

ECMA

EUROPEAN COMPUTER MANUFACTURERS ASSOCIATION

STANDARD ECMA-13

FOR

MAGNETIC TAPE LABELLING
AND FILE STRUCTURE FOR
INFORMATION INTERCHANGE

2nd Edition-August 1973

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

ECMA

EUROPEAN COMPUTER MANUFACTURERS ASSOCIATION

STANDARD ECMA-13

FOR

MAGNETIC TAPE LABELLING
AND FILE STRUCTURE FOR
INFORMATION INTERCHANGE

2nd Edition-August 1973

F O R E W O R D

The aim of this Standard is to make possible the interchange of information recorded on magnetic tapes between different users and different computers. This is accomplished by means of magnetically recorded labels to identify and structure files, and by providing a means of describing the basic characteristics of the blocks containing the records that constitute a file.

The features provided by this Standard allow the user to consider only the logical structure of his files.

In most implementations of the Standard a general purpose operating system will be in use, but in other cases there may only be installation or user written input/output routines which may form part of a special purpose operating system. The Standard has allowed for this by avoiding the term "operating system" and using instead "label handling routines". However, for proper implementation of the Standard the installation or user-written input/output routines are expected to provide the same minimum facilities as a general purpose operating system within the area defined by this Standard.

This second edition of the Standard differs technically from the first particularly in the introduction of the record spanning technique, which allows a record to span more than one block and even more than one reel of tape. A detailed description of these differences is given in Appendix B, together with reasons for making the changes included in this revision.

In addition, the Standard has been subject to a major editorial revision.

In the whole of the Standard, the usage of the Standard ECMA-6 "7-bit Input/Output Coded Character Set" is implied.

This second Edition of the Standard ECMA-13 has been adopted by the General Assembly of ECMA on June 14, 1973.

This Second Edition supersedes the Edition dated November 1967.

TABLE OF CONTENTS

	<u>Page</u>
1. <u>SCOPE AND FIELD OF APPLICATION</u>	1
2. <u>REFERENCES</u>	1
3. <u>DEFINITIONS</u>	1
4. <u>FORMATS AND CONTENTS OF LABELS</u>	6
4.1 Volume Header Label (VOL1)	7
4.2 First File Header Label (HDR1)	7
4.3 Second File Header Label (HDR2) (optional)	9
4.4 First End of Volume Label (EOV1)	10
4.5 Second End of Volume Label (EOV2) (optional)	10
4.6 First End of File Label (EOF1)	11
4.7 Second End of File Label (EOF2) (optional)	11
4.8 Other optional Labels (HDR3-HDR9, EOV3-EOV9, EOF3-EOF9)	11
4.9 User Labels (UVLn, UHLa, UTLa) (optional)	12
5. <u>PROCESSING OF LABEL FIELDS</u>	13
5.1 Fields in required Labels	13
5.2 Fields in optional Labels	13
5.3 Use of data in Label Fields	13
5.4 Volume Header Label (VOL1)	13
5.5 First File Header Label (HDR1)	13
5.6 Second File Header Label (HDR2)	14
6. <u>ARRANGEMENTS OF LABELS AND DATA</u>	15
6.1 Labels	15
6.2 Placement of Tape Marks	15
6.3 Volume Header Labels	15
6.4 File Header Labels	15
6.5 Data	15
6.6 End of File Labels	15
6.7 Tape Mark after End of File Label Group	16
6.8 End of Volume Labels	16
6.9 Empty File or File Section	16
6.10 Continuation File Section in a Multi-Volume File	16
6.11 Coincidence of End of File and End-of-Tape Marker	16
6.12 Coincidence of Beginning of File and End-of-Tape Marker	17

TABLE OF CONTENTS (cont'd)

	<u>Page</u>
7. <u>STRUCTURING THE FILES</u>	18
7.1 Configurations of Files	18
7.2 Label Groups and Tape Marks	19
7.3 Label Groups and Volumes	19
7.4 Optional Labels	20
7.5 Examples of the Grouping of Labels	20
8. <u>BLOCK STRUCTURE</u>	21
8.1 Grouping Records into Blocks	21
8.2 Recording Density	22
8.3 Size of Data Blocks	22
9. <u>PADDING</u>	24
9.1 Use of Padding	24
9.2 Fixed Block Length	24
9.3 Word-oriented Computers	24
9.4 Padding of Label Blocks	24
9.5 Padding of Data Blocks	24
Table 1 Classification of Labels	5
Figure 1 Structure of Magnetic Tape Files	18
Figure 2 Empty File Section at Beginning of Volume	18
Figure 3 Empty File Section at End of Volume	19
Figure 4 Spanned Records, unblocked	23
Figure 5 Spanned Records, blocked	23
Figure 6 One spanned Record, unblocked	23
Figure 7 Two spanned Records, blocked	23
<u>Appendix A - NOTES ON IMPLEMENTATION</u>	
A.1 <u>INTRODUCTION</u>	25
A.2 <u>REQUIRED AND OPTIONAL LABELS</u>	25
A.2.1 Required Labels	25
A.2.2 Optional Labels	25

TABLE OF CONTENTS (cont'd)

	<u>Page</u>
A.3 <u>USE OF FIELDS IN THE LABELS</u>	27
A.4.1 Limitation on "a" Characters	27
A.4.2 Volume Header Label (VOL1)	27
A.4.3 First File Header Label (HDR1)	28
A.4.4 Second File Header Label (HDR2)	29
A.4.5 First End of Volume Label (EOV1) and First End of File Label (EOF1)	29
A.4 <u>BLOCK SEQUENCE INDICATOR</u>	30
A.5.1 Use	30
A.5.2 Optionality	30
A.5.3 Position	30
A.5.4 Influence on Block Length	30
A.5.5 Counting Technique	30
A.5.6 Counted Blocks	31
Appendix B - <u>DIFFERENCES BETWEEN THE FIRST AND SECOND VERSION OF THIS STANDARD</u>	32
Appendix C - <u>FUTURE WORK ON LEVELS OF LABELLING</u>	35

1. SCOPE AND FIELD OF APPLICATION

This Standard ECMA-13 specifies the file structure for data interchange on magnetic tape and specifies magnetically recorded labels to identify the files, the file sections and the reels of tape.

