

109042 **NAME**

109043 pax — portable archive interchange

109044 **SYNOPSIS**109045 pax [-dv] [-c|-n] [-H|-L] [-o options] [-f archive] [-s replstr]...
109046 [pattern...]109047 pax -r[-c|-n] [-dikuv] [-H|-L] [-f archive] [-o options]... [-p string]...
109048 [-s replstr]... [pattern...]109049 pax -w [-dituvX] [-H|-L] [-b blocksize] [[-a] [-f archive]] [-o options]...
109050 [-s replstr]... [-x format] [file...]109051 pax -r -w [-dikltuvX] [-H|-L] [-o options]... [-p string]... |
109052 [-s replstr]... [file...] directory109053 **DESCRIPTION**109054 The *pax* utility shall read, write, and write lists of the members of archive files and copy
109055 directory hierarchies. A variety of archive formats shall be supported; see the *-x format* option.109056 The action to be taken depends on the presence of the *-r* and *-w* options. The four combinations
109057 of *-r* and *-w* are referred to as the four modes of operation: **list**, **read**, **write**, and **copy** modes,
109058 corresponding respectively to the four forms shown in the SYNOPSIS section.109059 **list** In **list** mode (when neither *-r* nor *-w* are specified), *pax* shall write the names of
109060 the members of the archive file read from the standard input, with pathnames
109061 matching the specified patterns, to standard output. If a named file is of type
109062 directory, the file hierarchy rooted at that file shall be listed as well.109063 **read** In **read** mode (when *-r* is specified, but *-w* is not), *pax* shall extract the members of
109064 the archive file read from the standard input, with pathnames matching the
109065 specified patterns. If an extracted file is of type directory, the file hierarchy rooted
109066 at that file shall be extracted as well. The extracted files shall be created performing
109067 pathname resolution with the directory in which *pax* was invoked as the current
109068 working directory.109069 If an attempt is made to extract a directory when the directory already exists, this
109070 shall not be considered an error. If an attempt is made to extract a FIFO when the
109071 FIFO already exists, this shall not be considered an error.109072 The ownership, access, and modification times, and file mode of the restored files
109073 are discussed under the *-p* option.109074 **write** In **write** mode (when *-w* is specified, but *-r* is not), *pax* shall write the contents of
109075 the *file* operands to the standard output in an archive format. If no *file* operands are
109076 specified, a list of files to copy, one per line, shall be read from the standard input
109077 and each entry in this list shall be processed as if it had been a *file* operand on the
109078 command line. A file of type directory shall include all of the files in the file
109079 hierarchy rooted at the file.109080 **copy** In **copy** mode (when both *-r* and *-w* are specified), *pax* shall copy the *file* operands
109081 to the destination directory.109082 If no *file* operands are specified, a list of files to copy, one per line, shall be read
109083 from the standard input. A file of type directory shall include all of the files in the
109084 file hierarchy rooted at the file.109085 The effect of the **copy** shall be as if the copied files were written to a *pax* format
109086 archive file and then subsequently extracted, except that copying of sockets may be

109087 supported even if archiving them in write mode is not supported, and that there
 109088 may be hard links between the original and the copied files. If the destination
 109089 directory is a subdirectory of one of the files to be copied, the results are
 109090 unspecified. If the destination directory is a file of a type not defined by the System
 109091 Interfaces volume of POSIX.1-202x, the results are implementation-defined;
 109092 otherwise, it shall be an error for the file named by the *directory* operand not to
 109093 exist, not be writable by the user, or not be a file of type directory.

109094 In **read** or **copy** modes, if intermediate directories are necessary to extract an archive member,
 109095 *pax* shall perform actions equivalent to the *mkdir()* function defined in the System Interfaces
 109096 volume of POSIX.1-202x, called with the following arguments:

- 109097 • The intermediate directory used as the *path* argument
- 109098 • The value of the bitwise-inclusive OR of S_IRWXU, S_IRWXG, and S_IRWXO as the *mode*
 109099 argument

109100 If any specified *pattern* or *file* operands are not matched by at least one file or archive member,
 109101 *pax* shall write a diagnostic message to standard error for each one that did not match and exit
 109102 with a non-zero exit status.

109103 The archive formats described in the EXTENDED DESCRIPTION section shall be automatically
 109104 detected on input. The default output archive format shall be implementation-defined.

109105 A single archive can span multiple files. The *pax* utility shall determine, in an implementation-
 109106 defined manner, what file to read or write as the next file.

109107 If the selected archive format supports the specification of linked files, it shall be an error if these
 109108 files cannot be linked when the archive is extracted. For archive formats that do not store file
 109109 contents with each name that causes a hard link, if the file that contains the data is not extracted
 109110 during this *pax* session, either the data shall be restored from the original file, or a diagnostic
 109111 message shall be displayed with the name of a file that can be used to extract the data. In
 109112 traversing directories, *pax* shall detect infinite loops; that is, entering a previously visited
 109113 directory that is an ancestor of the last file visited. When it detects an infinite loop, *pax* shall
 109114 write a diagnostic message to standard error and shall terminate.

109115 OPTIONS

109116 The *pax* utility shall conform to XBD [Section 12.2](#) (on page 215), except that the order of
 109117 presentation of the **-o**, **-p**, and **-s** options is significant.

109118 The following options shall be supported:

- 109119 **-r** Read an archive file from standard input.
- 109120 **-w** Write files to the standard output in the specified archive format.
- 109121 **-a** Append files to the end of the archive. It is implementation-defined which devices
 109122 on the system support appending. Additional file formats unspecified by this
 109123 volume of POSIX.1-202x may impose restrictions on appending.
- 109124 **-b** *blocksize* Block the output at a positive decimal integer number of bytes per write to the
 109125 archive file. Devices and archive formats may impose restrictions on blocking.
 109126 Blocking shall be automatically determined on input. Conforming applications
 109127 shall not specify a *blocksize* value larger than 32256. Default blocking when
 109128 creating archives depends on the archive format. (See the **-x** option below.)
- 109129 **-c** Match all file or archive members except those specified by the *pattern* or *file*
 109130 operands.

109131	-d	Cause files of type directory being copied or archived or archive members of type directory being extracted or listed to match only the file or archive member itself and not the file hierarchy rooted at the file.
109132		
109133		
109134	-f <i>archive</i>	Specify the pathname of the input or output archive, overriding the default standard input (in list or read modes) or standard output (write mode).
109135		
109136	-H	If a symbolic link referencing a file of type directory is specified on the command line, <i>pax</i> shall archive the file hierarchy rooted in the file referenced by the link, using the name of the link as the root of the file hierarchy. Otherwise, if a symbolic link referencing a file of any other file type which <i>pax</i> can normally archive is specified on the command line, then <i>pax</i> shall archive the file referenced by the link, using the name of the link. The default behavior, when neither -H or -L are specified, shall be to archive the symbolic link itself.
109137		
109138		
109139		
109140		
109141		
109142		
109143		
109144	-i	Interactively rename files or archive members. For each archive member matching a <i>pattern</i> operand or file matching a <i>file</i> operand, a prompt shall be written to the file /dev/tty . The prompt shall contain the name of the file or archive member, but the format is otherwise unspecified. A line shall then be read from /dev/tty . If this line is blank, the file or archive member shall be skipped. If this line consists of a single period, the file or archive member shall be processed with no modification to its name. Otherwise, its name shall be replaced with the contents of the line. The <i>pax</i> utility shall immediately exit with a non-zero exit status if end-of-file is encountered when reading a response or if /dev/tty cannot be opened for reading and writing.
109145		
109146		
109147		
109148		
109149		
109150		
109151		
109152		
109153		
109154		
109155	-k	Prevent the overwriting of existing files.
109156	-l	(The letter ell.) In copy mode, hard links shall be made between the source and destination file hierarchies whenever possible. If specified in conjunction with -H or -L, when a symbolic link is encountered, the hard link created in the destination file hierarchy shall be to the file referenced by the symbolic link. If specified when neither -H nor -L is specified, when a symbolic link is encountered, the implementation shall create a hard link to the symbolic link in the source file hierarchy or copy the symbolic link to the destination.
109157		
109158		
109159		
109160		
109161		
109162		
109163	-L	If a symbolic link referencing a file of type directory is specified on the command line or encountered during the traversal of a file hierarchy, <i>pax</i> shall archive the file hierarchy rooted in the file referenced by the link, using the name of the link as the root of the file hierarchy. Otherwise, if a symbolic link referencing a file of any other file type which <i>pax</i> can normally archive is specified on the command line or encountered during the traversal of a file hierarchy, <i>pax</i> shall archive the file referenced by the link, using the name of the link. The default behavior, when neither -H or -L are specified, shall be to archive the symbolic link itself.
109164		
109165		
109166		
109167		
109168		
109169	-n	Select the first archive member that matches each <i>pattern</i> operand. No more than one archive member shall be matched for each pattern (although members of type directory shall still match the file hierarchy rooted at that file).
109170		
109171		
109172	-o <i>options</i>	Provide information to the implementation to modify the algorithm for extracting or writing files. The value of <i>options</i> shall consist of one or more <comma>-separated keywords of the form:
109173		
109174		
109175		
109176		
109177		<code>keyword[[:]=value] [, keyword[[:]=value], ...]</code>

109178 Some keywords apply only to certain file formats, as indicated with each
 109179 description. Use of keywords that are inapplicable to the file format being
 109180 processed produces undefined results.

109181 Keywords in the *options* argument shall be a string that would be a valid portable
 109182 filename as described in XBD [Section 3.264](#) (on page 70).

109183 **Note:** Keywords are not expected to be filenames, merely to follow the same character
 109184 composition rules as portable filenames.

109185 Keywords can be preceded with white space. The *value* field shall consist of zero or
 109186 more characters; within *value*, the application shall precede any literal <comma>
 109187 with a <backslash>, which shall be ignored, but preserves the <comma> as part of
 109188 *value*. A <comma> as the final character, or a <comma> followed solely by white
 109189 space as the final characters, in *options* shall be ignored. Multiple **-o** options can be
 109190 specified; if keywords given to these multiple **-o** options conflict, the keywords
 109191 and values appearing later in command line sequence shall take precedence and
 109192 the earlier shall be silently ignored. The following keyword values of *options* shall
 109193 be supported for the file formats as indicated:

109194 **delete=pattern**

109195 (Applicable only to the **-x pax** format.) When used in **write** or **copy** mode, *pax*
 109196 shall omit from extended header records that it produces any keywords
 109197 matching the string pattern. When used in **read** or **list** mode, *pax* shall ignore
 109198 any keywords matching the string pattern in the extended header records. In
 109199 both cases, matching shall be performed using the pattern matching notation
 109200 described in [Section 2.14.1](#) (on page 2506) and [Section 2.14.2](#) (on page 2507).
 109201 For example:

109202 **-o delete=security.***

109203 would suppress security-related information. See [pax Extended Header](#) (on
 109204 page 3236) for extended header record keyword usage.

109205 When multiple **-odelete=pattern** options are specified, the patterns shall be
 109206 additive; all keywords matching the specified string patterns shall be omitted
 109207 from extended header records that *pax* produces.

109208 **exthdr.name=string**

109209 (Applicable only to the **-x pax** format.) This keyword allows user control over
 109210 the name that is written into the **ustar** header blocks for the extended header
 109211 produced under the circumstances described in [pax Header Block](#) (on page
 109212 3235). The name shall be the contents of *string*, after the following character
 109213 substitutions have been made:

<i>string</i> Includes:	Replaced by:
%d	The directory name of the file, equivalent to the result of the <i>dirname</i> utility on the translated pathname.
%f	The filename of the file, equivalent to the result of the <i>basename</i> utility on the translated pathname.
%p	The process ID of the <i>pax</i> process.
%%	A '%' character.

109222 Any other '%' characters in *string* produce undefined results.

109223 If no **-o exthdr.name=string** is specified, *pax* shall use the following default

109224 value:
 109225 %d/PaxHeaders.%p/%f

globexthdr.name=string

(Applicable only to the **-x pax** format.) When used in **write** or **copy** mode with the appropriate options, *pax* shall create global extended header records with **ustar** header blocks that are treated as regular files by previous versions of *pax*. This keyword allows user control over the name that is written into the **ustar** header blocks for global extended header records. The name shall be the contents of *string*, after the following character substitutions have been made:

<i>string</i> Includes:	Replaced by:
%n	An integer that represents the sequence number of the global extended header record in the archive, starting at 1.
%p	The process ID of the <i>pax</i> process.
%%	A '%' character.

109239 Any other '%' characters in *string* produce undefined results.

109240 If no **-o globexthdr.name=string** is specified, *pax* shall use the following
 109241 default value:

109242 \$TMPDIR/GlobalHead.%p.%n

109243 where *\$TMPDIR* represents the value of the *TMPDIR* environment variable. If
 109244 *TMPDIR* is not set, *pax* shall use **/tmp**.

invalid=action

(Applicable only to the **-x pax** format.) This keyword allows user control over the action *pax* takes upon encountering values in an extended header record that, in **read** or **copy** mode, are invalid in the destination hierarchy or, in **list** mode, cannot be written in the codeset and current locale of the implementation. The following are invalid values that shall be recognized by *pax*:

- 109252 — In **read** or **copy** mode, a filename or link name that contains character
 109253 encodings invalid in the destination hierarchy. (For example, the name
 109254 may contain embedded NULs.)
- 109255 — In **read** or **copy** mode, a filename or link name that is longer than the
 109256 maximum allowed in the destination hierarchy (for either a pathname
 109257 component or the entire pathname).
- 109258 — In **list** mode, any character string value (filename, link name, user name,
 109259 and so on) that cannot be written in the codeset and current locale of the
 109260 implementation.

