

ECMA

EUROPEAN COMPUTER MANUFACTURERS ASSOCIATION

ECMA STANDARD
for
MAGNETIC TAPE LABELLING

November 1967

Free copies of this standard **ECMA-13** are available from
ECMA European Computer Manufacturers Association
114, Rue du Rhône - 1204 Geneva (Switzerland)

ECMA

EUROPEAN COMPUTER MANUFACTURERS ASSOCIATION

ECMA STANDARD
for
MAGNETIC TAPE LABELLING

November 1967

BRIEF HISTORY

ECMA Technical Committee TC 1 have produced in April 1965 their Standard ECMA-6 (revised and re-issued in May 1967) for a 7 bit Input/Output Character Code. The work of the Committee led then to the publication of other standards directed to the implementation of the first one on different media, in particular on magnetic tape (ECMA-5 and ECMA-12). Furthermore it was felt that a standard way of labelling the magnetic tapes was required. A Task Group was set up to produce this Standard and comprised representatives of following companies:

CII, Compagnie Internationale d'Informatique
Compagnie des Machines Bull
English Electric Computers Ltd
IBM-WTEC
ICT, International Computers and Tabulators Ltd
ITT Europe Inc.
NCR, The National Cash Register Company Ltd
Olivetti General Electric
Siemens Aktiengesellschaft
UNIVAC Computers (Europe) Ltd

In addition the Committee collaborated with following organizations:

International Organization for Standardization (ISO)
United States of America Standards Institute (USASI)

Adopted as Standard ECMA-13 by the General Assembly of Nov.28,1967

1. Scope

To establish a standard for information interchange utilizing magnetic tape, by providing magnetically recorded labels to identify and structure files, and by providing a standard structure for the blocks containing the records that constitute a file.

2. Definition of Terms

As used in this Standard, the following terms have the meanings indicated. When a term is in common use in a context related to the subject-matter of this Standard, the definition conforms to common usage; when a term has been used in related standards, the definition conforms to its usage there.

Record. A record is a collection of related items of data, which for operating systems logic purposes is treated as a unit of information. Conceptually, a record corresponds (in the context of business data) to a transaction, a customer's account, etc. In other contexts, the delineation of a record may be relatively arbitrary, and be determined by the designer of the information formats.

Block. A block is a group of contiguous characters recorded on and read from magnetic tape as a unit. A block may contain one or more complete records.

File. A file is a major collection of information, consisting of all the records pertaining to a general subject. Conceptually, this term (in the context of business data) relates to such collections as a Payroll File, an Inventory File, etc. In other contexts, the delineation of a file may be relatively arbitrary.

File Set. A file set is a collection of one or more related files, recorded on one or more volumes. A file set may consist of:

One file, recorded on a single volume.

More than one file, recorded on a single volume.

One file, recorded on more than one volume.

More than one file, recorded on more than one volume.

Volume. A volume is a physical unit of storage media. The word volume, as used in this Standard, is completely synonymous with 'reel of magnetic tape'.

Label. A label is a block at the beginning or end of a volume or a file, and serves to identify and/or delimit that volume or file.

Label Group. A label group is a collection of contiguous labels of the same type (see 3.1.3 for label types).

Tape Mark. A tape mark is a special configuration recorded on magnetic tape, essentially indicating the boundary between files and labels, and also between certain label groups. The tape mark configuration is defined in the relevant Recorded Magnetic Tape Standards.

Operating System. The operating system is a program or set of programs, usually provided by the manufacturer, which, among other things, handles the functions of reading from and writing blocks on to tape, label handling and related functions.

3. Labelling

3.1 Elements of the Labelling Scheme.

3.1.1 Each label shall be an 80-character block.

3.1.2 Labels are divided into two general classes: Operating System Labels and User Labels. In addition, labels are classified as required or optional.

3.1.3 There are two types of labels: Volume Labels and File Labels.

3.1.4 The identifiers and numbers (4 characters) for the various classes and types of labels are shown in the following chart. The labels are also classified to show which labels are required by this Standard and which additional labels are permissible (optional).

