INTERNATIONAL STANDARD

ISO/IEC 9594-1

Eighth edition
2017-05

Information technology — Open Systems Interconnection — The Directory —

Part 1: Overview of concepts, models and services

Technologies de l'information — Interconnexion de systèmes ouverts (OSI) — L'annuaire —

Partie 1: Aperçu général des concepts, modèles et services

Reference number
ISO/IEC 9594-1:2017(E)

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

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see the following URL: www.iso.org/iso/foreword.html.

This eighth edition cancels and replaces the seventh edition (ISO/IEC 9594-1:2014), which has been technically revised.

This document was prepared by ISO/IEC JTC 1, *Information technology*, SC 6, *Telecommunications and information exchange between systems*, in collaboration with ITU-T. The identical text is published as ITU-T X.500 (10/2016).

A list of all parts in the ISO/IEC 9594 series, published under the general title *Information technology — Open Systems Interconnection — The Directory*, can be found on the ISO website.

CONTENTS

Introduction

This Recommendation | International Standard together with other Recommendations | International Standards, has been produced to facilitate the interconnection of information processing systems to provide directory services. A set of such systems, together with the directory information that they hold, can be viewed as an integrated whole, called the *Directory*. The information held by the Directory, collectively known as the Directory Information Base (DIB), is typically used to facilitate communication between, with or about objects such as application entities, people, terminals and distribution lists.

The Directory plays a significant role in Open Systems Interconnection, whose aim is to allow, with a minimum of technical agreement outside of the interconnection standards themselves, the interconnection of information processing systems:

– from different manufacturers;

– under different managements;

– of different levels of complexity; and

– of different ages.

This Recommendation | International Standard introduces and models the concepts of the Directory and of the DIB and overviews the services and capabilities which they provide. Other Recommendations | International Standards make use of these models in defining the abstract service provided by the Directory, and in specifying the protocols through which this service can be obtained or propagated.

This Recommendation | International Standard provides the foundation frameworks upon which industry profiles can be defined by other standards groups and industry forums. Many of the features defined as optional in these frameworks, may be mandated for use in certain environments through profiles. This eighth edition technically revises and enhances, the seventh edition of this Recommendation | International Standard.

This eighth edition specifies versions 1 and 2 of the Directory protocols.

The first and second editions specified only version 1. Most of the services and protocols specified in this edition are designed to function under version 1. However some enhanced services and protocols, e.g., signed errors, will not function unless all Directory entities involved in the operation have negotiated version 2. Whichever version has been negotiated, differences between the services and between the protocols defined in the seven editions, except for those specifically assigned to version 2, are accommodated using the rules of extensibility defined in Rec. ITU-T X.519 | ISO/IEC 9594-5.

Annex A, which is an integral part of this Recommendation | International Standard, describes the types of use to which the Directory can be applied.

Annex B, which is not an integral part of this Recommendation | International Standard, lists the amendments and defect reports that have been incorporated to form this edition of this Recommendation | International Standard.

INTERNATIONAL STANDARD
ITU-T RECOMMENDATION

Information technology – Open Systems Interconnection –The Directory: Overview of concepts, models and services

1 Scope

The Directory provides the directory capabilities required by OSI applications, OSI management processes, other OSI layer entities, and telecommunications services. Among the capabilities which it provides are those of "user-friendly naming", whereby objects can be referred to by names which are suitable for citing by human users (though not all objects need have user-friendly names); and "name-to-address mapping" which allows the binding between objects and their locations to be dynamic. The latter capability allows OSI networks, for example, to be "self-configuring" in the sense that addition, removal and the changes of object location do not affect OSI network operation.

The Directory is not intended to be a general-purpose database system, although it may be built on such systems. It is assumed, for instance, that, as is typical with communications directories, there is a considerably higher frequency of "queries" than of updates. The rate of updates is expected to be governed by the dynamics of people and organizations, rather than, for example, the dynamics of networks. There is also no need for instantaneous global commitment of updates; transient conditions, where both old and new versions of the same information are available, are quite acceptable.