109261 The following mutually-exclusive values of the *action* argument are supported:

binary In **write** mode, *pax* shall generate a **hdrcharset=BINARY** extended header record for each file with a filename, link name, group name, owner name, or any other field in an extended header record that cannot be translated to the UTF-8 codeset, allowing the archive to contain the files with unencoded extended header record values. In **read** or **copy** mode, *pax* shall use the values specified in the header without translation,

109269		regardless of whether this may overwrite an existing file with a valid name. In list mode, <i>pax</i> shall behave identically to the bypass action.
109270		
109271		
109272	bypass	In read or copy mode, <i>pax</i> shall bypass the file, causing no change to the destination hierarchy. In list mode, <i>pax</i> shall write all requested valid values for the file, but its method for writing invalid values is unspecified.
109273		
109274		
109275		
109276	rename	In read or copy mode, <i>pax</i> shall act as if the -i option were in effect for each file with invalid filename or link name values, allowing the user to provide a replacement name interactively. In list mode, <i>pax</i> shall behave identically to the bypass action.
109277		
109278		
109279		
109280	UTF-8	When used in read , copy , or list mode and a filename, link name, owner name, or any other field in an extended header record cannot be translated from the pax UTF-8 codeset format to the codeset and current locale of the implementation, <i>pax</i> shall use the actual UTF-8 encoding for the name. If a hdrcharset extended header record is in effect for this file, the character set specified by that record shall be used instead of UTF-8. If a hdrcharset=BINARY extended header record is in effect for this file, no translation shall be performed.
109281		
109282		
109283		
109284		
109285		
109286		
109287		
109288		
109289	write	In read or copy mode, <i>pax</i> shall write the file, translating the name, regardless of whether this may overwrite an existing file with a valid name. In list mode, <i>pax</i> shall behave identically to the bypass action.
109290		
109291		
109292		
109293		If no -o invalid=option is specified, <i>pax</i> shall act as if -o invalid=bypass were specified. Any overwriting of existing files that may be allowed by the -o invalid= actions shall be subject to permission (-p) and modification time (-u) restrictions, and shall be suppressed if the -k option is also specified.
109294		
109295		
109296		
109297	linkdata	
109298		(Applicable only to the -x pax format.) In write mode, <i>pax</i> shall write the contents of a file to the archive even when that file is merely a hard link to a file whose contents have already been written to the archive.
109299		
109300		
109301	listopt=format	
109302		This keyword specifies the output format of the table of contents produced when the -v option is specified in list mode. See List Mode Format Specifications (on page 3230). To avoid ambiguity, the listopt=format shall be the only or final keyword=value pair in a -o option-argument; all characters in the remainder of the option-argument shall be considered part of the format string. When multiple -olistopt=format options are specified, the format strings shall be considered a single, concatenated string, evaluated in command line order.
109303		
109304		
109305		
109306		
109307		
109308		
109309		
109310	times	
109311		(Applicable only to the -x pax format.) When used in write or copy mode, <i>pax</i> shall include atime and mtime extended header records for each file. See pax Extended Header File Times (on page 3239).
109312		
109313		
109314		In addition to these keywords, if the -x pax format is specified, any of the keywords and values defined in pax Extended Header (on page 3236), including
109315		

109316 implementation extensions, can be used in **-o** option-arguments, in either of two
109317 modes:

109318 **keyword=***value*

109319 When used in **write** or **copy** mode, these keyword/value pairs shall be
109320 included at the beginning of the archive as **typeflag g** global extended header
109321 records. When used in **read** or **list** mode, these keyword/value pairs shall act
109322 as if they had been at the beginning of the archive as **typeflag g** global
109323 extended header records.

109324 **keyword:=***value*

109325 When used in **write** or **copy** mode, these keyword/value pairs shall be
109326 included as records at the beginning of a **typeflag x** extended header for each
109327 file. (This shall be equivalent to the <equals-sign> form except that it creates
109328 no **typeflag g** global extended header records.) When used in **read** or **list**
109329 mode, these keyword/value pairs shall act as if they were included as records
109330 at the end of each extended header; thus, they shall override any global or file-
109331 specific extended header record keywords of the same names. For example, in
109332 the command:

```
109333 pax -r -o "  
109334 gname:=mygroup,  
109335 " <archive
```

109336 the group name is forced to a new value for all files read from the archive. |

109337 The precedence of **-o** keywords over various fields in the archive is described in
109338 [pax Extended Header Keyword Precedence](#) (on page 3239). If the **-o**
109339 **delete=***pattern*, **-o keyword=***value*, or **-o keyword:=***value* options are used to
109340 override or remove any extended header data needed to find files in an archive
109341 (e.g., **-o delete=size** for a file whose size cannot be represented in a **ustar**
109342 header or **-o size=100** for a file whose size is not 100 bytes), the behavior is
109343 undefined.

109344 **-p** *string* Specify one or more file characteristic options (privileges). The *string* option-
109345 argument shall be a string specifying file characteristics to be retained or discarded
109346 on extraction. The string shall consist of the specification characters *a*, *e*, *m*, *o*, and
109347 *p*. Other implementation-defined characters can be included. Multiple
109348 characteristics can be concatenated within the same string and multiple **-p** options
109349 can be specified. The meaning of the specification characters are as follows:

- 109350 a Do not preserve file access times.
- 109351 e Preserve the user ID, group ID, file mode bits (see XBD [Section 3.145](#), on page
109352 52), access time, modification time, and any other implementation-defined file
109353 characteristics.
- 109354 m Do not preserve file modification times.
- 109355 o Preserve the user ID and group ID.
- 109356 p Preserve the file mode bits. Other implementation-defined file mode attributes
109357 may be preserved.

109358 In the preceding list, “preserve” indicates that an attribute stored in the archive
109359 shall be given to the extracted file, subject to the permissions of the invoking
109360 process. The access and modification times of the file shall be preserved unless
109361 otherwise specified with the **-p** option or not stored in the archive. All attributes

109362 that are not preserved shall be determined as part of the normal file creation action
109363 (see [Section 1.1.1.4](#), on page 2440).

109364 If neither the `e` nor the `o` specification character is specified, or the user ID and
109365 group ID are not preserved for any reason, *pax* shall not set the `S_ISUID` and
109366 `S_ISGID` bits of the file mode.

109367 If the preservation of any of these items fails for any reason, *pax* shall write a
109368 diagnostic message to standard error. Failure to preserve these items shall affect
109369 the final exit status, but shall not cause the extracted file to be deleted.

109370 If file characteristic letters in any of the *string* option-arguments are duplicated or
109371 conflict with each other, the ones given last shall take precedence. For example, if
109372 `-p eme` is specified, file modification times are preserved.

109373 `-s replstr` Modify file or archive member names named by *pattern* or *file* operands according
109374 to the substitution expression *replstr*, using the syntax of the *ed* utility. The concepts
109375 of “address” and “line” are meaningless in the context of the *pax* utility, and shall
109376 not be supplied. The format shall be:

109377 `-s /old/new/[gpsS]` |

109378 where as in *ed*, *old* is a basic regular expression and *new* can contain an
109379 `<ampersand>`, `'\n'` (where *n* is a digit) back-references, or subexpression
109380 matching. The *old* string shall also be permitted to contain `<newline>` characters.

109381 Any non-null character can be used as a delimiter (`'/'` shown here). Multiple `-s`
109382 expressions can be specified; the expressions shall be applied in the order
109383 specified, terminating with the first successful substitution. The optional trailing
109384 `'g'` is as defined in the *ed* utility. The optional trailing `'p'` shall cause successful
109385 substitutions to be written to standard error. The optional trailing `'s'` and `'S'`
109386 control whether the substitutions are applied to symbolic link contents: `'s'` shall
109387 cause them not to be applied; `'S'` shall cause them to be applied. If neither is
109388 present, it is unspecified which is the default. If both are present, the behavior is
109389 unspecified. File or archive member names that substitute to the empty string shall
109390 be ignored when reading and writing archives. Symbolic link contents that
109391 substitute to the empty string shall not be treated specially.

109392 `-t` When reading files from the file system, and if the user has the permissions
109393 required by *futimens()* to do so, set the access time of each file read to the access
109394 time that it had before being read by *pax*.

109395 `-u` Ignore files that are older (having a less recent file modification time) than a pre-
109396 existing file or archive member with the same name. In **read** mode, an archive
109397 member with the same name as a file in the file system shall be extracted if the
109398 archive member is newer than the file. In **write** mode, an archive file member with
109399 the same name as a file in the file system shall be superseded if the file is newer
109400 than the archive member. If `-a` is also specified, this is accomplished by appending
109401 to the archive; otherwise, it is unspecified whether this is accomplished by actual
109402 replacement in the archive or by appending to the archive. In **copy** mode, the file
109403 in the destination hierarchy shall be replaced if the file in the source hierarchy is
109404 newer.

109405 `-v` In **list** mode, produce a verbose table of contents (see the **STDOUT** section).
109406 Otherwise, write archive member pathnames to standard error (see the **STDERR**
109407 section).

109408	-x format	Specify the output archive format. The <i>pax</i> utility shall support the following formats:
109409		
109410	cpio	The cpio interchange format; see the EXTENDED DESCRIPTION section. The default <i>blocksize</i> for this format for character special archive files shall be 5120. Implementations shall support all <i>blocksize</i> values less than or equal to 32256 that are multiples of 512.
109411		
109412		
109413		
109414	pax	The pax interchange format; see the EXTENDED DESCRIPTION section. The default <i>blocksize</i> for this format for character special archive files shall be 5120. Implementations shall support all <i>blocksize</i> values less than or equal to 32256 that are multiples of 512.
109415		
109416		
109417		
109418	ustar	The tar interchange format; see the EXTENDED DESCRIPTION section. The default <i>blocksize</i> for this format for character special archive files shall be 10240. Implementations shall support all <i>blocksize</i> values less than or equal to 32256 that are multiples of 512.
109419		
109420		
109421		
109422		Implementation-defined formats shall specify a default block size as well as any other block sizes supported for character special archive files.
109423		
109424		Any attempt to append to an archive file in a format different from the existing archive format shall cause <i>pax</i> to exit immediately with a non-zero exit status.
109425		
109426	-X	When traversing the file hierarchy specified by a pathname, <i>pax</i> shall not descend below directories that have a different device ID (<i>st_dev</i> ; see XSH <i>fstatat()</i>) than the specified pathname; that is, when a directory with a different device ID is encountered, <i>pax</i> shall process (archive or copy) the directory itself but shall not process any files below the directory.
109427		
109428		
109429		
109430		
109431		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.
109432		
109433		The options that operate on the names of files or archive members (-c , -i , -n , -s , -u , and -v) shall interact as follows. In read mode, the archive members shall be selected based on the user-specified <i>pattern</i> operands as modified by the -c , -n , and -u options. Then, any -s and -i options shall modify, in that order, the names of the selected files. The -v option shall write names resulting from these modifications.
109434		
109435		
109436		
109437		
109438		In write mode, the files shall be selected based on the user-specified pathnames as modified by the -u option. Then, any -s and -i options shall modify, in that order, the names of these selected files. The -v option shall write names resulting from these modifications.
109439		
109440		
109441		If both the -u and -n options are specified, <i>pax</i> shall not consider a file selected unless it is newer than the file to which it is compared.
109442		
109443		List Mode Format Specifications
109444		In list mode with the -o listopt=format option, the <i>format</i> argument shall be applied for each selected file. The <i>pax</i> utility shall append a <newline> to the listopt output for each selected file. The <i>format</i> argument shall be used as the <i>format</i> string described in XBD Chapter 5 (on page 113), with the exceptions 1. through 6. defined in the EXTENDED DESCRIPTION section of <i>printf</i> , plus the following exceptions:
109445		
109446		
109447		
109448		
109449		7. The sequence (<i>keyword</i>) can occur before a format conversion specifier. The conversion argument is defined by the value of <i>keyword</i> . The implementation shall support the following keywords:
109450		
109451		

- 109452 — Any of the Field Name entries in Table 3-15 (on page 3240) and Table 3-17 (on page
 109453 3244). The implementation may support the *cpio* keywords without the leading *c_* in
 109454 addition to the form required by Table 3-17 (on page 3244).
- 109455 — Any keyword defined for the extended header in pax Extended Header (on page
 109456 3236).
- 109457 — Any keyword provided as an implementation-defined extension within the extended
 109458 header defined in pax Extended Header (on page 3236).

109459 For example, the sequence "%(charset)s" is the string value of the name of the character
 109460 set in the extended header.

109461 The result of the keyword conversion argument shall be the value from the applicable
 109462 header field or extended header, without any trailing NULs.

109463 All keyword values used as conversion arguments shall be translated from the UTF-8
 109464 encoding (or alternative encoding specified by any **hdrcharset** extended header record) to
 109465 the character set appropriate for the local file system, user database, and so on, as
 109466 applicable.

- 109467 8. An additional conversion specifier character, *T*, shall be used to specify time formats. The *T*
 109468 conversion specifier character can be preceded by the sequence (*keyword=subformat*), where
 109469 *subformat* is a date format as defined by *date* operands. The default *keyword* shall be **mtime**
 109470 and the default subformat shall be:

109471 %b %e %H:%M %Y

- 109472 9. An additional conversion specifier character, *M*, shall be used to specify the file mode string
 109473 as defined in *ls* Standard Output. If (*keyword*) is omitted, the **mode** keyword shall be used.
 109474 For example, %.1M writes the single character corresponding to the <entry type> field of the
 109475 *ls -l* command.

- 109476 10. An additional conversion specifier character, *D*, shall be used to specify the device for block
 109477 or special files, if applicable, in an implementation-defined format. If not applicable, and
 109478 (*keyword*) is specified, then this conversion shall be equivalent to %(*keyword*)u. If not
 109479 applicable, and (*keyword*) is omitted, then this conversion shall be equivalent to <space>.

- 109480 11. An additional conversion specifier character, *F*, shall be used to specify a pathname. The *F*
 109481 conversion character can be preceded by a sequence of <comma>-separated keywords:

109482 (*keyword*[,*keyword*] ...)

109483 The values for all the keywords that are non-null shall be concatenated together, each
 109484 separated by a '/'. The default shall be (**path**) if the keyword **path** is defined; otherwise,
 109485 the default shall be (**prefix,name**).

- 109486 12. An additional conversion specifier character, *L*, shall be used to specify a symbolic link
 109487 expansion. If the current file is a symbolic link, then %L shall expand to:

109488 "%s -> %s", <value of keyword>, <contents of link>

109489 Otherwise, the %L conversion specification shall be the equivalent of %F.

109490 OPERANDS

109491 The following operands shall be supported:

109492 *directory* The destination directory pathname for **copy** mode.

