--　--

IEEE P1003.2b Draft 2 − December 1991

`nolog`
: (Applicable to the `−x` *pax* format.) Suppress logging of ignored records and fields. See 4.48.7.2. Messages specifically described as error messages shall not be suppressed by this option.

`notrans` 2
: Suppress translation of character data or filenames, even 2
when the character set or charmap is known. This option 2
shall override the `charmap` option described previously. 2

`volid=`*id*
: (Applicable to the `−x` *pax* format.) Specify a volume identifier for the VOL1 label. See 4.48.7.5.1.2. If fewer than six octets are given for *id*, the value shall be padded at the end with `<space>`s.

⇒ **4.48.3 `pax` Options.** *Modify the description of the* `−x` *format option as follows: HLJ : The changes to this are pointers to the Extended Description instead of POSIX. 1 and the new* `pax` *format; these are diff-marked. The rest is the same as the 199x standard.*

`−x` *format* Specify the output archive format. The `pax` utility shall recognize the following formats:

`cpio` The extended `cpio` interchange format specified
in 4.48.8. The default *blocksize* for this format for 1
character special archive files shall be 5120. Implementations shall support all *blocksize* values less than or equal to 32 256 that are multiples of 512.

`pax` The interchange format specified in 4.48.7. The 1
default *blocksize* for this format for character spe- 1
cial archive files shall be 10 240. Implementa- 1
tions shall support all *blocksize* values less than 1
or equal to 32 256 that are multiples of 512. *HLJ :* 1
That was copied from cpio. Should it be 20 000 1
instead? What about multiples? 1

`ustar` The extended `tar` interchange format specified in 1
4.48.9. The default *blocksize* for this format for 1
character special archive files shall be 10 240. Implementations shall support all *blocksize* values less than or equal to 32 256 that are multiples of 512.

Implementation-defined formats shall specify a default block size as well as any other block sizes supported for character special archive files.

Any attempt to append to an archive file in a format different from the existing archive format shall cause `pax` to exit immediately with a nonzero exit status.

⇒ **4.48.6.1 `pax` Standard Output.** *Replace the first paragraph with:*

If the `−w` option is specified and neither the `−f` nor `−r` options are specified, the standard output shall be the archive formatted according to one of the specifications in the following subclauses or some other implementation-defined format. (See `−x` *format* under 4.48.3.)

⇒ **4.48.6.3 `pax` Output Files.** *Replace the second paragraph with:*

If the `−w` option is specified, but the `−r` option is not, the output file named by the `−f` option argument shall be a file formatted according to one of the specifications in the following subclauses or some other, implementation-defined, format.

⇒ **4.48.7 `pax` Extended Description.** *Replace this subclause with three subclauses, as follows. This will also cause the renumbering of 4.48.8 and 4.48.9 to be 4.48.10 and 4.48.11. HLJ : This was done for the editorial reason that we cannot go too deeply into the heading levels and if this were all in 4.48.7, we would have to start at the fourth level for each format, leaving little maneuvering room.*

4.48.7 Extended Description — **pax** Interchange Format 1

This format is based upon ISO 1001 {9}; a magnetic tape conforming to this standard also conforms to ISO 1001. The headers of these formats are defined (by ISO 1001) to use a subset of the invariant part of the characters and codings of ISO/IEC 646 {1}. There is an exception for certain names that would reasonably be in a different or larger character set.

4.48.7.1 Definitions

The following terms are used in this subclause (4.48.7).

4.48.7.1.1 a-character:

The 57 characters:

```
<space> ! " % & ′ ( ) * + , - . / : ; < = > ? _ 0-9 A-Z
```

as defined by ISO/IEC 646 {1}.

(The ISO 1001 definition is more elaborate, but equivalent.)

4.48.7.1.2 alternate character set:

A character set specified in the label records in which certain data fields may be encoded.

4.48.7.1.3 al-characters:

The a-characters (ISO 1001 {9}) plus `a-z`.

4.48.7.1.4 archive:

A file or medium recorded in the format specified in this subclause (4.48.7).

4.48.7.1.5 block:

When the medium is magnetic tape or another device that records information in discrete units that have a variable length: the meaning defined in ISO 1001 {9}.

> A group of octets recorded consecutively in accordance with the relevant standard for recorded magnetic tape. [ISO 1001.] 2

Otherwise, it is a quantity of data operated on as a single unit.

