American National Standard

Adopted for Use by the Federal Government

FIPS PUB 79

See Notice on Inside Front Cover

magnetic tape labels and file structure for information interchange

ANSI X3.27-1978

american national standards institute, inc.
1430 broadway, new york, new york 10018

This standard has been adopted for Federal Government use.

Details concerning its use within the Federal Government are contained in FIPS PUB 79, Magnetic Tape Labels and File Structure for Information Interchange. For a complete list of the publications available in the Federal Information Processing Standards Series, write to the Office of Standards Administration, Institute for Computer Sciences and Technology, National Bureau of Standards, Washington, D.C. 20234.

ANSI®
X3.27-1978
Revision of
X3.27-1969

American National Standard Magnetic Tape Labels and File Structure for Information Interchange

Secretariat

Computer and Business Equipment Manufacturers Association

Approved April 18, 1977

American National Standards Institute, Inc

Abstract

The requirements are specified for the labeling and block structure recorded on magnetic tape and for the data processing systems that process the labels and blocks. The arrangement of labels and data on the magnetic tape volume, the data code, block and record formats, and the padding to be used are included. Processing requirements for systems that utilize the standard are stated. Four well-nested subsets (levels) of the requirements on the media content and on the supporting systems are declared.

The standard is primarily directed toward information interchange between dissimilar computing systems, but consideration is afforded to local tape processing needs. Such parochial requirements are recognized and considered, with the intent that the standard serve as a guide for the designers of systems that also support a local environment that is more specialized than the interchange environment.

American National Standard

An American National Standard implies a consensus of those substantially concerned with its scope and provisions. An American National Standard is intended as a guide to aid the manufacturer, the consumer, and the general public. The existence of an American National Standard does not in any respect preclude anyone, whether he has approved the standard or not, from manufacturing, marketing, purchasing, or using products, processes, or procedures not conforming to the standard. American National Standards are subject to periodic review and users are cautioned to obtain the latest editions.

CAUTION NOTICE: This American National Standard may be revised or withdrawn at any time. The procedures of the American National Standards Institute require that action be taken to reaffirm, revise, or withdraw this standard no later than five years from the date of publication. Purchasers of American National Standards may receive current information on all standards by calling or writing the American National Standards Institute.

Published by

American National Standards Institute
1430 Broadway, New York, New York 10018

Printed in the United States of America

P2M978/1050

Foreword

(This Foreword is not a part of American National Standard Magnetic Tape Labels and File Structure for Information Interchange, ANSI X3.27-1978.)

The aim of this standard is to reduce the difficulty of interchange of information recorded on magnetic tapes between different users and different computing systems. This aim is accomplished by specifying the format, content, and arrangement of magnetically recorded labels to identify and structure files, and by specifying the format and arrangement 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 this standard there is a general-purpose operating system present, but in other cases there may be installation or user supplied label and file processing functions that may form part of a special-purpose operating system. For proper implementation of this standard, it is expected that the installation or user supplied label and file processing functions will provide the required facilities within the area defined by this standard.

This standard contains specifications for four levels of labeling. 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.

The capabilities for information interchange are more readily used if the implementor of a system providing these capabilities identifies the product, the level, maximum record or block size accommodated, other optional alternatives available, and how these capabilities may be exercised.

This edition of the standard differs technically from ANSI X3.27-1969, particularly in the introduction of the record spanning technique, which allows a record to span more than one block and even more than one volume. Processing requirements are included to ensure proper treatment and understanding of volumes and their contents in information interchange. Specifications for levels of labeling have been added.

A detailed description of these differences is given in Appendix C. together with reasons for making the changes included in this revision.

In addition, this standard has been subjected to a major editorial revision.

Throughout this standard, the usage of American National Standard Code for Information Interchange, ANSI X3.4-1977 (ASCII), is implied.

Suggestions for improvement of this standard will be welcome. They should be sent to the American National Standards Institute, 1430 Broadway, New York, N.Y. 10018.

This standard was processed and approved for submittal to ANSI by American National Standards Committee on Computers and Information Processing, X3. Committee approval of the standard does not necessarily imply that all members voted for its approval. At the time it approved this standard, the X3 Committee had the following members:

J. F. Auwaerter, Chairman
R. M. Brown, Vice-Chairman
W. F. Hanrahan, Secretary

Organization Represented	*Name of Representative*
Addressograph Multigraph Corporation	A. C. Brown
	D. D. Anderson (Alt)
Air Transport Association	F. C. White
	C. Hart (Alt)
American Library Association	J. R. Rizzolo
	J. C. Kountz (Alt)
	M. S. Malinconico (Alt)

Organization Represented	*Name of Representative*
American Nuclear Society	D. R. Vondy
	M. K. Butler (Alt)
American Society for Information Science	M. C. Kepplinger
Association for Computing Machinery	P. Skelly
	J. A. N. Lee (Alt)
	H. E. Thiess (Alt)
Association for Educational Data Systems	A. K. Swanson
Association of American Railroads	R. A. Petrash
Association of Computer Programmers and Analysts	L. A. Ruh
	M. A. Morris, Jr (Alt)
Association of Data Processing Service Organizations	J. B. Christiansen
Burroughs Corporation	E. Lohse
	J. S. Foley (Alt)
	J. F. Kalbach (Alt)
California Computer Products, Inc	R. C. Derby
Computer and Communications Industry Association	N. J. Ream
	A. G. W. Biddle (Alt)
Control Data Corporation	C. E. Cooper
	G. I. Williams (Alt)
Data Processing Management Association	A. E. Dubnow
	D. W. Sanford (Alt)
Datapoint Corporation	V. D. Poor
	H. W. Swanson (Alt)
Digital Equipment Corporation	P. W. White
	A. R. Kent (Alt)
Edison Electric Institute	S. P. Shrivastava
	J. L. Weiser (Alt)
General Services Administration	D. L. Shoemaker
	M. L. Burris (Alt)
GUIDE International	T. E. Wiese
	B. R. Nelson (Alt)
	D. Stanford (Alt)
Honeywell Information Systems, Inc	T. J. McNamara
	E. H. Clamons (Alt)
Institute of Electrical and Electronics Engineers, Communications Society	(Representation Vacant)
Institute of Electrical and Electronics Engineers, Computer Society	R. L. Curtis
	T. Feng (Alt)
International Business Machines Corporation	L. Robinson
	W. F. McClelland (Alt)
Joint Users Group	T. E. Wiese
	L. Rodgers (Alt)
Life Office Management Association	R. E. Ricketts
	J. F. Foley (Alt)
Litton Industries	I. Danowitz
National Association of State Information Systems	G. H. Roehm
	J. L. Lewis (Alt)
	G. I. Theis (Alt)
National Bureau of Standards	H. S. White, Jr
	R. E. Rountree (Alt)
National Communications System	M. L. Cain
	G. W. White (Alt)
National Machine Tool Builders' Association	O. A. Rodriques
NCR Corporation	R. J. Mindlin
	A. R. Daniels (Alt)
	T. W. Kern (Alt)
Olivetti Corporation of America	E. J. Almquist
Printing Industries of America	N. Scharpf
	E. Masten (Alt)
Recognition Equipment, Inc	H. F. Schantz
	W. E. Viering (Alt)
Sanders Associates, Inc	A. L. Goldstein
	T. H. Buchert (Alt)
Scientific Apparatus Makers Association	A. Savitsky
	J. E. French (Alt)
SHARE Inc	T. B. Steel, Jr
	E. Brubaker (Alt)
Society of Certified Data Processors	T. M. Kurihara
	A. E. Dubnow (Alt)

Organization Represented	*Name of Representative*
Telephone Group	V. N. Vaughan, Jr
	J. C. Nelson (Alt)
	S. M. Garland (Alt)
3M Company	R. C. Smith
UNIVAC, Division of Sperry Rand Corporation	M. W. Bass
	C. D. Card (Alt)
U.S. Department of Defense	W. L. McGreer
	W. C. Rinehuls (Alt)
	W. B. Robertson (Alt)
U.S. Department of Health, Education, and Welfare	W. R. McPherson, Jr
	M. A. Evans (Alt)
VIM (CDC 6000 Users Group)	S. W. White
	M. R. Speers (Alt)
Xerox Corporation	J. L. Wheeler
	A. R. Machell (Alt)

Subcommittee X3L5 on Data Formats and Labeling, which developed this standard, had the following members:

J. G. Smith, Vice-Chairperson
J. N. Dunn, Former Chairperson
W. D. Kasner, Former Chairperson

S. E. Blount
R. L. Brandt
M. Burris
M. Davis
N. Eiden
S. W. Godsell
F. E. Holberton
J. Kent
H. G. Madsen
C. J. Masterson, Jr
H. S. Meltzer
G. Rittenburg
G. J. Rosenkrantz
D. L. Shoemaker
W. Unke

Contents

SECTION PAGE

SECTION PAGE

American National Standard
Magnetic Tape Labels
and File Structure
for Information Interchange

1. Scope and Field of Application

1.1 Scope

1.1.1 This standard establishes four levels of label formats, blocking structure, and tape-mark relationships on magnetically recorded tapes so that these volumes can be used for information interchange.

1.1.2 This standard contains specifications 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.

1.1.3 It is not the intention of this standard that every instance of its implementation should necessarily include all of its provisions; however, to accommodate this standard, each implementation is expected to be able to produce and accept volumes that correspond to a level selected by the implementors.

1.1.4 Any volume (set) can be processed correctly by any system of equal or higher level; and any system can process correctly any volume (set) of equal or lower level.

1.2 Field of Application

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

1.2.2 If the first record on a volume is 80 or more characters in length, and contains the ASCII characters "VOL1" in character positions 1-4, and the ASCII digit "3" in character position 80, then the volume is recognized as conforming to this standard, and is subsequently processed according to this standard. If the first record on a volume fails any of these conditions, the volume is either unlabeled (but not unrecorded), or it is labeled as being not in conformance with this standard, and need not be processed according to this standard.

1.2.3 Failure of a volume to conform to this standard may result in loss of ability to interchange information effectively.

1.3 Exclusions

1.3.1 This standard assumes the existence of, but does not specify:

(1) A mechanism for recording labels and files upon a volume and reading labels and files from a volume; a code in which the characters of the labels and files are represented, and the representation of such a code upon a volume

(2) A mechanism for a system to obtain information from a program, operator, installation, or user, as appropriate, to initialize, create, or verify labels (for example, systems generation, operating system control statements, utility program control statements, common programming language declarations, messages from a terminal, etc)

(3) A mechanism for a system to communicate information to a program, operator, installation, or user, as appropriate, with respect to errors or unusual conditions

(4) A mechanism for a program, operator, installation, or user, as appropriate, to choose among the alternatives the system makes available

(5) A mechanism for a program, operator, installation, or user, as appropriate, to invoke a facility a system makes available

(6) A mechanism to pass control to an installation volume label processing routine to process User Volume Labels, or to an application program routine to process user file labels

(7) Common programming language declarations and processing statements to define the attributes of output files, to identify input files, to process user file labels, and to process records within the file

(8) A method for an originator to communicate with a recipient a description of the file, including, but not limited to:

(a) The identification of the file

(b) The content, layout, and format of user labels

(c) The content, layout, and format of the fields in each of the record types

(d) The indicator for each of the record types

(e) The keys for each of the records

(f) The structure or sequence of the records in the file

There are in existence a variety of programming language, industry-oriented, application-specific conventions that may apply.

(9) Standards for representation of numeric values, and of common data elements (such as dates and states)

1.3.2 This standard admits the existence of, but does not specify:

(1) A mechanism for an operator, systems administrator, security officer, or user, as appropriate, to be informed of errors and exceptional conditions, and to take corrective action

(2) A mechanism for an operator, systems administrator, or security officer, as appropriate, to override controls specified in this standard to avoid undue delay in processing

2. Relation to Parochial Practice

This standard does not conflict with parochial practice (for example, within an installation, an area, an industry group, or a user's group) that provides all, some, or none of these facilities. However, it is good practice to distinguish in the Volume-Header Label (VOL1) between a volume conforming to this standard and one conforming to a parochial practice (see 4.3, Label-Standard Version (VOL1 CP 80)).

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 a related standard 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
record. Related data treated as a unit of information. **Examples.** In the context of business data; a transaction record, a customer's account record. NOTE 1: The delineation of a record may be arbitrary and is determined by the designer of the information format. NOTE 2: A record may be recorded in all or part of a block or in more than one block.	**block.** A collection of characters written or read as a unit. NOTE 1: A block may contain one or more complete records. NOTE 2: A block may contain segments of one or more spanned records. A single block does not contain multiple segments of the same spanned record. NOTE 3: For the purposes of this standard, blocks are separated by an interblock gap.
file. A collection of information consisting of records pertaining to a single subject. **Examples.** In the context of business data; a payroll file, an inventory file. NOTE 1: The description, content or organization of a file may be arbitrary. NOTE 2: 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, that is, a reel of magnetic tape. NOTE 1: A volume may contain part of a file, a complete file, or more than one file. NOTE 2: A volume may contain sections of one or more files. A volume does not contain multiple sections of the same file.
file section. That part of a file that is recorded on any one volume. NOTE: The sections of a file do 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 one and only one file set is recorded.
unspanned record. A record contained in a file in which each record by design ends in the same 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. NOTE: 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.	

Logical	Physical
record segment. That part of a spanned record that is contained in any one block. NOTE: The segments of a record do not have segments of other records interspersed.	
unblocked record. A record contained in a file in which each block by design contains only one record or record segment.	
blocked record. A record contained 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 records by design have the same length.	
variable-length record. A record contained in a file in which the records may have different lengths.	

application program. For the purposes of this standard, a program that processes the contents of records and, under control of an operating system, uses the label and file processing functions of that system.

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-Section Label Group and the tape mark preceding the End-of-File-Section Label Group or End-of-File Label Group of that file section or file.

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

label and file processing functions. A set of routines that read, write and process labels; and that read and write files. Label and file processing functions are an integral part of a system's software, that may be an operating system provided by a supplier.

NOTE: When an installation or user supplies the functions ordinarily associated with a general-purpose operating system, it is expected that the installation or user supplied label and file processing functions provide the required facilities within the area defined by this standard.

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.

label set. A collection of one or more contiguous labels with the same three initial characters (Label Identifier). See Table 1.

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 a general-purpose operating system, used in many installations, as is frequently the case when it is provided by a supplier; or it may be a special-purpose operating system, used in a single installation, and provided by or for that installation.

tape mark. A delimiter used to indicate the boundary between file data and label groups and also between certain label groups.

NOTE: The configuration of a tape mark is specified in the relevant recorded magnetic tape standard.

user. For the purposes of this standard, the creator of application programs.

Table 1
Classification of Labels

Label Group Name	Label Set Name	Label Identifier	Number
Beginning-of-Volume	Volume-Header	VOL	1
	User Volume	UVL	1 to 9
	File-Header	HDR	1 to 9
	User Header	UHL	a
Beginning-of-File or of File-Section	File-Header	HDR	1 to 9
	User Header	UHL	a
End-of First or of Intermediate File-Section	End-of-Volume	EOV	1 to 9
	User Trailer	UTL	a
End-of-File or of Last File-Section	End-of-File	EOF	1 to 9
	User Trailer	UTL	a

NOTES:
(1) The Beginning-of-Volume Label Group also includes a File-Header and may include a User Header Label Set.

(2) The meaning of "a" under Number is explained in 4.2.1.

4. Formats and Contents of Labels

4.1 General. In 4.3 through 4.12, the meanings of the headings are as follows:

CP: character position in the label
Field Name: reference name of the field
L: length of the field (number of characters)
Content: content of the field

Throughout this standard, references to contents of fields are in the form:

Field Name (LBLx CP yy-zz)

where LBL is the Label Identifier of the referenced label, x is the Label Number of that label and yy-zz are the first and last character positions of the field in that label.

4.2 General Characteristics of Label Fields

4.2.1 In this standard an "n" means any numeric character from 0 to 9. An "a" means any numeric, alphabetic or special character of the center four columns of the code table specified in ANSI X3.4-1977 except for position 5/15 and those positions where there is a provision for alternative graphic representation.

4.2.2 If a numeric (n) value is shorter than the field, then the value is right-adjusted and unused positions filled with zeros. If an alphanumeric (a) value is shorter than the field, then the value is left-adjusted and unused positions filled with spaces.

4.2.3 If any of the alphanumeric (a) fields, the name of which includes "Identifier" is all spaces, then functions of this labeling system may fail (see 8.2.4, 8.3.4, 8.4.4, and 8.5.4 for allowed use of spaces).

4.3 Volume-Header Label (VOL1)

CP	Field Name	L	Content
1 to 3	Label Identifier	3	VOL
4	Label Number	1	1
5 to 10	Volume Identifier	6	"a" characters. Assigned permanently by the owner to identify this volume.
11	Accessibility	1	"a" character. Indicates restrictions on access to the information in this volume. Space means no restrictions. The numeric characters and all special characters of the "a" set are reserved for operating system definition. "A" through "Z" are reserved for installation assignment.
12 to 37	Reserved for Future Standardization	26	Spaces.
38 to 51	Owner Identifier	14	"a" characters. Identifies the owner of the volume. Identifiers are not specified in this standard.
52 to 79	Reserved for Future Standardization	28	Spaces.
80	Label-Standard Version	1	Indicates the version of this standard to which the labels and data formats on this volume conform. "3" means this version.

4.4 User Volume Labels (UVL1-UVL9)

CP	Field Name	L	Content
1 to 3	Label Identifier	3	UVL
4	Label Number	1	1, 2, 3, 4, 5, 6, 7, 8, or 9 as appropriate.
5 to 80	Reserved for Installation Use	76	"a" characters. Not intended for use in an interchange environment.

