

DRAFT INTERNATIONAL STANDARD

ISO/IEC DIS 20919

ISO/IEC JTC 1

Secretariat: ANSI

Voting begins on:
2020-04-14

Voting terminates on:
2020-07-07

Information technology — Linear tape file system (LTFS) Format specification

Technologies de l'information — Spécification du format de système de fichier à bande magnétique

ICS: 35.220.20

iTeh STANDARD PREVIEW (standards.iteh.ai)

[ISO/IEC DIS 20919](https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919)

<https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919>

THIS DOCUMENT IS A DRAFT CIRCULATED FOR COMMENT AND APPROVAL. IT IS THEREFORE SUBJECT TO CHANGE AND MAY NOT BE REFERRED TO AS AN INTERNATIONAL STANDARD UNTIL PUBLISHED AS SUCH.

IN ADDITION TO THEIR EVALUATION AS BEING ACCEPTABLE FOR INDUSTRIAL, TECHNOLOGICAL, COMMERCIAL AND USER PURPOSES, DRAFT INTERNATIONAL STANDARDS MAY ON OCCASION HAVE TO BE CONSIDERED IN THE LIGHT OF THEIR POTENTIAL TO BECOME STANDARDS TO WHICH REFERENCE MAY BE MADE IN NATIONAL REGULATIONS.

RECIPIENTS OF THIS DRAFT ARE INVITED TO SUBMIT, WITH THEIR COMMENTS, NOTIFICATION OF ANY RELEVANT PATENT RIGHTS OF WHICH THEY ARE AWARE AND TO PROVIDE SUPPORTING DOCUMENTATION.

This document is circulated as received from the committee secretariat.

FAST TRACK PROCEDURE



Reference number
ISO/IEC DIS 20919:2020(E)

© ISO/IEC 2020

iTeh STANDARD PREVIEW (standards.iteh.ai)

ISO/IEC DIS 20919

<https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919>



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2020

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Fax: +41 22 749 09 47
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Contents

1	Introduction	1
2	Scope	2
2.1	Versions	2
2.2	Conformance.....	3
3	Normative references	4
3.1	Approved references	4
3.2	References under development.....	4
3.3	Other references	4
4	Definitions and Acronyms	5
4.1	Definitions.....	5
4.2	Acronyms	7
5	Volume Layout	9
5.1	LTFS Partitions.....	9
5.2	LTFS Constructs	9
5.3	Partition Layout.....	10
5.4	Index Layout.....	11
6	Data Extents	14
6.1	Extent Lists.....	14
6.2	Extents Illustrated.....	14
6.3	Files Illustrated	16
7	Data Formats	20
7.1	Boolean format	20
7.2	Creator format	20
7.3	Extended attribute value format	20
7.4	Name format.....	21
7.5	Name pattern format	22
7.6	String format.....	22
7.7	Time stamp format	22
7.8	UUID format	23

iTeh STANDARD PREVIEW
(standards.iteh.ai)

<https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919>

8 Label Format	24
8.1 Label Construct	24
9 Index Format	27
9.1 Index Construct	27
9.2 Index.....	27
10 Medium Auxiliary Memory	42
10.1 Volume Change Reference	42
10.2 Volume Coherency Information.....	43
10.3 Use of Volume Coherency Information for LTFS	43
10.4 Use of Host-type Attributes for LTFS	45
10.5 Volume Advisory Locking	47
Annex A (normative) LTFS Label XML Schema	49
Annex B (normative) LTFS Index XML Schemas	51
B.1 LTFS Full Index XML Schema	51
B.2 LTFS Incremental Index XML Schema	53
Annex C (normative) Reserved Extended Attribute definitions	57
C.1 Software Metadata	57
C.2 Drive Metadata	57
C.3 Object Metadata	58
C.4 Volume Metadata	58
C.5 Media Metadata.....	60
Annex D (informative) Example of Valid Simple Complete LTFS Volume	63
Annex E (informative) Complete Example LTFS Full Index	64
Annex F (normative) Interoperability Recommendations	69
F.1 Spanning Files across Multiple Tape Volumes in LTFS	69
F.2 File Permissions in LTFS	74
F.3 Storing File Hash Values in LTFS	77
F.4 LTFS Media Pools.....	78
Annex G (informative) Character representations	80
Annex H (informative) Incremental Indexes	83

ITC STANDARD PREVIEW
 (standards.iteh.ai)

