ECMA

EUROPEAN COMPUTER MANUFACTURERS ASSOCIATION

STANDARD ECMA-13

MAGNETIC TAPE LABELLING AND FILE STRUCTURE FOR INFORMATION INTERCHANGE

3rd Edition-January 1978

ECMA

EUROPEAN COMPUTER MANUFACTURERS ASSOCIATION

STANDARD ECMA-13

MAGNETIC TAPE LABELLING AND FILE STRUCTURE FOR INFORMATION INTERCHANGE

3rd Edition-January 1978

BRIEF HISTORY

In November 1967 the General Assembly of ECMA adopted the first edition of Standard ECMA-13, prepared by a task group of TC 1, Input/Output Codes.

Subsequent developments led to the setting up of TC 15, Labelling and File Structure, and to the preparation of a second edition of the Standard, adopted by the General Assembly of ECMA in June 1973. Beside a major editorial revision, this new edition included provisions for record spanning, while a number of fields, optional in the previous edition, were made mandatory.

Work in the area of subsetting, in co-operation with ISO, led to the definition of four levels of labelling, providing explicit subsetting and facilitating the interchange of data among users.

This third edition incorporates the definition of levels and other minor changes in the areas of Record Format, Block Length and Block Sequence Indicator.

This third edition of Standard ECMA-13 has been adopted by the General Assembly of ECMA on Dec. 14, 1977.

THIS THIRD EDITION SUPERSEDES THE EDITION DATED AUGUST 1973.

TABLE OF CONTENTS

1 SCOPE AND FIELD OF APPLICATION

The purpose of this Standard ECMA-13 is to facilitate 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.

This Standard contains specifications for four levels of labelling. This provides a fully compatible, well nested system of labels for use of smallest and simplest, to largest and most sophisticated computing systems, and ensures the capability for interchange among them with fewest restrictions.

A recorded magnetic tape intended to be interchanged between systems of potentially different architecture is expected to correspond to one of these four levels. The constraints of this Standard may not be needed to apply to data not intended for interchange between systems of potentially different architecture.

It is the intention of this 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.

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

In most implementations of this 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. This has been allowed for in this Standard by avoiding the term "operating system" and using instead "label handling routines".

However, for proper implementation of this 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 third edition of Standard ECMA-13 differs technically from the second edition in the introduction of the concept of levels of labelling.

A detailed description of the differences is given in Appendix C, together with the reasons for making the change and, where applicable, some future work that may be required.

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 bpmm.

ECMA-35   Extensions of the 7-Bit Coded Character Set.

ECMA-36   Data Interchange on 9-Track Phase Encoded Magnetic Tape at 63 bpmm.

ECMA-43   8-Bit Coded Character Set.

3 DEFINITIONS

For the purpose 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.

LOGICAL	PHYSICAL
RECORD : Related data treated as a unit of information. - 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.	BLOCK : A group of characters written or read as a unit. - 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.
FILE : A collection of information consisting of records pertaining to a single subject. - 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.	VOLUME : A dismountable physical unit of storage media, viz. a reel of magnetic tape. - 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.

LOGICAL	PHYSICAL
For a file recorded on more than 1 Volume ; FILE SECTION : that part of a file that is recorded on any one volume. - The section of a file shall not have sections of other files interspersed.	
FILE SET : A collection of one or more related files, recorded consecutively on a volume set.	VOLUME SET : A collection of one or more volumes on which a file set is recorded.
UNSPANNED RECORD : A record contained in a file in which each record by design ends in the block in which it begins.	
SPANNED RECORD : A record contained in a file in which each record may begin in one block and end in another. - Each record consists of one or more segments. Records are contained in one or more consecutive blocks, such that only one segment of each record can appear in any one block (see Fig. 4 and 5).	
RECORD SEGMENT : That part of a spanned record that is contained in any one block. - The segments of a record shall not have segments of another record interspersed.	
UNBLOCKED RECORD : A record contained in a file in which each block by design contains only one record or record segment.	

LOGICAL	PHYSICAL
BLOCKED RECORD : A record con- tained in a file in which each block may contain more than one record or record segment.	
FIXED-LENGTH RECORD : A record contained in a file in which all the re-cords by design have the same length.	
VARIABLE-LENGTH RECORD : 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 label (see Table 1).

LABEL GROUP :   A collection of one or more contiguous label sets that delimit one end of a volume, of a file section or of a file (see Table 1).

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 2).*

DOUBLE TAPE
MARK :   A delimiter consisting of two consecutive tape marks that is used to indicate the end of a volume of a file s

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 beginnin*

of file label group and the tape mark preceding the
end of volume or end of file label group of that file
section of file.

OPERATING
SYSTEM :
Software that controls the execution of computer pro-
grams and that may provide scheduling, debugging, input/
output control, accounting, compilation, storage as-
signment, 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 systems's software.

NOTES: - The system 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.

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

Group Name	Set Name	Identifier
Beginning of Volume	Volume Header	VOL
	User Volume Header	UVL
Beginning of File or of File Section	File Header	HDR
	User File Header	UHL
End of First or of Intermediate File Section	End of Volume (Volume Trailer)	EOV
	User Trailer (User End of Volume)	UTL
End of File or of Last File Section	End of File (File Trailer)	EOF
	User Trailer (User End of File)	UTL

Table 1 - Classification of Labels

4. FORMATS AND CONTENTS OF LABELS

4.1 General

The characters permitted in the labels shall be the characters of the International Reference Version (IRV) of the ECMA 7-bit Coded Character Set contained in positions 2/0, 2/1, 2/2, 2/5 to 3/5 and 4/1 to 5/10.

In the following sections the meaning of the table headings shall be:

- CP : Character position within the label.

- Field Name : Reference name of the field.

- L : Length of the field in number of characters.

- Content : The content of the field as follows:

    a-character : any permitted character or combination of permitted characters.

    digit   : any digit or combination of digits from ZERO to NINE.