4.5 First File-Header Label (HDR1)

CP	Field Name	L	Content
1 to 3	Label Identifier	3	HDR
4	Label Number	1	1
5 to 21	File Identifier	17	"a" characters. Assigned by the originator to identify the file.
22 to 27	File-Set Identifier	6	"a" characters. Identifies this file set among other file sets.
28 to 31	File Section Number	4	"n" characters. Identifies this section among the sections of this file.
32 to 35	File Sequence Number	4	"n" characters. Identifies this file among the files of this file set.
36 to 39	Generation Number	4	"n" characters. Distinguishes among successive generations of this file.
40 to 41	Generation Version Number	2	"n" characters. Distinguishes among successive iterations of the same generation.
42 to 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. Indicates date on which this file was created. Space followed by five zeros means that no creation date is specified.
48 to 53	Expiration Date	6	Same format as Creation Date. Indicates date on which this file may be overwritten. Space followed by five zeros means that this file is expired.
54	Accessibility	1	"a" character. Indicates restrictions on access to the information in this file. Space means no restrictions. The numeric characters and all special characters of the "a" set are reserved for operating system definition (see System Code (CP 61-73)). "A" through "Z" are reserved for installation assignment.
55 to 60	Block Count	6	000000

CP	Field Name	L	Content
61 to 73	System Code	13	"a" characters. Identifies the system that recorded this file. Identifiers are not specified in this standard. Spaces means that system-defined controls on file access, buffer offset, Reserved for System Use, and Other System Labels should not be used.
74 to 80	Reserved for Future Standardization	7	Spaces

4.6 Second File Header Label (HDR2)

CP	Field Name	L	Content
1 to 3	Label Identifier	3	HDR
4	Label Number	1	2
5	Record Format	1	F = fixed length D = variable length S = spanned
6 to 10	Block Length	5	"n" characters. Specifies the maximum number of characters per block.
11 to 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. In this case, 00000 indicates that the maximum record length may be greater than 99999.
16 to 50	Reserved for System Use	35	"a" characters. Not intended for use in an interchange environment.

CP	Field Name	L	Content
51 to 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.
53 to 80	Reserved for Future Standardization	28	Spaces

4.7 First End-of-Volume Label (EOV1)

CP	Field Name	L	Content
1 to 3	Label Identifier	3	EOV
4	Label Number	1	1
5 to 54	Same as the corresponding fields in HDR1	50	Same as the corresponding fields in HDR1
55 to 60	Block Count	6	"n" characters. Denotes the number of data blocks since the preceding Beginning-of-File-Section Label Group.
61 to 80	Same as the corresponding fields in HDR1	20	Same as the corresponding fields in HDR1

4.8 Second End-of-Volume Label (EOV2)

CP	Field Name	L	Content
1 to 3	Label Identifier	3	EOV
4	Label Number	1	2
5 to 80	Same as the corresponding fields in HDR2	76	Same as the corresponding fields in HDR2

4.9 First End-of-File Label (EOF1)

CP	Field Name	L	Content
1 to 3	Label Identifier	3	EOF
4	Label Number	1	1
5 to 54	Same as the corresponding fields in HDR1	50	Same as the corresponding fields in HDR1
55 to 60	Block Count	6	"n" characters. Denotes the number of data blocks since the preceding Beginning-of-File-Section Label Group.

CP	Field Name	L	Content
61 to 80	Same as the corresponding fields in HDR1	20	Same as the corresponding fields in HDR1

4.10 **Second End-of-File Label (EOF2)**

CP	Field Name	L	Content
1 to 3	Label Identifier	3	EOF
4	Label Number	1	2
5 to 80	Same as the corresponding fields in HDR2	76	Same as the corresponding fields in HDR2

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

CP	Field Name	L	Content
1 to 3	Label Identifier	3	HDR, EOV, or EOF as appropriate
4	Label Number	1	3, 4, 5, 6, 7, 8, or 9 as appropriate
5 to 80	Reserved for System Use	76	"a" characters. Not intended for use in an interchange environment.

4.12 **User File Labels (UHLa, UTLa)**

CP	Field Name	L	Content
1 to 3	Label Identifier	3	UHL or UTL as appropriate
4	Label Number	1	"a" character
5 to 80	Reserved for User Application	76	"a" characters

5. Arrangement of Labels and Data

5.1 General. In Fig. 1 through 3 and in the examples of the label groups, the beginning of the volume is at the left, and the end of the volume is at the right. Labels are represented by the four characters of their identifiers and numbers, and tape marks are represented by asterisks (*). In Fig. 1 through 3, each label group is represented by the first (or only) label in it, with the exception that the Beginning-of-Volume Label Group is represented by a VOL1 label followed by an HDR1 label.

5.2 Labels

5.2.1 A label is an 80-character record, the character positions (CP) of which are numbered 1 to 80.

5.2.2 Each label is recorded in a separate block.

5.2.3 The block that contains the label record may be extended by padding, as specified in 6.3.3.

5.2.4 A label is not considered part of the file section it delimits.

5.3 Label Sets

5.3.1 A label set is one or more contiguous labels with the same three-letter label identifier.

5.3.2 System labels (HDR1-HDR9, EOV1-EOV9, EOF1-EOF9) are recorded in consecutive ascending order. The label numbers of the consecutive system labels are 1, 2, 3, 4, 5, 6, 7, 8, 9, respectively.

5.3.3 System label sets are symmetric about a file section. That is, corresponding labels are used in each File-Header Label Set, End-of-Volume Label Set, and End-of-File Label Set for the entire file.

5.3.4 User Volume Labels (UVL1-UVL9) are recorded in consecutive ascending order. The label numbers of the consecutive User Volume Labels are 1, 2, 3, 4, 5, 6, 7, 8, 9, respectively.

5.3.5 User file labels (UHLa, UTLa) have no constraints upon the order of label numbers or symmetry of occurrence, as recorded.

5.4 Label Groups

5.4.1 A label group is one or more contiguous label sets that together delimit one end of a volume, file section, or file.

5.4.2 Label groups are delimited by a tape mark. However, at the beginning of a volume, when the file sets that normally are included in a Beginning-of-Volume Label Group and Beginning-of-File-Section Label Group are contiguously recorded with no intervening tape mark, they are considered to constitute a Beginning-of-Volume Label Group.

5.4.3 A label group is completed on the same volume where the first label of the label group is recorded.

5.5 Tape Marks

5.5.1 Single tape marks separate the Beginning-of-File-Section Label Group from the file data, the file data from the End-of-File-Section Label Group or End-of-File Label Group, and an End-of-File Label Group from the Beginning-of-File-Section Label Group of the succeeding file.

5.5.2 Double tape marks terminate End-of-File-Section Label Groups, and the last End-of-File Label Group in the file set.

5.5.3 A tape mark is not used at any other place in the volume.

5.6 Beginning-of-Volume Label Group (VOL1, UVL1-UVL9)

5.6.1 The Volume-Header Label (VOL1) follows the beginning-of-tape marker. It is the first block on a volume and contains one record. This label is not used at any other place in the volume.

5.6.2 If User Volume Labels (UVL1-UVL9) are used, they immediately follow the Volume-Header Label (VOL1).

5.6.3 A Beginning-of-Volume Label Group includes a Volume-Header Label and a File-Header Label Set. It may also include a User-Volume Label Set and a User-Header Label Set.

5.7 Beginning-of-File-Section Label Group (HDR1-HDR9, UHLa)

5.7.1 Each file section is preceded by a Beginning-of-File-Section Label Group, the first label of which is HDR1.

5.7.2 If other file-header labels (HDR2-HDR9) are used, they immediately follow the first File-Header Label (HDR1).

5.7.3 If User Header Labels (UHLa) are used, they immediately follow the last File-Header Label (HDRn).

5.7.4 A tape mark immediately follows the last label of the Beginning-of-File-Section Label Group.

5.8 File Data

5.8.1 The first data block of that section of the file immediately follows the tape mark following the Beginning-of-File-Section Label Group.

5.8.2 A tape mark immediately follows the last data block of that section of the file.

5.8.3 If no data blocks occur in that section of the file, the two tape marks bounding that empty file section are recorded contiguously. That is, these two consecutive tape marks delimit an "empty" or "null" file section.

5.9 End-of-File-Section Label Group (EOV1-EOV9, UTLa)

5.9.1 If a file extends over the end of a volume, that intermediate file section is delimited by an End-of-File-Section Label Group, the first label of which is EOV1. The first End-of-Volume Label (EOV1) immediately follows the tape mark that follows the last data block of that section of the file.

5.9.2 If other End-of-Volume Labels (EOV2-EOV9) are used, they immediately follow the first End-of-Volume Label (EOV1).

5.9.3 If User Trailer Labels (UTLa) are used, they immediately follow the last End-of-Volume Label (EOVn).

5.9.4 A double tape mark immediately follows the last label of the End-of-File-Section Label Group.

5.9.5 No information is considered to appear following the double tape mark at the end of an intermediate file section.

5.10 End-of-File Label Group (EOF1-EOF9, UTLa)

5.10.1 If a file is completed on a volume, that terminal file section is delimited by an End-of-File Label Group, the first label of which is EOF1. The first End-of-File Label (EOF1) immediately follows the tape mark that follows the last data block of that section of the file.

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

5.10.3 If User Trailer Labels (UTLa) are used, they immediately follow the last End-of-File Label (EOFn).

5.10.4 If the file is an intermediate file of a file set, a tape mark immediately follows the last label of an End-of-File Label Group.

5.10.5 If the file is the last file of a file set, a double tape mark immediately follows the last label of the End-of-File Label Group.

5.10.6 No information is considered to appear following the double tape mark at the end of a file set.

5.11 Configurations of Files

5.11.1 File length is effectively unbounded in that there may be as many as 9999 sections in one file.

5.11.2 There is only one section of a file on a volume.

5.11.3 The sections of a file are written in consecutive order (that is, section N + 1 is written immediately following section N), and do not have sections of other files interspersed.

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

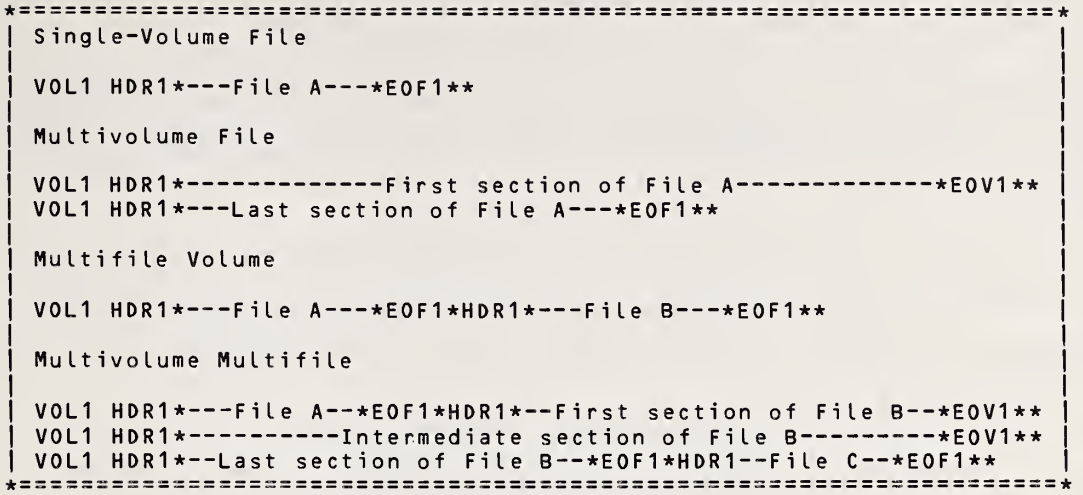

Fig. 1
Configurations of Magnetic Tape Files

5.12 Coincidence of Beginning of File and End-of-Tape Marker. If the end-of-tape marker is recognized while the Beginning-of-File-Section Label Group is being written, the Beginning-of-File-Section Label Group is completed, as described in 5.7. However, no data blocks are written in this file section, as described in 5.8.3. The volume is terminated, as described in 5.9. The file is continued on the next volume, as described in 5.6 and 5.7. See Fig. 2, File B. The File Section Number (HDR1 CP 28-31) is 1 on the original volume and 2 on the continuation volume.

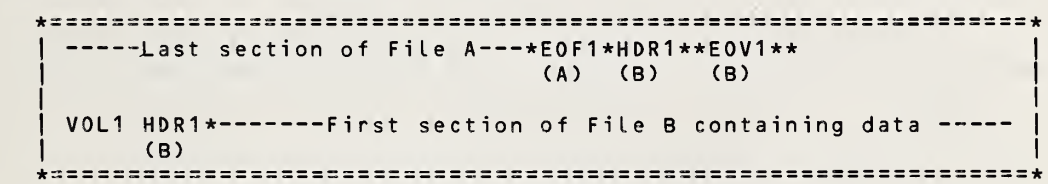

Fig. 2
Empty File Section at End of Volume

5.13 Coincidence of Data Block and End-of-Tape Marker. There are two possibilities for recognizing the end-of-tape marker while writing a file, as described in 5.13.1 and 5.13.2.

5.13.1 The end-of-tape marker is recognized while a data block is being written, and the data block is not the last data block in the file. In this case, the data block is completed. The volume is terminated, as described in 5.9. The file is continued on the next volume, as described in 5.6 and 5.7. See Fig. 1, Multivolume File.

5.13.2 The end-of-tape marker is recognized while the last data block in the file is being written. In this case, the data block is completed. The volume is terminated, as described in 5.9. The file is continued on the next volume, as described in 5.7. However, no data blocks are written on the continuation volume, and an End-of-File Label Group delimits the empty file section, as described in 5.8.3 and 5.10. See Fig. 3, File A.

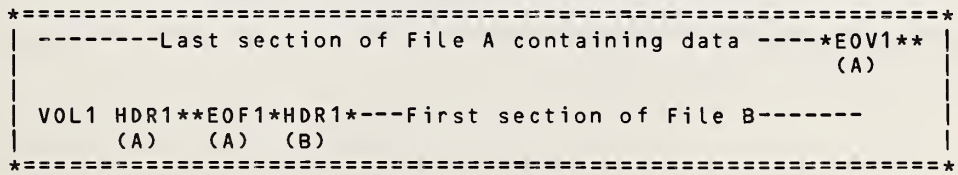

Fig. 3
Empty File Section at Beginning of Volume

5.14 Coincidence of End of File and End-of-Tape Marker. There are two possibilities for recognizing the end-of-tape marker while terminating a file, as described in 5.14.1 and 5.14.2.

5.14.1 The end-of-tape marker is recognized while the End-of-File Label Group is being written, and the file is not the last file in the file set. In this case, the End-of-File Label Group is completed, as described in 5.10. The Beginning-of-File-Section Label Group of the succeeding file is written, as described in 5.7. However, no data blocks are written in this file section, as described in 5.8.3. The volume is closed, as described in 5.9. The file is continued on the next volume, as described in 5.7. See Fig. 2, File B.

5.14.2 The end-of-tape marker is recognized while the End-of-File Label Group is being written, and the file is the last file in the file set. In this case, the End-of-File Label Group is completed, as described in 5.10. The file is terminated, as described in 5.10.5. See Fig. 1, Single-Volume File.

5.15 Examples of Label Groups. In each of the examples given in 5.15.1 through 5.15.6, there is only one file section on the volume.

5.15.1 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**

5.15.2 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*

5.15.3 Physical beginning of tape to end of file set:

VOL1 UVL1 ... UVLn HDR1 HDR2 ... HDRn UHLa ... UHLa* File Data *EOF1 ... EOFn UTLa ... UTLa**

5.15.4 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**

5.15.5 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*

5.15.6 Beginning of new file (not beginning of tape) to end of file set:

HDR1 HDR2 ... HDRn UHLa ... UHLa* File Data *EOF1 ... EOFn UTLa ... UTLa**

6. Structure of Blocks

6.1 Character Code in Data. The data in each record are recorded using only characters of the code table specified in ANSI X3.4-1977. Either a 7- or 8-bit code version, structured as specified in American National Standard Recorded Magnetic Tape for Information Interchange standards, may be used. These standards include ANSI X3.14-1973, ANSI X3.22-1973, ANSI X3.39-1973, and ANSI X3.54-1976 (see Section 9, Related Standards).

6.2 Block Formats

6.2.1 Blocking Records. No explicit indication of the boundaries between records is required. There is an integral number of records in a block for fixed-length records (F) and variable-length records (D). There is an integral number of segments in a block for spanned records (S). Blocks of fixed-length records may be shorter than Block Length (HDR2 CP 6-10). Blocks of varying length are permitted for variable-length records and for segments of spanned records. On input, it is assumed that the actual number of characters that have been read in can be determined by the system.

6.2.2 Fixed-Length Records (F). No indication of record length is required within the file. Examples of fixed-length records appear in Fig. 4 and 5.

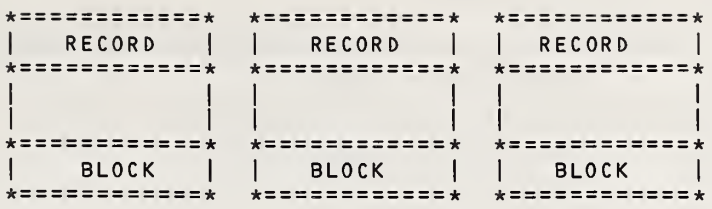

Fig. 4
Fixed-Length Records (F), Unblocked

Fig. 5
Fixed-Length Records (F), Blocked

6.2.3 Variable-Length Records (D). A record control word (RCW) precedes each record. The RCW consists of four characters. The record length (that is, the number of characters it contains) is expressed as a decimal numeral occupying the entire RCW. The record length includes the length of the RCW. Examples of variable-length records appear in Fig. 6 through 8.

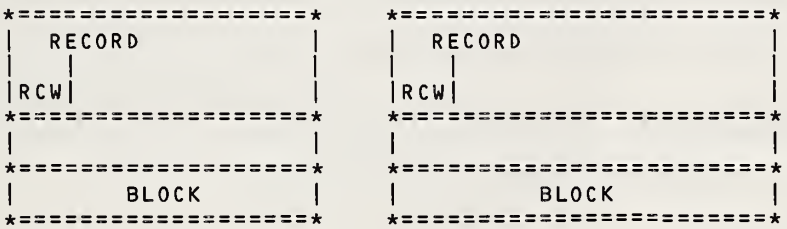

Fig. 6
Variable-Length Records (D), Unblocked

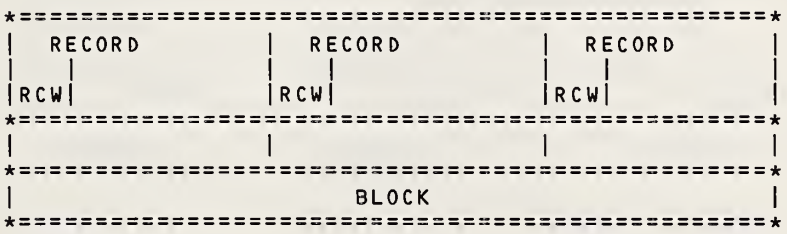

Fig. 7
Variable-Length Records (D), Blocked

NOTE: Each line represents a block.