It is the intention of this ECMA Standard, not that every instance of its implementation should necessarily include all of its provisions, but that the interchange of information should be facilitated by the use of a recorded magnetic tape conforming to this Standard to the following extent :

- it and every file upon it contains all of the required elements of the labelling system, formatted and placed as specified in this Standard, and containing an accurate description of the volume or the file to which they relate, and,
- if it or any file upon it contains any optional elements of the labelling system, they are formatted and placed as specified in this Standard, and contain an accurate description of the volume or file to which they relate, and,
- it and every file upon it contains only record/block formats and file configurations permitted by this Standard.

Failure to conform to this Standard may result in loss of the ability to interchange data effectively.

2. REFERENCES

- | | |
|---------|--|
| ECMA-5 | Data Interchange on 7-Track Magnetic Tape. |
| ECMA-6 | 7-Bit Input/Output Coded Character Set. |
| ECMA-12 | Data Interchange on 9-Track Magnetic Tape at 32 bits per mm (800 bpi). |
| ECMA-36 | Data Interchange on 9-Track Magnetic Tape at 63 bpmm (1600 bpi) Phase Encoded. |

3. DEFINITIONS

For the purposes of this Standard the following terms have the meanings indicated. For a better explanation, the concepts have, where appropriate, been listed separately as logical and physical. The definition of a term that is used in an ECMA Standard related to this subject conforms to its usage in that Standard; the definition of a term that is in common use in a context related to this Standard conforms to that common usage.

LOGICAL	PHYSICAL
<p><u>record</u></p> <p>Related data treated as a unit of information.</p> <p>EXAMPLES. In the context of business data, a transaction record, a customer's account record</p> <ul style="list-style-type: none">- the delineation of a record may be arbitrary and determined by the designer of the information format.- A record may be recorded in all or part of a block or in more than one block.	<p><u>block</u></p> <p>A group of characters written or read as a unit</p> <ul style="list-style-type: none">- A block may contain one or more complete records.- A block may contain segments of one or more spanned records. A single block shall not contain multiple segments of the same spanned record.
<p><u>file</u></p> <p>A collection of information consisting of records pertaining to a single subject.</p> <p>EXAMPLES. In the context of business data, a payroll file, an inventory file.</p> <ul style="list-style-type: none">- The delineation of a file may be arbitrary.- A file may be recorded on all or part of a volume, or on more than one volume.	<p><u>volume</u></p> <p>A dismountable physical unit of storage media, i.e. a reel of magnetic tape.</p> <ul style="list-style-type: none">- A volume may contain part of a file, a complete file or more than one file.- A volume may contain sections of one or more files but not multiple sections of the same file.
<p><u>file section</u></p> <p>That part of a file that is recorded on any one volume.</p> <ul style="list-style-type: none">- The sections of a file shall not have sections of other files interspersed.	
<p><u>file set</u></p> <p>A collection of one or more related files, recorded consecutively on a volume set.</p>	<p><u>volume set</u></p> <p>A collection of one or more volumes on which a file set is recorded.</p>

LOGICAL	PHYSICAL
<p><u>unspanned record</u></p> <p>A record contained in a file in which each record by design ends in the block in which it begins.</p>	
<p><u>spanned record</u></p> <p>A record contained in a file in which each record may begin in one block and end in another.</p> <ul style="list-style-type: none">- Each record consists of one or more segments, each segment being contained in a block, the blocks being written consecutively.	
<p><u>record segment</u></p> <p>That part of a spanned record that is contained in any one block.</p> <ul style="list-style-type: none">- The segments of a record shall not have segments of another record interspersed.	
<p><u>unblocked record</u></p> <p>A record contained in a file in which each block by design contains only one record or record segment.</p>	
<p><u>blocked record</u></p> <p>A record contained in a file in which each block may contain more than one record or record segment.</p>	
<p><u>fixed-length record</u></p> <p>A record contained in a file in which all the records by design have the same length.</p>	

LOGICAL	PHYSICAL
<u>variable-length record</u> A record contained in a file in which the records may have different lengths.	

Label : A record, at the beginning or at the end of a volume or of a file, that identifies, characterizes and/or delimits that volume or file. A label is not considered to be part of a file.

Label Set : A collection of contiguous labels with the same label identifier.

Label Identifier : Three characters recorded in the label to identify the label (see table 1).

Label Type : Labels are classified into four types as shown in table 1.

Label Group : A collection of contiguous label sets of the same label type.

Tape Mark : A delimiter used to indicate the boundary between file data and label groups and also between certain label groups.

Note : The tape mark configuration is specified in the relevant recorded magnetic tape Standards (see Section 2).

Double Tape Mark : A delimiter consisting of two consecutive tape marks that is used to indicate the end of a volume or of a file set.

Note : Two consecutive tape marks also occur when an empty file section or an empty file exists on a volume, in which case they are not interpreted as a double tape mark but rather as two single tape marks framing an empty file section. In this context "empty" means that no blocks are present between the tape mark following the beginning of file label group and the tape mark preceding the end of volume or end of file label group of that file section or file.