    Specific character(s), digit(s) or sign(s).

					b7	0	0	0	0	1	1	1	1
					b6	0	0	1	1	0	0	1	1
					b5	0	1	0	1	0	1	0	1
b4	b3	b2	b1			0	1	2	3	4	5	6	7
0	0	0	0	0		NUL	T.C.₇ (DLE)	SP	0	à	P	`	p
0	0	0	1	1		T.C.₁ (SOH)	DC₁	!	1	A	Q	a	q
0	0	1	0	2		T.C.₂ (STX)	DC₂	"	2	B	R	b	r
0	0	1	1	3		T.C.₃ (ETX)	DC₃	#	3	C	S	c	s
0	1	0	0	4		T.C.₄ (EOT)	DC₄	¤	4	D	T	d	t
0	1	0	1	5		T.C.₅ (ENQ)	T.C.₈ (NAK)	%	5	E	U	e	u
0	1	1	0	6		T.C.₆ (ACK)	T.C.₉ (SYN)	&	6	F	V	f	v
0	1	1	1	7		BEL	T.C.₁₀ (ETB)	'	7	G	W	g	w
1	0	0	0	8		FE₀ (BS)	CAN	(	8	H	X	h	x
1	0	0	1	9		FE₁ (HT)	EM	)	9	I	Y	i	y
1	0	1	0	10		FE₂ (LF)	SUB	*	:	J	Z	j	z
1	0	1	1	11		FE₃ (VT)	ESC	+	;	K	[	k	{
1	1	0	0	12		FE₄ (FF)	IS₄ (FS)	,	<	L	\	l	¦
1	1	0	1	13		FE₅ (CR)	IS₃ (GS)	-	=	M	]	m	}
1	1	1	0	14		SO	IS₂ (RS)	.	>	N	^	n	‾
1	1	1	1	15		SI	IS₁ (US)	/	?	O	_	o	DEL

TABLE 2

4.2 Volume Header Label (System Label VOL 1)

CP	Field name	L	Content
1-3	Label Identifier	3	VOL
4	Label Number	1	1
5-10	Volume Identifier	6	a-characters
11	Accessibility	1	a-character
12-37	(Reserved for Future Standardization)	26	SPACEs
38-51	Owner Identifier	14	a-characters
52-79	(Reserved for Future Standardization)	28	SPACEs
80	Labelling Standard Version	1	a-character

The Volume Header Label (VOL 1) shall identify the volume, its owner, its security level and the version of the Labelling Standard that has been used.

4.2.1 Volume Identifier (VOL 1 CP 5-10)

This field shall be permanently assigned by the owner to identify the volume.

4.2.2 Accessibility (VOL 1 CP 11)

This field shall indicate restrictions on access to the information in the volume. SPACE shall mean no restrictions.

4.2.3 Owner identifier (VOL 1 CP 38-51)

This field shall identify the owner of the volume.

4.2.4 Labelling Standard Version (VOL 1 CP 80)

This field shall indicate the version of this standard to which labels and data formats on this volume conform.

It shall contain one of the following characters :

- 3 meaning this version dated January 1978.

- 2 meaning the version dated August 1973

- 1 meaning the version dated November 1967

4.3 First File Header Label (System 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
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

CP	Field name	L	Content
40-41	Generation Version Number	2	Digits
42-47	Creation Date	6	SPACE followed by five digits
48-53	Expiration Date	6	SPACE followed by five digits
54	Accessibility	1	a-character
55-60	Block Count	6	ZEROes
61-73	System Code	13	a-characters
74-80	(Reserved for Future Standardization)	7	SPACEs

The First File Header Label (HDR1) shall identify the file, its parameters and its security level.

4.3.1 File Identifier (HDR1 CP 5-21)

This field shall be assigned by the originator to identify the file.

4.3.2 File Set Identifier (HDR1 CP 22-27)

This field shall identify the file set among file sets.

4.3.3 File Section Number (HDR1 CP 28-31)

This field shall identify the section among other sections of the file.

4.3.4 File Sequence Number (HDR1 CP 32-35)

This field shall identify the file among other files of the file set.

4.3.5 Generation Number (HDR1 CP 36-39)

This field shall distinguish among successive generations of the file.

4.3.6 Generation Version Number (HDR1 CP 40-41)

This field shall distinguish among successive iterations of the same generation.

4.3.7 Creation Date (HDR1 CP 42-47)

This date shall be expressed as follows :
- SPACE
- two digits for the last two digits of the year
- three digits for the day within the year (000 to 366).
A SPACE followed by five ZEROes means no significant Creation Date.

4.3.8 Expiration Date (HDR1 CP 48-53)

This date shall be expressed as follows :
- SPACE
- two digits for the last two digits of the year
- three digits for the day within the year (000 to 366).
A SPACE followed by five ZEROes shall indicate an expired file.

4.3.9 Accessibility (HDR1 CP 54)

This field shall indicate restrictions on access to the information in the file. SPACE shall mean no restrictions.

4.3.10 System Code (HDR1 CP 61-73)

This field shall identify the system that recorded the file. The identifiers are not specified in this ECMA Standard.

4.4 Second File Header Label (System Label HDR2)

CP	Field name	L	Content
1-3	Label Identifier	3	HDR
4	Label Number	1	2
5	Record Format	1	capital letter F or capital letter D or capital letter S
6-10	Block Length	5	Digits
11-15	Record Length	5	Digits
16-50	(Reserved for System Software Use)	35	a-characters
51-52	Buffer Offset Length	2	Digits
53-80	(Reserved for Future Standardization)	28	SPACEs

The Second File Header Label (HDR2) shall describe the physical characteristics of the file.

4.4.1 Record Formats (HDR2 CP 5)

This field shall identify the format of all the records of the file as follows :

- Capital letter F, meaning fixed-length records
- Capital letter D, meaning variable-length records
- Capital letter S, meaning spanned records

4.4.2 Block length (HDR2 CP 6-10)

This field shall specify the maximum number of characters per block of the file. It shall not be in contradiction with the hardware standard of the medium used.

4.4.3 Record length (HDR2 CP 11-15)

This field shall specify the length in characters of the record in the file. Its meaning depends on the contents of the Record Format field (HDR2 CP 5) as follows :

- F (fixed length) :   the length of the record.
- D (variable length) :   the maximum length of the record including the count field.
- S (spanned) :   the maximum record length, excluding the segment control word(s). In this case 00000 shall mean that the maximum length may be greater than 99999 characters.

4.4.4 Buffer Offset Length (HDR2 CP 51-52)

This field shall specify the length in characters of any additional field inserted before the first record in a data block.