4.48.7.1.6 file section:

That part of a file that is recorded on any one volume. [ISO 1001.]

4.48.7.1.7 folded to a-characters:

A string of characters when converted to contain only a-characters, as defined in 4.48.7.2.

4.48.7.1.8 label:

A record that identifies and characterizes a volume, or a file section on a volume. [ISO 1001.]

4.48.7.1.9 log:

Write to `pax` standard error. 1

2

4.48.7.1.10 magnetic tape:

Media as defined by ISO 1863 {10}, ISO 3788 {11}, and ISO 5652 {12}. [ISO 1001.]

4.48.7.1.11 record (noun):

Related data treated as a unit of information. [ISO 1001.]

4.48.7.1.12 record (verb):

When discussing the content of an archive, to transfer data from the file hierarchy to the archive.

4.48.7.1.13 restore:

Transfer data from the archive to the file hierarchy.

4.48.7.1.14 tape mark:

A control block used as a delimiter. [ISO 1001.]

NOTE: The structure of tape marks is specified by the relevant standards for recorded magnetic tape.

4.48.7.1.15 volume:

When referring to magnetic tape, this shall have the same meaning as in ISO 1001 {9}:

> A dismountable reel of magnetic tape. [ISO 1001.]

When referring to other media, it shall refer to dismountable media. It shall refer to the content of a regular file when this format is recorded in a regular file.

The following notation is used in tables describing record formats:

- BP Octet position within the label 2
- L Length of the field in number of octet positions
- digit(s) Any digit from `0` through `9`.
- P When this field in a table is marked, it indicates that a more specific requirement than that of ISO 1001 {9} is made.

4.48.7.2 General Requirements

This subclause is an extension to ISO 1001 {9}; conflicts (as opposed to further specification) with ISO 1001 shall be resolved in favor of ISO 1001.

When restoring from an archive, the process invoking `pax` might not have the appropriate privileges to create files with all the characteristics described in the archive. In such a case, the file shall be created with as many of the characteristics that it had in the archive as the privileges of the process and the system implementation permit. In the case of a process that has no additional privileges, the protection information (ownership and access permissions) shall be set in the same fashion that the POSIX.1 {8} *open*() function would when given the *mode* argument matching the file permissions supplied by the mode field supplied in the File Information File (see 4.48.7.6). A process with permissions to do so shall be capable of restoring the ownership, permissions, and other file characteristics exactly as recorded on the medium.

The `pax` utility has various options that systematically change the names of files on the archive or as restored to the file hierarchy. For simplicity, the rest of this description is written as if such name changes were not made; `pax` shall operate in such a fashion that the relationship between files specified by links and symbolic links is retained in the presence of such changes, to the extent permitted by the organization of the file hierarchy.

It is implementation defined what protections are applied to an archive based upon protection fields recorded in it. Interpretation of the protection fields need not be applied until the information has been restored.

When a device that records data in discrete units of variable length is used, a *block* may be longer than the record or records it contains. The `pax` utility shall deal with the extra information in such a way that the extra octets do not participate in the data recorded on the archive.

When a device that does not support discrete units of variable length is used, a *block* shall be exactly the specified length.

When logging of information is required by this standard, the name of the file (as it appears in the FIF) shall be included in the log in such a way that it can be associated with the other logged information for that file.

Throughout this subclause, it is stated that certain fields, records, and labels ''shall be ignored.'' Such items (if present) may be recorded with any value consistent with ISO 1001 {9}. Nothing in this standard shall be construed to require that these fields shall be ignored in the presence of options to `pax` enabling interpretation of some or all of these fields. However, such fields shall be ignored when such options are not used. The existence and content of such items may be logged. Logging of ignored information can be suppressed by the `pax –o nolog` option.

A string shall be folded to a-characters to fit in a specific field in a label subject to the following constraints:

(1) The character set is converted to ISO/IEC 646 IRV {1}; characters not in ISO/IEC 646 {1} IRV shall be converted in an implementation-defined manner.

(2) The characters `a-z` are converted to `A-Z`.

(3) The string is truncated to the field width, or padded with <space>s on the right, as needed to make it exactly the width of the label field.

4.48.7.3 Media

The originator and recipient of the volume used for interchange shall agree on the media upon which this format is to be recorded. If that media is agreed to be magnetic tape, a tape mark shall be recorded and expected on the medium in the locations defined below. Where an analogous concept exists for other media, it should be used.