	Class		
Type	Operating System		User
Volume	VOL 1 EOV 1	(None) EOV 2 - 9	UVL 1 - 9 UTLx
File	HDR 1 EOF 1	HDR 2 - 9 EOF 2 - 9	UHLx UTLx
	Required	Optional	

x = any 'a' character
as defined in 3.3

3.1.5 Tape Marks shall be used only where specified in this Standard.

3.2 Structuring the File.

3.2.1 Required Labels and Tape Marks shall be used to establish the file structure according to the following rules, as illustrated in Figure 1. In that figure, the beginning of the tape is at the left, and the end of the tape is at the right. Required Labels are indicated by the three characters of their identifiers, and Tape Marks are indicated by asterisks (*).

The rules and the figure are presented as though there were no Optional Operating System Labels or User Labels. Rules for using these optional classes of labels are set forth in 3.2.3 to 3.2.8.

Figure 1: Structure of Magnetic Tape Files
(* means Tape Mark)

```
Single-Volume File  
VOL HDR* --- Data Blocks --- * EOF **  
Multi-Volume File  
VOL HDR * --- First Volume Data --- * EOF **  
VOL HDR * --- Last Volume Data --- * EOF **  
Multi-File Volume  
VOL HDR * -- File A -- * EOF * HDR * --- File B --- * EOF **  
Multi-Volume Multi-File  
VOL HDR * -- File A -- * EOF * HDR * -- File B --- * EOF **  
VOL HDR * --- Continuation of File B ----- * EOF **  
VOL HDR * -- End of File B -- * EOF * HDR * -- File C -- * EOF **
```

Volume Header Label. Every volume shall have a Volume Header Label as the first block in the volume. The Volume Header Label shall not be used at any other place in the volume.

File Header Label. Every file shall be preceded by a File Header Label. Whenever a volume ends within a file, the continuation of that file in the next volume shall also be preceded by a File Header Label. The File Header Label shall be followed by a Tape Mark (see also 3.2.6).

End of File Label. The last block of every file shall be followed by an End of File Label. A Tape Mark shall immediately precede, and another Tape Mark shall follow, the End of File Label. The End of File Label that appears at the end of the last (or only) file in a set shall be followed by two Tape Marks, rather than one (see also 3.2.6).


End of Volume Label. Whenever a volume ends within a file, the last block of the file in that volume shall be followed by an End of Volume Label. One Tape Mark shall immediately precede, and two Tape Marks shall follow, the End of Volume Label (see also 3.2.6). File sets shall not be terminated by an End of Volume Label.

3.2.2 Whenever end-of-volume and end-of-file coincide, the labelling configuration shall be one of the following:

Figure 2



Figure 3



3.2.2.1 By agreement between the interchange parties it is allowable to arrange that any file of a set may start at the beginning of a volume. If this procedure is used the labelling configuration shall be as shown in the Appendix A.2.6.

3.2.3 Optional Operating System Labels and User Labels shall be fitted into the file structure as described in 3.2.4 to 3.2.8, without otherwise modifying the relationship between the Required Labels and the files. When these optional classes of labels are used, any reference to a label within 3.2 shall be construed to mean the entire label group.

3.2.4 Optional Operating System Labels of a given type, when used, shall directly follow a Required Label of the same type. The fourth character of consecutive Optional Operating System Labels in a group shall be the numbers 2, 3, 4, 5, 6, 7, 8, and 9 respectively.

3.2.5 User Labels of a given type, when used, shall directly follow a consecutive group of Operating System Labels of the same type. When no Optional Operating System Labels are used in the label group, the User Labels shall directly follow a Required Label of the same type.