Operating System :

Software that controls the execution of computer programs and that may provide scheduling, debugging, input/output control, accounting, compilation, storage assignment, data management, and related services.

NOTE: An operating system may be used in a single installation or it may be used in many installations, as is frequently the case when it is provided by a supplier.

Label handling routines : A set of routines that process labels and that are an integral part of a system's software.

NOTE 1. The system's software may be an operating system provided by a supplier or it may be provided by an installation or a user. In either case it must provide at least the label facilities required by this Standard.

NOTE 2. This Standard has been written as if the label handling routines were not contained in the user program.

TABLE 1 Classification of Labels

TYPE	NAME	IDENTIFIER
BEGINNING OF VOLUME	VOLUME HEADER LABEL	VOL
	USER VOLUME HEADER LABEL	UVL
BEGINNING OF FILE OR OF FILE SECTION	FILE HEADER LABEL	HDR
	USER FILE HEADER LABEL	UHL

TYPE	NAME	IDENTIFIER
END OF FIRST OR OF INTERMEDIATE FILE SECTION	END OF VOLUME LABEL (VOLUME TRAILER LABEL)	EOV
	USER TRAILER LABEL (USER END OF VOLUME LABEL)	UTL
END OF FILE OR OF LAST FILE SECTION	END OF FILE LABEL (FILE TRAILER LABEL)	EOF
	USER TRAILER LABEL (USER END OF FILE LABEL)	UTL

4. FORMATS AND CONTENTS OF LABELS

In this standard "n" means any numeric character from 0 to 9. An "a" means any numeric, alphabetic or special character of the centre four columns of the ECMA-7-bit code table except position 5/15 and those positions where there is provision for alternative graphic representation. (See Section 2).

In subsections 4.1 to 4.9 the meaning of the table's headings is the following one :

CP - character position in the label
FIELD NAME - reference name of the field
L - length of the field
CONTENT - content of the field

4.1 Volume Header Label (VOL1)

CP	FIELD NAME	L	CONTENTS
1-3	Label Identifier	3	VOL
4	Label Number	1	1
5-10	Volume Identifier	6	"a" characters. Permanently assigned by the owner to identify the volume.
11	Accessibility	1	"a" character. Indicates restrictions on access to the information on the volume. Space means no restrictions.
12-37	Reserved for future standardization	26	Spaces
38-51	Owner Identifier	14	"a" characters. Indicates the owner of the volume.
52-79	Reserved for future standardization	28	Spaces
80	Label Standard Version	1	Indicates the version of the Standard to which the labels and data formats in the volume conform. 2 means this version 1 means the first version (see Appendix B).

4.2 First File Header Label (HDR1)

CP	FIELD NAME	L	CONTENT
1-3	Label Identifier	3	HDR
4	Label Number	1	1
5-21	File Identifier	17	"a" characters. Assigned by the originator to identify the file.

CP	FIELD NAME	L	CONTENT
22-27	File Set Identifier	6	"a" characters. Identifies the file set among file sets
28-31	File Section Number	4	"n" characters. Identifies the section among other sections of the file.
32-35	File Sequence Number	4	"n" characters. Identifies the file among other files of the file set
36-39	Generation Number	4	"n" characters. Distinguishes among successive generations of the file.
40-41	Generation Version Number	2	"n" characters. Distinguishes among successive iterations of the same generation.
42-47	Creation Date	6	One space followed by two "n" characters for the year followed by three "n" characters for the day (001 to 366) within the year.
48-53	Expiration Date	6	One space followed by two "n" characters for the year followed by three "n" characters for the day (001 to 366) within the year.
54	Accessibility	1	"a" character. Indicates restrictions on access to the information in this file. "Space" means no restrictions.
55-60	Block Count	6	000000
61-73	System Code	13	"a" characters. Identifies the system that recorded the file. The identifiers are not specified in this Standard.
74-80	Reserved for future standardization	7	Spaces

4.3 Second File Header Label (HDR2) (optional)

CP	FIELD NAME	L	CONTENT
1-3	Label Identifier	3	HDR
4	Label Number	1	2
5	Record Format	1	F = Fixed-length D = Variable-length S = Spanned U = Undefined
6-10	Block Length	5	"n" characters. Specifies the maximum number of characters per block.
11-15	Record Length	5	"n" characters. Specifies the record length, in conjunction with Record Format (CP 5). - If Record Format is F, this field contains the actual record length. - If Record Format is D, this field contains the maximum record length including the count field. - If Record Format is S, this field contains the maximum record length, excluding all the Segment Control Words. - If Record Format is S, 00000 indicates that the maximum length may be greater than 99999. - If Record Format is U, the content of this field is undefined.
16-50	Reserved for Systems Software Use	35	"a" characters. Not intended for use in an interchange environment.
51-52	Buffer Offset Length	2	"n" characters. Specifies the length in characters of any additional field inserted before the first record in a data block.

CP	FIELD NAME	L	CONTENT
53-80	Reserved for future standardization	28	Spaces

4.4. First End of Volume Label (EOV1)

CP	FIELD NAME	L	CONTENT
1-3	Label Identifier	3	EOV
4	Label Number	1	1
5-54	Same as the corresponding fields in HDR1	total 50	Same as the corresponding fields in HDR1.
55-60	Block Count	6	"n" characters. Denotes the number of data blocks since the preceding beginning of file label group. This count excludes label blocks and tape mark blocks.
61-80	Same as the corresponding fields in HDR1	total 20	Same as the corresponding fields in HDR1.

