NAME

pax – portable archive interchange

SYNOPSIS

pax [–**cdnv**] [–**f** *archive*] [–**s** *replstr*] . . . [*pattern* . . .]
pax –**r** [–**cdiknuv**] [–**f** *archive*] [–**o** *options*] . . . [–**p** *string*] . . . [–**s** *replstr*] . . .
[*pattern* . . .]
pax –**w** [–**dituvX**] [–**b** *blocksize*] [–**a**] [–**f** *archive*] [–**o** *options*] . . .
[–**s** *replstr*] . . . [–**x** *format*] [*file* . . .]
pax –**r** –**w** [–**diklntuvX**] [–**p** *string*] . . . [–**s** *replstr*] . . . [*file* . . .] *directory*

AVAILABILITY

SUNWcsu

DESCRIPTION

The **pax** command reads, writes and writes lists of the members of archive files and copy directory hierarchies. A variety of archive formats are supported; see the –**x** *format* option.

Modes of Operations

The action to be taken depends on the presence of the –**r** and –**w** options. The four combinations of –**r** and –**w** are referred to as the four modes of operation: list, read, write, and copy modes, corresponding respectively to the four forms shown in the **SYNOPSIS**.

list
In list mode (when neither –**r** nor –**w** are specified), **pax** writes the names of the members of the archive file read from the standard input, with path names matching the specified patterns, to standard output. If a named file is of type directory, the file hierarchy rooted at that file will be written out as well.

read
In read mode (when –**r** is specified, but –**w** is not), **pax** extracts the members of the archive file read from the standard input, with path names matching the specified patterns. If an extracted file is of type directory, the file hierarchy rooted at that file will be extracted as well. The extracted files is created relative to the current file hierarchy.

The ownership, access and modification times, and file mode of the restored files are discussed under the –**p** option.

write
In write mode (when –**w** is specified, but –**r** is not), **pax** writes the contents of the file operands to the standard output in an archive format. If no *file* operands are specified, a list of files to copy, one per line, will be read from the standard input. A file of type directory will include all of the files in the file hierarchy rooted at the file.

copy
In copy mode (when both –**r** and –**w** are specified), **pax** copies the file operands to the destination directory.

If no *file* operands are specified, a list of files to copy, one per line, will be read from the standard input. A file of type directory will include all of the files in the file hierarchy rooted at the file.

The effect of the copy is as if the copied files were written to an archive file and then subsequently extracted, except that there may be hard links between the original and the copied files. If the destination directory is a subdirectory of one of the files to be copied, the results are unspecified. It is an error if *directory* doesn't to exist, is not writable by the user, or is not a directory.

In read or copy modes, if intermediate directories are necessary to extract an archive member, **pax** will perform actions equivalent to the **mkdir**(2) function, called with the following arguments:

- the intermediate directory used as the *path* argument
- the octal value of **777** or **rwx** (read, write, and execute permissions) as the *mode* argument (see **chmod**(1)).

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

The supported archive formats are automatically detected on input. The default output archive format is **tar**(1).

If the selected archive format supports the specification of linked files, it is an error if these files cannot be linked when the archive is extracted. Any of the various names in the archive that represent a file can be used to select the file for extraction.

OPTIONS

The following options are supported:

–r Read an archive file from standard input.

–w Write files to the standard output in the specified archive format.

–a Append files to the end of the archive. This option will not work for some archive devices, such as 1/4-inch streaming tapes and 8mm tapes.

–b *blocksize* Block the output at a positive decimal integer number of bytes per write to the archive file. Devices and archive formats may impose restrictions on blocking. Blocking is automatically determined on input. Portable applications must not specify a *blocksize* value larger than **32256**. Default blocking when creating archives depends on the archive format. (See the **–x** option below.)

–c Match all file or archive members except those specified by the *pattern* or *file* operands.

–d Cause files of type directory being copied or archived or archive members of type directory being extracted to match only the file or archive member itself and not the file hierarchy rooted at the file.

–f *archive* Specify the path name of the input or output archive, overriding the default standard input (in list or read modes) or standard output (write mode).