- 3.2.6 There shall be no Tape Mark within a group of labels. Wherever Figures 1, 2, 3, and the descriptions in 3.2.1 and 3.2.2 indicate a Tape Mark following a Required Label, that Tape Mark shall actually follow the last label of the entire group.
- 3.2.7 Every label group shall be completed on the volume where the first label of the group was recorded.
- 3.2.8 Examples of the grouping of optional labels.
- 3.2.8.1 Physical beginning of tape to physical end of tape (not end of file)
- ```
VOL1 UVL1 ... UVLn HDR1 HDR2 ... HDRn UHL1 ...
UHLa* File Data* EOV1 ... EOVn UTL1 ... UTLa**
```
- 3.2.8.2 Physical beginning of tape to end of file
- ```
VOL1 UVL1 ...UVLn HDR1 HDR2 ...HDRn UHL1 ...  
UHLa* File Data* EOF1 ...EOFn UTL1 ... UTLa*
```
- 3.2.8.3 Physical beginning of tape to end of file set
- ```
VOL1 UVL1 ... UVLn HDR1 HDR2 ...HDRn UHL1 ...
UHLa* File Data* EOF1 ...EOFn UTL1 ... UTLa**
```
- 3.2.8.4 Beginning of new file (not beginning of tape) to physical end of tape (not end of file)
- ```
*HDR1 HDR2 ... HDRn UHL1 ... UHLa*  
File Data * EOV1 ... EOVn UTL1 ... UTLa**
```
- 3.2.8.5 Beginning of new file (not beginning of tape) to end of file
- ```
HDR1 HDR2 ... HDRn UHL1 ... UHLa
File Data* EOF1 ... EOFn UTL1 ... UTLa*
```
- 3.2.8.6 Beginning of new file (not beginning of tape) to end of file set
- ```
*HDR1 HDR2 ... HDRn UHL1 ... UHLa*  
File Data* EOF1 ... EOFn UTL1 ... UTLa**
```

3.3 Formats and Contents of Labels

In this Standard 'n' means any numeric digit from 0 to 9. An 'a' means any of:

The characters occupying the central four columns of the ECMA 7 bit code table (2/0 through 5/15) except position 5/15 and those positions where there is provision for alternative graphic representation.

The limitation on 'a' characters is intended as a guide to provide maximum interchangeability and consistent printing especially when international interchange is a possibility. Checking for conformity to this limitation is not implied.

The word 'optional' is sometimes used in the name or description of fields in this Standard. When used, 'optional' means that the field may, but need not contain the information described. If an optional

field does not contain the designated information, it shall contain spaces (or zeroes if so indicated).

Fields which are not described as 'optional' are considered to be 'mandatory'. 'Mandatory' fields must be written as specified.

Although this Standard does not require any particular label processing on reading, certain desirable processing is implied.

3.3.1 Volume Header Label

<u>Field</u>	<u>Name</u>	<u>Length</u>	<u>Description</u>
1	Label Identifier	3	Must be VOL
2	Label Number	1	Must be 1
3	Volume Serial Number	6	Six 'a' characters permanently assigned by the owner to identify this physical volume (i.e., reel of tape).
4	Accessibility	1	An 'a' character which indicates any restrictions on who may have access to the information in this volume. A 'space' means unlimited access; any other character means special handling, in a manner agreed on between the interchange parties.
5	Reserved for operating systems	20	Reserved for operating systems use. Any 'a' characters.
6	Reserved for future standardization	6	Must be 'spaces'.
7	Owner Identification	14	Any 'a' characters, identifying the owner of the physical volume.
8	Reserved for future standardization	28	Must be 'spaces'.
9	Label Standard Level	1	'1' means that the labels and data formats on this volume conform to the requirements of this Standard. 'Space' means that the labels and data formats on this volume require the agreement of the interchange parties.

3.3.1.1 Anyone recording on a magnetic tape that he does not own must preserve the information content of the entire Volume Header Label, except as authorized by the owner of the tape. This is not intended to preclude the rewriting of the Volume Header Label.

3.3.2 First File Header Label

<u>Field</u>	<u>Name</u>	<u>Length</u>	<u>Description</u>
1	Label Identifier	3	Must be HDR
2	Label Number	1	Must be 1
3	File Identifier	17	Any 'a' characters agreed on between the interchange parties.
4	Set Identification	6	Any 'a' characters to identify the set of files of which this is one. This identification must be the same for all files of a multifile set.
5	File Section Number	4	The file section number of the first Header Label of each file is 0001. This applies to the first or only file on a volume and to subsequent files on a multi-file volume. This field is incremented by 1 on each subsequent volume of the file.
6	File Sequence Number	4	Four 'n' characters denoting the sequence (i.e., 0001, 0002, etc.) of files within the volume or set of volumes. In all the labels for a given file, this field will contain the same number.
7	Generation Number (optional)	4	Four 'n' characters denoting the current stage in the succession of one file-generation by the next. When a file is first created its generation number is 0001.
8	Generation Version Number (optional)	2	Two 'n' characters distinguishing successive iterations of the same generation. The generation version number of the first attempt to produce this file is 00.