4.5. Second End of Volume Label (EOV2) (optional)

CP	FIELD NAME	L	CONTENT
1-3	Label Identifier	3	EOV
4	Label Number	1	2
5-80	Same as the corresponding fields in HDR2	total 76	Same as the corresponding fields in HDR2.

4.6 First End of File Label (EOF1)

CP	FIELD NAME	L	CONTENT
1-3	Label Identifier	3	EOF
4	Label Number	1	1
5-54	Same as the corresponding fields in HDR1	total 50	Same as the corresponding fields in HDR1.
55-60	Block Count	6	"n" characters. Denotes the number of data blocks since the preceding beginning of file label group. This count excludes label blocks and tape mark blocks.
61-80	Same as the corresponding fields in HDR1	total 20	Same as the corresponding fields in HDR1.

4.7 Second End of File Label (EOF2) (optional)

CP	FIELD NAME	L	CONTENT
1-3	Label Identifier	3	EOF
4	Label Number	1	2
5-80	Same as the corresponding fields in HDR2	total 76	Same as the corresponding fields in HDR2.

4.8 Other optional labels (HDR3 - HDR9, EOVS - EOVS9, EOF3 - EOF9)

CP	FIELD NAME	L	CONTENT
1-3	Label Identifier	3	HDR, EOVS or EOF as appropriate
4	Label Number	1	3,4,5,6,7,8 or 9
5-80	Reserved for System Software Use	76	"a" characters.

4.9 User Labels (UVLn, UHLa, UTLa) (optional)

CP	FIELD NAME	L	CONTENT
1-3	Label Identifier	3	UVL, UHL or UTL as appropriate.
4	Label Number	1	"n" character in UVL Labels, "a" character in UHL or UTL labels
5-80	Reserved for user application	76	"a" characters.

5. PROCESSING OF LABEL FIELDS

5.1 Fields in required labels

All fields must be written with content as specified, on reading they may be treated as desired.

5.2 Fields in optional labels

If the System Software elects to write an optional label, then the fields must be written as specified in the standard.

If the System Software elects to read an optional label, then the contents may be treated as needed.

5.3 Use of data in label fields

On input, the system may override data found in labels being processed by that system with new values of that data provided from other sources. The new values may be supplied before the file is processed (e.g., compiled values) or after the processing has begun (e.g., system control statements), at the option of the system implementors. However, data found in VOL1 is not overlaid or overridden.

5.4 Volume Header Label (VOL1)

The Volume Header Label must be preserved except as specified below. This does not preclude the rewriting of the label with the contents unchanged.

The Volume Header Label may be changed only if authorized by the owner, and then only as prescribed by the owner.

5.5 First File Header Label (HDR1)

5.5.1 File Set Identifier (CP 22-27)

The identification shall be the same for all files of a set.

5.5.2 File Section Number (CP 28-31)

The number of the first section of a file is 0001. This number is increased by 1 for each successive volume of the file.

5.5.3 File Sequence Number (CP 32-35)

The File Sequence Number of the first file in a file set is 0001. This number is increased by 1 for each successive file of the set. In all the labels for a given file, whether that file be single or multi-volume, this field contains the same number.

5.5.4 Generation Number (CP 36-39)

The Generation Number of the first generation of a file is 0001.

If subsequent generations of a file are noted, this number is increased by 1 for each successive generation of the file.

5.5.5 Generation Version Number (CP 40-41)

The Generation Version Number of the first attempt to produce a generation of a file is 00. If the number of subsequent attempts to produce that generation of the file is noted, this number is increased by 1 at each subsequent attempt.

The Generation Version Number is reset to 00 when the Generation Number is increased by 1.

5.5.6 Expiration Date (CP 48-53)

A file is regarded as "expired" on a day whose date is equal to, or later than the date given in this field. When this condition is satisfied, the remainder of the volume set 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 dates of all previous files on the volume set.

5.6 Second File Header Label (HDR2)

5.6.1 Block Length (CP 6-10)

The count in this field includes buffer offset and padding. The actual maximum block capacity for data is reduced by the content of the Buffer Offset Length field and by the number of padding characters, in order that the block should not exceed the maximum length specified in the applicable recorded magnetic tape standards.

5.6.2 Reserved for System Software Use (CP 16-50)

On input and/or output, this field may be used by those System Software routines which recognize the System Code in HDR1, which will identify the particular System Software routine that created the file. In interchange the contents of this field are ignored.

5.6.3 Buffer Offset Length (CP 51-52)

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 will be specified in this field. If no such information is included, the contents of this field are 00.

6. ARRANGEMENTS OF LABELS AND DATA

6.1 Labels

A label is an 80 character block, the character positions (CP) of which are numbered 1 to 80. However, this may be extended by padding characters (see section 9).

Numbered labels shall be recorded in consecutive ascending order. The first label shall be numbered "1". User header labels and user trailer labels have no such restrictions.

6.2 Placement of tape marks

Label groups are always delimited by a tape mark except at the beginning of a volume, where the beginning of volume and beginning of file label groups are contiguously recorded with no intervening tape mark.

6.3 Volume header labels

The first block on a volume shall be the Volume Header Label (VOL1). This label shall not be used at any other place in the volume.

If User Volume Labels (UVL1-UVL9) are used they shall immediately follow the VOL1 label.

6.4 File header labels

Each file shall be preceded by Header Labels, the first of which shall be the first File Header Label (HDR1). If other file header labels are used, they shall immediately follow the HDR1 label.

If User File Header Labels (UHLa) are used, they shall immediately follow the last HDRn label.

6.5 Data

File data shall follow the last label of the beginning of file label group and be separated from the label group by a tape mark.