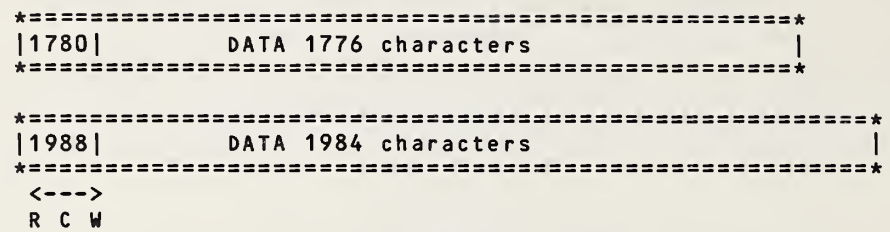

Fig. 8
Example of Unblocked Variable-Length Records

6.2.4 Spanned Records (S). A segment control word (SCW) is the first five characters of each segment. The first character of the SCW is called the Spanning Indicator. It has the following meanings:

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

The segment length (that is, the number of characters it contains) is expressed as a decimal numeral occupying the next four character positions of the SCW. The segment length includes the length of the SCW.

6.2.4.1 For spanned records there is no explicit record control word that expresses the total record length.

6.2.4.2 Records may span volumes.

6.2.4.3 Record length is unbounded in that there is no limit to the number of segments in one record.

6.2.4.4 There is only one segment of the same record in a block.

6.2.4.5 The segments of a record are written in consecutive order (that is, segment N + 1 is written immediately following segment N), and do not have segments of other records interspersed.

6.2.4.6 Examples of spanned records appear in Fig. 9 through 12.

NOTE: First block showing the maximum block size.

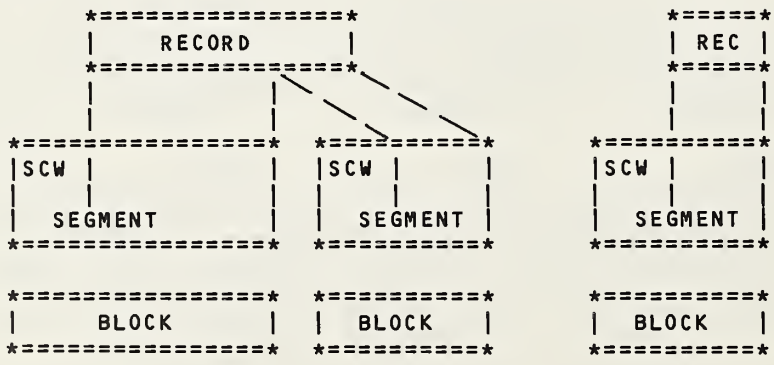

Fig. 9
Spanned Record (S), Unblocked

NOTES:
(1) All blocks showing the maximum block size.
(2) Last record continued in next subsequent block.

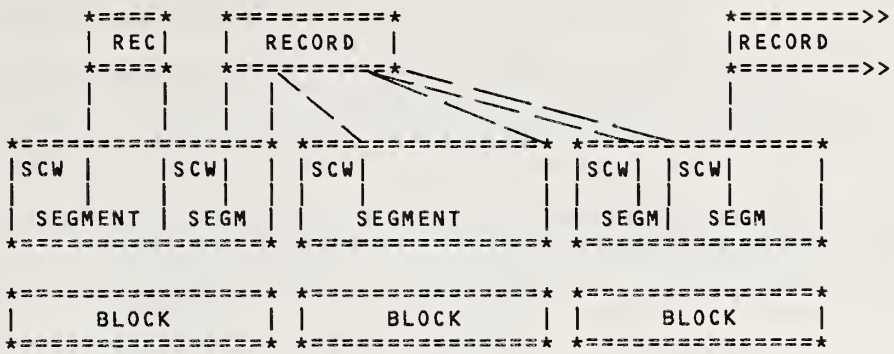

Fig. 10
Spanned Records (S), Blocked

NOTES:
(1) Length of record is 4241 characters.
(2) Each line represents a block.

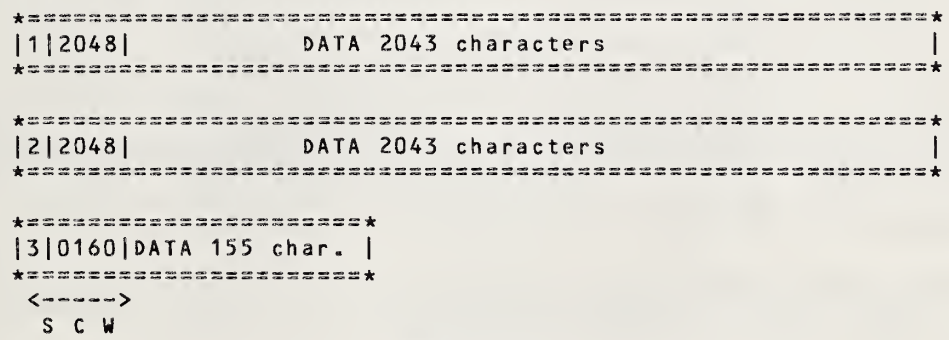

Fig. 11
Example of Unblocked Spanned Record

NOTES:
(1) Length of record 1 is 4231 characters. Length of record 2 is 5936 characters.
(2) Each line represents a block.

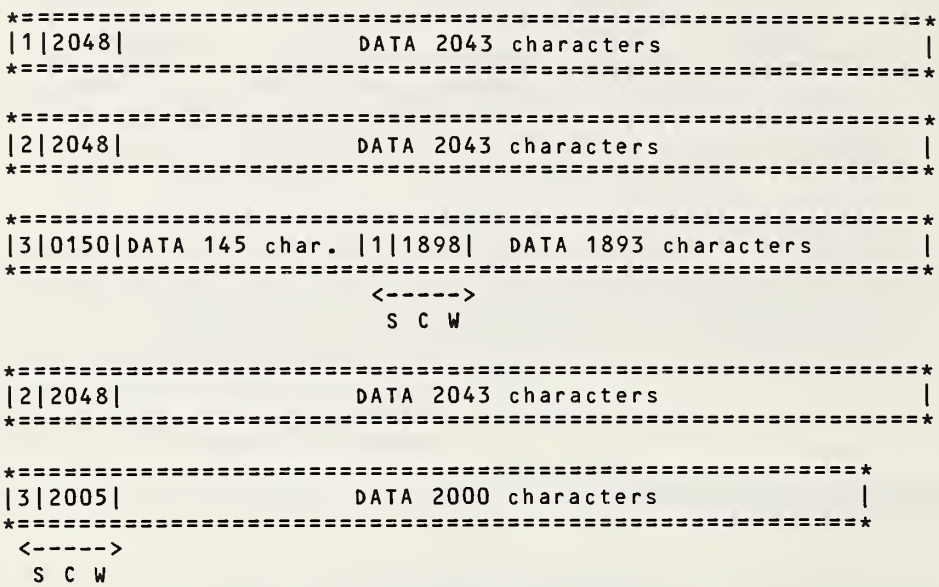

Fig. 12
Example of Blocked Spanned Records

6.2.5 Undefined Records. When records do not meet the specifications in 6.2.2, 6.2.3, or 6.2.4, they are undefined in format. The interchange of information that is undefined in format is not within the purview of this standard.

6.2.6 Record Format in a File. Records of only one format are recorded in a file.

6.2.7 Bypass or Checkpoint Records. Only data blocks containing information to be interchanged are written on a volume. Because bypass information or checkpoint records are considered to be extraneous to interchange, no standard means of identification is provided.

6.3 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 required length. Padding is not counted in the Record Control Word or the Segment Control Word.

6.3.1 Fixed Block Length. Whenever a volume 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.

6.3.2 Word-Oriented Computer. Whenever a volume 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. Note that blocks containing segments of spanned records may include padding beyond the end of the segment to the length required by a word-oriented computer.

6.3.3 Padding of Label Blocks. A block containing a label is padded out to the desired length using any desired padding characters.

6.3.4 Padding of Data Blocks. Blocks within a file are padded out to the desired length by the use of circumflex accent characters (position 5/14 of the ASCII code table). To ensure that padding after fixed-length blocked records can be distinguished from valid records, fixed-length records may not consist entirely of circumflex characters.

6.4 Recording Density. File sets and their associated labels are recorded on all volumes at the same density.

6.5 Block Length Requirements. The minimum and maximum size of data blocks are specified in the relevant recorded magnetic tape standard (see ANSI X3.14-1973, ANSI X3.22-1973, ANSI X3.39-1973, or ANSI X3.54-1976, as appropriate).

7. Processing of Label Fields, Labels, and Blocks

NOTE: In this section, any reference to processing of a field in the File-Header Label Set (HDR1-HDR9) is equally applicable to that field in the End-of-Volume Label Set (EOV1-EOV9) or in the End-of-File Label Set (EOF1-EOF9), when reading backward, unless otherwise noted.

7.1 Use of Information in Label Fields. On output, a system is provided file identification (for example, File Identifier (HDR1 CP 5-21), Generation Number (HDR1 CP 36-39)), and file attribute parameters (for example, Record Format (HDR2 CP 5), Block Length (HDR2 CP 6-10)) to use in labels. File statistics, including File Section Number (HDR1 CP 28-31), File Sequence Number (HDR1 CP 32-35), and Block Count (HDR1 CP 55-60) are developed and propagated by the system.

On input, a system is provided file identification to use to verify labels. A system may adopt and use, diagnose and reject, or ignore any file attribute parameters found in labels being processed by that system. A system may allow a user to claim different file attribute parameters from those claimed in the label. A system may override file attribute parameters found in labels being processed by that system with new values for these parameters provided from other sources. However, information found in the Volume-Header Label (VOL1) is neither overridden nor ignored.

7.2 Volume-Header Label (VOL1)

7.2.1 Volume Identifier (CP 5-10). Volume Identifier is not duplicated within an installation. It is assigned by an owner or installation when the volume is initialized.

7.2.2 Accessibility (CP 11). Additional controls over volume Accessibility may be applied, but they are not specified in this standard. Volume Accessibility is assigned by an owner or installation when the volume is initialized, to protect the volume from unauthorized access. Accessibility may be used on both input and output. A space means no access protection. Nonspace means additional qualification is required for access to the rest of the volume. As a minimum, an operating system distinguishes space and nonspace, and should not permit further volume access for nonspace. Other support the operating system may provide includes an interface to an installation security routine or an interface to the User Volume Label processing routine as described in 7.3. The installation security routine would then be considered as a part of the operating system in use at that installation, and would be provided with Accessibility and Owner Identifier. Accessibility may be processed separately from any User Volume Labels, and system support for both an installation security routine and a User Volume Label processing routine is not precluded.

Support associated with any installation-defined functions of Accessibility use the character values A through Z.

Systems that define functions related to numeric or special character values for Accessibility should specify identification procedures that ensure that the processing system can recognize the system that created the volume.

7.2.3 Owner Identifier (CP 38-51). Owner Identifier is assigned by an owner or installation when the volume is initialized, to identify the owner. Owner Identifier is supplied to the installation label processing routine by the system. If Owner Identifier is all spaces, then User Volume Labels should not be used.

7.2.4 Label-Standard Version (CP 80). Label-Standard Version is assigned by this standard. It is inserted when the volume is initialized.

7.3 User Volume Labels (UVL1-UVL9). User Volume Labels contain information about the installation use of that volume and will normally be retained permanently on the volume. On output or input, or both, these labels may be used by an installation label processing routine that recognizes Owner Identifier (VOL1 CP 38-51). These labels may be used independently, or in conjunction with Accessibility (VOL1 CP 11). User Volume Labels may be updated on the same volume by the installation label processing routines that recognize Owner Identifier (VOL1 CP 38-51), but only when all labels and data beyond these updated UVL labels are also written. In interchange with installations using different owner identification, the contents of these labels are ignored. If recognition of Owner Identifier (VOL1 CP 38-51) is not performed by the label processing routine, incorrect interpretation of the User Volume Labels may occur.

If Owner Identifier is not all spaces and there is at least one User Volume Label, the installation volume label processing routine is invoked and Owner Identifier is made available along with the User Volume Label set. Accessibility is made available to the installation volume label processing routine unless only the minimum operating system support is provided for this field.

7.4 Altering the Volume-Header Label (VOL1) and User Volume Labels (UVL1-UVL9). The Volume-Header Label (VOL1) and the User Volume Labels (UVL1-UVL9) may not be changed, added to, or deleted, except as specified in 7.3. This does not preclude the rewriting of the Volume-Header Label, along with the User Volume Labels, with the contents unchanged.

7.5 First File-Header Label (HDR1)

7.5.1 File Identifier (CP 5-21). File Identifier is not duplicated within a file set. It is assigned by a user.

7.5.2 File-Set Identifier (CP 22-27). File-Set Identifier is the same for all files of a file set. The user supplies this value in many systems. In some cases the system label routines supply the unique value, for example, Volume Identifier (VOL1 CP 5-10) of the first volume of the file set. It is propagated by a system.

7.5.3 File Section Number (CP 28-31). File Section Number of the first section of a file is 0001. This number is increased by 1 for each successive volume of the file. It is maintained by a system.

7.5.4 File Sequence Number (CP 32-35). 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 labels for a given file, whether that file be single or multivolume, this field contains the same number. It is maintained by a system.

7.5.5 Generation Number (CP 36-39). In some systems, Generation Number is supplied by the system. In other systems, it may be supplied by the user. The value of the first generation written for 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.

7.5.6 Generation Version Number (CP 40-41). Generation Version Number of the first attempt to produce a generation of a file is 00. If subsequent versions of a generation are noted, this number is increased by 1 for each successive attempt. This number is reset to 00 when the Generation Number (HDR1 CP 36-39) is increased by 1. The system normally supplies the properly updated value for this field. If a system is unable to provide a properly updated value, the user may be required to supply a usable value.

7.5.7 Creation Date (CP 42-47). In some systems creation date is assigned by a user. In other systems, it may be the current date maintained by a system. Space followed by five zeros is a valid value.

7.5.8 Expiration Date (CP 48-53). The file is regarded as expired on a day whose date is equal to or later than Expiration Date. The user supplies this value in many systems; in some cases the installation will provide a retention period that the system utilizes to calculate an expiration date, using the current date as the base. Or the system may provide a facility to calculate a date by utilizing a retention period supplied by the user in a control statement or common programming language declaration. In some systems, an expired value will be supplied if the user does not specify a value or if the other possible methods of setting the value are not available. Space followed by five zeros is a valid value.

7.5.9 Accessibility (CP 54). Accessibility is examined by the system and no additional controls over file access are applied if the value is a space. As a minimum, if a nonspace value is found, access to the file is denied. If a numeric or special character is found, a system that recognizes System Code (HDR1 CP 61-73) may utilize checks defined for that system and associated with a numeric or special character from the "a" code set. At the system implementors option, the values "A" through "Z" are made available to an installation-provided security routine. The installation may then apply additional controls over file access, but those controls are not defined in this standard. If Accessibility is "A" through "Z" and no installation processing routine is available, access to the file is denied by the system. File Accessibility is assigned by a user.

NOTE: It is strongly recommended that installations that process Accessibility perform an identification check, using Owner Identifier, to assure proper interpretation of the values in the Accessibility field.

7.5.10 Block Count (CP 55-60). A count of blocks written or read is maintained by a system so that Block Count can be written on output, or it can be verified that the correct number of blocks were read on input. Block Count is six zeros in the first File-Header Label (HDR1). This count is maintained so that if a sequence of blocks is read, backspaced, reread, etc, or written, backspaced, overwritten, etc, this count is incremented (decremented when reading backward) only once for each block that appears in each file section on the volume.

7.5.11 System Code (CP 61-73). System Code is a constant for a given system. It is inserted by the system that created the file, to identify itself. A system recognizes the particular System Code for that file before that system makes use of file Accessibility (HDR1 CP 54), of Reserved for System Use

(HDR2 CP 16-50), of buffer offset (see 7.6.5), or of any other system label (HDR3-HDR9, EOV3-EOV9, EOF3-EOF9). If System Code is all spaces, then these facilities should not be used.

7.6 Second File-Header Label (HDR2). All of the file attribute parameters are assigned by a user.

7.6.1 Record Format (CP 5). Record Format specifies the format of the records in the file.

7.6.2 Block Length (CP 6-10). Block Length includes buffer offset, RCW, SCW, data and padding. The actual maximum capacity for data within a block is reduced by Buffer-Offset Length (HDR2 CP 51-52) and by the number of padding characters so that the block does not exceed the maximum length specified in the applicable recorded magnetic tape standard.

7.6.3 Record Length (CP 11-15). If Record Format (HDR2 CP 5) is F, this field contains the actual length of each record. If Record Format (HDR2 CP 5) is D, then Record Length includes the length of the record control word. If Record Format (HDR2 CP 5) is S, then Record Length is the maximum record length, excluding all the Segment Control Words. If the maximum length of the spanned record may be more than 99999 characters, or if the maximum length is not known, then Record Length is 00000.

7.6.4 Reserved for System Use (CP 16-50). On output or input, or both, this field may be used by a system that recognizes System Code (HDR1 CP 61-73). In interchange the contents of this field are ignored.

7.6.5 Buffer-Offset Length (CP 51-52). The length of information that prefixes each data block is specified in Buffer-Offset Length. If the prefix information (called buffer offset) is not included, then Buffer-Offset Length is 00.

7.7 Other System Labels (HDR3-HDR9, EOV3-EOV9, EOF3-EOF9). On output or input, or both, the labels may be used by a system that recognizes System Code (HDR1 CP 61-73).

7.8 User File Labels (UHLa, UTLa). On output or input, or both, these labels may be used by an application program that processes the records in a file. The labels may be associated with the beginning or end of the file, or both, or they may be associated with the beginning or end of each file section, or with both the beginning and end of each file section.

Note that recorded tape standards limit the area available for the recording of information after the end-of-tape marker (EOT) to 10 feet of the 25-foot area beyond EOT.

7.9 Label and Block Processing. This section specifies the requirements for processing of labels and blocks in information interchange.

7.9.1 File Opening

7.9.1.1 On input, if the desired file is at load point, the entire Beginning-of-Volume Label Group (VOL1, UVL1-UVL9, HDR1-HDR9, UHLa) is verified. The system examines the Volume-Header Label (VOL1). Then the system support for the User Volume Label Set (UVL1-UVL9) is invoked. The processing after this point applies also if the desired file is other than the first file on the volume. The system positions the volume to process the File-Header Label Set (HDR1-HDR9), reads the File-Header Label Set (HDR1-HDR9), verifies file identification (see 7.5 and 7.10), and processes the File-Header Label Set (HDR1-HDR9). Then the system support for the User Header Label Set (UHLa) is invoked. Finally, at the tape mark, the system prepares to read the data in the file section.

7.9.1.2 On output, if the tape is at load point, the Beginning-of-Volume Label Group is verified. The system examines the Volume-Header Label (VOL1). Then the system support for the User Volume Label Set (UVL1-UVL9) is invoked. When a system has positioned the volume at the expected location of the File-Header Label Set (HDR1-HDR9), it expects to read an existing first File-Header Label (HDR1). The system verifies access to overwrite that file. The system then backspaces and overlays the existing Beginning-of-File-Section Label Group with the File-Header Label Set (HDR1-HDR9) of the new file. Then the system support for the User Header Label Set (UHLa) is invoked. Finally, the system writes a tape mark and prepares to write the data in the file section.