<u>Field</u>	<u>Name</u>	<u>Length</u>	<u>Description</u>
9	Creation Date	6	A 'space' followed by two 'n' characters for the year, followed by three 'n' characters for the day (001 to 366) within the year.
10	Expiration Date	6	Same format as field 9. This file is regarded as 'expired' when today's date is equal to, or later than the date given in this field. When this condition is satisfied, the remainder of this volume may be overwritten. To be effective on multi-file volumes therefore, the expiration date of a file must be less than, or equal to the expiration date of all previous files on the volume.
11	Accessibility	1	An 'a' character which indicates any restrictions on who may have access to the information in this file. A 'space' means unlimited access; any other character means special handling, in a manner agreed on between the interchange parties.
12	Block Count	6	Must be '000000'.
13	System Code (optional)	13	Thirteen 'a' characters identifying the operating system that recorded this file.
14	Reserved for future standardization	7	Must be 'spaces'.

3.3.3 Second File Header Label (optional).

(Further explanations concerning the fields in this label are given in section 4).

<u>Field</u>	<u>Name</u>	<u>Length</u>	<u>Description</u>
1	Label Identifier	3	Must be HDR
2	Label Number	1	Must be 2

<u>Field</u>	<u>Name</u>	<u>Length</u>	<u>Description</u>
3	Record Format	1	F = Fixed Length. D = Variable, with the number of characters in the record specified in decimal. V = Variable, with the number of characters in the record specified in binary. U = Undefined.
4	Block Length	5	Five 'n' characters specifying the maximum number of characters per block.
5	Record Length	5	Five 'n' characters specifying: - Record Length, if Record Format (field 3) is F. - Maximum Record Length including any count fields, if Record Format (field 3) is D or V. - Undefined, if Record Format (field 3) is U.
6	Reserved for operating systems	35	Reserved for operating systems use. Any 'a' characters.
7	Buffer Offset (optional)	2	Two 'n' characters specifying the length in characters of any additional field inserted before a data block. This length is included in the Block Length (field 4).
8	Reserved for future standardization	28	Must be 'spaces'.

3.3.4 First End of File Label.

<u>Field</u>	<u>Name</u>	<u>Length</u>	<u>Description</u>
1	Label Identifier	3	Must be EOF.
2	Label Number	1	Must be 1.
3 to 11	Same as corresponding fields in the first File Header Label (all optional)	Total 50	Same as corresponding fields in the first File Header Label.

<u>Field</u>	<u>Name</u>	<u>Length</u>	<u>Description</u>
12	Block Count	6	Six 'n' characters denoting the number of data blocks (exclusive of labels and Tape Marks) since the preceding HDR Label group.
13 and 14	Same as corresponding fields in the first File Header Label (optional)	Total 20	Same as corresponding fields in the first File Header Label.

3.3.5 Second End of File Label (optional).

A second End of File Label having EOF2 as label identifier and label number and containing the same information in fields 3 to 8 (all optional) as in HDR2.

3.3.6 First End of Volume Label.

<u>Field</u>	<u>Name</u>	<u>Length</u>	<u>Description</u>
1	Label Identifier	3	Must be EOv
2	Label Number	1	Must be 1
3 to 11	Same as corresponding fields in the first File Header Label (all optional)	Total 50	Same as corresponding fields in the first File Header Label.
12	Block Count	6	Six 'n' characters denoting the number of data blocks (exclusive of labels and Tape Marks) since the preceding HDR Label group.
13 and 14	Same as corresponding fields in the first File Header Label (optional)	Total 20	Same as corresponding fields in the first File Header Label.

3.3.7 Other Optional Operating System Labels

<u>Field</u>	<u>Name</u>	<u>Length</u>	<u>Description</u>
1	Label Identifier	3	Must be the appropriate 3-letter code, according to the type of Operating System Label.
2	Label Number	1	Must be numbered consecutively up to 9 within a group of Operating System Labels.
3	Operating System Option	76	Any 'a' characters.