4.5 First End of Volume Label (System 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)	50	Same as the corresponding fields in HDR1

CP	Field name	L	Content
55-60	Block Count	6	Digits
61-80	(Same as the corresponding fields in HDR1)	20	Same as the corresponding fields in HDR1

The First End of Volume Label (EOV1) shall be used when the last file in the volume extends over the end of the volume. It shall contain the same data as the HDR1 Label, with the exception of the Label Identifier field (CP 1-3) and the Block Count field (CP 55-60).

4.5.1 Block Count (EOV1 CP 55-60)

This field shall indicate the number of data blocks since the preceding Beginning of File Section Label Group. This count excludes label blocks and tape mark blocks.

4.6 Second End of Volume Label (System Label EOV2)

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)	76	Same as the corresponding fields in HDR2

4.7 First End of File Label (System 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)	50	Same as the corresponding fields in HDR1

CP	Field name	L	Content
55-60	Block Count	6	Digits
61-80	(Same as the corres-ponding fields in HDR1)	20	Same as the corresponding fields in HDR1

The First End of File Label (EOF1) shall contain the same data as the First File Header Label, with the exception of the Label Identifier field (CP 1-3) and the Block Count field (CP 55-60).

4.7.1 Block Count (EOF1 CP 55-60)

This field shall indicate the number of data blocks since the preceding Beginning of File Section Label Group . This count excludes label blocks and data blocks.

4.8 Second End of File Label (System Label EOF2)

CP	Field name	L	Content
1-3	Label Identifier	3	EOF
4	Label Number	1	2
5-80	(Same as the corres-ponding fields in HDR2)	76	Same as the corresponding fields in HDR2

4.9 Other System Labels (HDR3 - HDR9, EOV3 - EOV9, EOF3 - EOF9)

CP	Field name	L	Content
1-3	Label Identifier	3	HDR or EOV or EOF as appropriate
4	Label Number	1	Digit 3 to 9
5-80	(Reserved for System Software Use)	76	a-characters

4.10 User Volume Labels (UVLn)

CP	Field name	L	Content
1-3	Label Identifier	3	UVL
4	Label Number	1	Digit 1 to 9
5-80	(Reserved for User Application)	76	a-characters

4.11 User File Labels (UHLa, UTLa)

CP	Field name	L	Content
1-3	Label Identifier	3	UHL or UTL
4	Label Number	1	a-character
5-80	(Reserved for User Application)	76	a-characters

5 PROCESSING OF LABEL FIELDS

5.1 Fields in VOL1, HDR1, EOV1 and EOF1 labels

All fields shall be written with content as specified.

5.2 Fields in UVLn, HDR2 to HDR9, EOV2 to EOV9 and EOF2 to EOF9 labels

If the system software elects to write any of these labels, then the fields shall be written as specified in this Standard. If the system software elects to read any of these labels, 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 (for example compiled values) or after the processing has begun (for example system control statements), at the option of the system implementors. However, data found in VOL 1 is not overlaid or overridden.

5.4 Volume Header Label (VOL1)

The Volume Header Label shall be preserved, except as specified below. This does not preclude the rewriting of the label with the content 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 (HDR1 CP 22-27) : the identification shall be the same for all files of a set.

5.5.2 File Section Number (HDR1 CP 28-31) : the number of the first section of a file shall be 0001. This number shall be increased by 1 for each successive volume of the file.

5.5.3 File Sequence Number (HDR1 CP 32-35) : the File Sequence Number of the first file in a file set shall be 0001. This number shall be 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 shall contain the same number.

5.5.4 Generation Number (HDR1 CP 36-39) : the Generation Number of the first generation of a file shall be 0001. If subsequent generations of a file are noted, this number shall be increased by 1 for each successive generation of the file.

5.5.5 Generation Version Number (HDR1 CP 40-41) : the Generation Version Number of the first attempt to produce a generation of a file shall be 00. If the number of subsequent attempts to produce that generation of the file is noted, this number shall be increased by 1 at each subsequent attempt.

The Generation Version Number shall be reset to 00 when the Generation Number (HDR1 CP 36-39) is increased by 1.

5.5.6 Expiration Date (HDR1 CP 48-53) : a file shall be 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 shall 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 (HDR2 CP 6-10) : the count of this field shall include not only data but also buffer offset, Segment Control Words and padding. The actual maximum block capacity for data is thus reduced by the additional items, 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 (HDR2 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 shall be ignored.

5.6.3 Buffer Offset Length (HDR2 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 shall be specified in this field. If no such information is included, the contents of this field shall be 00.

6 ARRANGEMENTS OF LABELS AND DATA

6.1 Labels

A label shall be an 80 character block, the character positions (CP) of which are numbered 1 to 80. However, this may be extended by padding characters (see 9).
System Label Sets shall be symmetric about a file section. That is corresponding labels shall be recorded in each File Header Label Set, End of Volume Label Set and End of File Label Set for the entire file.
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 shall always be 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 the 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 the 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 that 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.

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

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 and a tape mark 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 (HDR1 CP 28-31) is increased by 1 (see 5.5.2).

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

If the end of a file and the end of a volume coincide then, unless the system avoids the situation by, for example, erasing the current block, three situations are possible :

- The end-of-tape marker is recognized while the system is writing the last data block of the file. In this case the system shall 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).

- The end-of-tape marker is recognized while 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 shall complete writing the End of File Label Group. The Beginning of File Label Group of the next file shall then be written, followed by an empty file section and the volume is then 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).

- The end-of-tape marker is recognized while 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 the beginning of a file and end-of-tape marker

If the end-of-tape marker is recognized while the system is writing the Beginning of File Label Group, then the system shall complete the writing of that group, terminate the volume by writing an empty file section followed by 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 shall be used to establish the file structure according to the rules illustrated in figures 1 to 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 Configuration of files

The various configurations of files that can be formed according to these rules are illustrated in figure 1. The rules and figures are presented including only VOL1, HDR1, EOV1 and EOF1. Additional labels, where used, shall be included as specified in 7.4.

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--*EOF

FIGURE 1 - Structure of magnetic tape files

7.1.1 As described in 6.11, first situation, the labelling configuration to be used when the end-of-tape marker is recognized while the last data block of the file is being written shall be as shown in figure 2.

```
    ---------- Last section of File A containing data ----*EOV1**
VOL1 HDR1** EOF1* HDR1* --------First section of File B ------(A)
    (A)      (A)     (B)
```

FIGURE 2 - Empty file section at beginning of volume

7.1.2 As described in 6.11, second situation, the labelling configu-
      ration to be used when the end-of-tape marker is recognized
      while the label handling routine is writing the End of File
      Label Group, shall be as shown in figure 3. There is one exception,
      described in 6.11, third situation,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 End of File Label Group
      followed by a double tape mark.

