdnv][-H|-L][-f archive][-s replstr]...[pattern...]"

To:
"pax [-cdnv][-H|-L][-o options][-f archive][-s replstr]...[pattern...]"
Rationale:
The SYNOPSIS was inconsistent with the normative text.

Change Number: XCU/TC1/21 [XCU ERN 24]
On Page: 702 Line: 27167-27168 Section: pax
Change From:
"If the selected archive format supports the specification of linked files, it shall be an error if these files cannot be linked when the archive is extracted."
To:
"If the selected archive format supports the specification of linked files, it shall be an error if these files cannot be linked when the archive is extracted, except that if the files to be linked are symbolic links and the system is not capable of making hard links to symbolic links then separate copies of the symbolic link shall be created instead."

Change Number: XCU/TC1/22 [XCU ERN 22]
On Page: 704 Line: 27266 Section: pax
In the OPTIONS section,
Add new text to the end of the subsection for "delete=pattern"
"When multiple -odelete=pattern options are specified, the patterns shall be additive; all keywords matching the specified string patterns shall be omitted from extended header records that pax produces."

Change Number: XCU/TC1/23 [XCU ERN 9]
On Page: 706 Line: 27336-27339 Section: pax
In the OPTIONS section
Change From:
"write   In read or copy mode, pax shall write the file, translating or truncating the name, regardless of whether this may overwrite an existing file with a valid name. In list mode, pax shall behave identically to the bypass action."
To:
"write   In read or copy mode, pax shall write the file, translating the name, regardless of whether this may overwrite an existing file with a valid name. In list mode, pax shall behave identically to the bypass action."
Rationale:
The text currently written in the standard allows the pax utility to truncate pathnames of files being extracted from an archive or copied from one place to another in arbitrary ways, even if this action overwrites existing files. We agree that translation may cause this to occur, but the pax utility should not truncate file names on its own. It should be capable of restoring files even if the pathname in the archive is longer than {PATH_MAX}. (This may require the pax utility to perform one or more change directory operations to get to a point where the file being extracted can be created using a pathname shorter than {PATH_MAX}.) If the underlying file system performs filename truncation (which is not allowed on file systems conforming to this standard), the pax utility need not attempt to prevent or detect the action. But, the pax utility itself should never truncate a pathname component.

Change Number: XCU/TC1/24 [XCU ERN 7]
On Page: 708 Line: 27469 Section: pax
In the OPTIONS section
On line 27469
add as a new paragraph before the paragraph beginning "The options that operate on the names of files..."
"Specifying more than one of the mutually-exclusive options -H and –L shall not be considered an error and the last option specified shall determine the behavior of the utility."

Change Number: XCU/TC1/25 [XCU ERN 8]
On Page: 715 Line: 27723-27727 Section: pax
In the EXTENDED DESCRIPTION section (Pax Extended Header)
Delete the paragraph at line 27723
"ctime   The file creation time for the following file(s), equivalent to the value of the st_ctime member of the stat structure for a file, as described by the stat() function. The creation time shall be restored if the process has the appropriate privilege required to do so. The format of the <value> shall be as described in pax Extended Header File Times (on page 717)."

Rationale:
There is a contradiction in the definition of the ctime keyword for pax's extended header, in that the st_ctime member of the stat structure does not refer to a file creation time. No field in the standard stat structure from <sys/stat.h> includes a file creation time.

Change Number: XCU/TC1/26 [XCU ERN 23]
On Page: 720 Line: 27939-27940,27961 Section: pax
In the EXTENDED DESCRIPTION section (ustar Interchange Format)
On lines 27939-27940
Change From:
"Such files are identified by each file having the same device and file serial number."
To:
"Such files are identified by having the same device and file serial numbers, and pathnames that refer to different directory entries. All such files shall be archived as linked files."
Before line 27961
Insert the following paragraph:
"It is unspecified whether files with pathnames that refer to the same directory entry are archived as linked files or as separate files. If they are archived as linked files this means that attempting to extract both pathnames from the resulting archive will always cause an error (unless the -u option is used) because the link cannot be created.
It is unspecified whether files with the same device and file serial numbers being appended to an archive are treated as linked files to members that were in the archive before the append."