Where a tape mark or other out-of-band end-of-file marker is not defined for the agreed-to media it shall be omitted from the media when the interchange format is recorded.

4.48.7.4 File Representation

An instance of this archive format shall consist of one or more volumes. Each volume shall contain zero or more file sections. A file shall be recorded in one or more file sections consistent with the rules of ISO 1001 {9}.

Where tape marks are available, multiple File Sections and the appropriate End of File or Section and End of Volume Label Groups shall be used. An archive written on magnetic tape shall use tape marks.

On media where both a tape mark or equivalent is not used, and where records are not delimited in hardware or where this fact is hidden by the underlying system, only a single logical volume shall be written, which may span several physical volumes. Both the labels and the file data shall be combined as a single stream, and the labels shall appear immediately adjacent to each other. If the archive must span more than one physical volume, the last octet from the archive that is recorded on the first physical volume shall be the one preceding the first octet recorded on the second physical volume.

On media where a tape mark or equivalent is not used, but where records are delimited in hardware, only a single logical volume may be written, which may span several physical volumes. Each label shall appear as a separate physical record. File data shall be treated as a single stream. If the archive must span more than one physical volume, the last octet from the archive that is recorded on the first physical volume shall be the one preceding the first octet recorded on the second physical volume.

Media recorded in conformance with this standard shall have volume header and trailer labels as described below. Such media shall represent each POSIX file in either one or two ISO 1001 {9} format files. The first such ISO 1001 file shall consist of labeling information as described for the File Information File (FIF). The second section, the File Data File (FDF), is optional—when the FIF describes a file that does contain data, it shall not be recorded.

4.48.7.5 Volume Label Records

This subclause specifies how each volume label field is defined in the context of this standard.

4.48.7.5.1 The VOL 1 Header Label

BP	Field Name	L	P	Content	
1−3	Label Identifier	3		`VOL`	
4	Label Number	1		`1`	
5−10	Volume identifier	6		a-characters	
11	Volume accessibility	1		a-character	
12−24	(Reserved for future standardization)	13		`<space>`s	
25−37	Implementation Identifier	13	*	a-characters	3
38−51	Owner Identifier	14	*	a-characters	
52−79	(Reserved for future standardization)	28		`<space>`s	
80	Label Standard Version	1		`4`	

4.48.7.5.1.1 Fields Reserved for Future Standardization

These fields shall be reserved for future standardization for ISO 1001 {9} and its successors. The characters in these fields shall be `<space>`s.

4.48.7.5.1.2 Volume Identifier

This field shall specify an identification of the volume. When `pax` creates a 1
volume, the `−o volid` option can be used to set this field. If this field is not set 1
explicitly it shall contain spaces. When `pax` interprets a volume, this field shall 1
be logged. 1

4.48.7.5.1.3 Volume Accessibility

HLJ : Help from 1003.6 is requested in this area.

It is unspecified whether any mechanism is present to control access to files as recorded in this format based upon any identifying information or security information recorded with the file.

4.48.7.5.1.4 Implementation Identifier

This is a constant for the archive format, as opposed to implementation defined
for ISO 1001 {9}. It is implementation defined whether `pax` will attempt to read a 1
file or other media in ISO 1001 format that does not contain the values specified
for these fields.

4.48.7.5.1.5 Owner Identifier

This field shall contain a representation of the user name associated with the *uid* of the process creating the archive file, as found in the user database. The name shall be folded to a-characters.

4.48.7.5.2 The VOL2 Header Label

This header label is optional; if omitted, the fields described in the record shall all take their default values.

BP	Field Name	L	P	Content	
1−3	Label Identifier	3		VOL	
4	Label Number	1		2	
5−9	POSIX Identification	5	*	POSIX	3
10	POSIX Version	1	*	1	3
11−15	Standards Body	5	*	a-characters	3
16−20	Number	5	*	a-characters	3
21−25	Variant	5	*	a-characters	3
26−29	Revision	4	*	a-characters	3
30−80	Reserved for POSIX use	51	*	<space>s	3

NOTE: The reservation of BP 30−80 is an extension from ISO 1001 {9}; in that standard, labels VOL2 3
through VOL9 are reserved for implementation use. 2

4.48.7.5.2.1 Character Set Identification

The four fields Standards Body, Number, Variant, and Revision are used to identify the alternate character set used in data files, file names, and user and group names. The following are defined to refer to known standards. Additional names may be agreed on between the originator and the recipient. If a name is not recognized, `pax` may take implementation-defined actions, but there shall always 1
be a way for the data to be transferred to the receiving system, possibly without translation.