7.1.3 As described in 6.12, the labelling configuration resulting when
      the end-of-tape marker is recognized while the label handling
      routine is writing the Beginning of File Label Group, shall be 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 containing data -----
     (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 HDR1, EOV1 and EOF1 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 Grouping of labels

As described in 7.4.1, 7.4.2 and 7.4.3, labels shall be fitted
into the file structure without otherwise modifying the relation-
ship between VOL1, HDR1, EOV1, EOF1 labels and the files.
Illustrations of the order of these labels are given in 7.5.

7.4.1 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 Labels HDR2 - HDR9, EOV2 - EOV9 and EOF2 - EOF9, when used,
shall directly follow an HDR1, EOV1 or EOF1 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 Labels UHLa and UTLa, when used, shall directly follow the
labels as described in 7.4.2 and of the same label group. If,
however,such labels do not exist, labels UHLa and UTLa will
directly follow anHDR1, EOV1 or EOF1 label of the same label
group.

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 ... EOVn 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 begin-
ning 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 shall be an integral number of records in a block
for formats F and D. There shall be an integral number of segments
in a block for format S. Padded blocks are permitted (see 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. This 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) : Segment Control Word (SCW) shall
be included as the first five characters of each record
segment.

The first character of the SCW is called the spanning indi-
cator. This indicator shall have one of 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 sgements in a record. This does not prohibit the
system software from limiting the size of a work area avail-
able 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 includes the length
of the SCW, and 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 Not more than one of the record formats described in 8.1.1
to 8.1.3 shall be used in any one file.

8.1.5 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 identi-
fication 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 physical recording density.

8.3 Size of data blocks

The minimum and maximum size of data blocks are specified in the relevant recorded magnetic tape standard (see 2).

8.4 Character code in data

The data in each record and buffer offset field shall be written in any ECMA standard coded character set for information interchange.

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 the character UPWARD ARROW HEAD (position 5/14 of table 2 of Standard ECMA-6). To ensure that padding after fixed-length blocked records can be distinguished from valid records fixed-length records shall not consist entirely of padding characters.
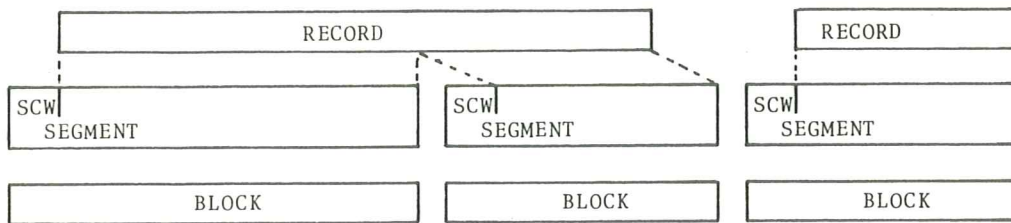
FIGURE 4 - Spanned records, unblocked
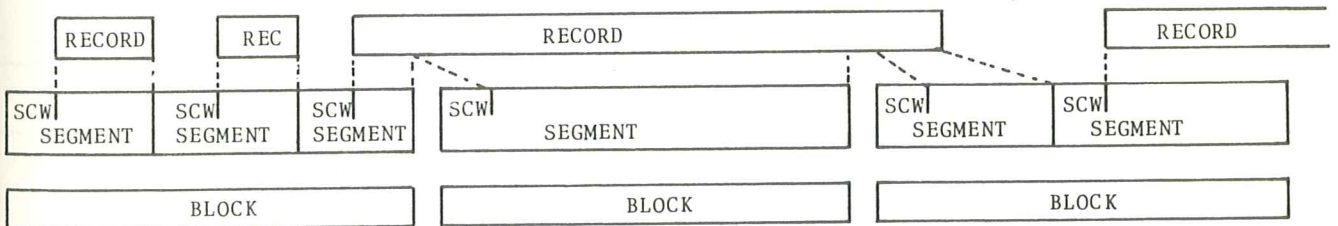
- first block showing the maximum block size

FIGURE 5 - Spanned records, blocked

- all blocks showing the maximum block size
- last record continues in next subsequent block

1	2048	DATA	2043 characters

2	2048	DATA	2043 characters

3	0160	DATA	155 characters

SCW

FIGURE 6 - One spanned record, unblocked.
Record length 4 241 characters (each line represents a block).

1	2048	DATA	2043 characters

2	2048	DATA	2043 characters

3	0150	DATA	145 char.	1	1898	DATA	1893 char

SCW   SCW

2	2048	DATA	2043 characters

3	2005	DATA	2000 characters

SCW

FIGURE 7 - Two spanned records, blocked.

Record 1 : record length 4 231 characters.
Record 2 : record length 5 936 characters (each line represents a block

NOTE : *It is not a requirement that blocks should all be of the
same size. In this figure the last block shown has been
truncated - a situation which could arise with the last
block of a file, or where it is essential to start the
following record at the beginning of a block.*

10 LEVELS OF LABELLING

The volume content as specified for each level described in this section acts as a ceiling on the volume content for each level. Note that UVL labels are not required at any of the levels described in 10.1 to 10.4.

10.1 Facilities available at level 1

10.1.1 File Sets : file sets shall be single-file single-volume or single-file multi-volume.

10.1.2 Labels : VOL1, HDR1, EOV1 and EOF1 shall appear on the volume. All other standard labels may appear, but a system shall ignore and bypass any additional labels it does not process.

10.1.3 Record/Block Structures : blocks shall consist of one or more fixed-length records.

10.1.4 Label fields : in a level 1 system, as a minimum the following basic label fields shall appear :

- VOL1

        Label Identifier
        Label Number
        Volume Identifier
        Accessibility
        Label Standard Version

- HDR1, EOV1, EOF1

        Label Identifier
        Label Number
        File Identifier
        File Section Number
        Expiration Date
        Block Count

However, all fields of the VOL1, HDR1, EOV1 and EOF1 labels shall contain meaningful information in accordance with this Standard. Therefore standard default values for certain fields are listed below. It should also be noted that fields which are Reserved for Future Standardization shall be filled with SPACEs.