Change Number: XCU/TC1/27 [XCU ERN 10]
On Page: 722,724 Line: 28039-28040,28099 Section: pax
In the EXTENDED DESCRIPTION section (cpio header)
Change From:
"c_nlink   Contains the number of links referencing the file at
the time the archive was created."
To:
"c_nlink   Contains a number greater than or equal to the number of links in the archive referencing the file. If the -a option is used to append to a cpio archive, then the pax utility need not account for the files in the existing part of the archive when calculating the c_nlink values for the appended part of the archive, and need not alter the c_nlink values in the existing part of the archive if additional files with the same c_dev and c_ino values are appended to the archive."
In the APPLICATION USAGE section
Insert before line 28099
"Caution is advised when using the -a option to append to a cpio format archive. If any of the files being appended happen to be given the same c_dev and c_ino values as a file in the existing part of the archive, then they may be treated as links to that file on extraction. Thus it is risky to use -a with cpio format except when it is done on the same system that the original archive was created on, and with the same pax utility, and in the knowledge that there has been little or no file system activity since the original archive was created that could lead to any of the files being appended being given the same c_dev and c_ino values as an unrelated file in the existing part of the archive. Also, when (intentionally) appending additional links to a file in the existing part of the archive, the c_nlink values in the modified archive can be smaller than the number of links to the file in the archive, which may mean that the links are not preserved on extraction."

Change Number: XCU/TC1/28 [XCU ERN 33]
On Page: 845 Line: 32699-32701,32892 Section: sed
In the EXTENDED DESCRIPTION section
Delete the sentence:
"If the size of a byte on the system is greater than nine bits, the format used for non-printable characters is implementation-defined."
Add to the RATIONALE section:
"Prior versions of this standard allowed for implementations with bytes other than eight bits, but this has been modified in this version."

Change Number: XCU/TC1/29 [XCU ERN 11]
On Page: 846 Line: 32770-32771 Section: sed
Change From:
"If the delimiter is not n, within string1 and string2, the delimiter itself can be used as a literal character if it preceeded by a backslash."
To:
"If the delimiter is not 'n', within string1 and string2, the delimiter itself can be used as a literal character if it preceeded by a backslash."
Note to editor the change is from italic n to a quoted character n.

Change Number: XCU/TC1/30 [XCU ERN 14]
On Page: 907-905 Line: 35124-35169 Section: test
In the OPERANDS section
Change From:
"-b file
   True if file exists and is a block special file.
-c file
   True if file exists and is a character special file.
-d file
   True if file exists and is a directory.
-e file
   True if file exists.
-f file
   True if file exists and is a regular file.
-g file
   True if file exists and its set-group-ID flag is set.
-h file
   True if file exists and is a symbolic link.
-L file
   True if file exists and is a symbolic link.
-n string
   True if the length of string is non-zero.
-p file
   True if file is a FIFO.
-r file
   True if file exists and is readable. True shall indicate that permission to read from file will be granted, as defined in Section 1.7.1.4 (on page 4).
-S file
   True if file exists and is a socket.
-s file
   True if file exists and has a size greater than zero.
-t file_descriptor
   True if the file whose file descriptor number is file_descriptor is open and is associated with a terminal.
-u file
   True if file exists and its set-user-ID flag is set.
-w file
   True if file exists and is writable. True shall indicate that permission to write from file will be granted, as defined in Section 1.7.1.4 (on page 4).
-x file
   True if file exists and is executable. True shall indicate that permission to execute file will be granted, as defined in Section 1.7.4 (on page 4). If file is a directory, true shall indicate that permission to search file will be granted.
-z string
   True if the length of string string is zero.
string
   True if the string string is not the null string.
s1 = s2
   True if the strings s1 and s2 are identical.
s1 != s2
   True if the strings s1 and s2 are not identical.