7.9.1.3 On output, when writing a file other than the first file on a volume, and the system has positioned the volume at the expected location of the Beginning-of- File-Section Label Group, it expects to read either an existing first File-Header Label (HDR1) or a tape mark.

If a first File-Header Label (HDR1) is found, the system verifies access to overwrite that file. The system then backspaces and overlays the existing Beginning-of-File-Section Label Group with the File-Header Label Set (HDR1-HDR9) of the new file. Then the system support for the User Header Label Set

(UHLa) is invoked. Finally, the system writes a tape mark and prepares to write the data in the file section.

If a tape mark is found, the system accepts the volume for output without further label verification. The system then backspaces and overlays the tape mark with the File-Header Label Set (HDR1-HDR9) of the new file. Then the system support for the User Header Label Set (UHLa) is invoked. Finally, the system writes a tape mark, and prepares to write the data in the file section.

The remainder of a file set is logically destroyed when an existing File-Header Label Set (HDR1-HDR9) is overwritten.

7.9.2 Processing Blocks of Records

7.9.2.1 The block processing techniques described in this section are compatible with the specifications of common programming languages. At the option of system implementors, additional techniques may be provided.

7.9.2.2 The system accepts from/delivers to the application program processing routines one record at each invocation. This may be performed by a system routine, by a common programming language library routine, by a routine automatically included in the user program by a common programming language processor, etc. The processing depends upon Record Format (HDR2 CP 5). Note that padding may appear following the last record or record segment in a block.

If the record is fixed length (F), the length of the record need not be made available at each invocation.

If the record is variable length (D), the length of each record is made available at each invocation. The interpretation of this datum is consistent with the specifications of the particular common programming language.

If the record is spanned (S), the length of each record is made available at each invocation. The interpretation of this datum is consistent with the specifications of the particular common programming language. The length of a segment is arbitrary, and a system may segment records to fit the blocks being written.

7.9.2.3 Blocks may be buffered. The system may maintain a backlog of records read, but not yet processed by an application program, or a backlog of records prepared by an application program, but not yet written. The disposition of these blocks at the end of a file section or at the end of a file (that is, the emptying of the buffer) is described in 7.9.3 and 7.9.4.

7.9.3 Volume Switching

7.9.3.1 On output, the end of a file section is determined either when an end-of-tape marker is sensed while writing the data in the file or when an application program requests a volume switch. If the output is buffered, the system empties any records in the buffers at a volume switch. The system then writes a tape mark and then the End-of-Volume Label Set (EOV1-EOV9).

On any single volume, the End-of-Volume Label Set (EOV1-EOV9) is an exact copy of the File-Header Label Set (HDR1-HDR9), except for Block Count (HDR1 CP 55-60). None of the other attributes or dates vary from beginning to end of a file section. Then the application program label processing routine is invoked to prepare the User Trailer Label Set (UTLa). The system invokes the routines once for each of the User Trailer Labels, and writes the labels.

The system writes a double tape mark following the End-of-File-Section Label Group and rewinds the volume.

The Beginning-of-Volume Label Group and the Beginning-of-File-Section Label Group on the next volume are processed as described in 7.9.1.

The File-Header Label Set (HDR1-HDR9) of each succeeding file section is an exact copy of the File-Header Label Set (HDR1-HDR9) of the preceding file section, except that File Section Number (HDR1 CP 28-31) is increased by 1 to properly reflect the sequence of the file sections. In case of rewrite of a portion of a file that is on a separate volume, Generation Version Number (HDR1 CP 40-41) may also vary, and in that case, or in the case of additional file sections that are made (for example, over several days), Creation Date (HDR1 CP 42-47) and Expiration Date (HDR1 CP 48-53) may also vary. However, any new value for Expiration Date is equal to or less (earlier) than the value in all preceding header and trailer labels of the file. None of the other attributes varies from file section to file section.

Finally, the system writes a tape mark and prepares to write the data in the file section.

7.9.3.2 On input, the end of a file section is determined either when the tape mark that precedes an End-of-File-Section Label Group is read or when an application program requests a volume switch.

If the tape mark that precedes an End-of-File-Section Label Group is read, the system reads and processes the End-of-Volume Label Set (EOV1-EOV9). Then the application program label processing routines are invoked to examine the User Trailer Label Set (UTLa). The system reads the labels and invokes the routine once for each of the user trailer labels.

If the application program requests a volume switch, processing of the End-of-File-Section Label Group is omitted.

In either case, the system rewinds the volume.

The Beginning-of-Volume Label Group and the Beginning-of-File-Section Label Group on the next volume are processed as described in 7.9.1. Note that verification for file Accessibility (HDR1 CP 54) is repeated for subsequent file sections of a file. Volume Identifier (VOL1 CP 5-10) is checked and File Section Number (HDR1 CP 28-31) is checked to verify that the expected file section is mounted.

Finally, after reading the tape mark, the system prepares to read the data in the file section.

7.9.4 File Closing

7.9.4.1 On output, the end of a file is determined when an application program closes the file. The system ensures that all the data records have been written, and then writes a tape mark and the End-of-File Label Set (EOF1-EOF9). On any single volume, the End-of-File Label Set (EOF1-EOF9) is an exact copy of the File-Header Label Set (HDR1-HDR9) except for Block Count (HDR1 CP 55-60). None of the other attributes vary from beginning to end of a file section.

In case of rewrite of a portion of a file that is on a separate volume, Generation Version Number (HDR1 CP 40-44) may also vary, and in that case, or in the case of additional file sections that are made (for example, over several days), Creation Date (HDR1 CP 42-47) and Expiration Date (HDR1 CP 48-53) may also vary. However, any new value for Expiration Date is equal to or less (earlier) than the value in all preceding header and trailer labels of the file. After the End-of-File Label Set is processed, the application program label processing routines are invoked to prepare the User Trailer Label Set (UTLa). The system invokes the routines once for each of the user trailer labels and writes the labels. The system writes a double tape mark following the End-of-File Label Group and may rewind the volume or leave it positioned between the two tape marks. When adding another file to the volume, the Beginning-of-File-Section Label Group of the new file overlays the second tape mark.

7.9.4.2 On input, the end of a file is determined either when an End-of-File Label Group is read or when an application program closes the file.

If the application program closes the file before the End-of-File Label Group is read, the End-of-File Label Group processing is omitted.

If an End-of-File Label Group is encountered, the system reads and processes the End-of-File Label Set (EOF1-EOF9). It then communicates the end-of-file condition to the application program, so that the file can be closed. After the file is closed, the system permits the existing User Trailer Label Set (UTLa) to be passed to the user. This is done by permitting the input user trailer label routine to be entered once for each of the User Trailer Labels (UTLa). The system reads the labels, and invokes the routine once for each user trailer label. After processing the End-of-File Label Group, the system may rewind the volume or it may leave the volume positioned before the next (or any) Beginning-of-File-Section Label Group or before the second tape mark of the double tape mark.

7.9.5 Reading Backward

7.9.5.1 Reading backward is not required for interchange. If a read backward facility is supported by the system and a file is opened for read backward, after recognizing the tape mark that precedes the End-of-File Label Group, the system prepares to read the data in the file section. The records are read in reverse sequence.

7.9.5.2 Blocked variable-length records (D) and segments of spanned records (S) are not designed to be read backward efficiently.

7.9.6 Restart Processing. If a system provides automatic restart from a checkpoint, it repositions volumes containing files that were open when the checkpoint was taken. Specifically:

(1) The system will ensure that the first record on the volume is a Volume-Header Label (VOL1) and verify Volume Identifier (VOL1 CP 5-10). The file labels need not be reprocessed since these labels were processed when the file was opened.

(2) The system repositions the volume at the proper record within the file.

7.10 Error Actions. Errors that may occur during reading, processing, and writing of labels and of data are of two kinds: hardware and programming. In case of hardware errors, the system-implementor-provided error recovery function can be used. In case of a programming error such as when a label or tape mark is not found where expected or a field in a label contains an improper or incorrect value, at the option of system implementors, a system may provide automatic restart or retry of the operation.

At the option of system implementors, a system may permit an operator or a systems administrator to manually correct an error and continue, to restart, or to continue despite the error.

At the option of system implementors, a system may permit an application program to accept a user file label or data within the file, and to continue despite the error.

When an error has occurred, after appropriate corrective actions have failed and the error is not overridden, then the file is closed, if necessary, and the volume rewound by the system.

8. Levels of Labeling

8.1 General. In order to facilitate interchange of information among systems of dissimilar capabilities, the concept of levels is established. These levels are intended to enable an implementor of a system to guide applications designers and users to ascertain the capability of the system with respect to the requirements of the application and in selection of facilities to ensure that the volume can be processed correctly by the receiving system. Such levels act as floors on system capabilities and ceilings on the volumes produced by such a system.

In these discussions a distinction is presumed between the operations necessary to assure that the content of a label is valid and the operations required to process labels. The former are not described except to fix responsibility for the checking of validity, either by the system, by the writer of operating system control statements, by equivalently used common programming language declarations, or by user-written validity checking programs. 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. Although 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 explicitly specified processing of each field defined in this standard.

Use of this standard for interchange of information does not require checking of system labels by user-written validity checking routines. However, if the system implementors have provided for such routines, the routines may be entered only after the system input validity checks are satisfied, and only before the system output validity checks are satisfied. Specification of interfaces to user-written validity checking routines is not within the purview of this standard. Different systems or different level implementations may have different interfaces for user-written validity checking routines, and the interchangeability of the programs that contain the checking routines is not assured.

The following terms are used herein in referring to the four levels of labeling systems:

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

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

extension. Any format or function of a labeling system not defined in this standard.

8.2 Facilities Available at Level 1

8.2.1 File Sets. File sets are single-file single-volume or single-file multivolume.

8.2.2 Labels. Labels written are VOL1, HDR1, EOV1, EOF1. It is allowable to write additional standard labels, but a system ignores and bypasses any additional labels it does not process.

8.2.3 Record/Block Structures. Blocks consist of one or more fixed-length records.

8.2.4 Label Fields. In a level 1 system, the following label fields are processed:

(1) VOL1
- Label Identifier
- Label Number
- Volume Identifier
- Accessibility
- Label-Standard Version

(2) HDR1, EOV1, EOF1
- Label Identifier
- Label Number
- File Identifier
- File Section Number
- Expiration Date
- Block Count

The processing of additional label fields is optional.

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

VOL1	CP 38-51 Owner Identifier	Spaces
HDR1, EOV1, EOF1		
	CP 22-27 File Set Identifier	Spaces
	CP 32-35 File Sequence Number	0001
	CP 36-39 Generation Number	0001
	CP 40-41 Generation Version Number	00
	CP 42-47 Creation Date	△00000
	CP 54 Accessibility	Space
	CP 61-73 System Code	Spaces

8.2.5 Validity of Labels. In a level 1 system, 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, the exercise of responsibility for validity of HDR1, EOV1 and EOF1 is as follows: The writer of operating system control statements (for example, Job Control Statements), equivalently used common programming language declarations, or user-written validity checking programs exercises responsibility for format, editing, and accuracy of all file system label fields except File Section Number (LBL1 CP 28-31) and Block Count (LBL1 CP 55-60). A system exercises responsibility for format, editing, and accuracy of File Section Number (LBL1 CP 28-31) and Block Count (LBL1 CP 55-60).

A level 1 system is not required to check or edit any fields in any manner, except those identified here as the validity checking responsibility of the level 1 system. The level 1 system may simply copy the fields from information supplied by the writer of operating system control statements or common programming language declarations. Additionally, the level 1 system may provide an interface for user-written validity checking programs to the system label fields not the responsibility of the level 1 system.

Additional diagnosis by a system is optional.

8.2.6 File Security. A level 1 system assumes that the volume is immediately accessible if Accessibility (VOL1 CP 11) is equal to a space. If the contents are not a 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.

8.3 Facilities Available at Level 2

8.3.1 File Sets. File sets are single-file single-volume, single-file multivolume, multifile single-volume, or multifile multivolume.

8.3.2 Labels. Labels written are VOL1, HDR1, EOV1, EOF1. It is allowable to write additional standard labels, but a system ignores and bypasses any additional labels it does not process.

8.3.3 Record/Block Structures. Blocks consist of one or more fixed-length records.

8.3.4 Label Fields. In a level 2 system, the following label fields are processed:

(1) VOL1
Label Identifier
Label Number
Volume Identifier
Accessibility
Label-Standard Version

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

The processing of additional label fields is optional.

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

VOL1 CP 38-51 Owner Identifier	Spaces
HDR1,EOV1,EOF1	
CP 36-39 Generation Number	0001
CP 40-41 Generation Version Number	00
CP 42-47 Creation Date	Δ00000
CP 61-73 System Code	Spaces

8.3.5 Validity of Labels. In a level 2 system, 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, the exercise of responsibility for validity of HDR1, EOV1, and EOF1 is as follows: The writer of operating system control statements, equivalently used common programming language declarations, or user-written validity checking programs exercises responsibility for format, editing, and accuracy of all file system label fields except File-Set Identifier (LBL1 CP 22-27), File Section Number (LBL1 CP 28-31), File Sequence Number (LBL1 CP 32-35) and Block Count (LBL1 CP 55-60). A system exercises responsibility for format, editing, and accuracy of File-Set Identifier (LBL1 CP 22-27), File Section Number (LBL1 CP 28-31), File Sequence Number (LBL1 CP 32-35), and Block Count (LBL1 CP 55-60).

A level 2 system is not required to check or edit any fields in any manner, except those identified here as the validity checking responsibility of the level 2 system, and may simply copy the fields from information supplied by the writer of operating system control statements or common programming language declarations, or may provide an interface for user-written validity checking programs to the system label fields not the responsibility of the level 2 system.

Additional diagnosis by a system is optional.

8.3.6 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 a space. If the contents of either or both fields are not spaces, access is denied unless the level 2 system provides additional controls. Such additional controls are not specified in this standard.

8.4 Facilities Available at Level 3

8.4.1 File Sets. File sets are single-file single-volume, single-file multivolume, multifile single-volume, or multifile multivolume.

8.4.2 Labels. Labels written are VOL1, HDR1, HDR2, UHLa, EOV1, EOV2, EOF1, EOF2, UTLa. It is allowable to write additional standard labels, but a system ignores and bypasses any additional labels it does not process.

8.4.3 Record/Block Structures. Blocks consist of one or more fixed-length records, or one or more variable-length records.

8.4.4 Label Fields. In a level 3 system, the following label fields are processed:

(1) VOL1
- Label Identifier
- Label Number
- Volume Identifier
- Accessibility
- Label-Standard Version

(2) 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

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

The processing of additional label fields is optional.

However, all fields of the VOL1, HDR1, HDR2, EOV1, EOV2, EOF1, and EOF2 labels contain meaningful information in accordance with this standard. Therefore, to assist implementation, standard default values for certain fields are listed below. It should also be noted that fields that are Reserved for Future Standardization should be space-filled.

VOL1 CP 38-51 Owner Identifier	Spaces
HDR1,EOV1,EOF1	
CP 36-39 Generation Number	0001
CP 40-41 Generation Version Number	00
CP 61-73 System Code	Spaces
HDR2,EOV2,EOF2	
CP 16-50 Reserved for System Software Use	Spaces

8.4.5 Validity of Labels. In a level 3 system, 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, the exercise of responsibility for validity of HDR1, EOV1, and EOF1 is as follows: The writer of operating system control statements, equivalently used common programming language declarations, or user-written validity checking programs exercises responsibility for format, editing, and accuracy of all file system label fields except Label Identifier (LBL1 CP 1-3), Label Number (LBL1 CP 4), File-Set Identifier (LBL1 CP 22-27), File Section Number (LBL1 CP 28-31), File Sequence Number (LBL1 CP 32-35), and Block Count (LBL1 CP 55-60). A system exercises responsibility for format, editing, and accuracy of Label Identifier (LBL1 CP 1-3), Label Number (LBL1 CP 4), File-Set Identifier (LBL1 CP 12-27), File Section Number (LBL1 CP 28-31), File Sequence Number (LBL1 CP 32-35), and Block Count (LBL1 CP 55-60).

In a level 3 system, the system exercises responsibility for the format of HDR2, EOV2, and EOF2 label fields: Label Identifier (LBL2 CP 1-3), Label Number (LBL2 CP 4), Record Format (LBL2 CP 5), Block Length (LBL2 CP 6-10), Record Length (LBL2 CP 11-15), and Buffer-Offset Length (LBL2 CP 51-52).

A level 3 system is not required to check or edit any other of the system label fields that are specified to be the validity checking responsibility of the writer of operating system control statements or equivalently used common programming language declarations. For validity checking, a system may provide an interface for user-written programs to the system label fields not the responsibility of the level 3 system.

Additional diagnosis by a system is optional.

8.4.6 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 a 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.

8.5 Facilities Available at Level 4

8.5.1 File Sets. File sets are single-file single-volume, single-file multivolume, multifile single-volume, or multifile multivolume.

8.5.2 Labels. Labels written are VOL1, HDR1, HDR2, UHLa, EOV1, EOV2, EOF1, EOF2, UTLa. It is allowable to write additional standard labels, but a system ignores and bypasses any additional labels it does not process.

8.5.3 Record/Block Structures. Blocks consist of one or more fixed-length records, one or more variable-length records, or one or more segments of spanned records.

8.5.4 Label Fields. In a level 4 system, the following label fields are processed:

(1) VOL1
- Label Identifier
- Label Number
- Volume Identifier
- Accessibility
- Label-Standard Version

(2) 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

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

The processing of additional label fields is optional.

However, all fields of the VOL1, HDR1, HDR2, EOV1, EOV2, EOF1, and EOF2 labels contain meaningful information in accordance with this standard. Therefore, to assist implementation, standard default values for certain fields are listed below. It should also be noted that fields that are Reserved for Future Standardization should be space-filled.

VOL1 CP 38-51 Owner Identifier — Spaces

HDR1,EOV1,EOF1
CP 61-73 System Code — Spaces

HDR2,EOV2,EOF2
CP 16-50 Reserved for System Software Use — Spaces

8.5.5 Validity of Labels. Level 4 systems attempt to ensure maximum correctness and consistency of files and their labels. In a level 4 system, the system exercises responsibility for the format of all system label fields in VOL1, HDR1, HDR2, EOV1, EOV2, EOF1, and EOF2. Operating system control statements and equivalently used common programming language declarations are edited to ascertain that the labels are properly formatted and 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 to maximize the consistency of the file and its labels.

8.5.6 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.