Body	Number	Variant	Revision	Formal Standard
<space>s	<space>s	<space>s	<space>s	ISO/IEC 646 IRV {1}
ISO	646	IRV	1990	Note 1
ISO	8859	1	1987	ISO 8859-1 {5}
ISO	8859	2	1987	ISO 8859-2 {6}
ISO/IEC	10646	Note 2	19xx	ISO/IEC 10646 {13}

NOTES:

(1) In this form of entry, when the value for variant is IRV, then ISO/IEC 646 IRV {1} is specified. Other values specify National Usage variants of this character set, and shall be agreed upon by sender and recipient.

(2) This value specifies the compaction method used.

HLJ : Working group input needed for the following. See my sample `−o charmap` 1
description. 1

The format-reading utility shall provide a means to control whether data files have character set translation applied. This may be done either on a per-file basis, or selecting the option for the whole archive. A utility that performs the same translation should be provided.

4.48.7.5.3 The VOL 3 Header Label 2

This label shall be used in those instances where a volume spans more than one 2
physical medium, and the normal end-of-section labels could not be written (due 2
to device/driver limitations, for example). 2

BP	Field Name	L	P	Content	
1–3	Label Identifier	3		`VOL`	2
4	Label Number	1		`3`	2
5–7	Media Continuation Identifier	3		`EOM`	2
8–11	Media Continuation Number	4		digits	2
12	FIF /FDF Identifier	1	*	(`H` or `D`)	2
13–29	File Identifier	17	*	4.48.7.6.1.1	2
30–33	File Section Number	4		4.48.7.6.1.3	2
34–38	POSIX Identification	5	*	`POSIX`	2
39	POSIX Version	1	*	`1`	2
40–45	Implementation Identification	6	*	`<space>`s	2
46–80	(Reserved for future standardization)	7		`<space>`s	2

4.48.7.5.3.1 Media Continuation Number 2

This field shall contain the count, starting from zero (0), of the physical media 2
used to write the volume. The first continuation tape shall be numbered 0001 2
(the initial tape being 0000). 2

4.48.7.5.4 The VOL 4–9 and UVL 1–9 Header Labels 2

These labels are free to be used for any implementation-defined purpose; they 1
shall be ignored on input. It is implementation-defined as to whether they are 1
included in output. 1

4.48.7.6 File Information File

This subclause specifies how the File Information File (FIF) is recorded; both its content and the header labels for it are defined. A FIF consists of ISO 1001 {9} header labels that define the content of the FIF ; the FIF itself contains the information necessary to recreate the POSIX file (except the data).

4.48.7.6.1 First FIF Header Label (HDR 1)

This record shall be recorded for all files in the archive.

BP	Field Name	L	P	Content	
1–3	Label Identifier	3		`HDR`	
4	Label Number	1		`1`	
5–21	File Identifier	17	*	a-characters	
22–27	File Set Identifier	6		a-characters	
28–31	File Section Number	4		digits	
32–35	File Sequence Number	4		digits	
36–39	Generation Number	4		digits	
40–41	Generation Version Number	2		digits	
42–47	Creation Date	6	*	See 4.48.7.6.1.7	1
48–53	Expiration Date	6	*	See 4.48.7.6.1.8	3
54	File Accessibility	1		a-characters	
55–60	Block Count	6		`000000`	1
61–73	Implementation Identification	13	*	<space>s	3
74–78	POSIX Identification	5	*	`POSIX`	3
79	POSIX Version	1	*	`1`	3
80	FIF Identifier	1	*	`H`	3

4.48.7.6.1.1 File Identifier

This shall be the name of the FIF file. It shall be derived from the name of the POSIX file being recorded by prefixing the string `INFO-` to the filename (not pathname) of the file being recorded, with the result folded to a-characters.

4.48.7.6.1.2 File Set Identifier

This field shall be ignored.

4.48.7.6.1.3 File Section Number

This field shall be as specified in ISO 1001 {9}.

4.48.7.6.1.4 File Sequence Number

This field shall be ignored.

4.48.7.6.1.5 Generation Number

This field shall be ignored.

4.48.7.6.1.6 Generation Version Number

This field shall be ignored.