It is a characteristic of the Directory that, except as a consequence of differing access rights or un-propagated updates, the results of directory queries will not be dependent on the identity or location of the inquirer. This characteristic renders the Directory unsuitable for some telecommunications applications, for example some types of routing. For cases where the results are dependent on the identity of the inquirer, access to directory information and updates of the Directory may be denied.

2 Normative references

The following Recommendations and International Standards contain provisions which, through reference in this text, constitute provisions of this Recommendation | International Standard. At the time of publication, the editions indicated were valid. All Recommendations and Standards are subject to revision, and parties to agreements based on this Recommendation | International Standard are encouraged to investigate the possibility of applying the most recent edition of the Recommendations and Standards listed below. Members of IEC and ISO maintain registers of currently valid International Standards. The Telecommunication Standardization Bureau of the ITU maintains a list of currently valid ITU-T Recommendations.

2.1 Identical Recommendations | International Standards

- Recommendation ITU-T X.200 (1994) | ISO/IEC 7498-1:1994, *Information technology – Open Systems Interconnection – Basic Reference Model: The basic model*.
- Recommendation ITU-T X.501 (2016) | ISO/IEC 9594-2:2017, *Information technology – Open Systems Interconnection – The Directory: Models*.
- Recommendation ITU-T X.509 (2016) | ISO/IEC 9594-8:2017, *Information technology – Open Systems Interconnection – The Directory: Public-key and attribute certificate frameworks*.
- Recommendation ITU-T X.511 (2016) | ISO/IEC 9594-3:2017, *Information technology – Open Systems Interconnection – The Directory: Abstract service definition*.
- Recommendation ITU-T X.518 (2016) | ISO/IEC 9594-4:2017, *Information technology – Open Systems Interconnection – The Directory: Procedures for distributed operation*.
- Recommendation ITU-T X.519 (2016) | ISO/IEC 9594-5:2017, *Information technology – Open Systems Interconnection – The Directory: Protocol specifications*.
- Recommendation ITU-T X.520 (2016) | ISO/IEC 9594-6:2017, *Information technology – Open Systems Interconnection – The Directory: Selected attribute types*.
- Recommendation ITU-T X.521 (2016) | ISO/IEC 9594-7:2017, *Information technology – Open Systems Interconnection – The Directory: Selected object classes*.
- Recommendation ITU-T X.525 (2016) | ISO/IEC 9594-9:2017, *Information technology – Open Systems Interconnection – The Directory: Replication*.

3 Definitions

For the purposes of this Recommendation | International Standard, the following definitions apply.

3.1 Communication model definitions

The following terms are defined in Rec. ITU-T X.519 | ISO/IEC 9594-5:

a) application-entity;

b) application layer;

c) application process.

3.2 Directory model definitions

The following terms are defined in Rec. ITU-T X.501 | ISO/IEC 9594-2:

a) access control;

b) Administration Directory Management Domain;

c) alias;

d) ancestor;

e) attribute;

f) attribute type;

g) attribute value;

h) authentication;

i) compound entry;

j) context;

k) Directory Information Tree (DIT);

l) Directory Management Domain (DMD);

m) Directory System Agent (DSA);

n) Directory User Agent (DUA);

o) distinguished name;

p) entry;

q) family (of entries);

r) hierarchical group;

s) LDAP client;

t) LDAP requester;

u) LDAP responder;

v) LDAP server;

w) name;

x) object (of interest);

y) Private Directory Management Domain;

z) related entries;

aa) relative distinguished name;

bb) root;

cc) schema;

dd) security policy;

ee) subordinate object;

ff) superior entry;

gg) superior object;

hh) tree.

3.3 Distributed Operation definitions

The following terms are defined in Rec. ITU-T X.518 | ISO/IEC 9594-4:

a) uni-chaining;

b) multi-chaining;

c) referral.

3.4 Replication definitions

The following terms are defined in Rec. ITU-T X.525 | ISO/IEC 9594-9:

a) caching;

b) cache copy;

c) entry copy;

d) master DSA;

e) replication;

f) shadow consumer;

g) shadow supplier;

h) shadowed information;

i) shadowing agreement.

3.5 Basic directory definitions

The following terms are defined in this Recommendation | International Standard:

3.5.1 **the Directory**: A collection of open systems cooperating to provide directory services.