6.6 End of file labels

The End of File Label (EOF1) shall follow the last data block of the file and be separated from the file by a tape mark.

If other End of File Labels (EOF2 - EOF9) are used, they shall immediately follow the EOF1 label.

If User End of File Labels (UTLa) are used, they shall immediately follow the last EOFn label.

6.7 Tape mark after end of file label group

If the file ends within a volume a tape mark shall immediately follow the last label of the end of file label group.

If the file is the last file on a file set a double tape mark shall immediately follow the last label of the end of file label group.

6.8 End of volume labels

If the file extends over the end of a volume, the End of Volume Label (EOV1) shall follow the last data block on that volume and be separated from that data block by a tape mark.

If other end of volume labels (EOV2-EOV9) are used they shall immediately follow the EOV1 label.

6.8.1 If User End of Volume Labels (UTLa) are used they shall immediately follow the last EOVn label.

6.8.2 The last label of an end of volume label group shall always be followed by a double tape mark.

6.9 Empty file or file section

When an empty file or empty file section is present, the rules stated in 6.5, 6.6 and 6.8 will cause two consecutive tape marks to be written between the beginning of file label group and the end of file or end of volume label group.

6.10 Continuation file section in a multi-volume file

The first data block of a continuation file section shall be preceded by a beginning of file label group as described in 6.4 and 6.5.

This includes an exact copy of the last File Header Label set on the previous volume, except that the File Section Number which is on HDR1, is increased by 1 (see 5.5.2).

6.11 Coincidence of end of file and end-of-tape marker

If the end of a file and the end of a volume coincide three situations are possible.

6.11.1 The end-of-tape marker is recognized whilst the system is writing the last data block of the file. In this case the system will complete writing the data block and close the volume as described in 6.8 and continue the file on the next volume as described in 6.10 except that no data blocks will be written on the next volume but an end of file label group as described in 6.6 (see 7.1.1).

6.11.2 The end-of-tape marker is recognized whilst the system is writing the end of file label group and the file is not the last file of a set. In this case the system will complete the writing of the end of file label group. The beginning of file label group of the next file will then be written, followed by an empty file section and the volume terminated by an end of volume label group. The beginning of file label group will then be rewritten at the start of the next volume (see 7.1.2).

6.11.3 The end-of-tape marker is recognized whilst the system is writing the end of file label group and the file is the last file of a set. In this case the configuration used to terminate the file set shall be as described in 7.1.2.

6.12 Coincidence of beginning of file and of end-of-tape marker

If the end-of-tape marker is recognized whilst the system is writing the beginning of file label group, then it will complete the writing of that group, terminate the volume by writing an end of volume label group and rewrite the beginning of file label group at the start of the next volume (see 7.1.3).

7. STRUCTURING THE FILES

Label sets and tape marks are used to establish the file structure according to the following rules, as illustrated in figures 1-3. In these figures the beginning of the tape is at the left, and the end of the tape is at the right. Labels are indicated by their first four characters, and a tape mark is represented by an asterisk (*).

7.1 Configurations of files

The various configurations of files that can be formed according to these rules are illustrated in fig. 1.

The rules and figures are presented including only the required labels.

Single-volume file

```
VOL1 HDR1*---File A---*EOF1**
```

Multi-volume File

```
VOL1 HDR1*---first section of File A---*EOV1**
```

```
VOL1 HDR1*---last section of File A---*EOF1**
```

Multi-file volume

```
VOL1 HDR1*--File A--*EOF1*HDR1*---File B---*EOF1**
```

Multi-volume multi-file

```
VOL1 HDR1*--File A--*EOF1*HDR1*--first section of File B---*EOV1**
```

```
VOL1 HDR1*---intermediate section of File B-----*EOV1**
```

```
VOL1 HDR1*--last section of File B--*EOF1*HDR1*--File C--*EOF1**
```

Figure 1 - Structure of magnetic tape files

7.1.1 As described in 6.11.1, the labelling configuration to be used when the end-of-tape marker is recognized whilst the last data block of the file is being written is as shown in figure 2.

```
----- last section of File A--*EOV1**  
VOL1 HDR1**EOF1*HDR1*- - - first section of File B----  
      (A)   (A)  (B)
```

Figure 2 - Empty File Section at beginning of volume

7.1.2 As described in 6.11.2, the labelling configuration to be used when the end-of-tape marker is recognized whilst the label handling routine is writing the end of file label group, is as shown in figure 3.

There is one exception, described in 6.11.3 : if the end of file label group belongs to the last file of a file set, then the file set will be terminated on that volume by the use of the configuration EOF1 followed by two tape marks.

7.1.3 As described in 6.12 the labelling configuration resulting when the end-of-tape marker is recognized whilst the label handling routine is writing the beginning of file label group, is as shown in figure 3. Note that the File Section Number (HDR1 CP 28-31) is 1 on the original volume and 2 on the continuation volume.

```
-----last section of File A --- *EOF1*HDR1**EOV1**  
                                     (A)  (B)  
VOL1 HDR1*---first section of File B-----  
                                     (B)
```

Figure 3 - Empty File Section at end of volume

7.2 Label groups and tape marks

There shall be no tape mark within a label group. Whenever figures 1, 2 and 3 indicate a tape mark following a required label, that tape mark shall actually follow the last label of the entire group.

7.3 Label groups and volumes

Every label group shall be completed on the volume where the first label of the group was recorded.

7.4 Optional labels

Optional labels shall be fitted into the file structure as described in 7.4.1, 7.4.2 and 7.4.3 without otherwise modifying the relationship between required labels and the files.

Illustrations of the order of these labels are given in 7.5.