4.48.7.6.1.7 Creation Date

This shall be the representation of the POSIX.1 {8} *st_ctime* field of the file, converted to a Julian day, according to the time zone active at the time `pax` was run. 1
The Julian date shall appear in the six octets as follows: 1

1 If the first two digits of the year are 19, this octet shall be a `<space>`; 1
otherwise, it shall be `0`. 1

2–3 These two octets shall be the two least significant digits of the year. 1

4–6 These three octets shall be the three-digit ordinal number of the day 1
from `001` to `366`. 1

4.48.7.6.1.8 Expiration Date

This shall be the same value as the creation date; see 4.48.7.6.1.7. 3

4.48.7.6.1.9 File Accessibility

HLJ : Help from 1003.6 is solicited.

The content of this field may be any permissible value; `pax` shall ignore this field. 1

4.48.7.6.1.10 Block Count

This field shall be as specified in ISO 1001 {9}.

4.48.7.6.2 Second File Header Label (HDR 2)

BP	Field Name	L	P	Content	
1–3	Label Identifier	3		`HDR`	
4	Label Number	1		`2`	
5	Record Format	1	*	`D`	
6–10	Block Length	5	*	`20000`	
11–15	Record Length	5	*	`09999`	
16–20	Standards Body	5	*	a-characters	
21–25	Number	5	*	a-characters	
26–30	Variant	5	*	a-characters	
31–34	Revision	4	*	a-characters	
35–50	(Reserved for Implementation)	16		not specified	1
51–52	Offset Length	2	*	`00`	2
53–80	(Reserved for Standard)	28		`<space>`s	1

4.48.7.6.2.1 Standards Body, Number, Variant, Revision

These fields record the same information, and are encoded in the same way, as the corresponding fields in the VOL2 Header Label. They specify the alternate character set to be used in the FIF file for those fields permitted to be in the alternate character set. If all these fields are spaces, instead of indicating ISO/IEC 646 {1}IRV, the alternate character set specified in the VOL2 label shall be used.

4.48.7.6.2.2 Record Format

Data is recorded as variable length records, blocked between 20 and 20K octets per block.

4.48.7.6.2.3 Block Length

This is the maximum length of a block, which is 20 000 octets.

4.48.7.6.2.4 Record Length

This is the maximum length of a record, which is limited to 9999 octets by the structure of variable-length records.

4.48.7.6.2.5 Offset Length

This field shall be zeroes. 2

4.48.7.6.3 Additional File Header Labels (HDR 3−9, UHL ?) 2

For a FIF, these labels are free to be used for any implementation purpose, and
shall be ignored on input. It is implementation-defined as to whether they are 1
written on output. On media that supports tape marks, the last header shall be
followed by a tape mark. On other media the tape mark shall be omitted.

4.48.7.6.4 FIF Contents

The FIF shall consist of several fixed record types of variable length, in the order 1
specified below. Each logical record shall be recorded in ISO/IEC 646 {1} IRV,
except that certain fields may be recorded in the alternate character set. Additional
optional record types are permitted in the format specified below. In the
record descriptions below, the content is described excluding the record length as
specified in ISO 1001 {9}.

Records present in a FIF, but which are not recognized by `pax`, shall be ignored. 1

4.48.7.6.4.1 FIF Filename Record

BP	Field Name	L	Content	
1	File Type	1	al-characters	
2−10	File Mode	9	al-characters	
11	Link Status	1	al-characters	
12−31	Modification time	20	digits (with possible decimal point)	2
32−51	Creation time	20	digits (with possible decimal point)	2
52−71	Access time	20	digits (with possible decimal point)	2
72−91	File Size	20	digits	1
92−	File Name		Alternate character set	1

File Type

This is the single character as follows:

- `-` Regular file
- `c` Character special file
- `b` Block special file

`d` Directory

`p` FIFO special file

`h` High-performance file

`l` Symbolic link

`A-Z` Reserved for implementations

other Reserved for future standardization.

The interpretation of the FIF and FDF for each of these file types shall be as specified below.

File Mode

The file protections shall consist of nine characters, presented as three groups of three characters each:

<*owner permissions* > Permissions for the file owner class.

<*group permissions* > Permissions for the file group class.

<*other permissions* > Permissions for the file other class.

Each field shall have three character positions:

(1) If `r`, the file is readable; if –, it is not readable.

(2) If `w`, the file is writable; if –, it is not writable.

(3) The first of the following that applies:

`S` If in <*owner permissions*> , the file is not executable and set-user- ID mode is set. If in <*group permissions*> , the file is not executable and set-group- ID mode is set.