109493	<i>file</i>	A pathname of a file to be copied or archived.
109494	<i>pattern</i>	A pattern matching one or more pathnames of archive members. A pattern needs to be given in the name-generating notation of the pattern matching notation in Section 2.14 (on page 2506), including the filename expansion rules in Section 2.14.3 (on page 2508). The default, if no <i>pattern</i> is specified, is to select all members in the archive.
109495		
109496		
109497		
109498		
109499	STDIN	
109500		In write mode, the standard input shall be used only if no <i>file</i> operands are specified. It shall be a file containing a list of pathnames, each terminated by a <newline> character.
109501		
109502		In list and read modes, if -f is not specified, the standard input shall be an archive file.
109503		Otherwise, the standard input shall not be used.
109504	INPUT FILES	
109505		The input file named by the <i>archive</i> option-argument, or standard input when the archive is read from there, shall be a file formatted according to one of the specifications in the EXTENDED DESCRIPTION section or some other implementation-defined format.
109506		
109507		
109508		The file <i>/dev/tty</i> shall be used to write prompts and read responses.
109509	ENVIRONMENT VARIABLES	
109510		The following environment variables shall affect the execution of <i>pax</i> :
109511	<i>LANG</i>	Provide a default value for the internationalization variables that are unset or null. (See XBD Section 8.2 (on page 169) the precedence of internationalization variables used to determine the values of locale categories.)
109512		
109513		
109514	<i>LC_ALL</i>	If set to a non-empty string value, override the values of all the other internationalization variables.
109515		
109516	<i>LC_COLLATE</i>	Determine the locale for the behavior of ranges, equivalence classes, and multi-character collating elements used in the pattern matching expressions for the <i>pattern</i> operand and the basic regular expression for the -s option.
109517		
109518		
109519		
109520	<i>LC_CTYPE</i>	Determine the locale for the interpretation of sequences of bytes of text data as characters (for example, single-byte as opposed to multi-byte characters in arguments and input files), and the behavior of character classes used in the pattern matching expressions for the <i>pattern</i> operand and the basic regular expression for the -s option.
109521		
109522		
109523		
109524		
109525	<i>LC_MESSAGES</i>	Determine the locale used to affect the format and contents of diagnostic messages and prompts written to standard error.
109526		
109527		
109528	<i>LC_TIME</i>	Determine the format and contents of date and time strings when the -v option is specified.
109529		
109530	XSI <i>NLSPATH</i>	Determine the location of messages objects and message catalogs.
109531	<i>TMPDIR</i>	Determine the pathname that provides part of the default global extended header record file, as described for the -o globexthdr= keyword in the OPTIONS section.
109532		
109533	<i>TZ</i>	Determine the timezone used to calculate date and time strings when the -v option is specified. If <i>TZ</i> is unset or null, an unspecified default timezone shall be used.
109534		

109535 **ASYNCHRONOUS EVENTS**

109536 Default.

109537 **STDOUT**

109538 In **write** mode, if **-f** is not specified, the standard output shall be the archive formatted
 109539 according to one of the specifications in the EXTENDED DESCRIPTION section, or some other
 109540 implementation-defined format (see **-x format**).

109541 In **list** mode, when the **-olistopt=format** has been specified, the selected archive members shall
 109542 be written to standard output using the format described under [List Mode Format Specifications](#)
 109543 (on page 3230). In **list** mode without the **-olistopt=format** option, the table of contents of the
 109544 selected archive members shall be written to standard output using the following format:

109545 "%s\n", <pathname>

109546 If the **-v** option is specified in **list** mode, the table of contents of the selected archive members
 109547 shall be written to standard output using the following formats.

109548 For pathnames representing hard links to previous members of the archive:

109549 "%sΔ==Δ%s\n", <ls -l listing>, <linkname>

109550 For all other pathnames:

109551 "%s\n", <ls -l listing>

109552 where <ls -l listing> shall be the format specified by the *ls* utility with the **-l** option. When
 109553 writing pathnames in this format, it is unspecified what is written for fields for which the
 109554 underlying archive format does not have the correct information, although the correct number of
 109555 <blank>-separated fields shall be written.

109556 In **list** mode, standard output shall not be buffered more than a pathname (plus any associated
 109557 information and a <newline> terminator) at a time.

109558 **STDERR**

109559 If **-v** is specified in **read**, **write**, or **copy** modes, *pax* shall write the pathnames it processes to the
 109560 standard error output using the following format:

109561 "%s\n", <pathname>

109562 These pathnames shall be written as soon as processing is begun on the file or archive member,
 109563 and shall be flushed to standard error. The trailing <newline>, which shall not be buffered, is
 109564 written when the file has been read or written.

109565 If the **-s** option is specified, and the replacement string has a trailing 'p', substitutions shall be
 109566 written to standard error in the following format:

109567 "%sΔ>>Δ%s\n", <original pathname>, <new pathname>

109568 In all operating modes of *pax*, optional messages of unspecified format concerning the input
 109569 archive format and volume number, the number of files, blocks, volumes, and media parts as
 109570 well as other diagnostic messages may be written to standard error.

109571 In all formats, for both standard output and standard error, it is unspecified how non-printable
 109572 characters in pathnames or link names are written.

109573 When using the **-xpax** archive format, if a filename, link name, group name, owner name, or any
 109574 other field in an extended header record cannot be translated between the codeset in use for that
 109575 extended header record and the character set of the current locale, *pax* shall write a diagnostic
 109576 message to standard error, shall process the file as described for the **-o invalid=** option, and then
 109577 shall continue processing with the next file.

109578 **OUTPUT FILES**

109579 In **read** mode, the extracted output files shall be of the archived file type. In **copy** mode, the
 109580 copied output files shall be the type of the file being copied. In either mode, existing files in the
 109581 destination hierarchy shall be overwritten only when all permission (**-p**), modification time (**-u**),
 109582 and invalid-value (**-oinvalid=**) tests allow it.

109583 In **write** mode, the output file named by the **-f** option-argument shall be a file formatted
 109584 according to one of the specifications in the EXTENDED DESCRIPTION section, or some other
 109585 implementation-defined format.

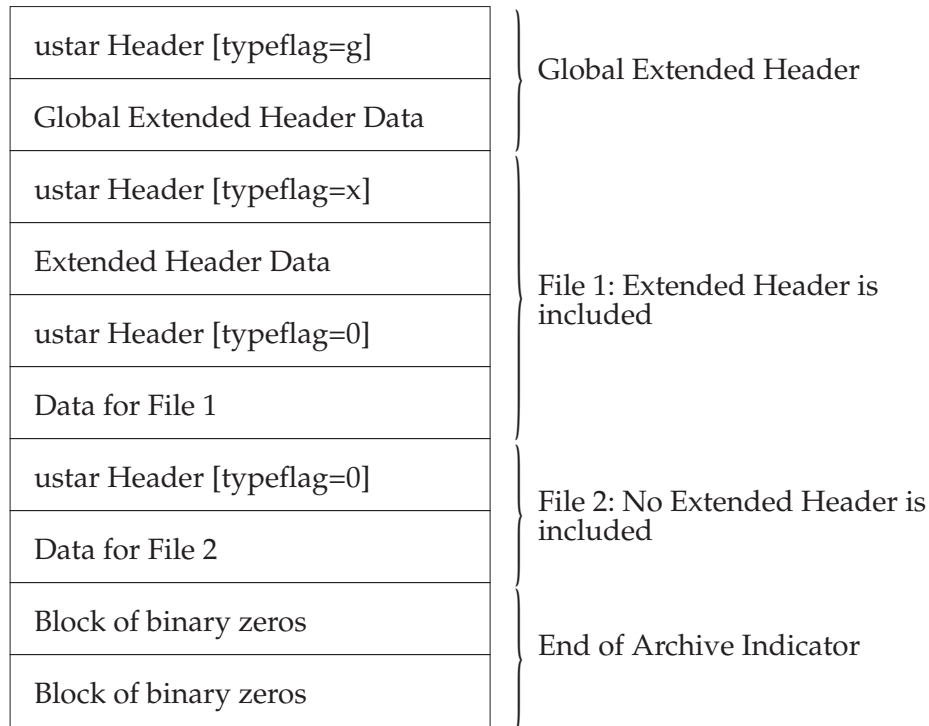
109586 **EXTENDED DESCRIPTION**109587 **pax Interchange Format**

109588 A *pax* archive tape or file produced in the **-xpax** format shall contain a series of blocks. The
 109589 physical layout of the archive shall be identical to the **ustar** format described in **ustar**
 109590 **Interchange Format** (on page 3240). Each file archived shall be represented by the following
 109591 sequence:

- 109592 • An optional header block with extended header records. This header block is of the form
 109593 described in **pax Header Block** (on page 3235), with a *typeflag* value of **x** or **g**. The
 109594 extended header records, described in **pax Extended Header** (on page 3236), shall be
 109595 included as the data for this header block.
- 109596 • A header block that describes the file. Any fields in the preceding optional extended
 109597 header shall override the associated fields in this header block for this file.
- 109598 • Zero or more blocks that contain the contents of the file.

109599 At the end of the archive file there shall be two 512-byte blocks filled with binary zeros,
 109600 interpreted as an end-of-archive indicator.

109601 A schematic of an example archive with global extended header records and two actual files is
 109602 shown in **Figure 3-1** (on page 3235). In the example, the second file in the archive has no
 109603 extended header preceding it, presumably because it has no need for extended attributes.



109604

Figure 3-1 pax Format Archive Example

109605

pax Header Block

109606

The **pax** header block shall be identical to the **ustar** header block described in [ustar Interchange Format](#) (on page 3240), except that two additional *typeflag* values are defined:

109607

109608

× Represents extended header records for the following file in the archive (which shall have its own **ustar** header block). The format of these extended header records shall be as described in [pax Extended Header](#) (on page 3236).

109609

109610

109611

g Represents global extended header records for the following files in the archive. The format of these extended header records shall be as described in [pax Extended Header](#) (on page 3236). Each value shall affect all subsequent files that do not override that value in their own extended header record and until another global extended header record is reached that provides another value for the same field. The *typeflag* g global headers should not be used with interchange media that could suffer partial data loss in transporting the archive.

109612

109613

109614

109615

109616

109617

For both of these types, the *size* field shall be the size of the extended header records in octets. The other fields in the header block are not meaningful to this version of the *pax* utility. However, if this archive is read by a *pax* utility conforming to the ISO POSIX-2:1993 standard, the header block fields are used to create a regular file that contains the extended header records as data. Therefore, header block field values should be selected to provide reasonable file access to this regular file.

109618

109619

109620

109621

109622

109623

A further difference from the **ustar** header block is that data blocks for files of *typeflag* 1 (the digit one) (hard link) may be included, which means that the size field may be greater than zero. Archives created by *pax -o linkdata* shall include these data blocks with the hard links.

109624

109625

109626 **pax Extended Header**

109627 A **pax** extended header contains values that are inappropriate for the **ustar** header block because
 109628 of limitations in that format: fields requiring a character encoding other than that described in
 109629 the ISO/IEC 646:1991 standard, fields representing file attributes not described in the **ustar**
 109630 header, and fields whose format or length do not fit the requirements of the **ustar** header. The
 109631 values in an extended header add attributes to the following file (or files; see the description of
 109632 the *typeflag* **g** header block) or override values in the following header block(s), as indicated in
 109633 the following list of keywords.

109634 An extended header shall consist of one or more records, each constructed as follows:

109635 "`%d %s=%s\n`", *<length>*, *<keyword>*, *<value>*

109636 The extended header records shall be encoded according to the ISO/IEC 10646-1:2000 standard
 109637 UTF-8 encoding. The *<length>* field, *<blank>*, *<equals-sign>*, and *<newline>* shown shall be
 109638 limited to the portable character set, as encoded in UTF-8. The *<keyword>* fields can be any
 109639 UTF-8 characters. The *<length>* field shall be the decimal length of the extended header record
 109640 in octets, including the trailing *<newline>*. If there is a **hdrcharset** extended header in effect for
 109641 a file, the *value* field for any **gname**, **linkpath**, **path**, and **uname** extended header records shall be
 109642 encoded using the character set specified by the **hdrcharset** extended header record; otherwise,
 109643 the *value* field shall be encoded using UTF-8. The *value* field for all other keywords specified by
 109644 POSIX.1-202x shall be encoded using UTF-8.

109645 The *<keyword>* field shall be one of the entries from the following list or a keyword provided as
 109646 an implementation extension. Keywords consisting entirely of lowercase letters, digits, and
 109647 periods are reserved for future standardization. A keyword shall not include an *<equals-sign>*.
 109648 (In the following list, the notations *file(s)* or *block(s)* is used to acknowledge that a keyword
 109649 affects the following single file after a *typeflag* **x** extended header, but possibly multiple files after
 109650 *typeflag* **g**. Any requirements in the list for *pax* to include a record when in **write** or **copy** mode
 109651 shall apply only when such a record has not already been provided through the use of the **-o**
 109652 option. When used in **copy** mode, *pax* shall behave as if an archive had been created with
 109653 applicable extended header records and then extracted.)

109654 **atime** The file access time for the following file(s), equivalent to the value of the *st_atim* |
 109655 member of the **stat** structure for a file, as described by the *stat()* function. The
 109656 access time shall be restored if the process has appropriate privileges required to
 109657 do so. The format of the *<value>* shall be as described in [pax Extended Header File](#)
 109658 [Times](#) (on page 3239).

109659 **charset** The name of the character set used to encode the data in the following file(s). The
 109660 entries in the following table are defined to refer to known standards; additional
 109661 names may be agreed on between the originator and recipient.

	<value>	Formal Standard
109662		
109663	ISO-IRΔ646Δ1990	ISO/IEC 646:1990
109664	ISO-IRΔ8859Δ1Δ1998	ISO/IEC 8859-1:1998
109665	ISO-IRΔ8859Δ2Δ1999	ISO/IEC 8859-2:1999
109666	ISO-IRΔ8859Δ3Δ1999	ISO/IEC 8859-3:1999
109667	ISO-IRΔ8859Δ4Δ1998	ISO/IEC 8859-4:1998
109668	ISO-IRΔ8859Δ5Δ1999	ISO/IEC 8859-5:1999
109669	ISO-IRΔ8859Δ6Δ1999	ISO/IEC 8859-6:1999
109670	ISO-IRΔ8859Δ7Δ1987	ISO/IEC 8859-7:1987
109671	ISO-IRΔ8859Δ8Δ1999	ISO/IEC 8859-8:1999
109672	ISO-IRΔ8859Δ9Δ1999	ISO/IEC 8859-9:1999
109673	ISO-IRΔ8859Δ10Δ1998	ISO/IEC 8859-10:1998
109674	ISO-IRΔ8859Δ13Δ1998	ISO/IEC 8859-13:1998
109675	ISO-IRΔ8859Δ14Δ1998	ISO/IEC 8859-14:1998
109676	ISO-IRΔ8859Δ15Δ1999	ISO/IEC 8859-15:1999
109677	ISO-IRΔ10646Δ2000	ISO/IEC 10646:2000
109678	ISO-IRΔ10646Δ2000ΔUTF-8	ISO/IEC 10646, UTF-8 encoding
109679	BINARY	None.