8.6 Implementation of Levels

8.6.1 Responsibility for System Labels. System labels (VOL1, HDR1-HDR9, EOV1-EOV9, EOF1-EOF9) are processed entirely by the label handling routines. The label and file processing functions 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 levels 3 and 4.

8.6.2 Responsibility for User Volume Labels. User Volume Labels (UVL1-UVL9) are processed partly by the label and file processing functions, 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 and file processing functions 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 and file processing functions and the resident installation routine are provided, except that no transfer will occur if no User Volume Labels appear.

These facilities are optional at each level.

8.6.3 Responsibility for User File Labels. User File Labels (UHLa and UTLa) are processed partly by the label and file processing functions, to the extent that they are accepted from user application program routines and written on a volume on output; and read from a volume, recognized, and passed to user application program routines on input. The label and file processing functions 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 and file processing functions and user routines are provided.

These facilities are optional at levels 1 and 2.

8.6.4 Allocation of Output Options. Selection of the level of labeling and provision of the facility for creation of labels optional at that level, or fields not required at that level, are 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.

8.6.5 Allocation of 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 processing facilities are not provided, or if provided are not invoked, are bypassed, and the information ordinarily obtained from them is lost.

8.7 Description of Media

8.7.1 Domain. This section 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 is not considered part of the information in interchange and is not considered in determining the level of the volume.

8.7.2 Conditions. A volume (set) corresponds to a level under all of the following conditions:

(1) It and every file upon it contains all of the elements of labeling required at that level, 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.

(2) If it or any file upon it contains any elements of labeling 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.

(3) It and every file upon it contains only file sets or record/block formats permitted at that level.

(4) It and every file upon it does not contain any element that is at variance with this standard, or any extension to elements of labeling defined in this standard.

8.7.3 Nonstandard Volumes. Applications that record nonstandard labels and files on volumes produce nonstandard volumes. This may result in a subsequent failure of information interchange.

8.8 Description of Implementations

8.8.1 Domain. This section applies only to the label and file processing functions of a system. Any other facilities are not considered in determining the level of an individual implementation.

8.8.2 Conditions. An implementation of label and file processing functions corresponds to a level under all of the following conditions:

(1) Except at the option of a user, it places upon the volume (set) all of the elements of labeling required at that level, 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.

(2) If it places upon the volume (set) any of the elements of labeling 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.

(3) Except at the option of a user, it places upon the volume (set) only file sets and record/block formats permitted at that level.

(4) Except at the option of a user, it does not place upon the volume, nor does it require the user to place upon the volume, any element that is at variance with this standard, or any extension to elements of labeling defined in this standard.

8.8.3 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 labeling defined in this standard. If a user opts such an omission, variance, or extension, then the resulting volume is nonstandard. See 8.6.3.

9. Related Standards

9.1 Revision of American National Standards Referred to in This Document

When any of the following American National Standards referred to in this document is superseded by a revision approved by the American National Standards Institute, Inc, the revision shall apply:

American National Standard Code for Information Interchange, ANSI X3.4-1977

American National Standard Recorded Magnetic Tape for Information Interchange (200 CPI, NRZI), ANSI X3.14-1973

American National Standard Recorded Magnetic Tape for Information Interchange (800 CPI, NRZI), ANSI X3.22-1973

American National Standard Programming Language COBOL, ANSI X3.23-1974

American National Standard Representation for Calendar Date and Ordinal Date for Information Interchange, ANSI X3.30-1971

American National Standard Recorded Magnetic Tape for Information Interchange (1600 CPI, PE), ANSI X3.39-1973

American National Standard Unrecorded Tape for Information Interchange (9-Track 200 and 800 CPI, NRZI, and 1600 CPI, PE), ANSI X3.40-1976

American National Standard Representation of Numeric Values in Character Strings for Information Interchange, ANSI X3.42-1975

American National Standard Recorded Magnetic Tape for Information Interchange (6250 CPI, Group-Coded Recording), ANSI X3.54-1976

American National Standard for Bibliographic Information Interchange on Magnetic Tape, ANSI Z39.2-1971

9.2 ISO Standards*

Magnetic Tape Labelling and File Structure for Information Interchange, ISO/DIS 1001.2

*Publications of the International Organization for Standardization are available from the American National Standards Institute, 1430 Broadway, New York, N.Y. 10018.

Appendixes

(These Appendixes are not a part of American National Standard Magnetic Tape Labels and File Structure for Information Interchange, ANSI X3.27-1977, but are included for information purposes only.)

Appendix A

Levels of Systems

A1. Systems of Levels of Labeling

In order to facilitate interchange of information among systems of dissimilar capabilities, the concept of levels of labeling 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 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 multivolume
	2	Single-file single-volume + single-file multivolume + multifile single-volume + multifile multivolume
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

There are four levels, numbered 1, 2, 3, and 4 in increasing order of complexity and made up of the option constituents as follows:

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

Level 1 is a minimum system.

Level 2 is effectively level 1 with multifile file sets added.

Level 3 provides most used facilities of the current standard, including variable-length records and HDR2, EOV2 and EOF2.

Level 4 is effectively level 3 with spanned records added.

The levels are so defined that systems software can process correctly all file sets in which the labeling level used is equal to or less than the system's nominated highest level. This applies equally to both reading and writing. Such required 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, 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 should contain meaningful information in accordance with this standard. Therefore, to assist the systems software writing the file set, default values have been

specified that meet this criterion.

Standard labels and user labels not included in a given level may be on the tape but bypassed by the systems software. Thus an 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 by adding other record formats without restriction, 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 for 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.

A2. Levels of Labeling

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 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 will be the amount of information provided external to the volumes interchanged.

It is not intended that every implementation 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 a equal or higher level; and any system can process correctly any volume (set) of equal or lower level.

A3. Distinguishing Characteristics of Levels of Implementation

A3.1 General. This section is intended to provide characterizations of each level of implementation in terms of the grouping of facilities. It is not intended to specify the internal working relationships of an individual implementation of a system. A system may identify the grouping of facilities that it processes correctly by a level number, without assuming any other characteristics of that level number. These levels are not intended to constrain or measure an individual implementation of a system. The provision in addition of any facility optional at that level precisely as defined in this standard, or any extensions to facilities otherwise defined in this standard, or both, does not in itself affect the definition of the level. The facilities cited in Section 8, Levels of Labeling, are desirable and consistent with the characterization of each level.

A3.2 Level 1 System. It is characteristic of a level 1 system that its label and file processing functions are not included in a resident monitor system. Typically, application source programs are assembled or compiled and are combined with their label and file processing functions by manual or automatic inclusion of source or object code of these routines. Thus, although the application program source code does not explicitly include the label and file processing algorithms, the application program object code effectively includes the label and file processing routines.

Typically, few options are available and they are expressed by operating system control statements or common programming language declarations submitted with the application program source code. Typically, the entire program is resident in limited main storage, and it is at assembly or compilation that a selection is made from among the options available; little or no adaptation to the attributes of the data is possible on execution. Level 1 systems frequently accept an (80-column) operating system control statement that is character for character copied as a file label or compared to an existing file label.

A3.3 Level 2 System. It is characteristic of a level 2 system that its label and file processing functions are not included in a resident monitor system, except to the extent necessary to locate, properly identify and validate access to specific files of a multifile file set on input, and to produce multifile file sets on output. Typically, application source programs are assembled or compiled and are combined with their label and file processing functions by manual or automatic inclusion of source or object code of these routines. Thus, although the application program source code does not explicitly include the label and file processing algorithms, the application program object code effectively includes the label and file processing routines.

Typically, few options are available and they are expressed by operating system control statements or common programming language declarations submitted with the application program source code. Typically,

the entire program is resident in limited main storage, and it is at assembly or compilation that a selection is made from among the options available; little or no adaptation to the attributes of the data is possible on execution. Level 2 systems frequently accept an (80-column) operating system control statement that is character for character copied as a file label or compared to an existing file label. Although these characteristics are similiar for level 1 systems, the level 2 system contains logic to provide additional facilities for the small system user.

A3.4 Level 3 System. It is characteristic of a level 3 system that its label and file processing functions are included in a resident monitor system and that selection from among options precedes any operations on the file or its labels. That is, the operating system is permanently resident in main storage, or partly in main storage and partly on auxiliary storage. Common label and file processing routines are linked with the application program later than at assembly or compilation, or centralized label and file processing routines are invoked at execution by the application program code.

Typically, the label and file processing functions are parameterized, and significant options are available using execution time operating system control statements. Typically, these options can override operating system control statements or common programming language declarations submitted with the application program source statements.

A3.5 Level 4 System. It is characteristic of a level 4 system that it may have the capability of highly interpretive control functions or of late selection from among options available for control functions. A level 4 system uses the information in the second file label to adapt itself to the file (unless this information is supplied from other sources).

Appendix B

Utilization of This Standard

B1. Introduction

This Appendix is intended to provide a guide for the implementation and utilization of this standard. It explains how elements of the standard can be used. Thus, to a large extent, it explains the rationale for this particular design of labels and file structures.

B2. Use of Tape Marks

B2.1 Switching of Mode of Processing. The use of tape marks as delimiters between file data and label groups and between certain label groups implies, for example, the following processing:

(1) Tape mark read while processing Beginning-of-Volume Label Group:

Switch to processing file blocks.

(2) Tape mark read while processing Beginning-of-File- Section Label Group:

Switch to processing file blocks.

(3) Tape mark read while processing file blocks:

Read next block.
If EOV1: Switch to processing End-of-File-Section Label Group.

If EOF1: Switch to processing End-of-File Label Group.

If anything else: Error.

(4) Tape mark read while processing End-of-File-Section Label Group:

Switch volumes.
Read one block.
If VOL1: Switch to processing Beginning-of-Volume Label Group.

If anything else: Error.

(5) Tape mark read while processing End-of-File Label Group:

Read next block.
If HDR1: Switch to processing Beginning-of-File-Section Label Group.

If tape mark: Terminate processing file set.
If anything else: Error.

Because there is no tape mark at the end of the Volume-Header or the User Volume Label Sets, the appearance of a first File-Header Label (HDR1) signals the switch to processing of the Beginning-of-File-Section Label Set.

B2.2 Framing of an Empty File Section. Two consecutive tape marks appear preceding EOV1 in Fig. 2 and at the beginning of the second volume (after VOL1 and HDR1) in Fig. 3; yet they are not interpreted as a double tape mark, but rather as framing an empty file section.

Conventional processing can proceed as follows:

Read and process HDR1.
Read and process any additional labels, as required.
Read tape mark: Switch to processing of file blocks.
Read tape mark: Switch to processing of labels.
In Fig. 2, read EOV1 occurring within File-B.
In Fig. 3, read EOF1 occurring at end of File-A.

This is a special case of mode switching, described in B2.1.

B2.3 Use of Double Tape Mark at End of Reel. The double tape mark at the end of each volume (see Fig. 1 through 3) permits either of the following procedures to be used in accomplishing the operation "Forward-space File":

(1) If a system has prior knowledge of which volume ends the volume set:

```
Having read HDR1,
A: Index forward until three
    tape marks are passed.
   Read next block.
    If HDR1:                One file has been indexed.
    If tape mark:
      If end of set:        File does not exist.
      If not end of set:    Rewind volume.
                            Alternate volumes.
                            Read and verify VOL1 and
                            HDR1 on next volume.
                            Return to A.
```

(2) If a system does not have prior knowledge of which volume ends the volume set:

```
Having read HDR1,
A: Index forward until two
    tape marks are passed.
   Read next block.
    If EOV1:  Rewind volume.
                                Alternate volumes.
                                Read and verify VOL1 and
                                 HDR1 on next volume.
                                Return to A.
    If EOF1: Index forward until
              one tape mark passed.
             Read next block.
              If HDR1:          One file has been indexed.
              If tape mark:     End of set.
                                File does not exist.
```

Thus, the double tape mark prevents tape runaway in forward spacing.

B3. Use of Fields in the Labels

B3.1 General. In this Appendix, any reference to use of a field in the File-Header Label Set (HDR1-HDR9) is equally applicable to that field in the End-of-Volume Label Set (EOV1-EOV9) or in the End-of-File Label Set (EOF1-EOF9), when reading backward, unless otherwise noted.

B3.2 "a" Characters. An "a" character is one of the set of the digits 0, 1, ..., 9, the uppercase letters A, B, ..., Z, and the following special characters:

SP ! " % & ' ( ) * + , - . / : ; < = > ?

These characters were chosen from the center four columns of the code table specified in ANSI X3.4-1977, omitting position 5/15 and those positions where there is provision for alternate graphic representation.

This limitation on "a" characters is intended as a guide to provide maximum interchangeability and consistent printing especially during international interchange. Checking for this limitation is not implied.

B3.3 Volume-Header Label (VOL1). The Volume-Header Label (VOL1) identifies the physical reel of magnetic tape, and the contents of that label relate to the identity of the volume.

B3.3.1 Volume Identifier (CP 5-10). Volume Identifier may be used to ensure that the proper volume is mounted. However, Volume Identifier alone may not be sufficient to establish unique identification in interchange. Positive identification can be secured by the combination of Volume Identifier and Owner Identifier (VOL1 CP 38-51).

B3.3.2 Accessibility (CP 11). Accessibility 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 would probably be accommodated in a government-specified User Volume Label (UVLn), but this field might be used as an indicator in conjunction with such a User Volume Label. For that reason, it is made available to the installation label processing routine by the system.

An Accessibility field appears in both the Volume-Header Label (VOL1) and first File-Header Label (HDR1), so that this function can be exercised either for the entire volume, or for each individual file, as appropriate. When this function is exercised, the system is responsible for specification of the processing of Accessibility (HDR1 CP 54). That is, the system implementors determine what support will be available beyond the minimum support required by this standard. If such support is allowed under a system, it may include processing of system-specified values or installation-specified values, or both. See 7.2.2 and 7.5.9 for the requirements on the system when additional support beyond the minimum is provided.

B3.3.3 Owner Identifier (CP 38-51). Owner Identifier identifies the owner of the volume. The name selected by the owner can also serve to identify the installation to which the volume belongs. Recognition of the Owner Identifier by the installation security routine and installation label processing routine will permit correct interpretation of Accessibility and of the User Volume Label Set (UVL1-UVL9). At the option of the installation the value may represent an individual user (owner) and be supplied by the user. It is anticipated that a standard method of identifying an owner will be defined. In the absence of such a standard, the parties should agree among themselves to choose identifiers so that each party is identified uniquely within the specific interchange environment.

B3.3.4 Label-Standard Version (CP 80). Label-Standard Version is used to indicate the version of this standard to which the labels and data formats on this volume conform. It also provides a means of extending this standard in the future, should the need arise, with minimum conflict between the future version of the standard and parochial practice that may develop in the meantime. It is intended to distinguish among future versions of this standard by the use of numerals in this field, rather than letters, to the extent possible. Letters are available for parochial practice.

B3.4 First File-Header Label (HDR1)

B3.4.1 File Identifier (CP 5-21). File Identifier may be used to ensure that the correct file within a multifile set is read. This field can be used for a search of a file set if the File Sequence Number (HDR1 CP 32-35) is not known, or it can be used to verify that the correct file is located at a given File Sequence Number (HDR1 CP 32-35).

B3.4.2 File-Set Identifier (CP 22-27). It is desirable that a unique identification within an installation be established. In most cases, this objective may be satisfied by duplicating Volume Identifier (VOL1 CP 5-10) of the first or only volume of the set as File-Set Identifier for that set.

B3.4.3 File Section Number (CP 28-31). File Section Number may be used to ensure that the file sections within a multivolume file set are read in the correct sequence. If it is not necessary to read the entire file, the File Section Number can also be used to identify a particular file section. The actual beginning of a file may be identified by "0001" in this field. Subsequent file sections are numbered sequentially on subsequent volumes.

B3.4.4 File Sequence Number (CP 32-35). File Sequence Number can be used by a system to position a volume to a particular file in conjunction with a requested file sequence number. It is most frequently used in directories or catalogs of files in file sets.

B3.4.5 Generation Number (CP 36-39). Generation Number is used to differentiate between different issues of a file produced periodically and bearing the same File Identifier (HDR1 CP 5-21). For example, it may be used to distinguish between different weekly files of a payroll program, between successive updates of a text file, or between different production runs of a design automation file.

B3.4.6 Generation Version Number (CP 40-41). Generation Version Number is used to differentiate between output data produced by repeated processing or writing operations that 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 a return to a rescue point, or between a file processed with input in error and the new copy of the same information with input corrected. If the file is partially re-created, for example, after a rescue point at the beginning of a file section, then Generation Version Number, and possibly Creation Date (HDR1 CP 42-47) and Expiration Date (HDR1 CP 48-53), may vary from the preceding file sections.

B3.4.7 Creation Date (CP 42-47). Creation Date may be used as a qualification of the identification of a file in interchange. If the possibility exists of creating, on the same date, multiple generations or generation versions of a file, then the Creation Date will not be a unique qualifier.

B3.4.8 Expiration Date (CP 48-53). The file is regarded as "expired" on a day whose date is equal to or later than Expiration Date. When this condition is satisfied, the remainder of this volume set may be overwritten. To be effective on multifile volumes, therefore, Expiration Date of a file is earlier than or equal to the expiration dates of all previous files on the volume set.

B3.4.9 Accessibility (CP 54). See B3.3.2.

B3.4.10 Block Count (CP 55-60). Block Count is most useful at the end of a file section. See B3.6.1.

B3.4.11 System Code (CP 61-73). A system verifies System Code if it makes use of system controls on file Accessibility (HDR1 CP 54), Reserved for System Use (HDR2 CP 16-50), of buffer offset (see B3.5.5), or of any other system labels for the file (HDR3-HDR9, EOV3-EOV9, EOF3-EOF9). It may also be used simply to provide an identification of the system. It is likely that, in time, a standard method of identifying the system will be defined.

B3.5 Second File-Header Label(HDR2). When used, this label contains fields, the contents of which are specified in this standard, and thus can be used effectively in information interchange. In addition, this label contains Reserved for System Use (CP 16-50) and, as explained in B3.5.4, the contents of this field are ignored in interchange. System Code (HDR1 CP 61-73) is used to identify the particular system that created these labels.

B3.5.1 Record Format (CP 5). A system may use Record Format to verify that it is capable of accessing the records in the blocks.

B3.5.2 Block Length (CP 6-10). A system may use Block Length to determine the length of a buffer required for reading.