`s` If in <*owner permissions*> , the file is executable and set-user- ID mode is set. If in <*group permissions*> , the file is executable and set-group- ID mode is set.

`T` If in <*other permissions*> , the file is executable and an unspecified concept of shared-executable is applied to the file. It is implementation defined whether such a mode is recorded. It is implementation defined what action, if any, beyond setting the the other-execute mode, is performed when `T` is found in this position in an archive.

`x` The file is executable or the directory is searchable.

`-` None of the attributes of `S`, `s`, `T`, or `x` applies.

Implementations may add other characters to this list for the third character position. Such additions shall, however, be recorded in lowercase if the file is executable or searchable, and in uppercase if it is not. It is implementation defined whether, when such an extension is found when restoring an archive, any action beyond setting the group-execute mode according to the case of the character occurs.

DT: This is a fairly conservative approach to all this; given that this is new work, should the suid, sgid, and sticky bits be expressed separately?

Link Status

This field indicates whether the file is a link or not. An ordinary file entry is indicated for the first (or only) entry in the archive for that file. Subsequent entries for additional links to that file shall indicate whether a link with or without data
is recorded. The `pax -o linkdata` controls whether data is recorded; by default, 1
this linked data is not recorded. 1

When a link, with or without data, is recorded, a record containing the *target* of the link, recorded in the alternate character set, shall immediately follow the FIF File Group record, and preceding any information specific to the file type. For full portability, the length of the target shall not exceed {_POSIX_PATH_MAX} octets.

If the target of the link was restored by `pax` during the same invocation, then a 1
link to the target shall be created and the FDF associated with the link entry shall be ignored. If the link cannot be created, an error shall be reported and the link ignored.

If a file was not restored during the same run the following actions shall be taken:

	File Type Does Not Require Data (Including Zero Length)	**File Type Requires Data**
Data present	see 4.48.7.7	restore
Data absent	restore log	ignore entry log error

A link with data shall be recorded with the File Size found at the time that specific link was recorded. If the length is found to be different from the length as recorded for a previous link, that fact shall be logged.

A link without data shall have the File Size as found at the time the specific link was recorded, but no FDF shall be recorded for that file.

The link type information is recorded as the single character as specified as follows:

`<space>` Ordinary file entry

`L` Link (without data)

`D` Link (with data)

other uppercase letters Reserved for implementations 2

other reserved for future standardization.

File Modification, Access, and Creation Times

These are the decimal representation of these times, as seconds since the Epoch. 2
If a period (.) decimal point character is present, the digits to the right of the point 2
shall represent the units of an implementation-defined subsecond timing granu- 2
larity, which shall be ignored by implementations not supporting these units. 2
The access and modification times shall be restored if the process has the
appropriate privilege required to do so.

File Size

For all file types except those noted below, this is the size of the file, in octets, that is recorded on the media as the FDF; if the size of the file on the originating system should change during recording of the file, the file on the archive shall be truncated or padding of octets containing zeroes added as required to assure that the exact number of octets specified is actually recorded.

Name

The rest of the record shall consist of the Name, which shall be recorded in the
alternate character set. No padding of Name shall occur, so that the length of
Name can be calculated from the record length. The pax utility shall translate 1
the Name in an implementation-defined manner to be the name of the file on the 1
receiving system. *HLJ : That needs work.* In no case shall a file name that is not 1
supported on the receiving system, or that cannot be subsequently accessed, be
created. The length shall not exceed {_POSIX_PATH_MAX} without agreement
between the originator and the recipient that this can occur.

4.48.7.6.4.2 FIF File Owner Record 1

This shall be the name of the owner of the file as found in the user database on the originating system, represented in the alternate character set. The (variable-length) record shall consist solely of the the file owner name, with no padding. A zero-length record shall be recorded if the implementation creating the format does not support ownership on this file type. If a zero length owner is encountered, it shall be treated as if the entry could not be found in the data base.

After an implementation-defined translation to the character set of the receiving
system, the File Owner name shall be used as the name of the owner of the file. If
the same name appears in the user database of the receiving system (without
truncation of either name), and if pax has the appropriate permission to do so, it 1
shall set the owner-ID of the file it is restoring to be the corresponding user-ID. If
the owner name does not appear in the user database of the receiving system, the
file shall be created with the ownership of the user running pax, and the name 1
that was not found shall be logged.