Label	CP	Field Name	Default Value
VOL1	38 - 51	Owner Identifier	SPACEs
HDR1 EOV1 EOF1	22 - 27	File Set Identifier	SPACEs
	32 - 35	File Sequence Number	0001
	36 - 39	Generation Number	0001
	40 - 41	Generation Version Number	00
	42 - 47	Creation date	SPACE 00000
	54	Accessibility	SPACE
	61 - 73	System Code	SPACEs

10.2 Facilities available at level 2

10.2.1 File Sets : file sets shall be single-file single-volume, single-file multi-volume, multi-file single-volume, or multi-file multi-volume.

10.2.2 Labels : VOL1, HDR1, EOV1, and EOF1 shall appear on the volume. All other standard labels may appear, but a system ignores and bypasses any additional labels it does not process.

10.2.3 Record/Block Structures : blocks shall consist of one or more fixed-length records.

10.2.4 Label Fields : in a level 2 system, as a minimum the following basic label fields shall appear :

- VOL1
    Label Identifier
    Label Number
    Volume Identifier
    Accessibility
    Label-Standard Version

- HDR1, EOV1, EOF1
    Label Identifier
    Label Number
    File Identifier
    File Set Identifier
    File Section Number
    File Sequence Number
    Expiration Date
    Accessibility
    Block Count

However, all fields of the VOL1, HDR1, EOV1 and EOF1 labels must contain meaningful information in accordance with this Standard. Therefore standard default values for certain fields are listed below. It should also be noted that fields which are reserved for future standardization shall be filled with SPACEs.

Label	CP	Field Name	Default Value
VOL1	38 - 51	Owner Identifier	SPACEs
HDR1 EOV1 EOF1	36 - 39	Generation Number	0001
	40 - 41	Generation Version Number	00
	42 - 47	Creation Date	SPACE 00000
	61 - 73	System Code	SPACEs

10.3 Facilities available at level 3

10.3.1 File Sets : file sets shall be single-file single-volume, single-file multi-volume, multi-file single-volume, or multi-file multi-volume.

10.3.2 Labels : VOL1, HDR1, HDR2, EOV1, EOV2, EOF1, EOF2 shall appear on the volume. UHLa and UTLa may appear if the user specifies them. All other standard labels may appear but a system ignores and bypasses any additional labels it does not process.

10.3.3 Record/Block Structures : blocks shall consist of one or more fixed-length records, or one or more variable-length records.

10.3.4 Label Fields : in a level 3 system, as a minimum the following basic label fields shall appear :

- VOL1
        Label Identifier
        Label Number
        Volume Identifier
        Accessibility
        Label Standard Version

- HDR1, EOV1, EOF1
        Label Identifier
        Label Number
        File Identifier
        File Set Identifier
        File Section Number
        File Sequence Number
        Creation Date
        Expiration Date
        Accessibility
        Block Count

- HDR2, EOV2, EOF2
        Label Identifier
        Label Number
        Record Format
        Block Length
        Record Length
        Buffer Offset Length

However, all fields of the VOL1, HDR1, HDR2, EOV1, EOV2, EOF1 and EOF2 labels shall contain meaningful information in accordance with this Standard. Therefore standard default values for certain fields are listed below. It should also be noted that fields which are reserved for future standardization shall be filled with SPACEs.

Label	CP	Field Name	Default Value
VOL1	38 - 51	Owner Identifier	SPACEs
HDR1 EOV1 EOF1	36 - 39	Generation Number	0001
	40 - 41	Generation Version Number	00
	61 - 73	System Code	SPACEs

Label	CP	Field Name	Default Value
HDR2 EOV2 EOF2	16 - 50	Reserved for System Software Use	SPACEs

10.4 Facilities available at level 4

10.4.1 File Sets : file sets shall be single-file single-volume, single-file multi-volume, multi-file single-volume, or multi-file multi-volume.

10.4.2 Labels : VOL1, HDR1, HDR2, EOV1, EOV2, EOF1, EOF2 shall appear on the volume. UHLa and UTLa may appear if the user specifies them. All other standard labels may appear but a system ignores and bypasses any additional labels it does not process.

10.4.3 Record/Block Structures : blocks shall consist of one or more fixed-length records, one or more variable-length records, or one or more segments of spanned records.

10.4.4 Label Fields : in a level 4 system, as a minimum the following basic label fields shall appear :

- VOL1
        Label Identifier
        Label Number
        Volume Identifier
        Accessibility
        Label Standard Version

- HDR1 , EOV1, EOF1
        Label Identifier
        Label Number
        File Identifier
        File Set Identifier
        File Section Number
        File Sequence Number
        Generation Number
        Generation Version Number
        Creation Date
        Expiration Date
        Accessibility
        Block Count

- HDR2, EOV2, EOF2
        Label Identifier
        Label Number
        Record Format
        Block Length
        Record Length
        Buffer Offset Length

However, all fields of the VOL1, HDR1, HDR2, EOV1, EOV2, EOF1, EOF2 labels shall contain meaningful information in accordance with this Standard. Therefore standard default values for certain fields are listed below. It should also be noted that fields which are reserved for future standard- ization shall be filled with SPACEs.

Label	CP	Field Name	Default Value
VOL1	38 - 51	Owner Identifier	SPACEs
HDR1 EOV1 EOF1	61 - 73	System Code	SPACEs
HDR2 EOV2 EOF2	16 - 50	Reserved for System Software Use	SPACEs

10.5 Description of Media

10.5.1 Domain : this section 10 applies to the information recorded on a volume and not to the physical characteristics of the volume. For the purpose of this Standard, only the portion of the volume between the beginning-of-tape marker and the double tape mark following the End of File Section Label Group or the last End of File Label Group is assumed to exist. Any contents of the volume beyond the double tape mark shall not be considered part of the information in interchange and shall not be considered in determining the level of the volume.

10.5.2 Conditions : a volume (set) corresponds to a level under the following conditions :

- It and every file upon it contains all of the elements of labelling required at that level, and they are formatted and placed as specified in this Standard, and they contain an accurate description of the volume or the file to which they pertain, and

- if it or any file upon it contains any elements of labelling optional at that level, then they are formatted and placed as specified in this Standard, and contain an accurate description of the volume or file to which they pertain, and

- it and every file upon it contains only file sets or record/block formats permitted at that level, and

- it and every file upon it does not contain any element that is at variance with this Standard, or any extension to elements of labelling defined in this Standard.