H.1 Background	83
H.2 Backwards Compatibility	83
H.3 Traversing the Index Back Pointer Chain	84
H.4 Incremental Index Format	84
H.5 Processing Incremental Indexes	86
H.6 Miscellaneous.....	87

iTeh STANDARD PREVIEW
(standards.iteh.ai)

ISO/IEC DIS 20919

<https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919>

List of Figures

Figure 1 — LTFS Partition.....	9
Figure 2 — Label Construct	9
Figure 3 — Index Construct	10
Figure 4 — Partition Layout.....	10
Figure 5 — Complete partition containing data.....	11
Figure 6 — Back Pointer example.....	12
Figure 7 — Back Pointer example for Incremental Indexes	13
Figure 8 — Extent starting and ending with full block	15
Figure 9 — Extent starting with full block and ending with fractional block	15
Figure 10 — Extent starting and ending in mid-block	15
Figure 11 — File contained in a single Data Extent.....	16
Figure 12 — File contained in two Data Extents.....	16
Figure 13 — Shared Blocks example	17
Figure 14 — Sparse files example	18
Figure 15 — Shared data example.....	18
Figure 16 — Label construct.....	24
Figure 17 — Index Construct	27
Figure D.1 — Content of a simple LTFS volume.....	63
Figure H.1 — Processing an Incremental Index (flowchart).....	88

List of Tables

Table 1 — Version elements	2
Table 2 — Version comparisons	3
Table 3 — Extent list entry starting and ending with full block	15
Table 4 — Extent list entry starting with full block and ending with fractional block	15
Table 5 — Extent list entry starting and ending in mid-block	16
Table 6 — Extent list entry for file contained in a single Data Extent	16
Table 7 — Extent list entry for a file contained in two Data Extents	16
Table 8 — Extent lists for Shared Blocks example	17
Table 9 — Extent list for sparse files example	18
Table 10 — Extent lists for shared data example	19
Table 11 — Creator format definitions	20
Table 12 — Reserved characters for name format	21
Table 13 — Characters which should be avoided for name format	21
Table 14 — Name percent-encoding	22
Table 15 — Time stamp format	23
Table 16 — VOL1 Label Construct	24
Table 17 — Volume Coherency Information	43
Table 18 — ACSI format for LTFS	44
Table 19 — Relevant Host-type Attributes for LTFS	45
Table 20 — Example of Host-type Attributes	47
Table 21 — Volume Locked MAM Attribute	47
Table 22 — Volume Locked MAM Attribute Values	47
Table F.1 — Hash Types	77
Table G.1 — Character representations : version 2.3 or later	80
Table G.2 — Character representations : version 2.2 or earlier	81

iTeh STANDARD PREVIEW
(standards.iteh.ai)

ISO/IEC DIS 20919

<https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919>

1 Introduction

This document defines a Linear Tape File System (LTFS) Format separate from any implementation on data storage media. Using this format, data is stored in LTFS Volumes. An LTFS Volume holds data files and corresponding metadata to completely describe the directory and file structures stored on the volume.

The LTFS Format has these features:

- An LTFS Volume can be mounted and volume content accessed with full use of the data without the need to access other information sources.
- Data can be passed between sites and applications using only the information written to an LTFS Volume.
- Files can be written to, and read from, an LTFS Volume using standard POSIX file operations.

The LTFS Format is particularly suited to these usages:

- Data export and import.
- Data interchange and exchange.
- Direct file and partial file recall from sequential access media.
- Archival storage of files using a simplified, self-contained or “self-describing” format on sequential access media.

ITeH STANDARD PREVIEW
(standards.iteh.ai)

[ISO/IEC DIS 20919](https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919)

<https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919>

2 Scope

This document defines the LTFS Format requirements for interchanged media that claims LTFS compliance. Those requirements are specified as the size and sequence of data blocks and file marks on the media, the content and form of special data constructs (the LTFS Label and LTFS Index), and the content of the partition labels and use of MAM parameters.

The data content (not the physical media) of the LTFS format shall be interchangeable among all data storage systems claiming conformance to this format. Physical media interchange is dependent on compatibility of physical media and the media access devices in use.

NOTE: This document does not contain instructions or tape command sequences to build the LTFS structure.

2.1 Versions

This document describes version 2.5.0 of the Linear Tape File System (LTFS) Format Specification.

The version number for the LTFS Format Specification consists of three integer elements separated by period characters of the form $M.N.R$, where M , N and R are positive integers or zero. Differences in the version number between different revisions of this specification indicate the nature of the changes made between the two revisions. Each of the integers in the format specification are incremented according to Table 1.