B3.5.3 Record Length (CP 11-15). A system may use Record Length to determine the length of a work area required for application program processing of the record. If Record Format (HDR2 CP 5) is S, then Record Length contains the maximum record length, excluding all the Segment Control Words. Thus, Record Length does not reflect the number of characters actually recorded in all segments of a spanned record (the sum of segment lengths in the Segment Control Words) but rather the length of the record after the segments have been concatenated into a single record, without consideration for any Segment Control Words. If Record Format (HDR2 CP 5) is S and Record Length is greater than the length of a work area, or if Record Length is 00000 and the record received is greater than the length of a work area and the content of the record is such that it cannot be meaningfully processed except as a complete record, an error has occurred. This does not require an implementation to provide, nor does it preclude an implementation from providing, as an extension, the facility to deliver a spanned record to the application program resegmented to fit the work area. If such an extension is provided, the system is made aware that segments of spanned records can be meaningfully processed by the particular application program.

B3.5.4 Reserved for System Use (CP 16-50). A system may use this field to contain any information needed to increase the efficiency of label and file processing. This information is pertinent to label and file processing functions more specialized than those contemplated in this standard. These functions may be developed, defined, and implemented differently by designers of different systems. Thus the nature of this field makes it fundamentally incompatible between different systems, and the contents of this field are ignored in interchange. Such functions, the contents of this field, and provisions for processing these contents are defined by the implementors of a system. System Code (HDR1 CP 61-73) is used to identify the particular system that created this label.

B3.5.5 Buffer-Offset Length (CP 51-52). A system may use Buffer-Offset Length to determine the location of the first character of the first record in a block. Certain applications 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, date, time of transmission, identification of the program that generated the information, etc. The contents of buffer offset, and provision for processing these contents, are defined by the implementors of a system. The System Code (HDR1 61-73) is used to identify the particular system that created the buffer offset.

B3.6 First End-of-Volume/File Label (EOV1, EOF1). These labels contain the same fields as the first File-Header Label (HDR1) to enable identical label processing when reading a file backward. This is of most use to a system that provides for reading backward and is capable of adapting to the attributes of the file.

B3.6.1 Block Count (CP 55-60). Block Count is used by a system to determine, when a file section is read, whether any blocks were skipped or whether any spurious blocks were inserted, except for compensating errors of equal numbers of skipped and spurious blocks.

When reading forward, the count of blocks read is maintained and compared with Block Count at the end of the file section. When reading backward, Block Count is used to establish a count that is decremented by one for each block read, and is compared with zero at the beginning of the file section. This check is not effective if reading of a file section is discontinued before reaching the end or beginning of the file section.

If the Block Count check fails on input, the system discontinues further processing of the file. The system communicates this information to an operator, a system's administrator, or the application program at the option of the system implementors. At the option of system implementors, a system may permit the operator, the system's administrator, or the application program to ignore the failure and continue processing.

At the option of system implementors, a system may provide automatic rerun of the file section in error.

B3.7 Second End-of-Volume/File Label (EOV2, EOF2). These labels contain the same fields as the second File-Header Label (HDR2) to enable identical label processing when reading a file backward. This is of most use to a system that provides reading backward and is capable of adapting to the attributes of the file.

B4. Use of Additional Labels

B4.1 General. The previous sections of this Appendix have described the use of the contents of the fields in the first and second system labels (VOL1, HDR1, EOV1, EOF1, HDR2, EOV2, EOF2) in detail. The use of fields in the other labels is amplified in this section.

B4.2 Other System Labels (HDR3-HDR9, EOV3-EOV9, EOF3-EOF9). When used, these labels contain only Reserved for System Use (CP 5 - 80) and, as explained in B3.5.4, the contents of these labels are ignored in interchange. System Code (HDR1 CP 61-73) is used to identify the particular system that created these labels.

B4.3 User Volume Labels (UVL1-UVL9). It is often convenient to use these labels for security controls or to contain information about the volume, such as purchase date, vendor, date of last certification, condition, length, department assignment, and the like. This information is useful for accounting and assignment control of the volumes, a practice usually parochial to a specific installation. Thus the contents of these labels are ignored in interchange. Owner Identifier (VOL1 CP 38-51) is used to identify the particular installation that created these labels.

If an operating system provides support for UVL labels, it provides an interface, when Owner Identifier is not all spaces, for the installation to create and utilize the labels in conjunction with Owner Identifier.

If UVL labels are not supported by an operating system, they are not created on output or made available on input to the installations that use that operating system.

If UVL labels are not supported by an operating system, they are ignored and bypassed when they are found on a volume to be used for output.

Some examples of possible operating system interface support and the functions made possible by that support are as follows:

(1) Only one UVL label can be created by an installation on output, and that label cannot be updated later.

(2) Only one UVL label can be created by an installation on output, and it can be updated and rewritten when all following records on the tape volume are also written.

(3) More than one UVL label can be created by an installation on output, each of which is made available individually on input, and none of the labels can be updated later.

(4) More than one UVL label can be created by an installation on output, all of which are made available as a set on input, and none of which can be updated later.

(5) More than one UVL label can be created by an installation on output, each of which is made available individually on input, and only one label can be updated and rewritten later with all following records on the tape volume also written.

(6) More than one UVL label can be created by an installation on output, either through a single transfer of the entire label set and of control to the operating system or through multiple transfers with one UVL label provided to the operating system with each transfer, and all of the label set updated and rewritten later, with all following records in the tape volume also written.

The use of any of the different possible implementations of UVL in this nonexhaustive list does not restrict or impede general information interchange. This is because UVL labels are not utilized for interchange except between installations with mutually known installation volume handling practices.

B4.4 User File Labels (UHLa, UTLa). It is often convenient to use these labels to contain summary information about the file section, file, or file set being interchanged, such as control totals, statistical tabulations, and the like. In such a case, that information is quite useful to the recipient of the file, so user file labels are part of the information being interchanged. A description of their contents can be forwarded by the originator of the file to the recipient. In absence of such knowledge, the contents of user labels are ignored in interchange.

Note that the length of (number of records in) a file section may vary between copies of a file, because of error recovery (erase operations), change in file configuration on the volume set, change in recording density, change in the number of records per block, etc. User file labels that depend upon the physical configuration of a file section are subject to change when the physical configuration of a file section is changed. Users should consider this factor when designing these labels.

B5. Volume Initialization

B5.1 Volumes to Be Used. Blank tapes or previously initialized volumes to which access for reinitialization has been verified are initialized by either a utility program or a routine integral to the output function.

B5.2 Arrangements of Labels for Interchange. Volumes to be sent to another installation for use in interchange are initialized to contain at least:

(1) A Beginning-of-Volume Label Group, as specified in 5.6

(2) A Beginning-of-File-Section Label Group, as specified in 5.7

(3) Two tape marks framing an empty file, as specified in 5.8.3

(4) An End-of-File Label Group, as specifed in 5.10

(5) A double tape mark, as specified in 5.10.5

B5.3 Contents of Labels. A Beginning-of-Volume Label Group contains valid values in all fields of the Volume Header Label (VOL1) as specified in 4.3. If the installation does not specify a value for Owner Identifier (CP 38-51), spaces are placed in that field by the system.

A Beginning/End-of-File Label Group contains the following fields and values in the first File-Header Label (HDR1) and first End-of-File Label (EOF1):

File Identifier	(CP 5-21)	any valid value
File-Set Identifier	(CP 22-27)	any valid value
File Section Number	(CP 28-31)	0001
File Sequence Number	(CP 32-35)	0001
Generation Number	(CP 36-39)	0001
Generation Version Number	(CP 40-41)	00
Creation Date	(CP 42-47)	any valid value
Expiration Date	(CP 48-53)	any expired date
Accessibility	(CP 54)	space
Block Count	(CP 55-60)	000000
System Code	(CP 61-73)	system specified
Reserved for Future Standardization	(CP 74-80)	spaces

Note that an initialized volume is one that is intended to be used for output and will have its system header label rewritten with appropriate updates of fields. Additional system labels used by the system will then be written.

B5.4 Processing. The system prepares and writes the Volume-Header Label (VOL1). Then the installation volume label processing routine is invoked to prepare the User Volume Label Set (UVL1-UVL9). Finally, the system writes a labeled empty file and an End-of-File Label Group, and terminates the volume with a double tape mark.

B5.5 On-Line Operation. An output function for on-line initialization may be provided at the option of the implementors of the system. A volume that is to be written with new files and file labels and that will retain no previously written files or file labels (a "scratch" volume) may be initialized with a new Volume-Header Label Set and a new User Volume Label Set before the file labels and file are written.

B5.6 Changes to Volume Contents on Reinitialization. Reinitialization may change the character code and the format of the labels of the volume to the character code and the label formats as specified by this standard. Reinitialization may change the tape density of the volume. The system implementation that performs this reinitialization reads and properly recognizes a Volume-Header Label as specified in this standard, in an earlier version of this standard, or in any other installation or parochial system standard. The Volume-Header Label contains a unique volume identifier field.

The volume is written as specified in this standard. The source of the identifier to be placed in Volume Identifier (VOL1 CP 5-10) is the same or analogous field on the scratch volume, restructured if necessary to conform to this standard. Any restrictions on access to the volume imposed by this or an earlier version of this standard (that is, Accessibility (VOL1 CP 11)), or any analogous restrictions imposed in any other standard the installation follows, are satisfied before allowing the volume to be overwritten. Fields in labels defined in parochial standards that are applicable to the fields defined in the volume label of this standard may be restructured to conform to this standard and included in the volume label.

If any User Volume Labels or an analogous set of parochial labels exist on a volume, they may be restructured to conform to this standard and rewritten on the new volume.

The first file on a volume written under this or any earlier version of this standard, or under any standard that the installation follows that contains fields analogous to Accessibility (HDR1 CP 54) and Expiration Date (HDR1 CP 48-53), is verified to be an expired file to which access is permitted.

The requirements on the Volume-Header Label do not preclude any updating of the contents of the defined fields of the Volume-Header Label (for example, make Label-Standard Version a 3). Similarly, the requirements on the User Volume Labels do not preclude any updating of, or additions to the contents or the quantity of the User Volume Labels (such as a new date of initialization in a UVL label). The imposition of additional checks on security, etc by a system before permitting the volume to be overwritten are not precluded.

The system implementors may permit a system administrator, operator, or other supervisory personnel to override any of these checks in order to avoid undue delay in processing.

B6. Copying

B6.1 Files. A file to be interchanged is either contained on the original volume or is contained as a copy on another volume. In either case the interchanged volume contains a Volume-Header Label (VOL1) with a Volume Identifier (VOL1 CP 5-10) that is unique at least within the originating installation. Further qualification of the volume identification can be made using Owner Identifier (VOL1 CP 38-51), which may contain an installation-assigned value.

If the copy is one of multiple copies of the same file, or a backup copy of a file, then one set of specifications for establishing a copied file is as follows:

(1) The following fields are unchanged:

File Identifier	(HDR1 CP 5-21)
Generation Number	(HDR1 CP 36-39)
Generation Version Number	(HDR1 CP 40-41)
Creation Date	(HDR1 CP 42-47)
Record Length	(HDR2 CP 11-15)

(2) The following fields reflect the new configuration of the file set:

File Set Identifier	(HDR1 CP 22-27)
File Sequence Number	(HDR1 CP 32-35)

(3) The following fields reflect the new configuration of the file:

File Section Number	(HDR1 CP 28-31)
Block Count	(HDR1 CP 55-60)

(4) At the option of system implementors, a system may permit the following fields to vary:

Expiration Date	(HDR1 CP 48-53)
Accessibility	(HDR1 CP 54)
System Code	(HDR1 CP 61-73)
Record Format	(HDR2 CP 5)
Block Length	(HDR2 CP 6-10)
Buffer-Offset Length	(HDR2 CP 51-52)

(5) The system copies Reserved for System Use (HDR2 CP 16-50) and other system labels (HDR3-HDR9, EOV3-EOV9, EOF3-EOF9) unless the system changes System Code (HDR1 CP 61-73).

(6) The length of a file section is arbitrary, and a system copying a file may change sections to fit the volumes being written.

(7) The blocking factor is arbitrary, and a system copying a file of fixed-length records (F) or variable-length records (D) may reblock records to fit the blocks being written.

(8) The length of a record segment is arbitrary, and a system copying a file of spanned records (S) may resegment the records to fit the blocks being written.

(9) At the option of system implementors, a system may permit a user to avoid changes in file sectioning, blocking, or segmenting.

B6.2 Volumes. Duplicating a complete volume results in duplicated Volume Identifier fields (VOL1 CP 5-10) on the copied-to and copied-from volumes. Good installation control of tape volumes may be impacted by indiscriminate copying of full volumes. The practice should be avoided whenever possible for that reason.

If the installation requirements dictate, a modified copy of a full volume including the system volume labels and user volume labels may be created. When such a copy is produced, the Volume Identifier (VOL1 CP 5-10) field of the copied-to volume is retained on that volume. Some implementations of this standard utilize the Volume Identifier (VOL1 CP 5-10) of the first volume of a file set in File Set Identifier (HDR1 CP 22-27) of each file in the file set. When the first volume of a file set is copied using the modified-copy process, the copied-from volume should not be used subsequently to contain the first volume of a file set with any file identifiers that appear in the file set. User Volume Labels may contain information that is relevant only to the physical volume on which they are written. In such a case this information in these labels will not be valid when the volume is copied. Any errors in the writing of the copy will possibly result in skipped areas of erased tape that will not be present on the original volume and that will influence the length of tape required. Satisfactory completion of such a copy operation requires that the used portion of the copied-to tape be as long as or longer than the used portion of the copied-from tape.

B7. Use of Undefined Record Formats

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 record formats are interchanged, it is not within the purview of this standard.

B8. Verification of Access

B8.1 General. Before a system grants access to a volume to record new data, or to a file to read or overwrite existing data, verification of access is necessary to prevent inadvertent destruction or misuse of data.

Verification of access involves three factors: expiration of existing data to be overwritten, identification of existing data to be used, and establishment of authorization to use the resource.

If verification fails, the system discontinues further processing of that volume. Overrides of failure of verification are described in this section.

B8.2 Control over Access

B8.2.1 Expiration of Existing Data. One requirement for access to a volume set to write new data is that the Expiration Date (HDR1 CP 48-53) for the existing data on the volume set past that point, if any, be expired. The system may request that another volume be mounted when an unexpired file is encountered.

At the option of system implementors, a system may permit an override of an unexpired date to avoid undue delay in processing. This override capability is not available to the application program.

At the option of system implementors, a system may permit further checks of an expired file before the file is purged.

At the option of system implementors, more stringent rules for expiration (buffer period beyond expiration date, minimum retention period, physical deletion of nonaccessible volumes or files, etc) may be applied.

B8.2.2 Identification of Existing Data. One requirement for access to a file is to uniquely identify the file. The name alone of a file is not always sufficient to establish unique identification. Multiple generations or versions of some files may exist, each of which bears the same File Identifier (HDR1 CP 5-21). Unique file identification can almost certainly be attained by a combination of File Identifier, File-Set Identifier (HDR1 CP 22-27), Generation Number (HDR1 CP 36-39), Generation Version Number (HDR1 CP 40-41), and Creation Date (HDR1 CP 42-47). As a practical matter, positive identification can often be secured by either the combination of File Identifier and File-Set Identifier (HDR1 CP 22-27), or the combination of File Identifier, Generation Number, and Generation Version Number. The combination of File Identifier and Creation Date, while not always unique, can be used to identify a file in interchange; the assumption is that different versions will not be in distribution simultaneously. Volume Identifier (VOL1 CP 5-10) and File Sequence Number (HDR1 CP 32-35) may be used to locate a file, but neither is used to establish file identification.

At the option of system implementors, a system may maintain (for example, catalog) file identification, and associate the file identification with a symbol (such as a synonym, a program, or a user's name).

At the option of system implementors, a system may permit an operator, a system administrator, etc, but not a user, to override a failure in file identification to avoid undue delay in processing.

B8.2.3 Establishment of Authorization. One requirement for access to a volume set to write new data is that the volume Accessibility (VOL1 CP 11) be a space. If it is not a space, then additional controls are exercised.

One requirement for access to an existing file is that both volume Accessibility (VOL1 CP 11) and file Accessibility (HDR1 CP 54) be spaces. If either is not a space, then additional controls are exercised.

No explicit additional controls are specified in this standard. The following are examples of additional controls:

(1) The character encountered in Accessibility is examined by an installation-provided security routine to determine that the routine recognizes and accepts the code as a valid accessibility indicator. The user or program may provide a matching or encompassing authorization indicator.

(2) The file is associated with a string of characters, maintained in a system file. The user or program may provide a matching string of characters (password).

At the option of system implementors, a system may permit an override of the accessibility check to avoid undue delay in processing. This override capability is not available to the application program.

B8.2.4 Unauthorized Transfer of Control. In order to better ensure the security of data in a file, the system should prevent possible access to that file if an unauthorized transfer of program control is attempted. This may be done by forcing the closing of the file and rewinding and unloading the volume on which it resides.

B8.3 Verification Procedures

B8.3.1 Volume Initialization. Access to initialize a (new or degaussed) blank tape is always granted.

Access to rewrite unchanged an existing Beginning-of-Volume Label Group is granted if the existing first file has expired (See B8.2.1) and if authorization to access the volume, authorization to access the existing first file, and identification of the existing first file have been given.

Access to reinitialize an existing volume, that is, to overwrite and change an existing Beginning-of-Volume Label Group (including a change in recording mode or density) is granted if the existing first file has expired (See B8.2.1), the owner identified in Owner Identifier (VOL1 CP 38-51) has authorized the change, and if authorization to access the volume, authorization to access the existing first file, and identification of the existing first file have been given.

Access to rewrite or reinitialize an existing Beginning-of-Volume Label Group is not granted if the existing Beginning-of-Volume Label Group cannot be read because of hardware error, recording mode, density, parochial practice, or any other reason. This restriction does not apply to a utility program specifically designed to initialize or to reinitialize volumes.

At the option of system implementors, a system may permit an override of the access checks to avoid undue delay in volume initialization. This override capability is not available to an application program.

B8.3.2 New File. Access to a volume to record new data is granted if authorization is established to access the volume, if the existing file to be overwritten has expired (See B8.2.1), and if authorization to access and identification of the existing file to be overwritten have been given.

At the option of system implementors, a system may permit requests for specific Volume Identifier (VOL1 CP 5-10), same volume as another file of the same file set, any available (scratch) volume, or other specifications, or any combination of these.

At the option of system implementors, a system may permit an override of the access checks to avoid undue delay in processing. This override capability is not available to an application program.

B8.3.3 Existing File. Access to a file to read existing data is granted if the file is identified and authorization is established to access the file.

Access to a file to extend or overwrite a portion of the existing data, incidentally to read existing data, is granted if the file is identified, authorization is established to access the file, and if the existing data have expired.

At the option of system implementors, a system may permit an override so that a portion of an unexpired file may be overwritten or extended. This override capability is not available to an application program.