APPENDIX A

NOTES ON IMPLEMENTATION

A.1 INTRODUCTION

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

A.2 SYSTEM AND USER LABELS

A.2.1 System labels

The following system labels shall always be required :

VOL1
HDR1
EOV1
EOF1

As further defined in Appendix B (Levels of Labelling), the following system labels may be required or may be optional :

HDR2
EOV2
EOF2

The remaining system labels are always optional :

HDR3 - HDR9
EOV3 - EOV9
EOF3 - EOF9

When used, these optional system 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 system software.

The contents of these optional system labels can be ignored in interchange. The writing of these labels is optional, but where used they shall be written in conformity with the requirement of this Standard in relation to content and position. There is no requirement under this Standard to take action on the label content.

In all system labels, the actual interpretation and treatment of those fields designated "Reserved for System Software Use" can be the subject of special agreement between the interchange parties.

A.2.2 User Labels

The following user labels fall into two categories.

The first category consists of UVL1 - UVL9.

These labels are provided for the retention of data relative to the volume of magnetic tape. Such data as security inform-ation, purchase date, vendor, date of last certification, department assignment, etc., are illustrative of the type of items which can be contained in these labels. All of these data are parochial to a specific installation, or possibly a group of installations, and are to be ignored in general interchange.

User Volume Labels are intended primarily for installation use, but may be used for application data by agreement of the inter-change parties.

Label handling routines provide only the Label Identifier and Number fields of User Volume Labels on output, and recognize these fields on input. The remainder of these labels (CP5-80) are furnished to the label handling routines on output. On in-put, the user volume labels are passed by the label handling routines to the invoking facility.

The second category consists of UHLa and UTLa.

These labels are handled partly by the label handling routines of the system, to the extent that they are recognized and passed on to the user on input, and are written to the tape on output.

Users provide the information in user labels UHLa and UTLa beyond the Label Identifier (CP 1 - 3) on output and utilize it on input.

When these labels are used, they shall be written in conformity with the requirement of this 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 follwing 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 HDR 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 next block.

If EOV1 : rewind, alternate, verify HDR1 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 an empty 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 EOV1 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 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 (VOL1 CP 11) : this field is expected to refer to such categories of information as company confidential, proprietary, etc. This field is not intended to fulfil 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 (VOL1 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 (VOL1 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 digits in this field,
rather than letters, to the extent of possible.

A.4.3 First File Header Label (HDR1)

A.4.3.1 File Set Identifier (HDR1 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 (HDR1 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 of that
file will be numbered sequentially on subsequent volumes.

A.4.3.3 Generation Version Number (HDR1 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 (HDR1 CP 54)   : see A.4.2.1.

A.4.3.5 System Code (HDR1 CP 61-73) : on output this field is
inserted by the label handling routines to identify the
System Software that created the file. This field indicates
the system that may have supplied Reserved for System Soft-
ware Use (HDR2 CP 16-50), Other System Labels (HDR3 - HDR9,
EOV3 -   EOV9, EOF3 - EOF9), Buffer Offset Length (HDR2 CP 51-52)
and/or buffer offset. 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, however,
in the interim period, implementors are recommended to define
their own system codes.

The use of these facilities without identifying the system may
result in loss of ability to interchange system data effectively.

A.4.4 Second File Header Label (HDR2)

A.4.4.1 Record Format (HDR2 CP 5) : implementors of computing and
operating systems may find that they have a requirement to
accommodate records in a format that does not conform to one
specified in this Standard. If such records are required by
a system, they may be incorporated into a parochial version
of this standard and appear only on volumes that are identified
as other than this version of the standard in Label Standard
Version (VOL1 CP 80). If volumes containing such a parochial
record format are interchanged, it is by agreement between
the interchange parties.

A.4.4.2 Record Length (HDR2 CP 11-15) : the content of this field is
undefined if Record Format (HDR2 CP 5) is other than F, D
or S. Any other value is outside this version of the Standard
(see Label Standard Version VOL1 CP 80).

A.4.4.3 Reserved for System Software Use (HDR2 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 (EOV1 CP 55-60 or EOF1 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.

APPENDIX B

LEVELS OF LABELLING

B.1 INTRODUCTION

In practice interchange of file sets between systems would not be easy when many stages of implementation are possible. The number of stages is reduced to an optimum by the definition of a small number of levels. By limiting the facilities used in interchange to those of a standard level the risk of a mismatch of facilities used will be much reduced, as also will be the amount of information provided external to the volumes interchanged.

This Appendix describes recommended methods for processing volumes that correspond to a level of this Standard to ensure proper treatment and understanding of these volumes and their contents in information interchange.

It is not intended that every implementation should necessarily include all described provisions; however, each implementation should be able to produce and accept volumes that correspond to a level selected by the implementors.

It is intended that any volume (set) can be processed correctly by any system designed for an equal or higher level; and any system can process correctly any volume (set) of equal or lower level.

The processing of labels under such a system of levels is proposed as an Appendix to the Standard since it can only be advisory, and since the present Standard is only concerned with standardization of the media.

In order to facilitate interchange of information among systems of dissimilar capabilities, the concept of levels of labelling is established. The level applies to the minimum facilities available in a system and to the maximum facilities resulting in the file set. Thus a system can limit the facilities it uses and produce a file set equivalent to that which can be produced or read by a system of lower level.

There are four levels numbered 1, 2, 3 and 4 in increasing order of complexity. The levels are so defined that systems software can process correctly all file sets in which the labelling level used is equal to or less than the system's nominated highest level. This applies equally to both reading and writing. Such levels act as floors on system capabilities.

Within a given level certain label fields are considered to be a basic part of that level. However, so as to avoid difficulties when a file set conforming to that level is read by systems software capable of handling a higher level, all label fields must contain meaningful information in accordance with this standard. Therefore, to assist the systems software writing the file set, default values have been specified which meet this criterion.

Standard labels and user labels not included in a given level may be present but may be ignored or by-passed by the systems software. Thus a HDR2 could be present in a level 1 set but would be considered as outside the interchange. However, those facilities of levels concerned with file formats and record formats cannot be enhanced in the same way without moving to a higher level since the resulting file set may be unreadable by the receiving system - for example, the addition of spanned records to a level 1 set as an optional facility requires that the option of providing HDR2 also be added to level 1, since the receiving system might not otherwise be able to properly handle the spanned records. Therefore, since the selection of such special optional facilities at a given level may result in a reduction of general interchangeability, it is recommended that whenever possible no such optional facilities should be added to the defined levels.