3.3.8 User Labels (optional)

<u>Field</u>	<u>Name</u>	<u>Length</u>	<u>Description</u>
1	Label Identifier	3	Must be the appropriate 3-letter code, according to the type of User Label.
2	Label Number	1	Must be numbered 1 to 9 consecutively within the User Volume Labels. In other User Labels, any 'a' character.
3	User Option	76	Any 'a' characters.

4. Block Structure

4.1 Grouping Records into Blocks.

No explicit indication of the boundaries between records is required. There must be an integral number of records in a block. Any use of padding requires the agreement of the interchange parties. (Refer to A.5 in the Appendix). Truncated blocks are permitted.

4.1.1 Fixed-Length Records. When all the records in a file are of the same length, no indication of that uniform length is required within a file.

4.1.2 Variable-Length Records. When the records in a file are not all of the same length, the length of each record (i.e., the number of characters it contains) shall be recorded as the first field in each record. That field shall be counted as part of the record length. The record length shall be expressed as a decimal number, occupying the first four positions of each record.

4.1.2.1 Alternatively, by agreement between the interchange parties, the record length may be expressed as a binary number, occupying the first two positions of each record length field. In this case, the remaining two characters may be used by the operating system.

4.1.3 Undefined Records. When records do not meet the definitions of 4.1.1, 4.1.2 or 4.1.2.1 they are undefined in format. In general the interchange of information in the undefined format will require the prior agreement of the interchange parties.

4.1.4 By-pass or check-point records. Only relevant data blocks shall be written on a tape used for interchange. Since by-pass information or checkpoint records are considered to be extraneous to the interchange, no standard means of identification is provided.

4.2 Block Length

4.2.1 For general information interchange, the maximum block length shall be 2048 characters. However, with the agreement of the interchange parties, larger blocks may be used.

4.2.2 The minimum block length shall be as specified in the applicable Recorded Magnetic Tape Standards.

4.2.3 Block Sequence Indicator. This optional checking facility, the full details of which are given in the Appendix, may only be used with the agreement of the interchange parties. Block Length will not include the BSI if the latter is used.

4.3 Recording density. The blocks recorded on all volumes containing a file set shall be recorded at the same density, except with the prior agreement of the interchange parties.

APPENDIX

(This Appendix is not part of the Standard, it is included to facilitate its use).

A.1 Classes of Labels (3.1.2)

- A.1.1 This Standard defines two classes of labels. The definitions of this Standard distinguish between Operating System Labels and User Labels and, in addition, denote each as required or optional. In one sense, all labels are Operating System Labels, since operating systems normally provide for their reading and writing. The processing of Operating System Labels, however, is normally provided as a function of the operating system, whereas the contents of User Labels are prescribed by the user and their processing is provided by a user routine. The following paragraphs provide additional clarity to the definitions used in this Standard.
- A.1.2 Optional Operating System Labels. These labels, when used, can be expected to contain information pertinent to labelling functions more specialized than those contemplated in this Standard. Such functions may be developed, defined, and implemented differently by the designers of different operating systems. Thus the nature of these labels makes them fundamentally incompatible between different systems, and these labels will be ignored in interchange.
- A.1.3 User Labels (UVL, UHL, UTL). It will often be convenient to use these labels to contain summary information about a file being interchanged, such as control totals, statistical tabulations, and the like. In such a case, that information will be quite useful to the recipient of the file, so that the User Labels would then become part of the information being interchanged, with their contents agreed on between the parties. In the absence of such an agreement, User Labels will be ignored in interchange.

A.2 Structuring the File (3.2)

- A.2.1 The Tape Mark is used in the following way:
- A.2.1.1 Following file information, the Tape Mark indicates the boundary between file blocks and labels.
 - A.2.1.2 Following a label group, the Tape Mark indicates the end of that label group.
- A.2.2 File Header Labels in a Multi-Volume File Set.
- A.2.2.1 The Set Identification (field 4) is identical in all HDR1 Labels throughout the entire file set.

A.2.2.2 The HDR1 Label at the beginning of every volume except the first is an exact copy of the last HDR1 Label on the previous volume, except that the File Section Number (field 5) is augmented by 1. (Except as provided in Section A.2.6).