109680 The encoding is included in an extended header for information only; when *pax* is
 109681 used as described in POSIX.1-202x, it shall not translate the file data into any other
 109682 encoding. The **BINARY** entry indicates unencoded binary data.

109683 When used in **write** or **copy** mode, it is implementation-defined whether *pax*
 109684 includes a **charset** extended header record for a file.

109685 **comment** A series of characters used as a comment. All characters in the <value> field shall
 109686 be ignored by *pax*.

109687 **gid** The group ID of the group that owns the file, expressed as a decimal number using
 109688 digits from the ISO/IEC 646:1991 standard. This record shall override the *gid* field
 109689 in the following header block(s). When used in **write** or **copy** mode, *pax* shall
 109690 include a *gid* extended header record for each file whose group ID is greater than
 109691 2 097 151 (octal 7 777 777).

109692 **gname** The group of the file(s), formatted as a group name in the group database. This
 109693 record shall override the *gid* and *gname* fields in the following header block(s), and
 109694 any *gid* extended header record. When used in **read**, **copy**, or **list** mode, *pax* shall
 109695 translate the name from the encoding in the header record to the character set
 109696 appropriate for the group database on the receiving system. If any of the characters
 109697 cannot be translated, and if neither the **-oinvalid=UTF-8** option nor the
 109698 **-oinvalid=binary** option is specified, the results are implementation-defined.
 109699 When used in **write** or **copy** mode, *pax* shall include a **gname** extended header
 109700 record for each file whose group name cannot be represented entirely with the
 109701 letters and digits of the portable character set.

109702 **hdrcharset** The name of the character set used to encode the value field of the **gname**,
 109703 **linkpath**, **path**, and **uname** *pax* extended header records. The entries in the
 109704 following table are defined to refer to known standards; additional names may be
 109705 agreed between the originator and the recipient.

109706
109707
109708

109709
109710
109711

109712
109713

109714
109715
109716
109717
109718
109719
109720
109721
109722
109723
109724

109725
109726
109727
109728
109729
109730

109731
109732
109733
109734

109735
109736
109737

109738

109739

109740
109741
109742
109743
109744

109745
109746
109747
109748
109749

<value>	Formal Standard
ISO-IRΔ10646Δ2000ΔUTF-8	ISO/IEC 10646, UTF-8 encoding
BINARY	None.

If no **hdrcharset** extended header record is specified, the default character set used to encode all values in extended header records shall be the ISO/IEC 10646-1:2000 standard UTF-8 encoding.

The **BINARY** entry indicates that all values recorded in extended headers for affected files are unencoded binary data from the underlying system.

linkpath The pathname of a link being created to another file, of any type, previously archived. This record shall override the *linkname* field in the following **ustar** header block(s). The following **ustar** header block shall determine the type of link created. If *typeflag* of the following header block is 1, it shall be a hard link. If *typeflag* is 2, it shall be a symbolic link and the **linkpath** value shall be the contents of the symbolic link. The *pax* utility shall translate the name of the link (contents of the symbolic link) from the encoding in the header to the character set appropriate for the local file system. When used in **write** or **copy** mode, *pax* shall include a **linkpath** extended header record for each link whose pathname cannot be represented entirely with the members of the portable character set other than NUL.

mtime The file modification time of the following file(s), equivalent to the value of the *st_mtim* member of the **stat** structure for a file, as described in the *stat()* function. This record shall override the *mtime* field in the following header block(s). The modification time shall be restored if the process has appropriate privileges required to do so. The format of the <value> shall be as described in [pax Extended Header File Times](#) (on page 3239).

path The pathname of the following file(s). This record shall override the *name* and *prefix* fields in the following header block(s). The *pax* utility shall translate the pathname of the file from the encoding in the header to the character set appropriate for the local file system.

When used in **write** or **copy** mode, *pax* shall include a *path* extended header record for each file whose pathname cannot be represented entirely with the members of the portable character set other than NUL.

realtime.any The keywords prefixed by “realtime.” are reserved for future standardization.

security.any The keywords prefixed by “security.” are reserved for future standardization.

size The size of the file in octets, expressed as a decimal number using digits from the ISO/IEC 646:1991 standard. This record shall override the *size* field in the following header block(s). When used in **write** or **copy** mode, *pax* shall include a *size* extended header record for each file with a size value greater than 8 589 934 591 (octal 77 777 777 777).

uid The user ID of the file owner, expressed as a decimal number using digits from the ISO/IEC 646:1991 standard. This record shall override the *uid* field in the following header block(s). When used in **write** or **copy** mode, *pax* shall include a *uid* extended header record for each file whose owner ID is greater than 2 097 151 (octal 7 777 777).

109750 **uname** The owner of the following file(s), formatted as a user name in the user database.
 109751 This record shall override the *uid* and *uname* fields in the following header block(s),
 109752 and any *uid* extended header record. When used in **read**, **copy**, or **list** mode, *pax*
 109753 shall translate the name from the encoding in the header record to the character set
 109754 appropriate for the user database on the receiving system. If any of the characters
 109755 cannot be translated, and if neither the **-oinvalid=UTF-8** option nor the
 109756 **-oinvalid=binary** option is specified, the results are implementation-defined.
 109757 When used in **write** or **copy** mode, *pax* shall include a **uname** extended header
 109758 record for each file whose user name cannot be represented entirely with the letters
 109759 and digits of the portable character set.

109760 If the *<value>* field is zero length, it shall delete any header block field, previously entered
 109761 extended header value, or global extended header value of the same name.

109762 If a keyword in an extended header record (or in a **-o** option-argument) overrides or deletes a
 109763 corresponding field in the **ustar** header block, *pax* shall ignore the contents of that header block
 109764 field.

109765 Unlike the **ustar** header block fields, NULs shall not delimit *<value>*s; all characters within the
 109766 *<value>* field shall be considered data for the field. None of the length limitations of the **ustar**
 109767 header block fields in Table 3-15 (on page 3240) shall apply to the extended header records.

109768 **pax Extended Header Keyword Precedence**

109769 This section describes the precedence in which the various header records and fields and
 109770 command line options are selected to apply to a file in the archive. When *pax* is used in **read** or
 109771 **list** modes, it shall determine a file attribute in the following sequence:

- 109772 1. If **-odelete=keyword-prefix** is used, the affected attributes shall be determined from step
 109773 7., if applicable, or ignored otherwise.
- 109774 2. If **-okeyword:=** is used, the affected attributes shall be ignored.
- 109775 3. If **-okeyword:=value** is used, the affected attribute shall be assigned the value.
- 109776 4. If there is a *typeflag* **x** extended header record, the affected attribute shall be assigned the
 109777 *<value>*. When extended header records conflict, the last one given in the header shall
 109778 take precedence.
- 109779 5. If **-okeyword=value** is used, the affected attribute shall be assigned the value.
- 109780 6. If there is a *typeflag* **g** global extended header record, the affected attribute shall be
 109781 assigned the *<value>*. When global extended header records conflict, the last one given in
 109782 the global header shall take precedence.
- 109783 7. Otherwise, the attribute shall be determined from the **ustar** header block.

109784 **pax Extended Header File Times**

109785 The *pax* utility shall write an **mtime** record for each file in **write** or **copy** modes if the file's
 109786 modification time cannot be represented exactly in the **ustar** header logical record described in
 109787 **ustar Interchange Format** (on page 3240). This can occur if the time is out of **ustar** range, or if
 109788 the file system of the underlying implementation supports non-integer time granularities and
 109789 the time is not an integer. All of these time records shall be formatted as a decimal representation
 109790 of the time in seconds since the Epoch. If a *<period>* (' . ') decimal point character is present,
 109791 the digits to the right of the point shall represent the units of a subsecond timing granularity,
 109792 where the first digit is tenths of a second and each subsequent digit is a tenth of the previous
 109793 digit. In **read** or **copy** mode, the *pax* utility shall truncate the time of a file to the greatest value

109794 that is not greater than the input header file time. In **write** or **copy** mode, the *pax* utility shall
 109795 output a time exactly if it can be represented exactly as a decimal number, and otherwise shall
 109796 generate only enough digits so that the same time shall be recovered if the file is extracted on a
 109797 system whose underlying implementation supports the same time granularity.

109798 **ustar Interchange Format**

109799 A **ustar** archive tape or file shall contain a series of logical records. Each logical record shall be a
 109800 fixed-size logical record of 512 octets (see below). Although this format may be thought of as
 109801 being stored on 9-track industry-standard 12.7 mm (0.5 in) magnetic tape, other types of
 109802 transportable media are not excluded. Each file archived shall be represented by a header logical
 109803 record that describes the file, followed by zero or more logical records that give the contents of
 109804 the file. At the end of the archive file there shall be two 512-octet logical records filled with
 109805 binary zeros, interpreted as an end-of-archive indicator.

109806 The logical records may be grouped for physical I/O operations, as described under the
 109807 **-bblocksize** and **-x ustar** options. Each group of logical records may be written with a single
 109808 operation equivalent to the *write()* function. On magnetic tape, the result of this write shall be a
 109809 single tape physical block. The last physical block shall always be the full size, so logical records
 109810 after the two zero logical records may contain undefined data.

109811 The header logical record shall be structured as shown in the following table. All lengths and
 109812 offsets are in decimal.

109813 **Table 3-15** ustar Header Block

Field Name	Octet Offset	Length (in Octets)
<i>name</i>	0	100
<i>mode</i>	100	8
<i>uid</i>	108	8
<i>gid</i>	116	8
<i>size</i>	124	12
<i>mtime</i>	136	12
<i>chksum</i>	148	8
<i>typeflag</i>	156	1
<i>linkname</i>	157	100
<i>magic</i>	257	6
<i>version</i>	263	2
<i>uname</i>	265	32
<i>gname</i>	297	32
<i>devmajor</i>	329	8
<i>devminor</i>	337	8
<i>prefix</i>	345	155

109831 All characters in the header logical record shall be represented in the coded character set of the
 109832 ISO/IEC 646: 1991 standard. For maximum portability between implementations, names should
 109833 be selected from characters represented by the portable filename character set as octets with the
 109834 most significant bit zero. If an implementation supports the use of characters outside of <slash>
 109835 and the portable filename character set in names for files, users, and groups, one or more
 109836 implementation-defined encodings of these characters shall be provided for interchange
 109837 purposes.

109838 However, the *pax* utility shall never create filenames on the local system that cannot be accessed

109839 via the procedures described in POSIX.1-202x. If a filename is found on the medium that would
 109840 create an invalid filename, it is implementation-defined whether the data from the file is stored
 109841 on the file hierarchy and under what name it is stored. The *pax* utility may choose to ignore these
 109842 files as long as it produces an error indicating that the file is being ignored.

109843 Each field within the header logical record is contiguous; that is, there is no padding used. Each
 109844 character on the archive medium shall be stored contiguously.

109845 The fields *magic*, *uname*, and *gname* are character strings each terminated by a NUL character.
 109846 The fields *name*, *linkname*, and *prefix* are NUL-terminated character strings except when all
 109847 characters in the array contain non-NUL characters including the last character. The *version* field
 109848 is two octets containing the characters "00" (zero-zero). The *typeflag* contains a single character.
 109849 All other fields are leading zero-filled octal numbers using digits from the ISO/IEC 646:1991
 109850 standard IRV. Each numeric field is terminated by one or more <space> or NUL characters.

109851 The *name* and the *prefix* fields shall produce the pathname of the file. A new pathname shall be
 109852 formed, if *prefix* is not an empty string (its first character is not NUL), by concatenating *prefix* (up
 109853 to the first NUL character), a <slash> character, and *name*; otherwise, *name* is used alone. In
 109854 either case, *name* is terminated at the first NUL character. If *prefix* begins with a NUL character, it
 109855 shall be ignored. In this manner, pathnames of at most 256 characters can be supported. If a
 109856 pathname does not fit in the space provided, *pax* shall notify the user of the error, and shall not
 109857 store any part of the file—header or data—on the medium.

109858 The *linkname* field, described below, shall not use the *prefix* to produce a pathname. As such, a
 109859 *linkname* is limited to 100 characters. If the name does not fit in the space provided, *pax* shall
 109860 notify the user of the error, and shall not attempt to store the link on the medium.

109861 The *mode* field provides 12 bits encoded in the ISO/IEC 646:1991 standard octal digit
 109862 representation. The encoded bits shall represent the following values:

109863 **Table 3-16** *ustar mode* Field

Bit Value	POSIX.1-202x Bit	Description
04 000	S_ISUID	Set UID on execution.
02 000	S_ISGID	Set GID on execution.
01 000	<reserved>	Reserved for future standardization.
00 400	S_IRUSR	Read permission for file owner class.
00 200	S_IWUSR	Write permission for file owner class.
00 100	S_IXUSR	Execute/search permission for file owner class.
00 040	S_IRGRP	Read permission for file group class.
00 020	S_IWGRP	Write permission for file group class.
00 010	S_IXGRP	Execute/search permission for file group class.
00 004	S_IROTH	Read permission for file other class.
00 002	S_IWOTH	Write permission for file other class.
00 001	S_IXOTH	Execute/search permission for file other class.

109877 When appropriate privileges are required to set one of these mode bits, and the user restoring
 109878 the files from the archive does not have appropriate privileges, the mode bits for which the user
 109879 does not have appropriate privileges shall be ignored. Some of the mode bits in the archive
 109880 format are not mentioned elsewhere in this volume of POSIX.1-202x. If the implementation does
 109881 not support those bits, they may be ignored.