3.5.2 **directory information base (DIB)**: The set of information managed by the Directory.

3.5.3 **(directory) user**: The end user of the Directory, i.e., the entity or person which accesses the Directory.

4 Abbreviations

For the purposes of this Recommendation | International Standard, the following abbreviations apply:

ACI	Access Control Information
DAP	Directory Access Protocol
DIB	Directory Information Base
DISP	Directory Information Shadowing Protocol
DIT	Directory Information Tree
DMD	Directory Management Domain
DOP	Directory Operational Binding Management Protocol
DSA	Directory System Agent
DSP	Directory System Protocol
DUA	Directory User Agent
LDAP	Lightweight Directory Access Protocol
OSI	Open Systems Interconnection
RDN	Relative Distinguished Name

5 Conventions

With minor exceptions this Directory Specification has been prepared according to the November 2001 edition of the *Rules for presentation of ITU-T | ISO/IEC common text*.

The term "Directory Specification" (as in "this Directory Specification") shall be taken to mean Rec. ITU-T X.500 | ISO/IEC 9594-1. The term "Directory Specifications" shall be taken to mean the ITU-T X.500-series Recommendations, except for Rec. ITU-T X.509, and all parts of ISO/IEC 9594, except for ISO/IEC 9594-8.

This Directory Specification uses the term *first edition systems* to refer to systems conforming to the first edition of these Directory Specifications, i.e., the 1988 edition of the series of CCITT X.500 Recommendations and the ISO/IEC 9594:1990 edition.

This Directory Specification uses the term *second edition systems* to refer to systems conforming to the second edition of these Directory Specifications, i.e., the 1993 edition of the series of ITU-T X.500 Recommendations and the ISO/IEC 9594:1995 edition.

This Directory Specification uses the term *third edition systems* to refer to systems conforming to the third edition of these Directory Specifications, i.e., the 1997 edition of the series of ITU-T X.500 Recommendations and the ISO/IEC 9594:1998 edition.

This Directory Specification uses the term *fourth edition systems* to refer to systems conforming to the fourth edition of the Directory Specifications, i.e., the 2001 editions of Recs ITU-T X.500, ITU-T X.501, ITU-T X.511, ITU-T X.518, ITU-T X.519, ITU-T X.520, ITU-T X.521, ITU-T X.525, and ITU-T X.530, the 2000 edition of Rec. ITU-T X.509, and parts 1-10 of the ISO/IEC 9594:2001 edition.

This Directory Specification uses the term *fifth edition systems* to refer to systems conforming to the fifth edition of these Directory Specifications, i.e., the 2005 edition of the series of ITU-T X.500 Recommendations and the ISO/IEC 9594:2005 edition.

This Directory Specification uses the term *sixth edition systems* to refer to systems conforming to the sixth edition of these Directory Specifications, i.e., the 2008 edition of the series of ITU-T X.500 Recommendations and the ISO/IEC 9594:2008 edition.

This Directory Specification uses the term *seventh edition systems* to refer to systems conforming to the seventh edition of these Directory Specifications, i.e., the 2012 edition of the series of ITU-T X.500 Recommendations and the ISO/IEC 9594:2014 edition.

This Directory Specification uses the term *eighth edition systems* to refer to systems conforming to the eighth edition of these Directory Specifications, i.e., the 2016 edition of the series of ITU-T X.500 Recommendations and the ISO/IEC 9594:2016 edition.

This Directory Specification presents ASN.1 notation in the **`bold Courier New`** typeface. When ASN.1 types and values are referenced in normal text, they are differentiated from normal text by presenting them in the **`bold Courier New`** typeface. The names of procedures, typically referenced when specifying the semantics of processing, are differentiated from normal text by displaying them in **bold Times New Roman**. Access control permissions are presented in *italicized Times New Roman*.

6 Overview of the Directory

The *Directory* is a collection of open systems which cooperate to hold a logical database of information about a set of objects in the real world. The *users* of the Directory, including people and computer programs, can read or modify the information, or parts of it, subject to having permission to do so. Each user is represented in accessing the Directory by a Directory User Agent (DUA) or an LDAP client, each of which is considered to be an application-process. These concepts are illustrated in Figure 1.