Table 1 — Version elements

Element	Description
M	Incremented when a major update has been made to the LTFS Format Specification. Major updates are defined as any change to the on-media format or specification semantics that are expected to break compatibility with older versions of the specification.
N	Incremented when a minor update has been made to the LTFS Format Specification. Minor updates are defined as any change to the on-media format or specification semantics that is not expected to break compatibility with older versions of the specification that have the same value for M in the version number.
R	Incremented when textual revisions are made to the LTFS Format Specification. Textual revisions are defined as revisions that improve the clarity of the specification document <i>without</i> changing the intent of the document. By definition, minor changes do not alter the on-media format or specification semantics.

NOTE 1: When any element of the specification version number is incremented, all sub-ordinate elements to the right are reset to zero. For example, if the version is 1.0.12 and N is incremented to 1, then R is set to zero resulting in version 1.1.0.

NOTE 2: The first public version of this document used version number 1.0. This value should be interpreted as equivalent to 1.0.0 in the version numbering defined in this document.

The result of comparison between two LTFS version numbers $M_A.N_A.R_A$ and $M_B.N_B.R_B$ is defined in

Table 2.

Table 2 — Version comparisons

Conditional	Description
$M_A < M_B$	$M_A.N_A.R_A$ is an earlier version than $M_B.N_B.R_B$.
$M_A = M_B$ and $M_A < N_B$	$M_A.N_A.R_A$ is an earlier version than $M_B.N_B.R_B$.
$M_A = M_B$ and $N_A = N_B$ and $R_A < R_B$	$M_A.N_A.R_A$ is an earlier version than $M_B.N_B.R_B$. However, as defined above, changes that result only in a different R value are descriptive changes in the specification rather than on media changes.

2.2 Conformance

Recorded media claiming conformance to this format shall be in a consistent state when interchanged or stored. See Section 4.1.4.

Any implementation conforming to this specification should be able to correctly read Label and Index structures from all prior versions of this specification and write Label and Index structures conforming to the descriptions in this document. The current Label and Index structures are defined in Section 8 Label Format and in Section 9 Index Format.

NOTE: Where practical, any implementation supporting a given version value for M should endeavor to support LTFS volumes with version numbers containing higher values for N and R than those defined at the time of implementation.

iTeh STANDARD PREVIEW
(standards.iteh.ai)

ISO/IEC DIS 20919

<https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919>

3 Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

3.1 Approved references

ISO/IEC 14776-453, SCSI Primary Commands - 3 (SPC-3) [ANSI INCITS.408-2005]

SSC-4 SCSI Stream Commands – 4 [SSC-4] [ANSI INCITS 516-2013]

IETF RFC 4648, The Base16, Base32, and Base64 Data Encodings, <http://www.ietf.org/rfc/rfc4648.txt>

ISO 8601:2004 Data elements and interchange formats – Information interchange – Representation of dates and times – (UTC)

ISO/IEC 10646:2012: Information technology - Universal Coded Character Set (UCS) (UTF-8)

IETF RFC 4122, Universally Unique Identifier (UUID) URN Namespace <http://www.ietf.org/rfc/rfc4122.txt>

IETF RFC 3986, Uniform Resource Identifier (URI): Generic Syntax, <http://www.ietf.org/rfc/rfc3986.txt>

ANSI X3.27-1978 American National Standard Magnetic Tape Labels and File Structure for Information

3.2 References under development

SCSI Primary Commands - 4 (SPC-4) [ANSI INCITS 513:2014]

3.3 Other references

W3C - Extensible Markup Language (XML) <http://www.w3.org/XML>

NFC – Unicode Normalization Forms - Unicode Standard Annex - UAX#15

<http://www.unicode.org/reports/tr15>

Unicode Text Segmentation - Unicode Standard Annex - UAX#29 <http://www.unicode.org/reports/tr29>

OSF CDE 1.1, Remote Procedure Call – Universal Unique Identifier (UUID)

<http://pubs.opengroup.org/onlinepubs/9629399/toc.pdf>

4 Definitions and Acronyms

For the purposes of this document the following definitions and acronyms shall apply.

4.1 Definitions

4.1.1

Block Position

The position or location of a recorded block as specified by its LTFS Partition ID and logical block number within that partition.

The block position of an Index is the position of the first logical block for the Index.

4.1.2

Complete Partition

An LTFS partition that consists of an LTFS Label Construct and a Content Area, where the last construct in the Content Area is an Index Construct.

4.1.3

Content Area

A contiguous area in a partition, used to record Index Constructs and Data Extents.