109882 The *uid* and *gid* fields are the user and group ID of the owner and group of the file, respectively.

109883 The *size* field is the size of the file in octets. If the *typeflag* field is set to specify a file to be of type
 109884 1 (a hard link) or 2 (a symbolic link), the *size* field shall be specified as zero. If the *typeflag* field is

- 109885 set to specify a file of type 5 (directory), the *size* field shall be interpreted as described under the
 109886 definition of that record type. No data logical records are stored for types 1, 2, or 5. If the *typeflag*
 109887 field is set to 3 (character special file), 4 (block special file), or 6 (FIFO), the meaning of the *size*
 109888 field is unspecified by this volume of POSIX.1-202x, and no data logical records shall be stored
 109889 on the medium. Additionally, for type 6, the *size* field shall be ignored when reading. If the
 109890 *typeflag* field is set to any other value, the number of logical records written following the header
 109891 shall be $(size+511)/512$, ignoring any fraction in the result of the division.
- 109892 The *mtime* field shall be the modification time of the file at the time it was archived. It is the
 109893 ISO/IEC 646:1991 standard representation of the octal value of the modification time obtained
 109894 from the *stat()* function.
- 109895 The *chksum* field shall be the ISO/IEC 646:1991 standard IRV representation of the octal value of
 109896 the simple sum of all octets in the header logical record. Each octet in the header shall be treated
 109897 as an unsigned value. These values shall be added to an unsigned integer, initialized to zero, the
 109898 precision of which is not less than 17 bits. When calculating the checksum, the *chksum* field is
 109899 treated as if it were all <space> characters.
- 109900 The *typeflag* field specifies the type of file archived. If a particular implementation does not
 109901 recognize the type, or the user does not have appropriate privileges to create that type, the file
 109902 shall be extracted as if it were a regular file if the file type is defined to have a meaning for the
 109903 *size* field that could cause data logical records to be written on the medium (see the previous
 109904 description for *size*). If conversion to a regular file occurs, the *pax* utility shall produce an error
 109905 indicating that the conversion took place. All of the *typeflag* fields shall be coded in the
 109906 ISO/IEC 646:1991 standard IRV:
- | | | |
|--------|------|--|
| 109907 | 0 | Represents a regular file. For backwards-compatibility, a <i>typeflag</i> value of binary zero (' <code>\0</code> ') should be recognized as meaning a regular file when extracting files from the archive. Archives written with this version of the archive file format create regular files with a <i>typeflag</i> value of the ISO/IEC 646:1991 standard IRV ' <code>0</code> '. |
| 109908 | | |
| 109909 | | |
| 109910 | | |
| 109911 | 1 | Represents a file linked to another file, of any type, previously archived. 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. The linked-to name is specified in the <i>linkname</i> field with a NUL-character terminator if it is less than 100 octets in length. |
| 109912 | | |
| 109913 | | |
| 109914 | | |
| 109915 | | |
| 109916 | 2 | Represents a symbolic link. The contents of the symbolic link shall be stored in the <i>linkname</i> field. |
| 109917 | | |
| 109918 | 3, 4 | Represent character special files and block special files respectively. In this case the <i>devmajor</i> and <i>devminor</i> fields shall contain information defining the device, the format of which is unspecified by this volume of POSIX.1-202x. Implementations may map the device specifications to their own local specification or may ignore the entry. |
| 109919 | | |
| 109920 | | |
| 109921 | | |
| 109922 | 5 | Specifies a directory or subdirectory. On systems where disk allocation is performed on a directory basis, the <i>size</i> field shall contain the maximum number of octets (which may be rounded to the nearest disk block allocation unit) that the directory may hold. A <i>size</i> field of zero indicates no such limiting. Systems that do not support limiting in this manner should ignore the <i>size</i> field. |
| 109923 | | |
| 109924 | | |
| 109925 | | |
| 109926 | | |
| 109927 | 6 | Specifies a FIFO special file. Note that the archiving of a FIFO file archives the existence of this file and not its contents. |
| 109928 | | |
| 109929 | 7 | Reserved to represent a file to which an implementation has associated some high-performance attribute. Implementations without such extensions should treat this file as a regular file (type 0). |
| 109930 | | |
| 109931 | | |

109932 A–Z The letters 'A' to 'Z', inclusive, are reserved for custom implementations. All other
109933 values are reserved for future versions of this standard.

109934 It is unspecified whether files with pathnames that refer to the same directory entry are archived
109935 as linked files or as separate files. If they are archived as linked files, this means that attempting
109936 to extract both pathnames from the resulting archive always causes an error (unless the `-u`
109937 option is used) because the link cannot be created.

109938 It is unspecified whether files with the same device and file serial numbers being appended to
109939 an archive are treated as linked files to members that were in the archive before the append.

109940 Attempts to archive a socket shall produce a diagnostic message when **ustar** interchange format
109941 is used, but may be allowed when **pax** interchange format is used. Handling of other file types is
109942 implementation-defined.

109943 The *magic* field is the specification that this archive was output in this archive format. If this field
109944 contains **ustar** (the five characters from the ISO/IEC 646:1991 standard IRV shown followed by
109945 NUL), the *uname* and *gname* fields shall contain the ISO/IEC 646:1991 standard IRV
109946 representation of the owner and group of the file, respectively (truncated to fit, if necessary).
109947 When the file is restored by a privileged, protection-preserving version of the utility, the user
109948 and group databases shall be scanned for these names. If found, the user and group IDs
109949 contained within these files shall be used rather than the values contained within the *uid* and *gid*
109950 fields.

109951 **cpio Interchange Format**

109952 The octet-oriented **cpio** archive format shall be a series of entries, each comprising a header that
109953 describes the file, the name of the file, and then the contents of the file.

109954 An archive may be recorded as a series of fixed-size blocks of octets. This blocking shall be used
109955 only to make physical I/O more efficient. The last group of blocks shall always be at the full
109956 size.

109957 For the octet-oriented **cpio** archive format, the individual entry information shall be in the order
109958 indicated and described by the following table; see also the **<cpio.h>** header.

109959

Table 3-17 Octet-Oriented cpio Archive Entry

109960

Header Field Name	Length (in Octets)	Interpreted as
<i>c_magic</i>	6	Octal number
<i>c_dev</i>	6	Octal number
<i>c_ino</i>	6	Octal number
<i>c_mode</i>	6	Octal number
<i>c_uid</i>	6	Octal number
<i>c_gid</i>	6	Octal number
<i>c_nlink</i>	6	Octal number
<i>c_rdev</i>	6	Octal number
<i>c_mtime</i>	11	Octal number
<i>c_namesize</i>	6	Octal number
<i>c_filesize</i>	11	Octal number
Filename Field Name	Length	Interpreted as
<i>c_name</i>	<i>c_namesize</i>	Pathname string
File Data Field Name	Length	Interpreted as
<i>c_filedata</i>	<i>c_filesize</i>	Data

109961

109962

109963

109964

109965

109966

109967

109968

109969

109970

109971

109972

109973

109974

109975

109976

cpio Header

109977

109978

109979

109980

109981

109982

For each file in the archive, a header as defined previously shall be written. The information in the header fields is written as streams of the ISO/IEC 646: 1991 standard characters interpreted as octal numbers. The octal numbers shall be extended to the necessary length by appending the ISO/IEC 646: 1991 standard IRV zeros at the most-significant-digit end of the number; the result is written to the most-significant digit of the stream of octets first. The fields shall be interpreted as follows:

109983

109984

c_magic Identify the archive as being a transportable archive by containing the identifying value "070707".

109985

109986

109987

c_dev, c_ino Contains values that uniquely identify the file within the archive (that is, no files contain the same pair of *c_dev* and *c_ino* values unless they are links to the same file). The values shall be determined in an unspecified manner.

109988

c_mode Contains the file type and access permissions as defined in the following table.

109989

Table 3-18 Values for `cpio c_mode` Field

109990

109991

109992

109993

109994

109995

109996

109997

109998

109999

110000

110001

110002

110003

110004

110005

110006

110007

110008

110009

110010

110011

File Permissions Name	Value	Indicates
C_IRUSR	000 400	Read by owner
C_IWUSR	000 200	Write by owner
C_IXUSR	000 100	Execute by owner
C_IRGRP	000 040	Read by group
C_IWGRP	000 020	Write by group
C_IXGRP	000 010	Execute by group
C_IROTH	000 004	Read by others
C_IWOTH	000 002	Write by others
C_IXOTH	000 001	Execute by others
C_ISUID	004 000	Set <i>uid</i>
C_ISGID	002 000	Set <i>gid</i>
C_ISVTX	001 000	Reserved
File Type Name	Value	Indicates
C_ISDIR	040 000	Directory
C_ISFIFO	010 000	FIFO
C_ISREG	0100 000	Regular file
C_ISLNK	0120 000	Symbolic link
C_ISBLK	060 000	Block special file
C_ISCHR	020 000	Character special file
C_ISSOCK	0140 000	Socket
C_ISCTG	0110 000	Reserved

110012

110013

110014

110015

110016

Directories, FIFOs, symbolic links, and regular files shall be supported on a system conforming to this volume of POSIX.1-202x; additional values defined previously are reserved for compatibility with existing systems. Additional file types may be supported; however, such files should not be written to archives intended to be transported to other systems.

110017

c_uid

Contains the user ID of the owner.

110018

c_gid

Contains the group ID of the group.

110019

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.

110020

110021

110022

110023

110024

110025

c_rdev

Contains implementation-defined information for character or block special files.

110026

c_mtime

Contains the latest time of modification of the file at the time the archive was created.

110027

110028

c_namesize

Contains the length of the pathname, including the terminating NUL character.

110029

c_filesize

Contains the length in octets of the data section following the header structure.

110030 cpio Filename

110031 The *c_name* field shall contain the pathname of the file. The length of this field in octets is the
110032 value of *c_namesize*.

110033 If a filename is found on the medium that would create an invalid pathname, it is
110034 implementation-defined whether the data from the file is stored on the file hierarchy and under
110035 what name it is stored.

110036 All characters shall be represented in the ISO/IEC 646:1991 standard IRV. For maximum
110037 portability between implementations, names should be selected from characters represented by
110038 the portable filename character set as octets with the most significant bit zero. If an
110039 implementation supports the use of characters outside the portable filename character set in
110040 names for files, users, and groups, one or more implementation-defined encodings of these
110041 characters shall be provided for interchange purposes. However, the *pax* utility shall never create
110042 filenames on the local system that cannot be accessed via the procedures described previously in
110043 this volume of POSIX.1-202x. If a filename is found on the medium that would create an invalid
110044 filename, it is implementation-defined whether the data from the file is stored on the local file
110045 system and under what name it is stored. The *pax* utility may choose to ignore these files as long
110046 as it produces an error indicating that the file is being ignored.

110047 cpio File Data

110048 Following *c_name*, there shall be *c_filesiz*e octets of data. Interpretation of such data occurs in a
110049 manner dependent on the file. For regular files, the data shall consist of the contents of the file.
110050 For symbolic links, the data shall consist of the contents of the symbolic link. If *c_filesiz*e is zero,
110051 no data shall be contained in *c_filedat*a.

110052 When restoring from an archive:

- 110053 • If the user does not have appropriate privileges to create a file of the specified type, *pax*
110054 shall ignore the entry and write an error message to standard error.
- 110055 • Only regular files and symbolic links have data to be restored. Presuming a regular file
110056 meets any selection criteria that might be imposed on the format-reading utility by the
110057 user, such data shall be restored.
- 110058 • If a user does not have appropriate privileges to set a particular mode flag, the flag shall be
110059 ignored. Some of the mode flags in the archive format are not mentioned elsewhere in this
110060 volume of POSIX.1-202x. If the implementation does not support those flags, they may be
110061 ignored.

110062 cpio Special Entries

110063 FIFO special files, directories, and the trailer shall be recorded with *c_filesiz*e equal to zero.
110064 Symbolic links shall be recorded with *c_filesiz*e equal to the length of the contents of the symbolic
110065 link. For other special files, *c_filesiz*e is unspecified by this volume of POSIX.1-202x. The header
110066 for the next file entry in the archive shall be written directly after the last octet of the file entry
110067 preceding it. A header denoting the filename **TRAILER!!!** shall indicate the end of the archive;
110068 the contents of octets in the last block of the archive following such a header are undefined.

110069 EXIT STATUS

110070 The following exit values shall be returned:

- 110071 0 All files were processed successfully.

110072 >0 An error occurred.

110073 CONSEQUENCES OF ERRORS

110074 If *pax* cannot create a file or a link when reading an archive or cannot find a file when writing an
 110075 archive, or cannot preserve the user ID, group ID, or file mode when the **-p** option is specified, a
 110076 diagnostic message shall be written to standard error and a non-zero exit status shall be
 110077 returned, but processing shall continue. In the case where *pax* cannot create a hard link to a file,
 110078 *pax* shall not, by default, create a second copy of the file.

110079 If the extraction of a file from an archive is prematurely terminated by a signal or error, *pax* may
 110080 have only partially extracted the file or (if the **-n** option was not specified) may have extracted a
 110081 file of the same name as that specified by the user, but which is not the file the user wanted.
 110082 Additionally, the file modes of extracted directories may have additional bits from the S_IRWXU
 110083 mask set as well as incorrect modification and access times.

110084 APPLICATION USAGE

110085 Caution is advised when using the **-a** option to append to a *cpio* format archive. If any of the
 110086 files being appended happen to be given the same *c_dev* and *c_ino* values as a file in the existing
 110087 part of the archive, then they may be treated as links to that file on extraction. Thus, it is risky to
 110088 use **-a** with *cpio* format except when it is done on the same system that the original archive was
 110089 created on, and with the same *pax* utility, and in the knowledge that there has been little or no
 110090 file system activity since the original archive was created that could lead to any of the files
 110091 appended being given the same *c_dev* and *c_ino* values as an unrelated file in the existing part of
 110092 the archive. Also, when (intentionally) appending additional links to a file in the existing part of
 110093 the archive, the *c_nlink* values in the modified archive can be smaller than the number of links to
 110094 the file in the archive, which may mean that the links are not preserved on extraction.

110095 The **-p** (privileges) option was invented to reconcile differences between historical *tar* and *cpio*
 110096 implementations. In particular, the two utilities use **-m** in diametrically opposed ways. The **-p**
 110097 option also provides a consistent means of extending the ways in which future file attributes can
 110098 be addressed, such as for enhanced security systems or high-performance files. Although it may
 110099 seem complex, there are really two modes that are most commonly used:

110100 **-p e** “Preserve everything”. This would be used by the historical superuser, someone with
 110101 all appropriate privileges, to preserve all aspects of the files as they are recorded in the
 110102 archive. The **e** flag is the sum of **o** and **p**, and other implementation-defined attributes.