–i Interactively rename files or archive members. For each archive member matching a *pattern* operand or file matching a *file* operand, a prompt will be written to the file **/dev/tty**. The prompt will contain the name of the file or archive member. A line will then be read from **/dev/tty**. If this line is blank, the file or archive member will be skipped.

If this line consists of a single period, the file or archive member will be processed with no modification to its name. Otherwise, its name will be replaced with the contents of the line. The **pax** command will 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.

–k Prevent the overwriting of existing files.

–l Link files. In copy mode, hard links will be made between the source and destination file hierarchies whenever possible.

–n Select the first archive member that matches each *pattern* operand. No more than one archive member will be matched for each pattern (although members of type directory will still match the file hierarchy rooted at that file).

–o *options* Reserved for special format-specific options.

–p *string* Specify one or more file characteristic options (privileges). The *string* option-argument must be a string specifying file characteristics to be retained or discarded on extraction. The string consists of the specification characters **a**, **e**, **m**, **o** and **p**. Multiple characteristics can be concatenated within the same string and multiple –**p** options can be specified. The meaning of the specification characters are as follows:

- **a** Do not preserve file access times.
- **e** Preserve the user ID, group ID, file mode bits, access time, and modification time.
- **m** Do not preserve file modification times.
- **o** Preserve the user ID and group ID.
- **p** Preserve the file mode bits. Other, implementation-dependent file-mode attributes may be preserved.

In the preceding list, ''preserve'' indicates that an attribute stored in the archive will be given to the extracted file, subject to the permissions of the invoking process; otherwise, the attribute will be determined as part of the normal file creation action.

If neither the **e** nor the **o** specification character is specified, or the user ID and group ID are not preserved for any reason, **pax** will not set the **setuid** and **setgid** bits of the file mode.

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

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

−s *replstr*
Modify file or archive member names named by *pattern* or *file* operands according to the substitution expression *replstr*, which is based on the **ed**(1) **s** (substitution) command, using the regular expression syntax on the **regex**(5) manual page. The concepts of ''address'' and ''line'' are meaningless in the context of the **pax** command, and must not be supplied. The format is:

−**s /** *old***/***new***/ [gp]**

where, as in **ed**, *old* is a basic regular expression and *new* can contain an ampersand (&) or a *n* backreference, where *n* is a digit. The *old* string also is permitted to contain newline characters.

Any non-null character can be used as a delimiter **(/** shown here). Multiple −**s** expressions can be specified; the expressions will be applied in the order specified, terminating with the first successful substitution. The optional trailing **g** is as defined in the **ed** command. The optional trailing **p** causes successful substitutions to be written to standard error. File or archive member names that substitute to the empty string are ignored when reading and writing archives.

−t
Cause the access times of the archived files to be the same as they were before being read by **pax**.

−u
Ignore files that are older (having a less recent file modification time) than a pre-existing file or archive member with the same name.

read mode
an archive member with the same name as a file in the file system will be extracted if the archive member is newer than the file.

write mode
an archive file member with the same name as a file in the file system will be superseded if the file is newer than the archive member.

copy mode
the file in the destination hierarchy will be replaced by the file in the source hierarchy or by a link to the file in the source hierarchy if the file in the source hierarchy is newer.

−v
In list mode, produce a verbose table of contents (see **Standard Output**). Otherwise, write archive member path names to standard error (see **Standard Error**).

−x *format*
Specify the output archive format. The **pax** command recognizes the following formats:

cpio
The extended **cpio** interchange format; see the IEEE 1003.1(1990) specifications. The default *blocksize* for this format for character special archive files is **5120**. Implementations support all *blocksize* values less than or equal to **32256** that are multiples of **512**.

ustar The extended **tar** interchange format; see the IEEE 1003.1(1990) specifications. The default *blocksize* for this format for character special archive files is **10240**. Implementations support all *blocksize* values less than or equal to **32256** that are multiples of **512**.

Any attempt to append to an archive file in a format different from the existing archive format will cause **pax** to exit immediately with a non-zero exit status.