A.2.3 The double Tape Mark at the end of each reel (see Figure 1) permits the following procedure to be used in accomplishing the operation 'Forward-space File'.

Having read the HDR Label,

- Ⓐ - Index forward till three Tape Marks passed.
Read the next block.
If HDR Label: One file has been indexed.
If Tape Mark: End of Volume, rewind.
If not end of set: Alternate, verify HDR Label on next volume, return to Ⓐ.

An alternative procedure for that operation would be:

Having read the HDR Label,

- Ⓐ - Index forward till two Tape Marks passed.
Read the next block.
If EO V Label: Rewind, alternate, verify HDR Label on next volume, return to Ⓐ.
If EOF Label: Index forward till one Tape Mark passed.
Read next block.
If HDR Label: One file has been indexed.
If Tape Mark: End of set has been reached.

Thus, while the Standard requires that two Tape Marks be recorded at the end of each volume, it does not require that the second one be read (except at the end of the file set, if there is no other indication which file is the last in the set).

A.2.4 Two consecutive Tape Marks appear at the beginning (after VOL and HDR labels) of the second volume in Figure 2, and at the end of the first volume in Figure 3; yet they are not interpreted as double Tape Marks, but rather as framing a 'null' section of file. Conventional processing can proceed as follows:

- Read the HDR Label.
- Pass over any optional labels.
- Read Tape Mark - switch to processing of file information.
- Read second Tape Mark - switch to processing of labels.
- In Figure 2, read EOF Label of File A.
- In Figure 3, read EO V Label occurring within File B.

A.2.5 In general, the pattern shown in Figure 2 will occur when the end-of-tape warning mark is reached while writing the last block

of a file. Usually, the operating system will not yet know that this is the last, and the EOF Label group will be recorded at this time.

On the other hand, the pattern shown in Figure 3 will occur when the warning mark is reached after the EOF Label group has been started.

A.2.6 When it has been agreed between the interchange parties to start a file of a set at the beginning of a volume, the labelling configuration to be used is as follows:

```
----- File A ---- * EOF * EOF **
                        (A)
VOL HDR * ----- File B -----
                        (B)
```

This configuration permits the initiation of any file within a set at the beginning of a volume. When this configuration is used, the Block Count (field 12) of the EOF label is undefined.

File sets shall not be terminated by an EOF label group.

A.3 Use of the Fields in the Labels

A.3.1 Volume Header Label (3.3.1)

The Volume Header Label identifies the physical reel of magnetic tape, and the contents of that label relate to the identity of the volume. Since there is no Optional Operating System Volume Header Label, a User Volume Header Label will (if present) always be the second block on the tape.

A.3.1.1 Accessibility (field 4) is expected to refer to such categories of information as company confidential, proprietary, etc. This field is not intended to fulfill the requirements of national security (which will probably be accommodated in a government-specified User Label), but this field might be used as an indicator in conjunction with such a User Label.

An Accessibility field appears in both the Volume Header and File Header Labels, so that this function can be exercised either for the entire volume, or for each individual file, as desired.

A.3.1.2 Label Standard Level. This field is used to indicate whether or not the information recorded on this volume conforms to this Standard. It may also indicate a deviation from the Standard for a particular interchange situation.

In some cases operating systems may need to distinguish between this Standard labelling system and similar related systems used parochially. This field is intended to facilitate this distinction. It also provides a means for extending this Standard in the future, should the need arise, with minimum conflict between the future standard and parochial practice that may develop in the meantime. It is intended to distinguish among future standard levels by the use of numerals in this field, rather than letters, to the extent possible.

- A.3.1.3 Owner Identification (field 8). It is likely that, in time, a standard method of identifying the owner will be defined.

In the absence of such a standard, the parties should agree among themselves to choose identifiers so that each party will be identified uniquely within the specific interchange environment.

A.3.2 First File Header Label (3.3.2)

- A.3.2.1 Set Identification (field 4). It is desirable that a unique identification be established. In most cases, this objective may be satisfied by duplicating field 3 (Volume Serial Number) of the VOL1 Label of the first or only volume of the set.