n1 -eq n2
   True if the integers n1 and n2 are algebraically equal.
n1 -ne n2
   True if the integers n1 and n2 are not algebraically equal.
n1 -gt n2
   True if the integer n1 is algebraically greater than the integer n2.
n1 -ge n2
   True if the integer n1 is algebraically greater than or equal to
   the integer n2.
n1 -lt n2
   True if the integer n1 is algebraically less than the integer n2.
n1 -le n2
   True if the integer n1 is algebraically less than or equal to the integer n2.
expression1 -a expression2
   [XSI] True if both expression1 and expression2 are true. The –a binary primary is left associative. It has a higher precedence than
   -o. [/XSI]
expression1 -o expression2
   [XSI] True if either expression1 or expression2 is true. The -o binary primary is left associative. [/XSI]
With the exception of the -h file and -L file primaries, if a file argument is a symbolic link, test shall evaluate the expression by resolving the symbolic link and using the file referenced by the link.
These primaries can be combined with the following operators:
! expression
   True if expression is false.
(expression)
   [XSI] True if expression is true. The parentheses can be used to alter the normal precedence and associativity. [/XSI]"
To:
"-b pathname
  True if pathname resolves to a file that exists and is a block special file. False if pathname cannot be resolved, or if pathname resolves to a file that exists but is not a block special file.
-c pathname
  True if pathname resolves to a file that exists and is a character special file. False if pathname cannot be resolved, or if pathname resolves to a file that exists but is not a character special file.
-d pathname
  True if pathname resolves to a file that exists and is a directory. False if pathname cannot be resolved, or if pathname resolves to a file that exists but is not a directory.
-e pathname
  True if pathname resolves to a file that exists. False if pathname cannot be resolved.
-f pathname
  True if pathname resolves to a file that exists and is a regular file. False if pathname cannot be resolved, or if pathname resolves to a file that exists but is not a regular file.
-g pathname
  True if pathname resolves to a file that exists and has its set-group-ID flag set. False if pathname cannot be resolved, or if pathname resolves to a file that exists but does not have its set-group-ID flag set.
-h pathname
  True if pathname resolves to a file that exists and is a symbolic link. False if pathname cannot be resolved, or if pathname resolves to a file that exists but is not a symbolic link. If the final component of pathname is a symlink that symlink is not followed.
-L pathname
  True if pathname resolves to a file that exists and is a symbolic link. False if pathname cannot be resolved, or if pathname resolves to a file that exists but is not a symbolic link. If the final component of pathname is a symlink that symlink is not followed.
-n string
  True if the length of string is non-zero, otherwise false.
-p pathname
  True if pathname resolves to a file that exists and is a FIFO. False if pathname cannot be resolved, or if pathname resolves to a file that exists but is not a FIFO.
-r pathname

True if pathname resolves to a file that exists and for which permission to read from the file will be granted, as defined in Section 1.7.1.4 (on page 4). False if pathname cannot be resolved, or if pathname resolves to a file for which permission to read from the file will not be granted.

-S pathname

True if pathname resolves to a file that exists and is a socket. False if pathname cannot be resolved, or if pathname resolves to a file that exists but is not a socket.

-s pathname

True if pathname resolves to a file that exists and has a size greater than zero. False if pathname cannot be resolved, or if pathname resolves to a file that exists but does not have a size greater than zero.

-t file_descriptor

True if file descriptor number file_descriptor is open and is associated with a terminal. False if file_descriptor is not a valid file descriptor number, or if file descriptor number file_descriptor is not open, or if it is open but is not associated with a terminal.

-u pathname

True if pathname resolves to a file that exists and has its set-user-ID flag set. False if pathname cannot be resolved, or if pathname resolves to a file that exists but does not have its set-user-ID flag set.

-w pathname

True if pathname resolves to a file that exists and for which permission to write to the file will be granted, as defined in Section 1.7.1.4 (on page 4). False if pathname cannot be resolved, or if pathname resolves to a file for which permission to write to the file will not be granted.