110103 **-p p** “Preserve” the file mode bits. This would be used by the user with regular privileges
 110104 who wished to preserve aspects of the file other than the ownership. The file times are
 110105 preserved by default, but two other flags are offered to disable these and use the time of
 110106 extraction.

110107 The one pathname per line format of standard input precludes pathnames containing <newline>
 110108 characters. Although such pathnames violate the portable filename guidelines, they may exist
 110109 and their presence may inhibit usage of *pax* within shell scripts. This problem is inherited from
 110110 historical archive programs. The problem can be avoided by listing filename arguments on the
 110111 command line instead of on standard input.

110112 It is almost certain that appropriate privileges are required for *pax* to accomplish parts of this
 110113 volume of POSIX.1-202x. Specifically, creating files of type block special or character special,
 110114 restoring file access times unless the files are owned by the user (the **-t** option), or preserving file
 110115 owner, group, and mode (the **-p** option) all probably require appropriate privileges.

110116 In **read** mode, implementations are permitted to overwrite files when the archive has multiple
 110117 members with the same name. This may fail if permissions on the first version of the file do not
 110118 permit it to be overwritten.

110119 The **cpio** and **ustar** formats can only support files up to 8 589 934 592 bytes ($8 * 2^{30}$) in size.

110120 When archives containing binary header information are listed , the filenames printed may
110121 cause strange behavior on some terminals.

110122 When all of the following are true:

- 110123 1. A file of type directory is being placed into an archive.
- 110124 2. The **ustar** archive format is being used.
- 110125 3. The pathname of the directory is less than or equal to 155 bytes long (it will fit in the *prefix*
110126 field in the **ustar** header block).
- 110127 4. The last component of the pathname of the directory is longer than 100 bytes long (it will
110128 not fit in the *name* field in the **ustar** header block).

110129 some implementations of the *pax* utility will place the entire directory pathname in the *prefix*
110130 field, set the *name* field to an empty string, and place the directory in the archive. Other
110131 implementations of the *pax* utility will give an error under these conditions because the *name*
110132 field is not large enough to hold the last component of the directory name. This standard allows
110133 either behavior. However, when extracting a directory from a **ustar** format archive, this standard
110134 requires that all implementations be able to extract a directory even if the *name* field contains an
110135 empty string as long as the *prefix* field does not also contain an empty string.

110136 When restricting file hierarchy traversal to one file system, it can sometimes be desirable for the +
110137 crossing points themselves to be processed (archived or copied) and sometimes for them not to +
110138 be processed. (Crossing points are mount points and, if the **-L** option is specified, symbolic links +
110139 to directories on other file systems.) With the **-X** option *pax* processes them, but there is no +
110140 standard way to have *pax* not process them. However, this can be achieved by using *find* to do +
110141 the hierarchy traversal and piping the output of *find* to *pax* (with the **-d** option); see the +
110142 APPLICATION USAGE for *find*.

110143 EXAMPLES

110144 The following command:

```
110145 pax -w -f /dev/rmt/1m .
```

110146 copies the contents of the current directory to tape drive 1, medium density (assuming historical
110147 System V device naming procedures—the historical BSD device name would be **/dev/rmt9**).

110148 The following commands:

```
110149 mkdir newdir  
110150 pax -rw olddir newdir
```

110151 copy the *olddir* directory hierarchy to *newdir*.

```
110152 pax -r -s '^//*usr//*,,' -f a.pax
```

110153 reads the archive **a.pax**, with all files rooted in **/usr** in the archive extracted relative to the current
110154 directory.

110155 Using the option:

```
110156 -o listopt="%M %(atime)T %(size)D %(name)s"
```

110157 overrides the default output description in Standard Output and instead writes:

```
110158 -rw-rw--- Jan 12 15:53 2003 1492 /usr/foo/bar
```

110159 Using the options:

```
110160 -o listopt='%L\t%(size)D\n%.7' \
110161 -o listopt='(name)s\n%(atime)T\n%T'
```

110162 overrides the default output description in Standard Output and instead writes:

```
110163 /usr/foo/bar -> /tmp 1492
110164 /usr/fo
110165 Jan 12 15:53 1991
110166 Jan 31 15:53 2003
```

110167 RATIONALE

110168 The *pax* utility was new for the ISO POSIX-2:1993 standard. It represents a peaceful compromise
110169 between advocates of the historical *tar* and *cpio* utilities.

110170 A fundamental difference between *cpio* and *tar* was in the way directories were treated. The *cpio*
110171 utility did not treat directories differently from other files, and to select a directory and its
110172 contents required that each file in the hierarchy be explicitly specified. For *tar*, a directory
110173 matched every file in the file hierarchy it rooted.

110174 The *pax* utility offers both interfaces; by default, directories map into the file hierarchy they root.
110175 The `-d` option causes *pax* to skip any file not explicitly referenced, as *cpio* historically did. The *tar*
110176 `-style` behavior was chosen as the default because it was believed that this was the more
110177 common usage and because *tar* is the more commonly available interface, as it was historically
110178 provided on both System V and BSD implementations.

110179 The data interchange format specification in this volume of POSIX.1-202x requires that processes
110180 with “appropriate privileges” shall always restore the ownership and permissions of extracted
110181 files exactly as archived. If viewed from the historic equivalence between superuser and
110182 “appropriate privileges”, there are two problems with this requirement. First, users running as
110183 superusers may unknowingly set dangerous permissions on extracted files. Second, it is
110184 needlessly limiting, in that superusers cannot extract files and own them as superuser unless the
110185 archive was created by the superuser. (It should be noted that restoration of ownerships and
110186 permissions for the superuser, by default, is historical practice in *cpio*, but not in *tar*.) In order to
110187 avoid these two problems, the *pax* specification has an additional “privilege” mechanism, the `-p`
110188 option. Only a *pax* invocation with the privileges needed, and which has the `-p` option set using
110189 the `e` specification character, has appropriate privileges to restore full ownership and permission
110190 information.

110191 Note also that this volume of POSIX.1-202x requires that the file ownership and access
110192 permissions shall be set, on extraction, in the same fashion as the *creat()* function when provided
110193 with the mode stored in the archive. This means that the file creation mask of the user is applied
110194 to the file permissions.

110195 Users should note that directories may be created by *pax* while extracting files with permissions
110196 that are different from those that existed at the time the archive was created. When extracting
110197 sensitive information into a directory hierarchy that no longer exists, users are encouraged to set
110198 their file creation mask appropriately to protect these files during extraction.

110199 The table of contents output is written to standard output to facilitate pipeline processing.

110200 An early proposal had hard links displaying for all pathnames. This was removed because it
110201 complicates the output of the case where `-v` is not specified and does not match historical *cpio*
110202 usage. The hard-link information is available in the `-v` display.

110203 The description of the `-l` option allows implementations to make hard links to symbolic links.
110204 Earlier versions of this standard did not specify any way to create a hard link to a symbolic link,
110205 but many implementations provided this capability as an extension. If there are hard links to

110206 symbolic links when an archive is created, the implementation is required to archive the hard
110207 link in the archive (unless **-H** or **-L** is specified). When in **read** mode and in **copy** mode,
110208 implementations supporting hard links to symbolic links should use them when appropriate.

110209 The archive formats inherited from the POSIX.1-1990 standard have certain restrictions that have
110210 been brought along from historical usage. For example, there are restrictions on the length of
110211 pathnames stored in the archive. When *pax* is used in **copy(-rw)** mode (copying directory
110212 hierarchies), the ability to use extensions from the **-xpax** format overcomes these restrictions.

110213 The default *blocksize* value of 5 120 bytes for *cpio* was selected because it is one of the standard
110214 block-size values for *cpio*, set when the **-B** option is specified. (The other default block-size value
110215 for *cpio* is 512 bytes, and this was considered to be too small.) The default block value of 10 240
110216 bytes for *tar* was selected because that is the standard block-size value for BSD *tar*. The
110217 maximum block size of 32 256 bytes (2^{15} –512 bytes) is the largest multiple of 512 bytes that fits
110218 into a signed 16-bit tape controller transfer register. There are known limitations in some
110219 historical systems that would prevent larger blocks from being accepted. Historical values were
110220 chosen to improve compatibility with historical scripts using *dd* or similar utilities to manipulate
110221 archives. Also, default block sizes for any file type other than character special file has been
110222 deleted from this volume of POSIX.1-202x as unimportant and not likely to affect the structure of
110223 the resulting archive.

110224 Implementations are permitted to modify the block-size value based on the archive format or the
110225 device to which the archive is being written. This is to provide implementations with the
110226 opportunity to take advantage of special types of devices, and it should not be used without a
110227 great deal of consideration as it almost certainly decreases archive portability.

110228 The intended use of the **-n** option was to permit extraction of one or more files from the archive
110229 without processing the entire archive. This was viewed by the standard developers as offering
110230 significant performance advantages over historical implementations. The **-n** option in early
110231 proposals had three effects; the first was to cause special characters in patterns to not be treated
110232 specially. The second was to cause only the first file that matched a pattern to be extracted. The
110233 third was to cause *pax* to write a diagnostic message to standard error when no file was found
110234 matching a specified pattern. Only the second behavior is retained by this volume of
110235 POSIX.1-202x, for many reasons. First, it is in general not acceptable for a single option to have
110236 multiple effects. Second, the ability to make pattern matching characters act as normal characters
110237 is useful for parts of *pax* other than file extraction. Third, a finer degree of control over the
110238 special characters is useful because users may wish to normalize only a single special character
110239 in a single filename. Fourth, given a more general escape mechanism, the previous behavior of
110240 the **-n** option can be easily obtained using the **-s** option or a *sed* script. Finally, writing a
110241 diagnostic message when a pattern specified by the user is unmatched by any file is useful
110242 behavior in all cases.

110243 In this version, the **-n** was removed from the **copy** mode synopsis of *pax*; it is inapplicable
110244 because there are no pattern operands specified in this mode.

110245 There is another method than *pax* for copying subtrees in POSIX.1-202x described as part of the
110246 *cp* utility. Both methods are historical practice: *cp* provides a simpler, more intuitive interface,
110247 while *pax* offers a finer granularity of control. Each provides additional functionality to the
110248 other; in particular, *pax* maintains the hard-link structure of the hierarchy while *cp* does not. It is
110249 the intention of the standard developers that the results be similar (using appropriate option
110250 combinations in both utilities). The results are not required to be identical; there seemed
110251 insufficient gain to applications to balance the difficulty of implementations having to guarantee
110252 that the results would be exactly identical.

110253 A single archive may span more than one file. It is suggested that implementations provide

110254 informative messages to the user on standard error whenever the archive file is changed.

110255 The **-d** option (do not create intermediate directories not listed in the archive) found in early
110256 proposals was originally provided as a complement to the historic **-d** option of *cpio*. It has been
110257 deleted.

110258 The **-s** option in early proposals specified a subset of the substitution command from the *ed*
110259 utility. As there was no reason for only a subset to be supported, the **-s** option is now compatible
110260 with the current *ed* specification. Since the delimiter can be any non-null character, the following
110261 usage with single <space> characters is valid:

```
110262 pax -s " foo bar " ...
```

110263 The **-t** description is worded so as to note that this may cause the access time update caused by
110264 some other activity (which occurs while the file is being read) to be overwritten.

110265 The default behavior of *pax* with regard to file modification times is the same as historical
110266 implementations of *tar*. It is not the historical behavior of *cpio*.

110267 Because the **-i** option uses */dev/tty*, utilities without a controlling terminal are not able to use
110268 this option.

110269 The **-y** option, found in early proposals, has been deleted because a line containing a single
110270 <period> for the **-i** option has equivalent functionality. The special lines for the **-i** option (a
110271 single <period> and the empty line) are historical practice in *cpio*.

110272 In early drafts, a **-charmap** option was included to increase portability of files between systems
110273 using different coded character sets. This option was omitted because it was apparent that
110274 consensus could not be formed for it. In this version, the use of UTF-8 should be an adequate
110275 substitute.

110276 The ISO POSIX-2:1993 standard and ISO POSIX-1 standard requirements for *pax*, however,
110277 made it very difficult to create a single archive containing files created using extended characters
110278 provided by different locales. This version adds the **hdrcharset** keyword to make it possible to
110279 archive files in these cases without dropping files due to translation errors.

110280 Translating filenames and other attributes from a locale's encoding to UTF-8 and then back again
110281 can lose information, as the resulting filename might not be byte-for-byte equivalent to the
110282 original. To avoid this problem, users can specify the **-o hdrcharset=binary** option, which will
110283 cause the resulting archive to use binary format for all names and attributes. Such archives are
110284 not portable among hosts that use different native encodings (e.g., EBCDIC *versus* ASCII-based
110285 encodings), but they will allow interchange among the vast majority of POSIX file systems in
110286 practical use. Also, the **-o hdrcharset=binary** option will cause *pax* in **copy** mode to behave
110287 more like other standard utilities such as *cp*.

110288 If the values specified by the **-o exthdr.name=value**, **-o globexthdr.name=value**, or by
110289 **\$TMPDIR** (if **-o globexthdr.name** is not specified) require a character encoding other than that
110290 described in the ISO/IEC 646:1991 standard, a **path** extended header record will have to be
110291 created for the file. If a **hdrcharset** extended header record is active for such headers, it will
110292 determine the codeset used for the value field in these extended **path** header records. These **path**
110293 extended header records always need to be created when writing an archive even if
110294 **hdrcharset=binary** has been specified and would contain the same (binary) data that appears in
110295 the **ustar** header record prefix and *name* fields. (In other words, an extended header **path** record
110296 is always required to be generated if the *prefix* or *name* fields contain non-ASCII characters even
110297 when **hdrcharset=binary** is also in effect for that file.)

110298 The **-k** option was added to address international concerns about the dangers involved in the
110299 character set transformations of **-e** (if the target character set were different from the source, the

110300 filenames might be transformed into names matching existing files) and also was made more
110301 general to protect files transferred between file systems with different {NAME_MAX} values
110302 (truncating a filename on a smaller system might also inadvertently overwrite existing files). As
110303 stated, it prevents any overwriting, even if the target file is older than the source. This version
110304 adds more granularity of options to solve this problem by introducing the **-oinvalid=option**—
110305 specifically the **UTF-8** and **binary** actions. (Note that an existing file is still subject to overwriting
110306 in this case. The **-k** option closes that loophole.)