4.1.4

Consistent State

A volume is consistent when both partitions are complete and the last Index Construct in the Index Partition has a back pointer to the last Full Index Construct in the Data Partition.

4.1.5

Data Extent

A contiguous sequence of recorded blocks.

[ISO/IEC DIS 20919](https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919)

<https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919>

4.1.6

Data Partition

An LTFS partition primarily used for data files.

4.1.7

File

A group of logically related extents together with associated file metadata.

4.1.8

Filesystem sync

An operation during which all cached file data and metadata is flushed to the media.

4.1.9

Full Index

A data structure that describes all valid data files in an LTFS volume. The Full Index is an XML document conforming to the XML schema shown in [Annex B \(normative\) LTFS Index XML Schema](#).

4.1.10

Generation number

A positive decimal integer which shall indicate the specific generation of an Index within an LTFS volume.

4.1.11

Incremental Index

A data structure that describes changes made to the LTFS volume since the last index was written. The Incremental Index is an XML document conforming to the XML schema shown in [Annex B \(normative\) LTFS Index XML Schema](#).

4.1.12

Index

Either a Full Index or an Incremental Index.

4.1.13

Index Construct

A data construct comprised of an Index and file marks.

4.1.14

Index Partition

An LTFS partition primarily used to store Index Constructs and optionally data files.

4.1.15

Label Construct

A data construct comprised of an ANSI VOL1 tape label, LTFS Label, and tape file marks.

4.1.16

Linear Tape File System (LTFS) (standards.iteh.ai)

This document describes the Linear Tape File System Format.

[ISO/IEC DIS 20919](#)

4.1.17

<https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919>

LTFS Construct

Any of three defined constructs that are used in an LTFS partition. The LTFS constructs are: Label Construct, Index Construct, and Data Extent.

4.1.18

LTFS Label

A data structure that contains information about the LTFS partition on which the structure is stored. The LTFS Label is an XML document conforming to the XML schema shown in [Annex A \(normative\) LTFS Label XML Schema](#).

4.1.19

LTFS Partition

A tape partition that is part of an LTFS volume. The partition contains an LTFS Label Construct and a Content Area.

4.1.20

LTFS Volume

A pair of LTFS partitions, one Data Partition and one Index Partition, that contain a logical set of files and directories. The pair of partitions in an LTFS Volume shall have the same UUID. All LTFS partitions in an LTFS volume are *related partitions*.

4.1.21**Medium Auxiliary Memory**

An area of non-volatile storage that is part of an individual storage medium. The method of access to this non-volatile storage is standardized as described in the T10/SPC-4 standard.

4.1.22**Partition Identifier (Partition ID)**

The logical partition letter to which LTFS data files and Indexes are assigned.

The linkage between LTFS partition letter and physical SCSI partition number is determined by the SCSI partition in which the LTFS Label is recorded. The LTFS partition letter is recorded in the LTFS Label construct, and the SCSI partition number is known by the SCSI positional context where they were read/written.

4.1.23**Sparse file**

A file that has some number of empty (unwritten) data regions. These regions are not stored on the storage media and are implicitly filled with bytes containing the value zero (0x00).

4.1.24**UUID**

Universally unique identifier; an identifier use to bind a set of LTFS partitions into an LTFS volume.

4.1.25**Volume Change Reference (VCR) (standards.iteh.ai)**

A value that represents the state of all partitions on a medium.

4.1.26**Volume Advisory Locking**

An indication that the LTFS volume has been locked against future modifications. This is a form of write protection under the control of host software rather than physical hardware.

iTeh STANDARD PREVIEW

[ISO/IEC DIS 20919](https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919)

<https://standards.iteh.ai/catalog/standards/sist/23e08a25-8e60-4ead-88a8-81a9a5aca3ec/iso-iec-dis-20919>

4.2 Acronyms

ASCII	American Standard Code for Information Interchange
CM	Cartridge Memory
DCE	Distributed Computing Environment
ISO	International Organization for Standardization
LTFS	Linear Tape File System
MAM	Media Auxiliary Memory
NFC	Normalization Form Canonical Composition
OSF	Open Software Foundation
POSIX	Portable Operating System Interface for Unix
T10/SSC-4	ISO/IEC 14776-334, SCSI Stream Commands - 4 (SSC-4) [T10/2123-D]
UTC	Coordinated Universal Time
UTF-8	8-bit UCS/Unicode Transformation Format
UUID	Universally Unique Identifier
W3C	World Wide Web Consortium