-x pathname

True if pathname resolves to a file that exists and for which permission to execute the file (or search it, if it is a directory) will be granted, as defined in Section 1.7.1.4 (on page 4). False if pathname cannot be resolved, or if pathname resolves to a file for which permission to execute (or search) the file will not be granted.

-z string

True if the length of string string is zero, otherwise false.

string

True if the string string is not the null string, otherwise false.

s1 = s2

True if the strings s1 and s2 are identical, otherwise false.

s1 != s2

True if the strings s1 and s2 are not identical, otherwise false.

n1 -eq n2

True if the integers n1 and n2 are algebraically equal, otherwise false.

n1 -ne n2

True if the integers n1 and n2 are not algebraically equal, otherwise false.

n1 -gt n2

True if the integer n1 is algebraically greater than the integer n2, otherwise false.

n1 -ge n2

True if the integer n1 is algebraically greater than or equal to the integer n2, otherwise false.

n1 -lt n2

True if the integer n1 is algebraically less than the integer n2, otherwise false.

n1 -le n2

True if the integer n1 is algebraically less than or equal to the integer n2, otherwise false.

expression1 -a expression2

[XSI] True if both expression1 and expression2 are true, otherwise false. The -a binary primary is left associative. It has a higher precedence than -o. [/XSI]

expression1 -o expression2

[XSI] True if either expression1 or expression2 is true, otherwise false. The -o binary primary is left associative. [/XSI]

With the exception of the -h file and -L file primaries, if a file argument is a symbolic link, test shall evaluate the expression by following the symbolic link and using the file referenced by the link.

These primaries can be combined with the following operators:

! expression

True if expression is false. False if expression is true.

(expression)

[XSI] True if expression is true. False if expression is false. The arentheses can be used to alter the normal precedence and ssociativity. [/XSI]

Rationale:

The previous descriptions of the file existence primaries for the test utility have been clarified.

Change Number: XCU/TC1/31 [XCU ERN 33]
On Page: 927 Line: 35871-35872,36041 Section: tr
In the EXTENDED DESCRIPTION section
Delete the sentence:
"If the size of a byte on the system is greater than nine bits, the valid escape sequence to represent a byte is implementation-defined."
Add to the RATIONALE section:
"Prior versions of this standard allowed for implementations with bytes other than eight bits, but this has been modified in this version."

Change Number: XCU/TC1/32 [XCU ERN 12]
On Page: 930 Line: 35992 Section: tr
In the EXAMPLES section
Change From:
```
"tr "[=e=]" e <file1 >file2"
```

To:
```
"tr "[=e=]" "[e*]" <file1 >file2"
```

Rationale:
The previous example depended on unspecified behavior.

Change Number: XCU/TC1/33 [XCU ERN 20]
On Page: 931 Line: 36041 Section: tr
In the RATIONALE section,
Change From:
"IEEE Std 1003.1-2001 specifies that octal sequences always refer to single byte binary values."
To:
"IEEE Std 1003.1-2001 specifies that octal sequences always refer to single byte binary values when used to specify an endpoint of a range of collating elements."
Rationale:
This clarification is for consistency with the normative text.

Change Number: XCU/TC1/34 [XCU ERN 35]
On Page: 945 Line: 36539-36540 Section: umask
In the RATIONALE section
Change From:
"The default output style is implementation-defined to permit implementers to provide migration to the new symbolic style at the time most appropriate to their users."
To:
"The default output style is unspecified to permit implementers to provide migration to the new symbolic style at the time most appropriate to their users."
Rationale: The change is made for consistency with the normative text.

Change Number: XCU/TC1/35 [XCU ERN 16]
On Page: 969 Line: 37351 Section: uudecode
In the DESCRIPTION section
Add at the end of the first paragraph:
"If either of the op characters '+' and '-' (see chmod) are specified in symbolic mode, the initial mode on which those operations are based is unspecified."
Note to editors: "op" is in italics