If the implementation does not support ownership on the file type being restored, the owner information may be ignored.

4.48.7.6.4.3 FIF File Group Record 1

This shall be treated in the same way as the file owner entry except that the group database shall be consulted, and the file group owner set where appropriate.

4.48.7.6.4.4 File-type Dependent Information

The next several records in the FIF vary depending on the File Type recorded in
the first record. The following subclauses describe these additional record types 1
and the interpretation of certain of the fields in the preceding records.

Regular Files

No File-type Dependent Information is recorded. The File Size is interpreted as
the number of octets found in the FDF, except that the FDF shall be omitted if the 2
file size is zero.

Block or Character Special Files

The file size in the first FIF record shall be zero, and no FDF shall be recorded.

BP	Field Name	L	Content
1–15	System Identifier	15	al-characters
16–	Device Specification	--	al-characters

System Identifier

When written by `pax`, this field shall contain an implementation-defined value 1
representing the name of the implementation, the model, and submodel numbers
sufficient to distinguish it from similar implementations that use a different
representation of the Device Specification field. If `pax` recognizes this field, it 1
shall restore the special file according to the information appearing in the Device
Specification. If `pax` does not recognize the System Identifier, the content of the 1
FIF, including the Device Specification and System Identifier, shall be logged and 1
`pax` may otherwise ignore the archive entry. 1

Device Specification

This field (of variable length) is unspecified by this standard, except that it shall contain information sufficient to create a file identical in function to the one written by `pax`, when restored on the same implementation on which it was written.

Directories

The information in a directory FIF shall be used to restore file ownership and per-
missions only after `pax` has completed reading the archive contents for that direc- 2
tory. The meaning of the file size field is implementation defined. If the implementation has no special definition for this field, it shall be ignored. No FDF shall be recorded for this entry type.

FIFO s

No additional records shall be written. The file size shall be ignored, and no FDF shall be recorded.

Symbolic Links

A single record consisting solely of the pathname of the target of the symbolic link, in the alternate character set, shall be recorded. For full portability, the length of the name shall not exceed {_POSIX_PATH_MAX} octets. The link shall be created whether or not the target of the link exists.

The ownership and protection information recorded for a symbolic link may be ignored.

If a symbolic link entry is encountered on a system that does not support symbolic links, it is implementation defined whether, and under what conditions, the link is translated to a hard link. This translation shall be logged giving both names. If the link is not restored, that fact shall be logged.

High Performance Files

HLJ : Help is solicited from P1003.4.

High-performance files shall be restored as regular files if the implementation does not support the file type. In such a case, the information in the records which follow shall be logged.

Zero or more of the following records shall be recorded to reflect the information required to represent the high-performance file's characteristics.

BP	Field Name	L	Content	
1–10	Flag	10	`HIGH-PERF:`	
11–	Parameter Specification	--	al-characters	1

Other File Types

For other, implementation-defined, file types, additional information records may be provided as specified below. If the File Size in the first FIF record is nonzero, and if the implementation does not recognize the file type, the subsequent FDF shall be restored as an ordinary file.

NOTE: This implies that the File Size field of the FDF must be accurate for all file types provided as an extension.

BP	Field Name	L	Content	
1–10	Flag	10	a-characters	
11–	Parameter Specification	--	unspecified	1

The Flag may be any combination of a-characters except those that are specified by this standard to appear as the first ten octets of a FIF record. The remaining characters shall either be al-characters or in the alternate character set. If the record type is not recognized by the implementation, it shall be logged;

translation from the alternate character set shall not occur when a record is logged.

4.48.7.6.5 Security Records 1

Following the last file-type specific record, additional records of the following format may be recorded.

HLJ : Help from 1003.6 is solicited.

BP	Field Name	L	Content	
1–10	Flag	10	`SECURITY:<space>`	2
11–	Security Specification	--	unspecified	1

On systems that do not implement security extensions, these records may be 2
ignored or logged.

4.48.7.6.6 Other FIF Information

For all file types, additional information records may be provided as specified below.

BP	Field Name	L	Content	
1–10	Flag	10	a-characters	
11–	Parameter Specification	--	unspecified	1

The Flag may be any combination of a-characters except those that are specified by this standard to appear as the first ten octets of a FIF record. The remaining characters shall either be al-characters or in the alternate character set. If the record type is not recognized by the implementation, it shall be logged; translation from the alternate character set shall not occur when a record is logged.