NOTE – The Directory Specifications refer to the Directory in the singular, and reflects the intention to create, through a single, unified, name space, one logical directory composed of many systems and serving many applications. Whether or not these systems choose to interwork will depend on the needs of the applications they support. Applications dealing with non-intersecting worlds of objects may have no such need. The single name space facilitates later interworking should the needs change. For a variety of reasons, such as security, connectivity, or business decisions, it is likely that some portions of the Directory may be unreachable from other portions of the Directory using third edition operations. This results in differing views of the Directory. Such differing views may contain related entries about a given real world object. Such related entries may or may not have the same distinguished name. Using fourth or subsequent edition systems, it is possible to perform operations across multiple, differing views to provide an integrated response to the user. Specifically:

- DMD administrators (see 9.2) may have a need to publish their own view (or views) of some specific real-world object; a real-world object may thus be modelled by multiple independent entries in the directory. This may happen whether or not they need to interwork. Interworking using DSP may also be unsupported.
- Notwithstanding the last sentence of the Note, it is also possible that particular DMDs may choose to publish information about real-world objects within their own distinct directory name-spaces (i.e., in one of multiple DITs); in this case, it would be possible to have a specific real-world object modelled by entries in the same or different DIT namespaces, with the same or different distinguished names in each. Note that certain Directory facilities (e.g., the acquisition of certificates, and related

functions based on digital signatures) cannot be implemented when distinct objects are permitted to share distinguished names.

- The objective of related entries is to provide a means whereby users can access such entries, bringing the resulting information together, if possible. This would apply to the situation described by both of the preceding bullet points.

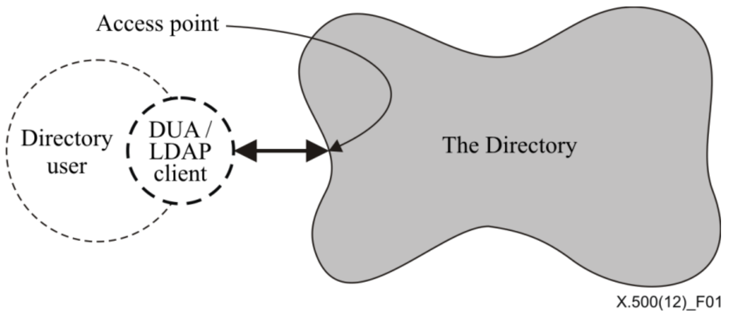

Figure 1 – Access to the Directory

The information held in the Directory is collectively known as the *Directory Information Base* (DIB). Clause 7 gives an overview of its structure.

The Directory provides a well-defined set of access capabilities, known as the abstract service of the Directory, to its users. This service, which is briefly described in clause 8, provides a simple modification and retrieval capability. This can be built on with local DUA functions to provide the capabilities required by the end-users.

The Directory is distributed, both along functional and organizational lines. Clause 9 gives an overview of the corresponding models of the Directory. These have been developed in order to provide a framework for the cooperation of the various components to provide an integrated whole.

The Directory exists in an environment where various administrative authorities control access to their portion of the information. Clause 10 gives an overview of access control.

When the Directory is distributed, it may be desirable to replicate information to improve performance and availability. Clause 11 gives an overview of the Directory replication mechanism.

The provision and consumption of the Directory services requires that the users (actually the DUAs and/or LDAP clients) and the various functional components of the Directory should cooperate with one another. In many cases this will require cooperation between application processes in different open systems, which in turn requires standardized application protocols, briefly described in clause 11, to govern this cooperation.

The Directory has been designed so as to support multiple applications, drawn from a wide range of possibilities. The nature of the applications supported governs which objects are listed in the Directory, which users access the information, and which kinds of access they carry out. Applications may be very specific, such as the provision of distribution lists for electronic mail, or generic, such as the 'inter-personal communications directory' application. The Directory provides the opportunity to exploit commonness among the applications:

- A single object may be relevant to more than one application: Perhaps even the same piece of information about the same object may be so relevant.
- To support this, a number of object classes and attribute types are defined, which are useful across a range of applications. These definitions are contained in Rec. ITU-T X.520 | ISO/IEC 9594-6 and Rec. ITU-T X.521 | ISO/IEC 9594-7.
- Certain patterns of use of the Directory are common across a range of applications: Annex A gives an overview of this area.