- 7.4.1 The optional labels UVL1 - UVL9, when used, shall directly follow the VOL1 label. The label number of the consecutive UVLn labels shall be 1, 2, 3, 4, 5, 6, 7, 8 and 9 respectively.
- 7.4.2 The optional labels HDR2 - HDR9, EOVS2 - EOVS9 and EOF2 - EOF9, when used, shall directly follow a required label with the same identifier. The label number of the consecutive optional labels shall be 2, 3, 4, 5, 6, 7, 8 and 9 respectively.
- 7.4.3 The remaining optional labels (UHLA, UTLA), when used, shall directly follow the optional labels as described in 7.4.2 and of the same type. If, however, such labels do not exist, the optional labels UHLA and UTLA will directly follow a required label of the same type.

7.5 Examples of the grouping of labels

- 7.5.1 Physical beginning of tape to physical end of tape (not end of file).

VOL1 UVL1 ... UVLn HDR1 HDR2 HDRn UHLA

UHLA*File Data*EOV1 ... EOVSn UTLA ... UTLA**

- 7.5.2 Physical beginning of tape to end of intermediate file of a file set.

VOL1 UVL1 ... UVLn HDR1 HDR2 ... HDRn UHLA

UHLA*File Data*EOF1 ... EOFn UTLA UTLA*

7.5.3 Physical beginning of tape to end of file set.

VOL1 UVL1 ... UVLn HDR1 HDR2 ... HDRn UHLA

UHLA*File Data*EOF1 ... EOFn UTLA ... UTLA**

7.5.4 Beginning of new file (not beginning of tape) to physical end of tape (not end of file).

HDR1 HDR2 ...HDRn UHLA ...UHLA*

File Data*EOV1 ... EOvn UTLA ... UTLA**

7.5.5 Beginning of any intermediate file of a file set (not beginning of tape) to end of file.

HDR1 HDR2 ...HDRn UHLA ... UHLA*

File Data*EOF1 ... EOFn UTLA ... UTLA*

7.5.6 Beginning of new file (not beginning of tape) to end of file set.

HDR1 HDR2 ... HDRn UHLA ... UHLA*

File Data*EOF1 ... EOFn UTLA ... UTLA**

8. BLOCK STRUCTURE

8.1 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 for formats F and D. There must be an integral number of segments in a block for format S. Padded blocks are permitted (see section 9). Truncated and varying length blocks are permitted.

8.1.1 Fixed-Length Records (F format). No indication of record length is required within a file.

8.1.2 Variable-Length Records (D format). 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 numeral, occupying the first four character positions of each record.

8.1.3 Spanned Records (S format). A segment control word (SCW) precedes each segment. The SCW consists of 5 characters. The first character of the SCW is called the Spanning indicator. This indicator may have the values 0, 1, 2 or 3 as follows :

- 0 - record begins and ends in this segment.
- 1 - record begins but does not end in this segment.
- 2 - record neither begins nor ends in this segment.
- 3 - record ends but does not begin in this segment.

Record length is unbounded in that there is no limit to the number of segments in a record. This does not prohibit the system software from limiting the size of a work area available to reconstruct a record.

The segment length is an arbitrary choice, and segments may be created and the lengths of existing segments changed when creating or copying a file in order to fit the size of the block to be written. The segment length shall be expressed as a decimal numeral occupying the last four character positions of the SCW.

There shall be only one segment of the same record in a block. The segments of a record shall not have segments of other records interspersed.

Records may span volumes.

Examples of the use of spanned records are given in figures 4 to 7.

8.1.4 Undefined Records (U format). When records do not meet the specifications of 8.1.1, 8.1.2 or 8.1.3, they are undefined in format. The interchange of information in the undefined format will require the prior agreement of the interchange parties.

8.1.5 Not more than one of the record formats described in 8.1.1 to 8.1.4 may be used in any one file.

8.1.6 By-pass or check-point records

Only relevant data blocks shall be written on a tape used for interchange. Since by-pass information or check-point records are considered to be extraneous to the interchange, and no standard means of identification is provided, the recording of by-pass and check-point information is not allowed on tapes for interchange.

8.2. Recording density

The blocks recorded on all volumes containing a file set shall be recorded at the same density.

8.3 Size of data blocks

The minimum and maximum size of data blocks are specified in the relevant recorded magnetic tape Standard. These limits are a minimum of 18 characters and a maximum of 2048 characters. (See section 2).




Fig. 4 Spanned records, unblocked .




Fig. 5 Spanned records, blocked .




Fig. 6 One spanned record, unblocked .
Record length 4241 characters
(each line represents a block)

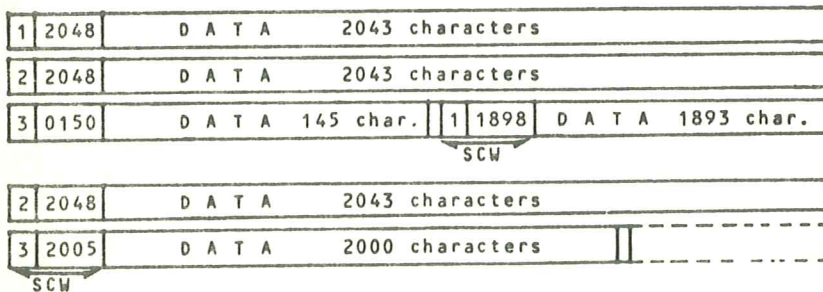


Fig. 7 Two spanned records, blocked .
Record #1 - Record length 4231 characters
Record #2 - Record length 5936 characters
(each line represents a block)

9. PADDING

9.1 Use of 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) to the desired length.

9.2 Fixed block length

Whenever a magnetic tape is recorded by a device or program that is restricted to a minimum or fixed block length, each data block and each label may be padded out to that minimum or fixed length.