110307 Some of the file characteristics referenced in this volume of POSIX.1-202x might not be
110308 supported by some archive formats. For example, neither the **tar** nor **cpio** formats contain the
110309 file access time. For this reason, the **e** specification character has been provided, intended to
110310 cause all file characteristics specified in the archive to be retained.

110311 It is required that extracted directories, by default, have their access and modification times and
110312 permissions set to the values specified in the archive. This has obvious problems in that the
110313 directories are almost certainly modified after being extracted and that directory permissions
110314 may not permit file creation. One possible solution is to create directories with the mode
110315 specified in the archive, as modified by the *umask* of the user, with sufficient permissions to
110316 allow file creation. After all files have been extracted, *pax* would then reset the access and
110317 modification times and permissions as necessary.

110318 The list-mode formatting description borrows heavily from the one defined by the *printf* utility.
110319 However, since there is no separate operand list to get conversion arguments, the format was
110320 extended to allow specifying the name of the conversion argument as part of the conversion
110321 specification.

110322 The **T** conversion specifier allows time fields to be displayed in any of the date formats. Unlike
110323 the *ls* utility, *pax* does not adjust the format when the date is less than six months in the past.
110324 This makes parsing the output more predictable.

110325 The **D** conversion specifier handles the ability to display the major/minor or file size, as with *ls*,
110326 by using **%-8(size)D**.

110327 The **L** conversion specifier handles the *ls* display for symbolic links.

110328 Conversion specifiers were added to generate existing known types used for *ls*.

110329 **pax Interchange Format**

110330 The new POSIX data interchange format was developed primarily to satisfy international
110331 concerns that the **ustar** and **cpio** formats did not provide for file, user, and group names encoded
110332 in characters outside a subset of the ISO/IEC 646:1991 standard. The standard developers
110333 realized that this new POSIX data interchange format should be very extensible because there
110334 were other requirements they foresaw in the near future:

- 110335 • Support international character encodings and locale information
- 110336 • Support security information (ACLs, and so on)
- 110337 • Support future file types, such as realtime or contiguous files
- 110338 • Include data areas for implementation use
- 110339 • Support systems with words larger than 32 bits and timers with subsecond granularity

110340 The following were not goals for this format because these are better handled by separate
110341 utilities or are inappropriate for a portable format:

- 110342 • Encryption
- 110343 • Compression
- 110344 • Data translation between locales and codesets
- 110345 • *inode* storage

110346 The format chosen to support the goals is an extension of the **ustar** format. Of the two formats
110347 previously available, only the **ustar** format was selected for extensions because:

- 110348 • It was easier to extend in an upwards-compatible way. It offered version flags and header
110349 block type fields with room for future standardization. The **cpio** format, while possessing a
110350 more flexible file naming methodology, could not be extended without breaking some
110351 theoretical implementation or using a dummy filename that could be a legitimate filename.
- 110352 • Industry experience since the original “*tar wars*” fought in developing the ISO POSIX-1
110353 standard has clearly been in favor of the **ustar** format, which is generally the default
110354 output format selected for *pax* implementations on new systems.

110355 The new format was designed with one additional goal in mind: reasonable behavior when an
110356 older *tar* or *pax* utility happened to read an archive. Since the POSIX.1-1990 standard mandated
110357 that a “format-reading utility” had to treat unrecognized *typeflag* values as regular files, this
110358 allowed the format to include all the extended information in a pseudo-regular file that
110359 preceded each real file. An option is given that allows the archive creator to set up reasonable
110360 names for these files on the older systems. Also, the normative text suggests that reasonable file
110361 access values be used for this **ustar** header block. Making these header files inaccessible for
110362 convenient reading and deleting would not be reasonable. File permissions of 600 or 700 are
110363 suggested.

110364 The **ustar** *typeflag* field was used to accommodate the additional functionality of the new format
110365 rather than magic or version because the POSIX.1-1990 standard (and, by reference, the previous
110366 version of *pax*), mandated the behavior of the format-reading utility when it encountered an
110367 unknown *typeflag*, but was silent about the other two fields.

110368 Early proposals for the first version of this standard contained a proposed archive format that
110369 was based on compatibility with the standard for tape files (ISO 1001, similar to the format used
110370 historically on many mainframes and minicomputers). This format was overly complex and
110371 required considerable overhead in volume and header records. Furthermore, the standard
110372 developers felt that it would not be acceptable to the community of POSIX developers, so it was
110373 later changed to be a format more closely related to historical practice on POSIX systems.

110374 The prefix and name split of pathnames in **ustar** was replaced by the single path extended
110375 header record for simplicity.

110376 The concept of a global extended header (*typeflag***g**) was controversial. If this were applied to an
110377 archive being recorded on magnetic tape, a few unreadable blocks at the beginning of the tape
110378 could be a serious problem; a utility attempting to extract as many files as possible from a
110379 damaged archive could lose a large percentage of file header information in this case. However,
110380 if the archive were on a reliable medium, such as a CD-ROM, the global extended header offers
110381 considerable potential size reductions by eliminating redundant information. Thus, the text
110382 warns against using the global method for unreliable media and provides a method for
110383 implanting global information in the extended header for each file, rather than in the *typeflag* **g**
110384 records.

110385 No facility for data translation or filtering on a per-file basis is included because the standard
110386 developers could not invent an interface that would allow this in an efficient manner. If a filter,
110387 such as encryption or compression, is to be applied to all the files, it is more efficient to apply the

110388 filter to the entire archive as a single file. The standard developers considered interfaces that
110389 would invoke a shell script for each file going into or out of the archive, but the system overhead
110390 in this approach was considered to be too high.

110391 One such approach would be to have **filter=** records that give a pathname for an executable.
110392 When the program is invoked, the file and archive would be open for standard input/output
110393 and all the header fields would be available as environment variables or command-line
110394 arguments. The standard developers did discuss such schemes, but they were omitted from
110395 POSIX.1-202x due to concerns about excessive overhead. Also, the program itself would need to
110396 be in the archive if it were to be used portably.

110397 There is currently no portable means of identifying the character set(s) used for a file in the file
110398 system. Therefore, *pax* has not been given a mechanism to generate charset records
110399 automatically. The only portable means of doing this is for the user to write the archive using the
110400 **-ocharset=string** command line option. This assumes that all of the files in the archive use the
110401 same encoding. The “implementation-defined” text is included to allow for a system that can
110402 identify the encodings used for each of its files.

110403 The table of standards that accompanies the charset record description is acknowledged to be
110404 very limited. Only a limited number of character set standards is reasonable for maximal
110405 interchange. Any character set is, of course, possible by prior agreement. It was suggested that
110406 EBCDIC be listed, but it was omitted because it is not defined by a formal standard. Formal
110407 standards, and then only those with reasonably large followings, can be included here, simply as
110408 a matter of practicality. The *<value>*s represent names of officially registered character sets in the
110409 format required by the ISO 2375:1985 standard.

110410 The normal *<comma>* or *<blank>*-separated list rules are not followed in the case of keyword
110411 options to allow ease of argument parsing for *getopts*.

110412 Further information on character encodings is in [pax Archive Character Set Encoding/Decoding](#)
110413 (on page 3256).

110414 The standard developers have reserved keyword name space for vendor extensions. It is
110415 suggested that the format to be used is:

110416 *VENDOR.keyword*

110417 where *VENDOR* is the name of the vendor or organization in all uppercase letters. It is further
110418 suggested that the keyword following the *<period>* be named differently than any of the
110419 standard keywords so that it could be used for future standardization, if appropriate, by
110420 omitting the *VENDOR* prefix.

110421 The *<length>* field in the extended header record was included to make it simpler to step
110422 through the records, even if a record contains an unknown format (to a particular *pax*) with
110423 complex interactions of special characters. It also provides a minor integrity checkpoint within
110424 the records to aid a program attempting to recover files from a damaged archive.

110425 There are no extended header versions of the *devmajor* and *devminor* fields because the
110426 unspecified format **ustar** header field should be sufficient. If they are not, vendor-specific
110427 extended keywords (such as *VENDOR.devmajor*) should be used.

110428 Device and *i*-number labeling of files was not adopted from *cpio*; files are interchanged strictly
110429 on a symbolic name basis, as in **ustar**.

110430 Just as with the **ustar** format descriptions, the new format makes no special arrangements for
110431 multi-volume archives. Each of the *pax* archive types is assumed to be inside a single POSIX file
110432 and splitting that file over multiple volumes (diskettes, tape cartridges, and so on), processing
110433 their labels, and mounting each in the proper sequence are considered to be implementation

110434 details that cannot be described portably.

110435 The **pax** format is intended for interchange, not only for backup on a single (family of) systems.
110436 It is not as densely packed as might be possible for backup:

- 110437 • It contains information as coded characters that could be coded in binary.
- 110438 • It identifies extended records with name fields that could be omitted in favor of a fixed-
110439 field layout.
- 110440 • It translates names into a portable character set and identifies locale-related information,
110441 both of which are probably unnecessary for backup.

110442 The requirements on restoring from an archive are slightly different from the historical wording,
110443 allowing for non-monolithic privilege to bring forward as much as possible. In particular,
110444 attributes such as “high performance file” might be broadly but not universally granted while
110445 set-user-ID or *chown()* might be much more restricted. There is no implication in POSIX.1-202x
110446 that the security information be honored after it is restored to the file hierarchy, in spite of what
110447 might be improperly inferred by the silence on that topic. That is a topic for another standard.

110448 Hard links are recorded in the fashion described here because a hard link can be to any file type. |
110449 It is desirable in general to be able to restore part of an archive selectively and restore all of those |
110450 files completely. If the data is not associated with each hard link, it is not possible to do this. |
110451 However, the data associated with a file can be large, and when selective restoration is not |
110452 needed, this can be a significant burden. The archive is structured so that files that have no |
110453 associated data can always be restored by the name of any link name of any hard link, and the |
110454 user can choose whether data is recorded with each instance of a file that contains data. The |
110455 format permits mixing of hard links with data and hard links without data in a single archive; |
110456 this can be done for special needs, and *pax* is expected to interpret such archives on input |
110457 properly, despite the fact that there is no *pax* option that would force this mixed case on output. |
110458 (When **-o linkdata** is used, the output must contain the duplicate data, but the implementation |
110459 is free to include it or omit it when **-o linkdata** is not used.)

110460 The time values are included as extended header records for those implementations needing
110461 more than the eleven octal digits allowed by the **ustar** format. Portable file timestamps cannot be
110462 negative. If *pax* encounters a file with a negative timestamp in **copy** or **write** mode, it can reject
110463 the file, substitute a non-negative timestamp, or generate a non-portable timestamp with a
110464 leading '-'. Even though some implementations can support finer file-time granularities than
110465 seconds, the normative text requires support only for seconds since the Epoch because the
110466 ISO POSIX-1 standard states them that way. The **ustar** format includes only *mtime*; the new
110467 format adds *atime* and *ctime* for symmetry. The *atime* access time restored to the file system will
110468 be affected by the **-p a** and **-p e** options. The *ctime* creation time (actually *inode* modification
110469 time) is described with appropriate privileges so that it can be ignored when writing to the file
110470 system. POSIX does not provide a portable means to change file creation time. Nothing is
110471 intended to prevent a non-portable implementation of *pax* from restoring the value.

110472 The *gid*, *size*, and *uid* extended header records were included to allow expansion beyond the
110473 sizes specified in the regular *tar* header. New file system architectures are emerging that will
110474 exhaust the 12-digit size field. There are probably not many systems requiring more than 8 digits
110475 for user and group IDs, but the extended header values were included for completeness,
110476 allowing overrides for all of the decimal values in the *tar* header.

110477 The standard developers intended to describe the effective results of *pax* with regard to file
110478 ownerships and permissions; implementations are not restricted in timing or sequencing the
110479 restoration of such, provided the results are as specified.

110480 Much of the text describing the extended headers refers to use in “**write** or **copy** modes”. The

110481 **copy** mode references are due to the normative text: “The effect of the copy shall be as if the
110482 copied files were written to an archive file and then subsequently extracted ...”. There is
110483 certainly no way to test whether *pax* is actually generating the extended headers in **copy** mode,
110484 but the effects must be as if it had.

110485 **pax Archive Character Set Encoding/Decoding**

110486 There is a need to exchange archives of files between systems of different native codesets.
110487 Filenames, group names, and user names must be preserved to the fullest extent possible when
110488 an archive is read on the receiving platform. Translation of the contents of files is not within the
110489 scope of the *pax* utility.

110490 There will also be the need to represent characters that are not available on the receiving
110491 platform. These unsupported characters cannot be automatically folded to the local set of
110492 characters due to the chance of collisions. This could result in overwriting previous extracted
110493 files from the archive or pre-existing files on the system.

110494 For these reasons, the codeset used to represent characters within the extended header records of
110495 the *pax* archive must be sufficiently rich to handle all commonly used character sets. The fields
110496 requiring translation include, at a minimum, filenames, user names, group names, and link
110497 pathnames. Implementations may wish to have localized extended keywords that use non-
110498 portable characters.

110499 The standard developers considered the following options:

- 110500 • The archive creator specifies the well-defined name of the source codeset. The receiver
110501 must then recognize the codeset name and perform the appropriate translations to the
110502 destination codeset.
- 110503 • The archive creator includes within the archive the character mapping table for the source
110504 codeset used to encode extended header records. The receiver must then read the
110505 character mapping table and perform the appropriate translations to the destination
110506 codeset.
- 110507 • The archive creator translates the extended header records in the source codeset into a
110508 canonical form. The receiver must then perform the appropriate translations to the
110509 destination codeset.

110510 The approach that incorporates the name of the source codeset poses the problem of codeset
110511 name registration, and makes the archive useless to *pax* archive decoders that do not recognize
110512 that codeset.

110513 Because parts of an archive may be corrupted, the standard developers felt that including the
110514 character map of the source codeset was too fragile. The loss of this one key component could
110515 result in making the entire archive useless. (The difference between this and the global extended
110516 header decision was that the latter has a workaround—duplicating extended header records on
110517 unreliable media—but this would be too burdensome for large character set maps.)

110518 Both of the above approaches also put an undue burden on the *pax* archive receiver to handle the
110519 cross-product of all source and destination codesets.

110520 To simplify the translation from the source codeset to the canonical form and from the canonical
110521 form to the destination codeset, the standard developers decided that the internal representation
110522 should be a stateless encoding. A stateless encoding is one where each codepoint has the same
110523 meaning, without regard to the decoder being in a specific state. An example of a stateful
110524 encoding would be the Japanese Shift-JIS; an example of a stateless encoding would be the
110525 ISO/IEC 646: 1991 standard (equivalent to 7-bit ASCII).