There are three main identifiable constituents of a level - the file formats, the labels used and the record format. The most useful options of these are as follows :

Constituent	Option	
FILES	1	Single-file single-volume + single-file multi-volume
	2	Single-file single-volume + single-file multi-volume
		+ multi-file single-volume + multi-file multi-volume
LABELS	1	VOL1 HDR1 EOV1 EOF1
	2	VOL1 HDR1 HDR2 EOV1 EOV2 EOF1 EOF2
RECORD FORMAT	1	Fixed
	2	Fixed + variable
	3	Fixed + variable + spanned

Levels of labelling founded on the content of the media are based on these constituents, and to enable a higher level system to read volumes produced by a lower level system the levels are nested.

These are two logics operating in these constituents. For maximum interchangeability it is essential that though a specific set of labels be defined for a given level, additional labels are allowed (though no guarantee of processing at the receiving end is given). However, the same approach cannot be applied to record format. If fixed length records are defined for level 1 then variable length formats should not be added as an option since

the receiving system might not be able to handle them.

This leads to the following four levels made up of the constituents as follows :

Level	Files	Labels	Record format
1	1	1	1
2	2	1	1
3	2	2	2
4	2	2	3

Level 1 is a minimum system.

Level 2 is effectively level 1 with multi-file file sets added.

Level 3 is equivalent to the basic and most used facilities of Standard ECMA-13 version 1, including variable length records and HDR2, EOV2 and EOF2.

Level 4 is effectively level 3 with spanned records added.

In these discussions a distinction is presumed between the operations necessary to assure that the content of a label is valid, and the operations necessary to process labels. The former are not described except to recommend responsibility for the checking of validity, either by the system or by the writer of operating system control statements or common programming language declarations. Examples of such checks are range checks on the values, alphabetic value checks of numeric-only fields, left justification of fields defined to be left justified, and any other checks that are feasible for the responsible party to perform. While system performance of such checks is, in reality, a form of processing of the fields being checked, that kind of processing is not to be confused with the specified recommended processing of each field defined in the standard.

The following terms are used below in referring to the four levels of labelling systems :

required facility : Format or function of the labelling system defined as constituting a level.

optional facility : Format or function of the labelling system not defined as required for that level but defined elsewhere in this Standard.

extension : Any format or function of a labelling system not defined in this Standard.

B.2.1 Level 1

B.2.1.1 Validity of labels : in a level 1 system, it is recommended that the system exercises responsibility for the format and content of VOL1 label fields : Label Identifier (VOL1 CP 1-3),

Label Number (VOL1 CP 4), Volume Identifier(VOL1 CP 5-10), Accessibility (VOL1 CP 11), and Label-Standard Version (VOL1 CP 80).

In a level 1 system, it is recommended that the writer of operating system control statements (e.g., Job Control Statements) or common programming language declaration exercises full responsibility for format, editing, and accuracy of all file system label fields except File Section Number (HDR1 CP 28-31) and Block Count (HDR1 CP 55-60). It is recommended that a system exercises responsibility for format, editing, and accuracy of File Section Number (HDR1 CP 28-31) and Block Count (HDR1 CP 55-60). A system is not required to check or edit any other fields in any manner, and may simply copy the fields from information supplied by the writer of operating system control statements or common programming language declarations.

B.2.1.2 File Security : a level 1 system assumes that the volume is immediately accessible if Accessibility (VOL1 CP 11) is equal to SPACE. If the contents are not SPACE, access is denied unless the system provides additional controls. Such additional controls are not specified in this standard.

Additional controls upon Accessibility (HDR1 CP 54) are optional.

B.2.2 Level 2

B.2.2.1 Validity of Labels : in a level 2 system, it is recommended that the system exercises responsibility for the format and content of VOL1 label fields : Label Identifier (VOL1 CP 1-3), Label Number (VOL1 CP 4), Volume Identifier (VOL1 CP 5-10), Accessibility (VOL1 CP 11), and Label-Standard Version (VOL1 CP 80).

In a level 2 system, it is recommended that the writer of operating system control statements or common programming language declarations exercises responsibility for format, editing, and accuracy of all file system label fields except File Set Identifier (HDR1 CP 22-27), File Section Number (HDR1 CP 28-31), File Sequence Number (HDR1 CP 32-35) and Block Count (HDR1 CP 55-60). It is recommended that a system exercises responsibility for format, editing, and accuracy of File Set Identifier (HDR1 CP 22-27), File Section Number (HDR1 CP 28-31), File Sequence Number (HDR1 CP 32-35), and Block Count(HDR1 CP 55-60). A system is not required to check or edit any other fields in any manner and may simply copy the fields from information supplied by the writer of operating system control statements or common programming language declarations.

B.2.2.2 File Security : a level 2 system assumes that a file is immediately accessible if both volume Accessibility (VOL1 CP 11) and file Accessibility (HDR1 CP 54) are equal to SPACEs. If the contents of either or both fields are not SPACE, access is denied unless the level 2 system provides additional

controls. Such additional controls are not specified in this Standard.

B.2.3 Level 3

B.2.3.1 Validity of Labels : in a level 3 system, it is recommended that the system exercises responsibility for the format and content of VOL1 label fields : Label Identifier (VOL1 CP 1-3), Label Number (VOL1 CP 4), Volume Identifier (VOL1 CP 5-10), Accessibility (VOL1 CP 11), and Label Standard Version (VOL1 CP 80).

In a level 3 system, it is recommended that the writer of operating system control statements or common programmming language declarations exercises responsibility for format, editing, and accuracy of all file system label fields except Label Identifier (HDR1 CP 1-3), Label Number (HDR1 CP 4), File Set Identifier (HDR1 CP 22-27), File Section Number (HDR1 CP 28-31), File Sequence Number (HDR1 CP 32-35), and Block Count (HDR1 CP 55-60). It is recommended that a system exercises responsibility for format, editing, and accuracy of Label Identifier (HDR1 CP 1-3), Label Number (HDR1 CP 4), File Set Identifier (HDR1 CP 22-27), File Section Number (HDR1 CP 28-31), File Sequence Number (HDR1 CP 32-35), and Block Count (HDR1 CP 55-60).