9.3 Word-oriented computers

Whenever a magnetic tape is recorded by a word-oriented computer, all data blocks and labels may be padded out to a multiple of the word length of that computer.

9.4 Padding of label blocks

Label blocks are padded out to the desired length, using any desired padding characters.

9.5 Padding of data blocks

Blocks within a file are padded out to the desired length using circumflex accent characters (position 5/14 of the ECMA 7-bit code table).

APPENDIX A

NOTES ON IMPLEMENTATION

A.1 INTRODUCTION

This Appendix provides a guide for those implementing the Standard, and explains and amplifies the information in the body of the Standard.

A.2 REQUIRED AND OPTIONAL LABELS

A.2.1 Required labels

For conformity with this standard the inclusion of the following labels in the correct position is required :

VOL1
HDR1
EOV1
EOF1

These labels are handled by the label handling routines.

A.2.2 Optional labels

The remaining labels are designated optional and fall into two categories.

The first category consists of :

HDR2 - HDR9
EOV2 - EOV9
EOF2 - EOF9

These labels are handled by the label handling routines. When used, these labels 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 systems software. The contents of these labels can be ignored in interchange. The writing of these labels is optional, but where used they must be written in conformity with the requirement of the standard in relation to content and position. Whether these labels are ignored or not when reading is at the discretion of the label handling routines. There is no requirement under this standard to take action on the label content.

In these labels, the actual interpretation and treatment of those fields designated "Reserved for Systems Software use" can be the subject of special agreement between the interchange parties.

The second category consists of the user labels :

UVLn
UHLa
UTLa

User labels (i.e. UVLn, UHLa and UTLa) are handled partly by the label handling routines of the system, to the extent that they are recognized and passed to the user on input, and are written to the tape on output.

Users provide the information in user labels beyond the Label Identifier and Number fields on Output, and utilize it on input.

It will often be convenient to use these labels to contain summary information about a volume or 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, the contents of user labels will be ignored in interchange.

When these labels are used, they must be written in conformity with the requirement of the Standard in relation to content and position.

A.3 USE OF TAPE MARKS

A.3.1 At end of volume

The double Tape Mark at the end of each volume (Figure 1) permits the following procedure to be used in accomplishing the operation "Forward-space File":

Having read HDR1,

A - Index forward till three Tape Marks passed. Read the next block.

if HDR1 : One file has been indexed.

if Tape Mark : End of volume, Rewind.

If not end of set : Alternate, verify HDR1 on next volume, return to A.

An alternative procedure for that operation would be as follows :

Having read HDR1,

A - Index forward till two Tape Marks passed. Read the next block.

If EOVL : Rewind, alternate, verify HDR1 on next volume, return to A.

If EOF1 : Index forward till one Tape Mark passed, Read next block.

If HDR1 : One file has been indexed.

If Tape Mark : End of set has been reached.

Thus, the double Tape Mark prevents tape runaway on forward spacing.

A.3.2 To frame an empty File Section

Two consecutive Tape Marks appear at the beginning 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 HDR1 Label.

Pass over or process any optional labels.

Read Tape Mark - switch to processing of file information.

Read Tape Mark - switch to processing of labels.

In figure 2, read File-A EOF1.

In figure 3, read EOVL occurring within File-B.

A.4 USE OF FIELDS IN THE LABELS

A.4.1 Limitation on "a" characters

The limitation on "a" characters as described in section 4 is intended as a guide to provide maximum interchangeability and consistent printing especially during international interchange. Checking for conformity to this limitation is not implied.

A.4.2 Volume Header Label (VOL1)

This label identifies the physical reel of magnetic tape and the contents of the label relate to the identity of the volume.

A.4.2.1 Accessibility (CP 11)

This field 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 Volume Label), but this field might be used as an indicator in conjunction with such a User Volume 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.4.2.2 Owner Identifier (CP 38-51)

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.4.2.3 Label Standard Version (CP 80)

This field is used to indicate whether or not the information recorded on this volume conforms to this Standard. 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 versions by the use of numerals in this field, rather than letters, to the extent possible.

A.4.3 First File Header Label (HDR1)

A.4.3.1 File Set Identifier (CP 22-27)

It is desirable that a unique identification be established. In most cases, this objective may be satisfied by duplicating the Volume Identifier (VOL1 CP 5-10) of the first or only volume of the set.

A.4.3.2 File Section Number (CP 28-31)

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 portion of the file held on each volume is termed a file section. The actual beginning of a file may be identified by "0001" in this field. Subsequent file sections will be numbered sequentially on subsequent volumes.

A.4.3.3 Generation Version Number (CP 40-41)

This field is used to differentiate output data which has been produced by repeated processing or writing operations and which in all other respects would bear the same identification. For example, it may be 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.

A.4.3.4 Accessibility (CP 54). See A.4.2.1

A.4.3.5 System Code (CP 61-73)

On output this field is inserted by the label handling routines to identify the System Software that created the file. The value to be entered is a constant for a given system.

It is likely that, in time, a standard method of identifying the system will be defined.

A.4.4 Second File Header Label (HDR2)

A.4.4.1 Record length (CP 11-15)

The content of this field is undefined if record format (HDR2 CP 5) is U. However, in normal usage it is expected that this field will be "space" filled.

A.4.4.2 Reserved for System Software Use (CP 16-50)

This field could contain any information needed to increase the efficiency of record processing by the System Software.

The System Code (HDR1 CP 61-73) may be used by label handling routines to determine the system software that created the file.