4.48.7.6.7 The FIF EOF 1 Label

The EOF 1 Record shall be identical to the HDR 1 record recorded for the FIF, except in the following fields.

BP	Field Name	L	P	Content
1–3	Label Identifier	3		`EOF`
55–60	Block Count	6		`<space>`, digits

4.48.7.6.7.1 Block Count

This field shall be as specified in ISO 1001 {9}. If the actual block count exceeds
the maximum value of a six-digit decimal number, it shall be recorded modulo 2
1 000 000. 2

4.48.7.6.8 The FIF EOF 2–9 and UTL 1–9 Labels

These records are optional, but if present shall be recorded in accordance with
ISO 1001. They shall be ignored on input, and it is implementation defined as to 1
whether they are written in output. 1

4.48.7.7 The File Data File

The file data file is optional, and when present shall contain the data contained in the file. If this file is omitted when it is expected, an empty file of the specified type shall be created, and an error logged. If a file of this type is found when it is not expected, the name from the HDR 1 record shall be used as the name of a file containing the data found in the FDF, in an implementation-defined directory. An error shall be logged.

4.48.7.7.1 The FDF HDR 1 Label

BP	Field Name	L	P	Content	
1–3	Label Identifier	3		HDR	
4	Label Number	1		1	
5–21	File Identifier	17	*	a-characters	
22–27	File Set Identifier	6		a-characters	
28–31	File Section Number	4		digits	
32–35	File Sequence Number	4		digits	
36–39	Generation Number	4		digits	
40–41	Generation Version Number	2		digits	
42–47	Creation Date	6	*	See 4.48.7.7.1.7	1
48–53	Expiration Date	6	*	See 4.48.7.7.1.8	3
54	File Accessibility	1		a-characters	
55–60	Block Count	6		000000	1
61–73	Implementation Identification	13	*	<space>s	3
74–78	POSIX Identification	5	*	POSIX	3
79	POSIX Version	1	*	1	3
80	FDF Identifier	1	*	D	3
74–80	(Reserved for future POSIX standardization)	7		<space>s	

4.48.7.7.1.1 File Identifier

This shall be the name of the FDF file. It shall be derived from the name of the POSIX file being recorded as the filename (not pathname) of the file being recorded, folded to a-characters.

4.48.7.7.1.2 File Set Identifier

This field shall be ignored.

4.48.7.7.1.3 File Section Number

This field shall be as specified in ISO 1001 {9}.

4.48.7.7.1.4 File Sequence Number

This field shall be ignored.

4.48.7.7.1.5 Generation Number

This field shall be ignored.

4.48.7.7.1.6 Generation Version Number

This field shall be ignored.

4.48.7.7.1.7 Creation Date

This shall be the representation of the POSIX.1 {8} *st_ctime* field of the file converted to a Julian day, according to the time zone active at the time `pax` was run. 1

4.48.7.7.1.8 Expiration Date

This shall be the same value as the creation date; see 4.48.7.7.1.7. 3

4.48.7.7.1.9 File Accessibility

HLJ : Help from 1003.6 is solicited.

The content of this field may be any permissible value; `pax` shall ignore this field. 1

4.48.7.7.1.10 Block Count

This field shall be as specified in ISO 1001 {9}.

4.48.7.7.2 Second FDF File Header Label (HDR 2)

BP	Field Name	L	P	Content	
1–3	Label Identifier	3		`HDR`	
4	Label Number	1		`2`	
5	Record Format	1	*	`F`	
6–10	Block Length	5	*	`16384`	1
11–15	Record Length	5	*	`00000`	
16–20	Standards Body	5	*	a-characters	
21–25	Number	5	*	a-characters	
26–30	Variant	5	*	a-characters	
31–34	Revision	4	*	a-characters	
35–50	(Reserved for Implementation)	16		not specified	1
51–52	Offset Length	2		`00`	2
53–80	(Reserved for Standard)	28		<space>s	2

4.48.7.7.2.1 Standards Body, Number, Variant, Revision

These fields record the same information, and are encoded in the same way, as the corresponding fields in the VOL2 Header Label. They specify the alternate character set to be used in the FDF file for those fields permitted to be in the alternate character set. If all these fields are spaces, instead of indicating ISO/IEC 646 {1} IRV, the alternate character set specified in the VOL2 label shall be used. If the Standards Body field is '`BINARY`'', the data recorded shall not be translated from the alternate character set.