−X When traversing the file hierarchy specified by a path name, **pax** will not descend into directories that have a different device ID (**st_dev**, see **stat**(2)).

The options that operate on the names of files or archive members (−**c**, −**i**, −**n**, −**s**, −**u** and −**v**) interact as follows. In read mode, the archive members are selected based on the user-specified *pattern* operands as modified by the −**c**, −**n** and −**u** options. Then, any −**s** and −**i** options will modify, in that order, the names of the selected files. The −**v** option will write names resulting from these modifications.

In write mode, the files are selected based on the user-specified path names as modified by the −**n** and −**u** options. Then, any −**s** and −**i** options will, in that order, modify the names of these selected files. The −**v** option will write names resulting from these modifications.

If both the −**u** and −**n** options are specified, **pax** does not consider a file selected unless it is newer than the file to which it is compared.

OPERANDS

The following operands are supported:

directory The destination directory path name for copy mode.

file A path name of a file to be copied or archived.

pattern A pattern matching one or more path names of archive members. A pattern must conform to the pattern matching notation found on the **fnmatch**(5) manual page. The default, if no *pattern* is specified, is to select all members in the archive.

OUTPUT

Standard Output

In write mode, if −**f** is not specified, the standard output will be the archive formatted according to **cpio** or **ustar**. (See −**x** *format*.)

In list mode, the table of contents of the selected archive members will be written to standard output using the following format:

"%s\n" **<***pathname***>**

If the −**v** option is specified in list mode, the table of contents of the selected archive members will be written to standard output using the following formats:

For path names representing hard links to previous members of the archive:

"%s∆==∆%s\n" <*ls –l listing*>, *linkname*

For all other path names:

<*pathname*> **"%s\n"** <*ls –l listing*>

where <*ls –l listing*> is the format specified by the **ls** command with the **–l** option. When writing path names in this format, it is unspecified what is written for fields for which the underlying archive format does not have the correct information, although the correct number of blank-character-separated fields will be written.

In list mode, standard output will not be buffered more than a line at a time.

Standard Error

If **–v** is specified in read, write or copy modes, **pax** will write the path names it processes to the standard error output using the following format:

These path names will be written as soon as processing is begun on the file or archive member, and will be flushed to standard error. The trailing newline character, which will not be buffered, will be written when the file has been read or written.

If the **–s** option is specified, and the replacement string has a trailing **p**, substitutions will be written to standard error in the following format:

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

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

ERRORS

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

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

USAGE

The **–p** (privileges) option was invented to reconcile differences between historical **tar**(1) and **cpio**(1) implementations. In particular, the two utilities use **–m** in diametrically opposed ways. The **–p** option also provides a consistent means of extending the ways in which future file attributes can be addressed, such as for enhanced security systems or high-performance files. Although it may seem complex, there are really two modes that will be most commonly used:

–p e ‘‘Preserve everything’’. This would be used by the historical superuser, someone with all the appropriate privileges, to preserve all aspects of the files as they are recorded in the archive. The **e** flag is the sum of **o**

and **p**, and other implementation-dependent attributes.

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

EXAMPLES

The following command:

example pax −w −f /dev/rmt/1m .

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

The following commands:

example% mkdir newdir example% pax −rw olddir newdir

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

example pax −r −s ',^//*usr//*,,' −f a.pax

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

ENVIRONMENT

See **environ**(5) for descriptions of the following environment variables that affect the execution of **pax**: **LC_CTYPE**, **LC_MESSAGES**, **LC_TIME**, and **NLSPATH**.

LC_COLLATE Determine the locale for the behaviour of ranges, equivalence classes, and multi-character collating elements used in the pattern matching expressions for the *pattern* operand, the basic regular expression for the −**s** option, and the extended regular expression defined for the **yesexpr** locale keyword in the **LC_MESSAGES** category.

EXIT STATUS

The following exit values are returned:

0 All files were processed successfully.

>0 An error occurred.

SEE ALSO

chmod(1), **cpio**(1), **ed**(1), **tar**(1), **mkdir**(2), **stat**(2), **environ**(5), **fnmatch**(5), **regex**(5)