A.4.5 First End of Volume Label (EOV1) and First End of File Label (EOF1)

A.4.5.1 Block Count CP 55-60

This field is provided in order that when a magnetic tape is read the label handling routine may ensure that no blocks have been skipped and no spurious blocks have been inserted. The particular error of equal numbers of skipped and spurious blocks may escape detection.

A.5 BLOCK SEQUENCE INDICATOR

A.5.1 Use

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 checking the Block Count. On detecting, 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:

- Immediate identification of any error involving added or missed blocks, including most compensating errors of equal numbers of skipped and spurious blocks, and
- 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.5.2 Optionality

The implementation and usage of the Block Sequence Indicator are both optional and require the agreement of the interchange parties. There is no implication that systems software should make provision to recognize this indicator solely for conformance with this feature. The definitions and rules in this Standard have been written as if the Block Sequence Indicator option were not implemented.

A.5.3 Position

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.5.4 Influence on Block Length

The Block Length shall not include the Block Sequence Indicator if the latter is used.

A.5.5 Counting Technique

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.5.6 Counted Blocks

This same count continues in unbroken sequence through all files and file sections on a given volume and applies to all labels, tape mark blocks (where it is not actually written) and data blocks.

APPENDIX B

DIFFERENCES BETWEEN THE FIRST AND SECOND VERSION
OF THIS STANDARD

When the Label Standard Version (VOL1 CP 80) has the value 1 it means that the volume was written according to the rules and options of the first version of the Standard (published in November 1967). The differences between that version and the present one are set out below.

In general, the reasons for making the changes fall into three categories :

- Correction of errors and ambiguities in the first version
- Making some of the options mandatory. This reduces the number of possibilities that the Systems Software has to deal with and reduces ambiguity, e.g. the problem of how the system writing the volume and the system reading it should deal with optional fields in mandatory labels, and optional fields in optional labels, is almost removed.
- Enhancing the technical content of the standard.

<u>Field</u>	<u>Version 1</u>	<u>Version 2</u>	<u>Reason for change</u>
VOL1 CP 12-31	Field name is "Reserved for operating systems"	Field name is "Reserved for future standardization"	It was logically impossible for the operating system to use this field while adhering to the rest of the standard

<u>Field</u>	<u>Version 1</u>	<u>Version 2</u>	<u>Reason for change</u>
HDR1 CP 36-39 Generation Number and HDR1 CP 40-41 Generation Version Number	Optional	Mandatory	To reduce the number of options and ambiguities resulting.
HDR2 CP 5 Record Format	Includes format V (variable) with the number of characters in the record specified in binary	No provision for format V	It was not possible to implement this format in a general interchange environment.
HDR2 CP 5 Record Format	-	Addition of format S to indicate record spanning. Further described in section 8	Enhancement of the Standard to meet the demand for such a facility.
HDR2 CP 51-52 Buffer Offset Length	Optional	Mandatory. Length added to the field name. Must contain 00 if not used.	This field must contain the length of the buffer offset which implies 00 if the buffer offset facility is not used. The original wording was ambiguous in that it permitted either 00 or two space characters.
EOF1 CP 5-54 and EOF1 CP 61-80	Optional	Mandatory	To reduce the number of options and ambiguities resulting.
EOF2 CP 5-80	Optional	Mandatory (in cases where this optional label is used)	To reduce the number of options and ambiguities resulting.

<u>Field</u>	<u>Version 1</u>	<u>Version 2</u>	<u>Reason for change</u>
EOV1 CP 5-54 and EOV1 CP 61/80	Optional	Mandatory	To reduce the number of options and ambiguities resulting.
EOV2 CP 5-80	Operating system option	Same as corresponding fields in HDR2.	For consistency with the approach adopted with EOVI for reading backwards.

<u>Section</u>	<u>Version 1</u>	<u>Version 2</u>	<u>Reason for change</u>
9 PADDING	Not part of the Standard (in the Appendix)	Part of the Standard	Enhancement of the Standard to meet the demand for such a facility.
A.2.6 Special arrangement for starting a file at beginning of volume	In Appendix A	Deleted	To reduce the number of options and ambiguities resulting.

APPENDIX C

FUTURE WORK ON LEVELS OF LABELLING

In order to facilitate interchange of information on systems of dissimilar capabilities it is proposed that studies be undertaken of a system of levels of labelling, with a view towards eventual inclusion in a later version of this standard. The scheme of levels would be compatible with the facilities defined in this ECMA Standard. It should provide a fully compatible, well-nested system of labels for use of smallest and simplest, to largest and most sophisticated computing systems, and ensure the capability for interchange among them with fewest restrictions. Thus any volume set could be processed correctly by any implementation of equal or higher level ; and any implementation could process correctly any volume set of equal or lower level.

Rather than to constrain or measure an implementation of a system it is intended that facilities excluded in this standard and facilities included in this standard as optional at a given level may be included without reflection upon the conformance of any implementation. These levels would be designed to enable implementers of systems to guide applications designers and users to ascertain the capability of a system with respect to the requirements of the application and in selection of facilities to ensure that the medium can be processed correctly by the receiving system.

The following are the areas of study.

- File sets (single/multi-file, single/multi-volume)
- Labels (VOL1,HDR1,HDR2,UHLA,EOV1,EOV2,EOF1,EOF2,UTLa)
- Record/Block Structures (fixed/variable/spanned/undefined, blocked/unblocked)

For each of the levels to be specified, the requirements for the supporting software would need to be specified :

- Characterization of a particular level of software
- Label fields processed
- Responsibility for the validity of labels
- File security.