Appendix C

Differences Between Versions of This Standard

C1. 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 Version 1 of this standard, ANSI X3.27-1969. When the Label-Standard Version (VOL1 CP 80) has the value 3, it means that the volume is written according to the rules and options of Version 3 of this standard, the present edition.

C2. Differences and Reasons for Differences

C2.1 General. The reasons for the differences between the versions fall into three categories:

(1) Correction of technical errors and ambiguities in Version 1.

(2) Clarification of requirements. This reduces the number of possibilities the system has to deal with and reduces ambiguity; for example, how the system writing the label and the system reading the label should deal with optional fields in mandatory labels and optional fields in optional labels is made more exact.

(3) Enhancing the technical content of this standard.

C2.2 Specific Differences. The differences between Version 1 and Version 3 are outlined in this section. In this Appendix, the meanings of the flags are as follows:

* Increase in requirements by addition of new material, or by moving material from the Appendix to the body of the standard.

\+ Change to specification of format or content of an element recorded on a volume.

/ Clarification or correction of error in specification.

= Substantive editorial change.

(1) Levels of labeling defined:

* In Version 1, subsetting was neither implied nor precluded. In Version 3, specifications for levels of labeled volumes and descriptions of levels of label and file processing functions are added.

This provides explicit subsets, to avoid confusion as to the requirements of this standard, and to ensure that products implementing only a subset can successfully interchange standard volumes with each other.

(2) Scope expanded:

* (a) In Version 3, requirements on implementations and volumes are added to the scope in Version 1.

This is to clarify how to implement this standard.

* (b) In Version 3, field of application is added to the scope of Version 1.

This is to clarify the domain in which this standard can be useful.

= (c) In Version 3, exclusions are added to the scope of Version 1.

This is to clarify the bounds of this standard, and is consistent with the latest style guides.

(3) References added:

= In Version 3, cross-references to other standards pertinent to the interchange of information recorded on magnetic tape are added.

This is to simplify the task of an implementor or user to determine where the specifications of which he should be aware can be found.

(4) Definitions expanded:

= In Version 3, the following definitions are added to those in Version 1, with differences as noted after the list:

application program
blocked record
double tape mark
file section
fixed-length record
label and file processing functions
label set
record segment
spanned record
unblocked record
unspanned record
user
variable-length record
volume set

All definitions were modified, but none substantively.

This is to enhance the understandability of this standard.

(5) Format of strings of characters in fields defined:

+ In Version 3, the requirement on adjustment of values in fields and on fill characters is added to the requirements of Version 1.

This is to ensure matches on comparisons of equal values.

(6) Column Headings for 4.3 through 4.12 changed:

= (a) In Version 1, "description" included format, processing, and intent, but not always consistently. In Version 3, "content" includes only format and intent, consistently. Processing appears elsewhere in Version 3.

This is to improve the organization of this standard.

= (b) In Version 1, reference was to field number. In Version 3, reference is to character position.

This is to permit addition or deletion of field specifications without upsetting the reference numbers of the fields following the changed field in the label.

* (c) In Version 1, the word "optional" appeared adjacent to the specification for the following fields:

Generation Number	(HDR1 CP 36-39)
Generation Version Number	(HDR1 CP 40-41)
System Code	(HDR1 CP 61-73)
Buffer-Offset Length	(HDR2 CP 51-52)
	(EOV1 CP 5-54)
	(EOV1 CP 61-80)
	(EOF1 CP 5-54)
	(EOF1 CP 61-80)
	(EOF2 CP 5-80)

In Version 3, this word does not appear.

This is to ensure that the labels accurately describe the file.

(7) Volume Identifier (VOL1 CP 5-10) changed:

= In Version 1, the name was "Volume Serial Number." In Version 3, the name is "Volume Identifier."

This is to reflect that the field may contain nonnumeric characters, and to achieve consistency in the names of all identifiers.

(8) Owner Identifier (VOL1 CP 38-51) changed:

= In Version 1, the name was "Owner Identification." In Version 3, the name is "Owner Identifier."

This is to achieve consistency in the names of all identifiers.

(9) Label-Standard Version (VOL1 CP 80) changed:

= (a) In Version 1, the name was "Label Standard Level." In Version 3, the name is "Label-Standard Version."

This is to avoid the implication of subsets or supersets.

+ (b) In Version 1, the value was 1. In Version 3, the value is 3.

This is to distinguish between the versions.

+ (c) In Version 1, the description included the specification, "'space' means the labels and data formats on this volume require the agreement of the interchange parties." In Version 3, this specification does not appear.

This is to avoid stating specifications for nonusers of this standard.

(10) File-Set Identifier (HDR1 CP 22-27) changed:

= In Version 1, the name was "Set Identification." In Version 3, the name is "File-Set Identifier."

This is to avoid ambiguity as to what kind of set is being identified and to achieve consistency in the names of all identifiers.

(11) Record Format (HDR2 CP 5) changed:

/ (a) In Version 1, V (variable length with binary count field) appeared. In Version 3, this format does not appear.

This is because it is not feasible to attempt to interchange mixed codes between systems of possibly different architecture.

/ (b) In Version 1, U (undefined) appeared. In Version 3, this format does not appear.

This is because it is not feasible to utilize undefined format records.

+ (c) In Version 3, S (spanned records) is added to the record formats in Version 1.

This is to enhance this standard, and to satisfy a stated requirement of American National Standard for Bibliographic Information Interchange on Magnetic Tape, ANSI Z39.2-1971.

(12) Reserved for System Use (HDR2 CP 16-50) changed:

= In Version 1, the name was "Reserved for Operating Systems." In Version 3, the name is "Reserved for System Use."

(13) Buffer-Offset Length (HDR2 CP 51-52) changed:

= (a) In Version 1, the name was "Buffer Offset." In Version 3, the name is "Buffer-Offset Length."

This is to more accurately describe the field that specifies the length of a buffer offset.

= (b) In Version 1, the description stated, "inserted before a data block." In Version 3, content states, "inserted before the first record in a data block."

This is more consistent with the specification of Block Length (HDR2 CP 6-10).

(14) Sequence of End of Volume, End of File changed:

= In Version 1, lists of labels were in the sequence VOL1, HDR1, EOF1, EOV1. In Version 3, lists of labels are in the sequence VOL1, HDR1, EOV1, EOF1.

This is to clarify the concept that EOV1 is an end-of- file-section label, whereas EOF1 is a label at the end of the entire file.

(15) Second End-of-Volume Label added:

/ In Version 3, the specifications for second End-of-Volume Label (EOV2) are added to those in Version 1.

This is to correct an oversight in Version 1.

(16) Name of additional labels changed:

= In Version 1, additional labels were called "Other Optional Operating System Labels," and the name was "Operating System Option." In Version 3, additional labels are called "Other System Labels," and the field name is "Reserved for System Use."

(17) Reserved for Installation Use (UVLn) changed:

= In Version 1, the field name was "User Option." In Version 3, the field name is "Reserved for Installation Use."

(18) Requirement that system label sets be symmetric added:

+ In Version 3, the requirement that system label sets be symmetric about file sections is added.

This is to permit equal facility in reading a file forward or backward.

(19) Consideration of a volume beyond the end of data added:

+ In Version 3, the specification that the contents of the volume beyond the double tape mark at the end of the volume not be considered in information interchange is added.

This is to avoid any implicit requirement for examining or erasing volumes beyond the end of data recorded according to this standard.

(20) Alternate method for ending volumes deleted:

* In Version 1, the alternate method for ending volumes by agreement of the interchange parties appeared in the Appendix. In Version 3, the alternate method for ending volumes does not appear in the standard.

(21) Processing of Label Fields added:

= In Version 1, some of the processing information appeared in "Format and Content of Labels" and some in the Appendix. In Version 3, all of the information is collected into one section. The source of the contents of label fields is clarified.

* (a) Use of information in label fields: In Version 3, this section is added to reduce the ambiguity in how to process the contents of label fields.

+ (b) Volume Identifier (VOL1 CP 5-10): In Version 3, the requirement that Volume Identifier be unique in an installation is added.

This is to ensure more precise identification of a file.

/ (c) User Volume Labels (UVL1-UVL9): In Version 1, no indication appeared as to how to determine if a particular label was significant to a particular installation. In Version 3, these labels are correlated with Owner Identifier (VOL1 CP 38-51).

This ensures positive identification of the creator of these labels.

/ (d) Owner Identifier (VOL1 CP 38-51): In Version 3, the recommendation that if this field is all spaces User Volume Labels should not be used is added to Version 1.

This assists in positive identification of the creator of these labels.

+ (e) File Identifier (HDR1 CP 5-21): In Version 3, the requirement that File Identifier be unique in a file set is added.

This is to ensure a more precise identification of a file. This also enhances the capability to search a file set for a specific file by name.

/ (f) Generation Version Number (HDR1 CP 40-41): In the Appendix of Version 1, this number was specified to be zero for the first attempt to produce a file. In the body of Version 3, this number is specified to be zero for the first attempt to produce a generation of a file.

This clarifies that Generation Version Number is reset to zero for each generation.

/ (g) Expiration Date (HDR1 CP 48-53): In Version 1, this field controlled the remainder of a volume. In Version 3, this field controls the remainder of a volume set.

This reflects that all of a file set is unusable beyond any part that is overwritten.

/ (h) Block Count (HDR1 CP 55-60): In Version 1, this was stated to be the count of blocks written. In Version 3, this is the count of blocks appearing in the file section.

This clarifies the case that a block may be written, backspaced, and rewritten (blocks written, 2; block count, 1).

/ (i) System Code (HDR1 CP 61-73): In Version 3, the recommendation that if this field is all spaces unique system facilities should not be used is added to Version 1.

This assists in positive identification of the creator of unique system facilities.

/ (j) Block Length (HDR2 CP 6-10): In Version 1, it was not stated whether Block Length included padding. In Version 3, it is stated that Block Length includes padding.

/ (k) Reserved for System Use (HDR2 CP 16-50): In Version 1, no indication appeared as to how to determine whether a particular field was significant to a particular system. In Version 3, this field is correlated with System Code (HDR1 CP 61-73).

This clarifies an ambiguity in Version 1.

= (l) Buffer-Offset Length (HDR2 CP 51-52): In Version 3, the requirement that this field is zero if buffer offset is not used is added.

This is to ensure that the labels accurately describe the file. In a sense, this is not a new requirement, because it is unambiguously derivable from other specifications in Version 1.

(22) Processing of labels added:

= In Version 1, some of the processing information appeared in "Format and Content of Labels," and some in the Appendix. In Version 3, all of the information is collected into one section. Various processing situations are described in more detail.

+ (a) In Version 3, the timing of the interfaces with regard to processing of user labels is added.

This is consistent with American National Standard Programming Language COBOL, ANSI X3.23-1974, and is not inconsistent with any other standards for common programming languages. This is intended to avoid inconsistencies in future additions to language specifications of common programming languages.

+ (b) In Version 3, the requirement that User Volume Labels (UVL1-UVL9) be accorded much the same treatment as the Volume-Header Label (VOL1) is added.

This is to preserve and protect accounting or other data pertinent to the volume.

/ (c) In Version 1, the requirement that the File-Header Label Set of a subsequent volume be an exact copy of its predecessor, except for File Section Number (HDR1 CP 28-31) appeared in the Appendix. In Version 3, the requirement that the File-Header Label Set of a subsequent volume be an exact copy of its predecessor, except for File Section Number (HDR1 CP 28-31), Generation Version Number (HDR1 CP 40-41), Creation Date (HDR1 CP 42-47), and Expiration Date (HDR1 CP 48-53), appears in the body of the standard.

This is to be consistent with the definition of Generation Version Number to distinguish among successive iterations, and to permit partial iterations; for example, a rescue point at other than the beginning of the entire file. Furthermore, this is to ensure that a file is homogeneous on all of its volumes.

+ (d) In Version 3, the treatment of errors is expanded.

This is to further safeguard volumes and information recorded on them in information interchange without unduly delaying processing.

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

In Version 1, no indication appeared as to how to determine if a particular label was significant to a particular system. In Version 3, these labels are correlated with System Code (HDR1 CP 61-73).

This clarifies an ambiguity in Version 1.

(23) Additional figures added:

= In Version 3, Fig. 4 through 8, showing unblocked and blocked fixed-length and variable-length records, are added to the figures in Version 1.

This is to clarify this standard. No additional requirements appear.

(24) Padding moved:

* In Version 1, specifications for padding appeared in the Appendix, its use required the prior agreement of the interchange parties, and the recorder required knowledge of the architecture of the recipient. In Version 3, the specification appears in the body of the standard, its use does not require the prior agreement of the interchange parties, and the recorder no longer requires knowledge of the architecture of the recipient.

This is to permit more free interchange from a word-oriented or a block-oriented computing system.

(25) Requirement for file consistency added:

= In Version 3, the requirement that all records in a file be recorded in the same format is added.

This is to ensure that Record Format (HDR2 CP 5), that is constant for a file, accurately describes the file. In a sense, this is not a new requirement, because it is unambiguously derivable from other specifications in Version 1.

(26) Requirement for constant recording density added:

+ In Version 3, the requirement that the blocks on all volumes of a file set be recorded in the same density is added.

This is to ensure that if a system can process any part of a file set, it can process the entire file set.

(27) Block length limits deleted:

= In Version 1, specific maximum block length was specified. In Version 3, a cross-reference to another standard appears. Maximum and minimum block length limits appear in the other standard.

This is to avoid redundant, possibly inconsistent, specifications.

(28) Block Sequence Indicator description deleted:

= In Version 1, the Block Sequence Indicator description appeared in the Appendix. In Version 3, it does not appear.

This is to clarify the actual requirements and to eliminate an option found to be unsatisfactory.

(29) Explanatory material on diagnosis of "a" characters moved:

= In Version 1, explanatory material on the diagnosis of "a" characters appeared in the body of the standard. In Version 3, this material appears in the Appendix.

This is to remove motivational material from the body and to limit the body to specification.

(30) Additional explanatory material on the use of this standard is included in the Appendix of Version 3. This assists implementors and users to take full advantage of the capabilities inherent in the particular design specified in this standard. No additional specifications or requirements are added.

Appendix D. Considerations Associated with Changes and Additions to the Earlier Version

D1. Introduction

D1.1 Role of Standardization in Interchange of Information Recorded on Magnetic Tape. With the completion of standards on the physical aspects of magnetic tape, one might well ask whether or not such standards guarantee the ability to exchange information recorded on magnetic tape between the equipment of various manufacturers. Presuming that they do in the hardware sense, such interchanges are possible.

Probably nothing but a physical standard is required. Once the volume is mounted and the bit patterns are in main storage, it is possible, with some labor, to program an understanding of the information. However, it has been recognized for some time that an environment in which a single character code is used is more satisfactory than where several representations are used. That is, significant economies in recognition and interpretation become available. Hence, a character code for information interchange was specified. A standard character code is desirable to reduce the probability of error and misunderstanding, to reduce programming labor, and to provide a more satisfactory exchange.

There is a programming cost involved in performing the various label and file processing functions. However, that programming is not really germane to the solution of the application programmer's problem. Although many of the attributes of data to be processed vary by application and therefore defy generalization, there are others that permit, with a degree of standardization, the performance of certain functions on a generalized basis rather than as a procedural function of an application program. With a minimum of standardization, it is possible to supply a set of programmed routines that perform these functions almost without knowledge of the applications programmer. Such a set of programmed routines are referred to as generalized routines, in the sense that a single set of routines can work with a wide variety of applications programs. Such routines have become part of a larger set, sometimes referred to as an operating system.

The degree of standardization that permits the creation of many operating systems, each capable of providing the label and file handling functions applicable to sending and receiving volumes in the interchange environment, is included in this standard.

D1.2 Functions of Labels. The basic functions of labels include:

(1) Resource control: Identification of the owner and of the volume provide a tool for control of the resource within an installation and for its disposition after an interchange.

(2) Data protection: A protective measure beyond that afforded by conventional file protect rings is provided through the expiration date. This date may be checked against the current date prior to writing on a volume to determine the availability of the volume for output. This checking capability can prevent the inadvertent destruction of data caused by an error in set-up procedure.

NOTE: In addition, minimal protection against inadvertent disclosure of data is provided through the accessibility indicators. The effectiveness of these protective measures depends upon the specifications of the operating system to which the interchange volumes are sent, and to the fidelity of the recipients in accessing the volumes through an operating system providing protection.

(3) Data identification: Positive identification of the input data is provided. This ensures that the application processes the correct data. The requirements for positive identification vary by application, and several techniques are included.

(4) Data characterization: Formats for records and segments of records are provided. In addition, specifications on the use of each format, the permitted aggregations of records into blocks, blocks into file sections, file sections into files, and files into file sets are included. Certain of the fields in the labels contain file attribute parameters, specifying these and other similar characteristics. All these parameters are useful to an operating system for diagnosis of its capabilities with respect to the file or for adapting itself (dynamically) to process the particular configuration of the file being interchanged.

(5) File positioning: Operating systems are aware of the relative positioning of the information. In particular, they are capable of detecting the beginning and ending of file sections and of files. The discrete label identifiers EOVn and EOFn distinguish between the end of a file section and the end of a file. The definition of the occurrence of all label groups and the specification of the positions of tape marks permit an operating system to easily maintain label processing and file processing modes.

(6) Error control: A minimum of audit control on the reading of data is provided by means of the block count. Other fields enable control of error in sequencing through the sections of files, and the files of the file set.

NOTE: Additional capabilities for parochial use are provided; identification of the system that created the file enables the system that reads the file to determine how to respond to these capabilities.

D1.3 Functions of Label and File Processing. Processing consists of six basic functions:

(1) Positioning the volume

(2) Checking input volumes to verify that the correct volume is mounted and to verify file identification, obtain additional characteristics of the file, and ascertain that access to the file is warranted

(3) Checking output volumes for existence and certain content of labels in order to ensure that correct volumes are mounted and to prevent premature overwriting of important data

(4) Writing new labels on output

(5) Reading and writing installation and user labels, passing the information to/from an installation volume label processing routine or to application program file label processing routines

(6) Deblocking and blocking records, passing them to/from an application program processing routine

D1.4 Restriction of Data Codes to ASCII. In 6.1, it is specified that only the ASCII 7-bit code or its 8-bit version as defined in the American National Standard Recorded Magnetic Tape for Information Interchange standards (ANSI X3.14-1973, ANSI X3.22-1973, ANSI X3.39-1973, and ANSI X3.54-1976) is used in records within the files to be interchanged. This is a reasonable requirement to enhance interchange by specifying one common code all systems can recognize, rather than requiring systems to recognize an ever-increasing number of codes. This reduces the probability of error and misunderstanding, reduces programming labor, and provides a more satisfactory exchange.