7 The Directory Information Base (DIB)

NOTE 1 – The DIB, and its structure, are defined in Rec. ITU-T X.501 | ISO/IEC 9594-2.

The DIB is made up of information about objects. It is composed of (*Directory*) *entries*, each of which consists of a collection of information on one object. An entry may be an aggregate of member entries each holding information about a particular aspect of an object. Such an aggregate entry is called a compound entry. Each entry is made up of *attributes*, each with a type and one or more values. The types of attribute which are present in a particular entry are dependent on the *class* of object which the entry describes. Each *value* of an attribute may be tagged with one or more *contexts* that specify information about a value that can be used to determine the applicability of the value.

The entries of the DIB are arranged in the form of a tree, the Directory Information Tree (DIT) where the vertices represent the entries. Entries higher in the tree (nearer the root) will often represent objects such as countries or organizations, while entries lower in the tree will represent people or application processes.

NOTE 2 – The services defined in the Directory Specifications operate only on a tree-structured DIT. The Directory Specifications do not preclude the existence in the future of other structures (as the need arises).

Every entry has a distinguished name, which uniquely and unambiguously identifies the entry. These properties of the distinguished name are derived from the tree structure of the information. The distinguished name of an entry is made up of the distinguished name of its superior entry, together with specially nominated attribute values (the distinguished values) from the entry.

Some of the entries at the leaves of the tree are alias entries, while other entries are object entries and compound entries. Alias entries point to object entries, and provide the basis for alternative names for the corresponding objects.

A compound entry is an entry representing a single object and it is an aggregate of member entries each representing a part of the information about the object.

The Directory enforces a set of rules to ensure that the DIB remains well-formed in the face of modifications over time. These rules, known as the *Directory schema*, prevent entries having the wrong types of attributes for its object class, attribute values being of the wrong form for the attribute type, and even entries having subordinate entries of the wrong class.

Figure 2 illustrates the above concepts of the DIT and its components.

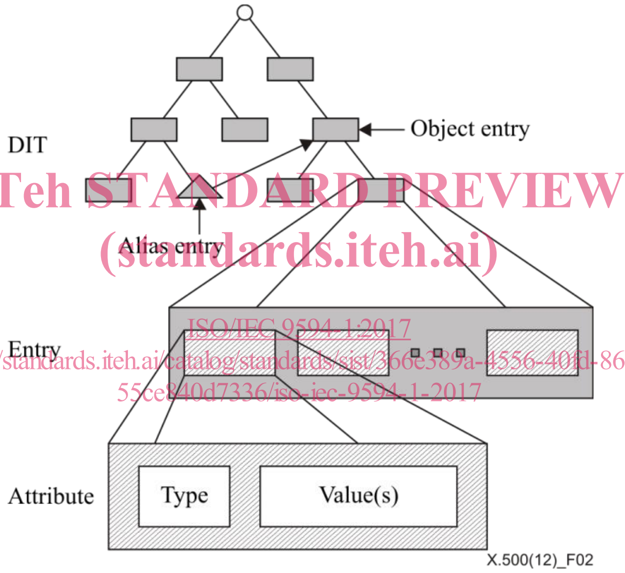

Figure 2 – Structure of the DIT and of entries

Figure 3 gives a hypothetical example of a DIT. The tree provides examples of some of the types of attributes used to identify different objects. For example the name:

{C=GB, L=Winslow, O=Graphic Services, CN=Laser Printer}

identifies the application entity, "Laser Printer", which has in its distinguished name the geographical attribute of Locality.

The residential person, John Jones, whose name is {C=GB, L=Winslow, CN=John Jones}, has the same geographical attribute in his distinguished name.

The growth and form of the DIT, the definition of the Directory schema, and the selection of distinguished names for entries as they are added, is the responsibility of various authorities, whose hierarchical relationship is reflected in the shape of the tree. The authorities shall ensure, for example, that all of the entries in their jurisdiction have unambiguous distinguished names, by carefully managing the attribute types and values which appear in those names. Responsibility is passed down the tree from superior to subordinate authorities, with control being exercised by means of the schema.