In a level 3 system, it is recommended that the system exercises responsibility for the format and content of HDR2, EOV2, EOF2 label fields : Label Identifier (CP 1-3), Label Number (CP 4), Record Format (CP 5), Block Length (CP 6-10), Record Length (CP 11-15) and Buffer Offset Length (CP 51-52).

B.2.3.2 File Security : a level 3 system assumes that a file is immediately accessible if both volume Accessibility (VOL1 CP 11) and file Accessibility (HDR1 CP 54) are equal to SPACE. If the contents of either or both fields are not SPACEs, access is denied unless the level 3 system provides additional controls. Such additional controls are not specified in this standard.

B.2.4 Level 4

B.2.4.1 Validity of Labels : in a level 4 system, it is recommended that the system exercises responsibility for the format and content of all system label fields in VOL1, HDR1, HDR2, EOV1, EOV2, EOF1 and EOF2. Operating system control statements and common programming language declarations are edited to ensure that the labels are properly formatted and are self-consistent. On output, the file is constructed to be consistent with the label declarations. On input the system is responsible for interpreting the labels and ensuring the self-consistency of the file.

B.2.4.2 File Security : a level 4 system assumes that a file is immediately accessible if both volume Accessibility (VOL1 CP 11) and file Accessibility (HDR1 CP 54) are equal to a SPACE. If

the contents of either or both fields are not SPACEs, access
is denied unless the level 4 system provides additional
controls. Such additional controls are not specified in this
standard.

B.2.5 Responsibility for implementation of levels

B.2.5.1 System Labels : VOL1, HDR1 - HDR9, EOV1 - EOV9, EOF1 - EOF9
are processed entirely by the label handling routines. The
label handling routines write and read those labels and use
the information in those labels.

Facilities for second system labels (HDR2, EOV2, EOF2) and for
other system labels (HDR3 - HDR9, EOV3 - EOV9, EOF3 - EOF9)
are optional at levels 1 and 2. Facilities for other system
labels (HDR3 - HDR9, EOV3 - EOV9, EOF3 - EOF9) are optional
at level 3 and 4.

B.2.5.2 User Volume Labels : UVL1 - UVL 9 are processed partly by the
label handling routines to the extent that they are accepted
from an installation label processing routine and written on
a volume on output; and read from a volume, recognized, and
passed to an installation label processing routine on input.
The label handling routines supply Label Identifier (UVLn CP 1-3)
and Label Number (UVLn CP 4). The installation routine provides
the information in the remainder of these labels (UVLn CP 5-80)
on output and utilizes it on input. Interfaces to permit
transfer of control between the label handling routines and the
resident installation routine are provided.

Alternatively the UVL can be used to contain application
program information. This use, which is expected to be compara-
tively rare, can be implemented by agreement between the inter-
change parties.

These facilities are optional at each level.

B.2.5.3 User File Labels : UHLa and UTLa are processed partly by the
label handling routines, to the extent that they are accepted
from user application object program routines and written on
a volume on output; and read from a volume, recognized, and
passed to user application object program routines on input.
The label handling routines supply Label Identifier (UHLa CP 1-3
or UTLa CP 1-3). The user routines provide Label Number (UHLa
CP 4 or UTLa CP 4) and the information in the remainder of these
labels (UHLa CP 5-80 or UTLa CP 5-80) on output and utilize it on
input. Interfaces to permit transfer of control between the label
handling routines and user routines are provided.

These facilities are optional at level 1 and level 2.

B.2.5.4 Output Options : selection of the level of labelling and pro-
vision of the facility for creation of labels optional at that
level, or fields not required at that level, is at the option
of the system implementors. Invocation of the facilities for
creation of user labels is at the option of the user, if such

facilities are provided. Invocation of the facilities for
creation of other system labels (HDR3 - HDR9, EOV3 - EOV9,
EOF3 - EOF9) is at the option of the system.

B.2.5.5 Input Options : provision of the facility for processing of
labels optional at that level, or fields not required at that
level, is at the option of system implementors. Invocation of
the facilities for accessing of user labels is at the option
of the user, if such facilities are provided. Invocation of
the facilities for processing of other system labels (HDR3-
HDR9, EOV3 - EOV9, EOF3 - EOF9) is at the option of the system.
Labels for which facilities are not provided or if provided
are not invoked, are bypassed and the information ordinarily
obtained from them is lost.

B.2.5.6 Non-standard facilities : at the option of system implementors,
an implementation of label and file processing functions may
provide a capability for a user to cause to be omitted from a
volume an element that is required, or to cause to be included
upon a volume an element that is at variance with this standard,
or an extension to elements of labelling designed in this
Standard. If a user opts for such an omission, variance, or
extension, then the resulting volume is non-standard.

APPENDIX C

VERSIONS OF THIS STANDARD

C.1 IDENTIFICATION OF THE VERSIONS

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 this Standard (November 1967). When the Label Standard Version has the value 2 it means that the volume was written according to the rules and options of the second version of this Standard (August 1973). When it has the value 3 it means that the volume was written according to the rules and options of this third version (January 1978).

C.2 DIFFERENCES BETWEEN VERSIONS

The major differences between the second and third version are as follows :

- Definition of the levels of labelling. In the previous versions subsetting was neither implied nor precluded.

  In this version specifications for levels of labelled volumes, and descriptions of levels of label and file processing functions, are added.
  This provides explicit subsets, to avoid confusion as to the re-quirements of this version, and to ensure that products implementing only a subset can successfully interchange standard volumes with each other.
  Differences between this Standard and previous versions will vary depending upon the level of labelling considered. The whole Appendix B is devoted to the subject of levels and it is not dealt with further here.

- Block sequence indicator description deleted. In the previous versions the block sequence indicator description appeared in Annex B. In this version it does not appear.
  This is to clarify the actual requirements and to eliminate an option found to be unsatisfactory.

- Record Format (HDR2 CP 5) changed. In the previous versions, U (undefined) appeared. In this version, it does not appear.
  This is because it is not feasible to utilize undefined format records in general interchange. Parochial formats can only be used in conjunction with a Label Standard Version (VOL1 CP 80) other than 3.

- Block length limits specified indirectly. In the previous version, specific maximum block length was specified. In this version a cross reference to another standard appear. Maximum and minimum block length limits appear in this other standard.
  This is to avoid redundant, possibly inconsistent specifications.