D1.5 Impact on Operating Systems. This standard provides, through the specification of levels of labeling, for participation in the interchange environment of minimal operating system. Although some operating systems will tend to be more sophisticated even than one following the highest level specified in this standard, certain fundamental capabilities are provided for both the interchange and the local environment. On the other hand, this standard has been intentionally limited in scope so that future dynamic advancement of operating system theory and capabilities will not be hampered. In fact, such advancements have been anticipated. Portions of the labels have been reserved for operating system use, allowing for unilateral extensions, and portions of the labels have been reserved for future standardization, allowing for common advancement.

D2. Automatic Label and File Processing Functions

D2.1 General. This standard provides sufficient specifications for implementation of the generalized functions described in D2.2 through D2.6 in a dynamic, fully automatic, operating system.

D2.2 Opening and Closing Files. There are many routine details involved in beginning and concluding the processing of files in an application program. These details can be accomplished by a generalized function. This function is frequently referred to as opening and closing, or label processing. This standard contains sufficient specifications for file positioning, system label processing, and user label processing that this function can be performed automatically and consistently.

D2.3 Blocking and Deblocking of Records. In most cases the accumulation of several or many records into a single block will produce a great improvement in utilization of the medium as well as in control over the data. This can be accomplished by a generalized blocking and deblocking, segmenting and desegmenting, function. This standard includes sufficient specification that this function can be performed automatically for records that are invariant in length, for records that may vary in length, for records that are shorter than blocks, for records that are equal in length to blocks, and for records that are longer than blocks. There is no significant limit to the length of a record accommodated in this standard.

D2.4 Volume Switching. The determination of end of available space on a volume on output, end of data on a volume on input, and switching to start using the next volume are bothersome chores. Especially when the data are buffered (a backlog of records already processed but not yet written or a backlog of records already read but not yet processed) is the synchronization of volume switching and file processing a nuisance. Since there should be no application dependency on volumes (especially for applications that benefit from device independence), this can be performed entirely automatically, including use of alternate drives when available. This standard contains sufficient specification to ensure that all needed file sections of a file are processed, and in the correct sequence.

D2.5 Error Detection. The amount of programmed error checking required to ensure proper hardware reading and writing of magnetic tape was given substantial consideration. Consideration was given in preparation of Version 1 to checksumming each block, sequentially numbering each block, and periodically recording a block or record count. The extent to which each of these checks were then used was discussed. At that time it was considered that only sequentially numbering blocks had merit. The argument for this facility was that it was additional assurance to detect omitted or spurious blocks. Furthermore, detecting such malfunction immediately could provide additional information to users and maintainers of the system. This could permit immediate programming corrective action, avoiding unnecessary processing and reprocessing of data in error. This also could lead to more rapid achievement of acceptable results.

An important argument against the one-character block sequence indicator was the difficulty to optionally allow for it in software, relative to the anticipated benefits. It was also speculated that hardware accuracy would improve. The conclusion reached was that block counts recorded in the End-of-Volume and End-of-File Labels were necessary and sufficient. The Block Sequence Indicator was described in the Appendix for international accord. The specification of block counts is sufficient for detection of an error in processing a file section and for triggering manual or automatic corrective operations.

This facility was reviewed in preparation of Version 3. The anticipated improvement in hardware was noted. Since it is anticipated that phase encoded tape will essentially supplant NRZI tape during the life of this version, the committee concluded that no relative benefit would be obtained by retaining this facility, and accordingly it has been removed from the Appendix.

D2.6 Backward Reading. In general, the labeling system is expected to meet the needs of backward reading capability in all areas where it is likely to be implemented. Backward reading of blocked variable-length records or of segments of spanned records is not expected to be commonly used. Although such a facility could be provided using this labeling system, no special assistance has been provided in this labeling system to accommodate such implementation.

D3. Labeling System Structure

D3.1 General. The considerations given in D3.2 through D3.8 led to adoption of the specified labeling system structure.

D3.2 Volume-Header Label (VOL1). In an interchange situation wherein many different installations prepare and submit data to a single installation, the recipient should have a way to identify the sender. The external labeling of the volume is not always reliable or sufficient. The Volume-Header Label (VOL1) on each volume is expected to satisfy identification requirements. Any restrictions as to access to the information recorded on the volume is indicated in the label.

D3.3 User Volume Labels (UVL1-UVL9). The task group realized that government security requirements could not be satisfied by a one-character accessibility indicator. This same consideration could also apply to industry, not only for security classification but also for other information about the volume. In order to permit the solution of this problem by installations, User Volume Labels are defined and included in this standard.

D3.4 First File-Header Label (HDR1). Protection of a file is accommodated in the first label of the file. Thus unnecessary or unauthorized exposure is limited. Sufficient data to protect, secure, identify, and audit the file are on one label. Thus a system that does not dynamically adapt itself to the characteristics of the file can limit its attention to one label.

D3.5 Second File-Header Label (HDR2). It is increasingly apparent that certain file attribute descriptions are desirable to be able to self-initialize an operating system dynamically to process a file. One method of doing this is to append such descriptions to the file. The second File-Header Label contains information such as record format, record length, block length, etc. This information is also of value to self-initialize general utility programs such as tape dumps and file reformatters.

D3.6 Additional System Labels (HDR3-HDR9). In operating systems with specialized labeling functions, more information may be required than is provided in the first and second File-Header Labels (HDR1, HDR2). To provide for this information, space is reserved for system use in the second File-Header Label, and seven more labels are reserved for system use. The first four characters of these labels are defined in this standard. The remainder of these labels are defined by the designer or implementor of an operating system. When used, these labels are incompatible between different systems and can be ignored in interchange.

D3.7 User File Labels (UHLa, UTLa). Users have required descriptive information about the file to enable their more generalized application programs to be self-initializing and have required summary information about the file for audit and control. In many designs of labeling systems space in the system labels was reserved for application programs, and these labels were examined by application

programs. The attempt to contain more and more information led to expanding labels and conflict over the contents. Furthermore, system labels contain accessibility and other information that should be examined by the system and not by the program being controlled. To resolve this conflict over resource and this conflict of interest, separate label sets are provided for users. These sets appear at the beginning of each file section for initialization and at the end of each file section for summarization. In some cases the labels at the beginning and end of file sections are ignored, and information is significant only at the beginning or end of the entire file or at both the beginning and end of the entire file.

D3.8 Tape Marks. The tape mark is specified in this standard to have but one function: to delimit label groups and data. It is a trigger that causes an operating system to alternate between a label processing mode and a data mode and to determine the end of the file set. The actual configuration recorded upon and read from magnetic tape is not included in this standard, but it can be found in the appropriate recorded magnetic tape standards for the particular recording technique. (See ANSI X3.14-1973, ANSI X3.22-1973, ANSI X3.39-1973, or ANSI X3.54-1976, as appropriate.)

However, it should be noted that this standard does state exactly where tape marks occur, and that they do not occur at any other place. The use of different kinds of tape marks or the inclusion of tape marks other than where specified in this standard may cause misprocessing of a volume, and thus it is not valid in information interchange.

D4. Label Fields

D4.1 General. The considerations given in D4.2 through D4.6 led to the adoption of the specified label field formats.

D4.2 Identification. Fields are provided for identification of owners, systems, files, volumes, etc. By nature these are variable-length identifiers. Fixed-length fields are assigned for simplicity in label processing. It is felt that the length of each field assigned is sufficient to establish unique identification.

D4.3 Date Format. The date format specified for use in labels is the ordinal date (YYDDD). It is recognized that calendar date format (YYMMDD) has also been standardized. The former format is specified for the following reasons:

(1) The ordinal format permits better efficiency in operating system date manipulation functions, in terms of storage and performance. Typically a user determines the retention period for the file, that is expressed as the number of days following creation date that the file should be retained. The expiration date can be calculated by adding the retention period to the creation date. The YYDDD format forces fewer carries, and a less complex algorithm, than the YYMMDD.

(2) This is a specialized use of date, manipulated only by an operating system.

(3) This format is compatible with that appearing in volumes written according to Version 1 of this standard.

D4.4 Accessibility. "Accessibility" as used in this standard refers to such categories of information as company confidential, proprietary, payroll, social security, or tax information. This field may be used as a "lock" on the accessibility of the information. In some applications, a simple "lock and key" may not provide the desired degree of security. Any additional security may be implemented through the use of User Volume Labels (UVL1-UVL9).

Volume accessibility can be interpreted as a factored control for all the files in the volume set; satisfying volume accessibility implies satisfying file accessibility. It can be interpreted as control over the disposition of a volume. Because an earlier recording can be detected by sensitive instruments, even though overwritten, a volume accessibility may act as a floor for control over the volume or as a ceiling for the sensitivity of the material written on the volume. It can be interpreted as an additional control; satisfying volume accessibility is a prerequisite to satisfying file accessibility. This standard assumes the last interpretation.

Selective security of individual records may be useful in some environments; for example, a file containing records of several departments, in which a manager can access only records of his own department. This standard contains no facilities for selective security of individual records.

D4.5 Maximum Block Length. In Version 1 of this standard, block length was constrained to 2048 characters to ensure capability to process an interchange volume on small systems and to ensure accommodation by reasonably sized buffers. This provision is repeated in recorded magnetic tape standards. To avoid redundancy, the provision was deleted from this version; however, by reference to recorded magnetic tape standards, the provision is yet effective.

D4.6 Block Sequence Indicator. The requirement for block sequence indicator was reconsidered. While the use of block sequence indicator may be of value in alleviating a problem in 800-bpi NRZI tape, the problem does not exist when a phase-encoding technique is used. As it is anticipated that phase-encoded tape will essentially supplant NRZI technology during the period of development of this revision, the X3L5 Committee found that the requirement for block sequence indicator in this revision would be too late for current equipment. The committee concurred with the conclusions in D2.5 that no relative benefit would be obtained by retaining this facility.

D5. Other Major Changes

D5.1 General. The considerations given in D5.2 through D5.6 led to the adoption of other major changes.

D5.2 Elimination of Optional Fields. ANSI X3.27-1969 called certain fields "optional." These fields had to occupy space, and they had to contain valid characters. In that sense the fields were required and not optional.

These fields included the following:

HDR1:	Generation Number
	Generation Version Number
	System Code
HDR2:	Positions 5-80
EOF1, EOV1:	Positions 5-54
	Positions 61-80
EOF2:	Positions 5-80

In Version 3 the notation "optional" has been removed. Specific requirements for format, content, and processing, as appropriate, appear within the body of this standard.

D5.3 Format V Records. ANSI X3.27-1969 included an additional record format (V) in which the record length was recorded as a binary number. It was later realized that mixed character and noncharacter representations (or even mixed characters from different codes) cannot be handled in interchange between independent systems of potentially different architecture that may require the data to be translated to/from the interchange code from/to the internal processing code. Therefore, the constraint was imposed that all data on the volume are recorded in characters. This is consistent with the constraint that the interchange code is ASCII (see D1.4). The record format V, though used in certain systems as a parochial format, cannot be used in interchange, and therefore the specification was deleted from Version 3 of this standard.

D5.4 Spanned Record Format. Compatibility with the work of American National Standards Committee on Library Work, Documentation, and Related Publishing Practices, Z39, has been an important consideration in the decision to include spanned records in Version 3 of the standard. The selection of the format (Segment Indicator followed by length) for the Segment Control Word was made, in Version 3, to assure compatibility with ANSI Z39.2-1971 for text on magnetic tape. Neither blocked variable-length records (D) nor segments of spanned records (S) were designed for convenience in reading backward.

D5.5 Format U Records. ANSI X3.27-1969 included an additional record format (U) for use when the records used did not meet the definitions that were included elsewhere in the standard for record formats. This undefined record format required that Record Length not be specified. This record format did not prove useful to parochial systems and required agreement between interchange parties for use between systems. Version 3 has removed U format definition from the standard. Parochial versions of the standard may, of course, define other formats that are useful to those systems in their local environment.

D5.6 Levels of Labeling. Three principal reasons for the definition and inclusion of levels of labeling in Version 3 are:

(1) It was observed that implementations of Version 1 of this standard chose subsets for implementation and that those subsets served the purposes of interchange for those implementations.

(2) It was observed that a number of other complex software standards had successfully established subsets or levels of those standards.

(3) It was clear that the new technical requirement of spanned record support would not be of general use to all systems.

These facts, along with the committee's recognition that it should provide for orderly conversion to the new version, furnished justification for definition of levels.

The committee realized that the content of the levels defined, as well as the number of such levels, might not be optimal for every situation. In the interest of promoting maximum interchange it was decided, however, that no more than four levels, containing the specifications described in this version, should be introduced at this time. It is understood that additional levels could be defined in later versions without obsoleting the implementations that are produced before the new version is adopted.

American National Standards for Information Processing

X3.1-1976 Synchronous Signaling Rates for Data Transmission
X3.2-1970 (R1976) Print Specifications for Magnetic Ink Character Recognition
X3.3-1970 (R1976) Bank Check Specifications for Magnetic Ink Character Recognition
X3.4-1977 Code for Information Interchange
X3.5-1970 Flowchart Symbols and Their Usage in Information Processing
X3.6-1965 (R1973) Perforated Tape Code for Information Interchange
X3.9-1978 FORTRAN
X3.11-1969 Specification for General Purpose Paper Cards for Information Processing
X3.14-1973 Recorded Magnetic Tape for Information Interchange (200 CPI, NRZI)
X3.15-1976 Bit Sequencing of the American National Standard Code for Information Interchange in Serial-by-Bit Data Transmission
X3.16-1976 Character Structure and Character Parity Sense for Serial-by-Bit Data Communication in the American National Standard Code for Information Interchange
X3.17-1977 Character Set and Print Quality for Optical Character Recognition (OCR-A)
X3.18-1974 One-Inch Perforated Paper Tape for Information Interchange
X3.19-1974 Eleven-Sixteenths-Inch Perforated Paper Tape for Information Interchange
X3.20-1967 (R1974) Take-Up Reels for One-Inch Perforated Tape for Information Interchange
X3.21-1967 Rectangular Holes in Twelve-Row Punched Cards
X3.22-1973 Recorded Magnetic Tape for Information Interchange (800 CPI, NRZI)
X3.23-1974 Programming Language COBOL
X3.24-1968 Signal Quality at Interface between Data Processing Terminal Equipment and Synchronous Data Communication Equipment for Serial Data Transmission
X3.25-1976 Character Structure and Character Parity Sense for Parallel-by-Bit Data Communication in the American National Standard Code for Information Interchange
X3.26-1980 Hollerith Punched Card Code
X3.27-1978 Magnetic Tape Labels and File Structure for Information Interchange
X3.28-1976 Procedures for the Use of the Communication Control Characters of American National Standard Code for Information Interchange in Specified Data Communication Links
X3.29-1971 Specifications for Properties of Unpunched Oiled Paper Perforator Tape
X3.30-1971 Representation for Calendar Date and Ordinal Date for Information Interchange
X3.31-1973 Structure for the Identification of the Counties of the United States for Information Interchange
X3.32-1973 Graphic Representation of the Control Characters of American National Standard Code for Information Interchange
X3.34-1972 Interchange Rolls of Perforated Tape for Information Interchange
X3.36-1975 Synchronous High-Speed Data Signaling Rates between Data Terminal Equipment and Data Communication Equipment
X3.37-1980 Programming Language APT
X3.38-1972 (R1977) Identification of States of the United States (Including the District of Columbia) for Information Interchange
X3.39-1973 Recorded Magnetic Tape for Information Interchange (1600 CPI, PE)
X3.40-1976 Unrecorded Magnetic Tape for Information Interchange (9-Track 200 and 800 CPI, NRZI, and 1600 CPI, PE)
X3.41-1974 Code Extension Techniques for Use with the 7-Bit Coded Character Set of American National Standard Code for Information Interchange
X3.42-1975 Representation of Numeric Values in Character Strings for Information Interchange
X3.43-1977 Representations of Local Time of the Day for Information Interchange
X3.44-1974 Determination of the Performance of Data Communication Systems
X3.45-1974 Character Set for Handprinting
X3.46-1974 Unrecorded Magnetic Six-Disk Pack (General, Physical, and Magnetic Characteristics)
X3.47-1977 Structure for the Identification of Named Populated Places and Related Entities of the States of the United States for Information Interchange
X3.48-1977 Magnetic Tape Cassettes for Information Interchange (3.810-mm [0.150-in] Tape at 32 bpmm [800 bpi], PE)
X3.49-1975 Character Set for Optical Character Recognition (OCR-B)
X3.50-1976 Representations for U.S. Customary, SI, and Other Units to Be Used in Systems with Limited Character Sets
X3.51-1975 Representations of Universal Time, Local Time Differentials, and United States Time Zone References for Information Interchange
X3.52-1976 Unrecorded Single-Disk Cartridge (Front Loading, 2200 BPI), General, Physical, and Magnetic Requirements
X3.53-1976 Programming Language PL/I
X3.54-1976 Recorded Magnetic Tape for Information Interchange (6250 CPI, Group Coded Recording)
X3.55-1977 Unrecorded Magnetic Tape Cartridge for Information Interchange, 0.250 Inch (6.30 mm), 1600 bpi (63 bpmm), Phase Encoded
X3.56-1977 Recorded Magnetic Tape Cartridge for Information Interchange, 4 Track, 0.250 Inch (6.30 mm), 1600 bpi (63 bpmm), Phase Encoded
X3.57-1977 Structure for Formatting Message Headings for Information Interchange Using the American National Standard Code for Information Interchange for Data Communication Systems Control
X3.58-1977 Unrecorded Eleven-Disk Pack, General, Physical, and Magnetic Requirements
X3.60-1978 Programming Language Minimal BASIC
X3.61-1978 Representation of Geographic Point Locations for Information Interchange
X3.62-1979 Paper Used in Optical Character Recognition (OCR) Systems
X3.64-1979 Additional Controls for Use with American National Standard Code for Information Interchange
X3.66-1979 Advanced Data Communication Control Procedures (ADCCP)
X3.73-1980 Single-Sided Unformatted Flexible Disk Cartridge (for 6631-BPR Use)
X3.77-1980 Representation of Pocket Select Characters in Information Interchange
X3.82-1980 One-Sided Single-Density Unformatted 5.25-Inch Flexible Disk Cartridge (for 3979-BPR Use)
X3.83-1980 ANSI Sponsorship Procedures for ISO Registration According to ISO 2375
X3.86-1980 Optical Character Recognition (OCR) Inks

X3/TR1-77 Dictionary for Information Processing (Technical Report)

American National Standards Institute, Inc
1430 Broadway
New York, N.Y. 10018