110526 For these reasons, the standard developers decided to adopt a canonical format for the
 110527 representation of file information strings. The obvious, well-endorsed candidate is the
 110528 ISO/IEC 10646-1:2000 standard (based in part on Unicode), which can be used to represent the
 110529 characters of virtually all standardized character sets. The standard developers initially agreed
 110530 upon using UCS2 (16-bit Unicode) as the internal representation. This repertoire of characters
 110531 provides a sufficiently rich set to represent all commonly-used codesets.

110532 However, the standard developers found that the 16-bit Unicode representation had some
 110533 problems. It forced the issue of standardizing byte ordering. The 2-byte length of each character
 110534 made the extended header records twice as long for the case of strings coded entirely from
 110535 historical 7-bit ASCII. For these reasons, the standard developers chose the UTF-8 defined in the
 110536 ISO/IEC 10646-1:2000 standard. This multi-byte representation encodes UCS2 or UCS4
 110537 characters reliably and deterministically, eliminating the need for a canonical byte ordering. In
 110538 addition, NUL octets and other characters possibly confusing to POSIX file systems do not
 110539 appear, except to represent themselves. It was realized that certain national codesets take up
 110540 more space after the encoding, due to their placement within the UCS range; it was felt that the
 110541 usefulness of the encoding of the names outweighs the disadvantage of size increase for file,
 110542 user, and group names.

110543 The encoding of UTF-8 is as follows:

110544	UCS4 Hex Encoding	UTF-8 Binary Encoding
110545	00000000–0000007F	0xxxxxxx
110546	00000080–000007FF	110xxxxx 10xxxxxx
110547	00000800–0000FFFF	1110xxxx 10xxxxxx 10xxxxxx
110548	00010000–001FFFFFFF	11110xxx 10xxxxxx 10xxxxxx 10xxxxxx
110549	00200000–03FFFFFFF	111110xx 10xxxxxx 10xxxxxx 10xxxxxx 10xxxxxx
110550	04000000–7FFFFFFF	1111110x 10xxxxxx 10xxxxxx 10xxxxxx 10xxxxxx 10xxxxxx

110551 where each 'x' represents a bit value from the character being translated.

110552 **ustar Interchange Format**

110553 The description of the **ustar** format reflects numerous enhancements over pre-1988 versions of
 110554 the historical *tar* utility. The goal of these changes was not only to provide the functional
 110555 enhancements desired, but also to retain compatibility between new and old versions. This
 110556 compatibility has been retained. Archives written using the old archive format are compatible
 110557 with the new format.

110558 Implementors should be aware that the previous file format did not include a mechanism to
 110559 archive directory type files. For this reason, the convention of using a filename ending with
 110560 <slash> was adopted to specify a directory on the archive.

110561 The total size of the *name* and *prefix* fields have been set to meet the minimum requirements for
 110562 {PATH_MAX}. If a pathname will fit within the *name* field, it is recommended that the pathname
 110563 be stored there without the use of the *prefix* field. Although the name field is known to be too
 110564 small to contain {PATH_MAX} characters, the value was not changed in this version of the
 110565 archive file format to retain backwards-compatibility, and instead the prefix was introduced.
 110566 Also, because of the earlier version of the format, there is no way to remove the restriction on the
 110567 *linkname* field being limited in size to just that of the *name* field.

110568 The *size* field is required to be meaningful in all implementation extensions, although it could be
 110569 zero. This is required so that the data blocks can always be properly counted.

110570 It is suggested that if device special files need to be represented that cannot be represented in the
 110571 standard format, that one of the extension types (A-Z) be used, and that the additional

- 110572 information for the special file be represented as data and be reflected in the *size* field.
- 110573 Attempting to restore a special file type, where it is converted to ordinary data and conflicts with
110574 an existing filename, need not be specially detected by the utility. If run as an ordinary user, *pax*
110575 should not be able to overwrite the entries in, for example, */dev* in any case (whether the file is
110576 converted to another type or not). If run as a privileged user, it should be able to do so, and it
110577 would be considered a bug if it did not. The same is true of ordinary data files and similarly
110578 named special files; it is impossible to anticipate the needs of the user (who could really intend
110579 to overwrite the file), so the behavior should be predictable (and thus regular) and rely on the
110580 protection system as required.
- 110581 The value 7 in the *typeflag* field is intended to define how contiguous files can be stored in a
110582 **ustar** archive. POSIX.1-202x does not require the contiguous file extension, but does define a
110583 standard way of archiving such files so that all conforming systems can interpret these file types
110584 in a meaningful and consistent manner. On a system that does not support extended file types,
110585 the *pax* utility should do the best it can with the file and go on to the next.
- 110586 The file protection modes are those conventionally used by the *ls* utility. This is extended beyond
110587 the usage in the ISO POSIX-2 standard to support the “shared text” or “sticky” bit. It is intended
110588 that the conformance document should not document anything beyond the existence of and
110589 support of such a mode. Further extensions are expected to these bits, particularly with
110590 overloading the set-user-ID and set-group-ID flags.
- 110591 **cpio Interchange Format**
- 110592 The reference to appropriate privileges in the **cpio** format refers to an error on standard output;
110593 the **ustar** format does not make comparable statements.
- 110594 The model for this format was the historical System V *cpio-c* data interchange format. This
110595 model documents the portable version of the **cpio** format and not the binary version. It has the
110596 flexibility to transfer data of any type described within POSIX.1-202x, yet is extensible to transfer
110597 data types specific to extensions beyond POSIX.1-202x (for example, contiguous files). Because it
110598 describes existing practice, there is no question of maintaining upwards-compatibility.
- 110599 **cpio Header**
- 110600 There has been some concern that the size of the *c_ino* field of the header is too small to handle
110601 those systems that have very large *inode* numbers. However, the *c_ino* field in the header is used
110602 strictly as a hard-link resolution mechanism for archives. It is not necessarily the same value as
110603 the *inode* number of the file in the location from which that file is extracted.
- 110604 The name *c_magic* is based on historical usage.
- 110605 **cpio Filename**
- 110606 For most historical implementations of the *cpio* utility, {PATH_MAX} octets can be used to
110607 describe the pathname without the addition of any other header fields (the NUL character
110608 would be included in this count). {PATH_MAX} is the minimum value for pathname size,
110609 documented as 256 bytes. However, an implementation may use *c_namesize* to determine the
110610 exact length of the pathname. With the current description of the **<cpio.h>** header, this
110611 pathname size can be as large as a number that is described in six octal digits.
- 110612 Two values are documented under the *c_mode* field values to provide for extensibility for known
110613 file types:

110614 **0110 000** Reserved for contiguous files. The implementation may treat the rest of the
 110615 information for this archive like a regular file. If this file type is undefined, the
 110616 implementation may create the file as a regular file.

110617 This provides for extensibility of the **cpio** format while allowing for the ability to read old
 110618 archives. Files of an unknown type may be read as “regular files” on some implementations. On
 110619 a system that does not support extended file types, the *pax* utility should do the best it can with
 110620 the file and go on to the next.

110621 FUTURE DIRECTIONS

110622 None.

110623 SEE ALSO

110624 [Chapter 2](#) (on page 2457), *cp*, *ed*, *getopts*, *ls*, *printf*

110625 XBD [Section 3.145](#) (on page 52), [Chapter 5](#) (on page 113), [Chapter 8](#) (on page 167), [Section 12.2](#)
 110626 (on page 215), [<cpio.h>](#), [<tar.h>](#)

110627 XSH *chown()*, *creat()*, *fstatat()*, *futimens()*, *mkdir()*, *mkfifo()*, *write()*

110628 CHANGE HISTORY

110629 First released in Issue 4.

110630 Issue 5

110631 A note is added to the APPLICATION USAGE indicating that the **cpio** and **tar** formats can only
 110632 support files up to 8 gigabytes in size.

110633 Issue 6

110634 The *pax* utility is aligned with the IEEE P1003.2b draft standard:

- 110635 • Support has been added for symbolic links in the options and interchange formats.
- 110636 • A new format has been devised, based on extensions to **ustar**.
- 110637 • References to the “extended” **tar** and **cpio** formats derived from the POSIX.1-1990
 110638 standard have been changed to remove the “extended” adjective because this could cause
 110639 confusion with the extended **tar** header added in this version. (All references to **tar** are
 110640 actually to **ustar**.)

110641 The *TZ* entry is added to the ENVIRONMENT VARIABLES section.

110642 IEEE PASC Interpretation 1003.2 #168 is applied, clarifying that *mkdir()* and *mkfifo()* calls can
 110643 ignore an [EEXIST] error when extracting an archive.

110644 IEEE PASC Interpretation 1003.2 #180 is applied, clarifying how extracted files are created when
 110645 in **read** mode.

110646 IEEE PASC Interpretation 1003.2 #181 is applied, clarifying the description of the **-t** option.

110647 IEEE PASC Interpretation 1003.2 #195 is applied.

110648 IEEE PASC Interpretation 1003.2 #206 is applied, clarifying the handling of links for the **-H**, **-L**,
 110649 and **-I** options.

110650 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/35 is applied, adding the process ID of
 110651 the *pax* process into certain fields. This change provides a method for the implementation to
 110652 ensure that different instances of *pax* extracting a file named **/a/b/foo** will not collide when
 110653 processing the extended header information associated with **foo**.

110654 IEEE Std 1003.1-2001/Cor 1-2002, item XCU/TC1/D6/36 is applied, changing **-x B** to **-x pax** in
 110655 the OPTIONS section.

- 110656 IEEE Std 1003.1-2001/Cor 2-2004, item XCU/TC2/D6/20 is applied, updating the SYNOPSIS to
110657 be consistent with the normative text.
- 110658 IEEE Std 1003.1-2001/Cor 2-2004, item XCU/TC2/D6/21 is applied, updating the
110659 DESCRIPTION to describe the behavior when files to be linked are symbolic links and the
110660 system is not capable of making hard links to symbolic links.
- 110661 IEEE Std 1003.1-2001/Cor 2-2004, item XCU/TC2/D6/22 is applied, updating the OPTIONS
110662 section to describe the behavior for how multiple **-odelete=pattern** options are to be handled.
- 110663 IEEE Std 1003.1-2001/Cor 2-2004, item XCU/TC2/D6/23 is applied, updating the **write** option
110664 within the OPTIONS section.
- 110665 IEEE Std 1003.1-2001/Cor 2-2004, item XCU/TC2/D6/24 is applied, adding a paragraph into
110666 the OPTIONS section that states that specifying more than one of the mutually-exclusive options
110667 (**-H** and **-L**) is not considered an error and that the last option specified will determine the
110668 behavior of the utility.
- 110669 IEEE Std 1003.1-2001/Cor 2-2004, item XCU/TC2/D6/25 is applied, removing the *ctime*
110670 paragraph within the EXTENDED DESCRIPTION. There is a contradiction in the definition of
110671 the *ctime* keyword for the *pax* extended header, in that the *st_ctime* member of the **stat** structure
110672 does not refer to a file creation time. No field in the standard **stat** structure from **<sys/stat.h>**
110673 includes a file creation time.
- 110674 IEEE Std 1003.1-2001/Cor 2-2004, item XCU/TC2/D6/26 is applied, making it clear that *typeflag*
110675 1 (**ustar** Interchange Format) applies not only to files that are hard-linked, but also to files that
110676 are aliased via symbolic links.
- 110677 IEEE Std 1003.1-2001/Cor 2-2004, item XCU/TC2/D6/27 is applied, clarifying the *cpio c_nlink*
110678 field.
- 110679 **Issue 7**
- 110680 Austin Group Interpretations 1003.1-2001 #011, #036, #086, and #109 are applied.
- 110681 Austin Group Interpretation 1003.1-2001 #126 is applied, changing the description of the
110682 *LC_MESSAGES* environment variable.
- 110683 SD5-XCU-ERN-2 is applied, making **-c** and **-n** mutually-exclusive in the SYNOPSIS.
- 110684 SD5-XCU-ERN-3 is applied, revising the default behavior of **-H** and **-L**.
- 110685 SD5-XCU-ERN-5, SD5-XCU-ERN-6, SD5-XCU-ERN-7, SD5-XCU-ERN-60 are applied.
- 110686 SD5-XCU-ERN-97 is applied, updating the SYNOPSIS.
- 110687 The *pax* utility is no longer allowed to create separate identical symbolic links when extracting
110688 linked symbolic links from an archive.
- 110689 POSIX.1-2008, Technical Corrigendum 1, XCU/TC1-2008/0128 [260], XCU/TC1-2008/0129
110690 [261], XCU/TC1-2008/0130 [261], XCU/TC1-2008/0131 [313], and XCU/TC1-2008/0132 [233]
110691 are applied.
- 110692 POSIX.1-2008, Technical Corrigendum 2, XCU/TC2-2008/0152 [886], XCU/TC2-2008/0153
110693 [814], XCU/TC2-2008/0154 [886], and XCU/TC2-2008/0155 [707] are applied.
- 110694 **Issue 8**
- 110695 Austin Group Defect 1122 is applied, changing the description of *NLSPATH*. +
- 110696 Austin Group Defect 1133 is applied, clarifying the **-X** option and adding a paragraph to the +
110697 APPLICATION USAGE section. +

110698	Austin Group Defect 1270 is applied, removing the <code>-n</code> option from the copy mode SYNOPSIS	+
110699	line.	+
110700	Austin Group Defect 1278 is applied, removing mention of the <code>-n</code> option in connection with	+
110701	write mode.	+
110702	Austin Group Defect 1330 is applied, removing obsolescent interfaces.	+
110703	Austin Group Defect 1331 is applied, changing <code>``st_atime``</code> to <code>``st_atim``</code> and <code>``st_mtime``</code> to	+
110704	<code>``st_mtim``</code> .	+
110705	Austin Group Defect 1379 is applied, changing the ENVIRONMENT VARIABLES section.	+
110706	Austin Group Defect 1380 is applied, changing text using the term <code>``link``</code> in line with its	+
110707	updated definition and changing the description of the <code>-u</code> option.	+
110708	Austin Group Defect 1618 is applied, adding optional trailing <code>'s'</code> and <code>'S'</code> characters to the	+
110709	option-argument of the <code>-s</code> option.	+