- A.3.2.2 File Section Number (field 5). There are occasions when it is desired to read selectively one or more of the files within a multi-volume multi-file set. Through error, an attempt might be made to begin by reading an incorrect volume: one that happens to start in the middle of the file. The actual beginning of a file may be identified by '0001' in this field.

- A.3.2.3 Generation Version Number (field 8) denotes the successive attempts to produce a new generation of a file or part of a file. That is to say, each time it is necessary to reproduce data the Generation Version Number is increased by 1. The Generation Version Number of the first attempt to produce this file is 00.

The Generation Version Number is used to differentiate output data which has been produced by repeated processing operations and which in all other respects would bear the same identification. It is used to distinguish between a partial file recorded during an aborted run, and the new copy of the same information recorded after return to a rescue point.

For interchange purposes this field has no significance to the recipient.

A.3.2.4 Expiration Date (field 10). If - when recording a multifile volume - any file is given a later Expiration Date than the files preceding it, that further protection will not be achieved. Because of the way in which operating systems are designed, the Expiration Date of the first file in a volume is the latest date on which any part of the volume will be protected.

A.3.2.5 Accessibility (field 11). See A.3.1.1

A.3.2.6 Block Count (field 12). This field is provided in order that the installation reading a magnetic tape may assure itself that no blocks were skipped and no spurious blocks were inserted. The Block Count provides for this verification each time an EOY Label or EOF Label is reached in reading a tape.

A.3.3 Second File Header Label (3.3.3)

A.3.3.1 Record Format, Block Length, Record Length (fields 3-5). Provide the operating system with information about the format of the data records recorded on the tape. For those cases where dynamic blocking or deblocking is provided by the operating system, these fields provide the necessary information to perform these functions. This information could be used to allocate the areas to perform the blocking and deblocking.

A.3.3.2 Reserved for Operating Systems (field 6). Could contain any information needed to increase the efficiency of record processing by the operating system.

A.3.3.3 Buffer Offset (field 7). Certain applications will require additional information at the front of each data block. This could include block length, the block address of the last record in the block, initial padding for word machines, dates, times of transmission etc. The length of the added information may be specified in this field.

A.4 Block Sequence Indicator (optional)

A.4.1 Malfunction of a tape handling device or a fault on magnetic tape may occasionally result in missing a complete block of information or, in essence, cause a single block to be 'read' twice. Such failures would be detected at the end of a volume or file by comparison of the Block Count. On detection, the file, file section or volume may be reprocessed.

As a more immediate check, on a block-by block-basis, an optional Block Sequence Indicator is provided for. This indicator will provide:

- a) Immediate identification of any error involving added or missed blocks and
- b) Facility to take appropriate corrective action at once in the operating system rather than through the initiation of a re-run at the end of a file, file section or volume.

A.4.2 The Block Sequence Indicator is optional and its use requires the agreement of the interchange parties. There is no implication that operating systems should make provision to recognize this indicator solely for conformance with this feature.

If the Block Sequence Indicator is not used, there will be no additional character or character position on the front of the Volume Header Label or any of the other blocks in the volume. The Block Length shall not include the Block Sequence Indicator if the latter is used.

A.4.3 The Block Sequence Indicator is placed on the front of each block on tape as an additional character, except for Tape Mark blocks, where it is not so added but is counted as if it were added.

A.4.4 The Block Sequence Indicator is inserted as a '1' on the Volume Header Label and is incremented on each block as a one position revolving counter, i.e., 1, 2, 3 ... 8, 9, 0, 1, etc.

A.4.5 This same count continues through each file on a given volume and applies to all labels, Tape Mark blocks (where it is not actually written) and data blocks.

A.5 Padding

Whenever it becomes necessary or advisable to extend the recorded length of a block beyond the end of the last (or only) record in it, the block may be extended (padded) with the agreement of the interchange parties.

A.5.1 Whenever a magnetic tape is recorded by, or is expected to be read by a word-oriented computer, all data blocks and labels shall be padded out to a multiple of the word lengths of the computer.

A.5.2 Blocks shall be padded out to the desired length by the use of circumflex characters (position 5/14 of the ECMA 7 bit code table).

A.5.2.1 Label blocks shall be padded only to the next higher multiple of the word length, using any desired padding characters, and as few characters as needed.

