
Information technology —

Open Systems Interconnection —

**The Directory: Overview of Concepts, Models, and
Services**

Recommendation X.500

ISO/IEC 9594-1

Contents

Foreword	iii
Introduction	iv
1 Scope	1
2 Normative references	1
3 Definitions	2
4 Abbreviations	3
5 Conventions	3
6 Overview of the Directory	4
7 The Directory Information Base (DIB)	5
8 The Directory Service	6
9 The Distributed Directory	8
10 Access control in the Directory	10
11 Replication in the Directory	11
12 Directory protocols	13
Annex A — Applying the Directory	15
Annex B — Amendments & corrigenda	19

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. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75% of the national bodies casting a vote.

International Standard ISO/IEC 9594-1 was prepared by Joint Technical committee ISO/IEC JTC 1, Information technology, Subcommittee SC21, Information retrieval, transfer and management for open systems interconnection (OSI), in collaboration with ITU-T. The identical text is published as ITU-T Recommendation X.500.

This second edition technically revises and enhances, but does not replace, the first edition of ISO/IEC 9594. Implementations may still claim conformance to the first edition.

This second edition of ISO/IEC 9594 specifies version 1 of the Directory service and protocols. The first edition also specifies version 1. Differences between the service and between the protocols defined in the two editions are accommodated using the rules of extensibility defined in part 5 of this edition.

ISO/IEC 9594 consists of the following parts, under the general title Information technology — Open systems Interconnection — The Directory:

- Part 1: Overview of concepts, models, and services
- Part 2: Models
- Part 3: Abstract service definition
- Part 4: Procedures for distributed operation
- Part 5: Protocol specifications
- Part 6: Selected attribute types
- Part 7: Selected object classes
- Part 8: Authentication framework
- Part 9: Replication

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 which 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 second edition technically revises and enhances, but does not replace, the first edition of this Recommendation | International Standard. Implementations may still claim conformance to the first edition.

This second edition specifies version 1 of the Directory service and protocols. The first edition also specifies version 1. Differences between the services and between the protocols defined in the two editions are accommodated using the rules of extensibility defined in this edition of 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 unpropagated 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.

2 Normative references

The following Recommendations and International Standards contain provisions which, through reference in this text, constitute provisions of this Recommendation | International Standard part. 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 editions 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

- ITU-T Recommendation X.501 (1993) | ISO/IEC 9594-2:1993, *Information technology — Open Systems Interconnection — The Directory: Models.*
- ITU-T Recommendation X.511 (1993) | ISO/IEC 9594-3:1993, *Information technology — Open Systems Interconnection — The Directory: Abstract Service Definition.*
- ITU-T Recommendation X.518 (1993) | ISO/IEC 9594-4:1993, *Information technology — Open Systems Interconnection — The Directory: Procedures for Distributed Operation.*
- ITU-T Recommendation X.519 (1993) | ISO/IEC 9594-5:1993, *Information technology — Open Systems Interconnection — The Directory: Protocol Specifications.*
- ITU-T Recommendation X.520 (1993) | ISO/IEC 9594-6:1993, *Information technology — Open Systems Interconnection — The Directory: Selected Attribute Types.*
- ITU-T Recommendation X.521 (1993) | ISO/IEC 9594-7:1993, *Information technology — Open Systems Interconnection — The Directory: Selected Object Classes.*
- ITU-T Recommendation X.509 (1993) | ISO/IEC 9594-8:1993, *Information technology — Open Systems Interconnection — The Directory: Authentication Framework.*

- ITU-T Recommendation X.525 (1993) | ISO/IEC 9594-9:1993, *Information technology — Open Systems Interconnection — The Directory: Replication*
- ITU-T Recommendation X.880 (1994) | ISO/IEC 13712-1:1994, *Information technology — Remote Operations: Concepts, Model and Notation*

2.2 Paired Recommendations | International Standards equivalent in technical content

- CCITT Recommendation X.200 (1988) *Reference Model of Open Systems Interconnection for CCITT Applications*.
ISO 7498:1984, *Information Processing Systems — Open Systems Interconnection — Basic Reference Model*.

3 Definitions

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

3.1 OSI Reference Model definitions

The following terms are defined in Rec. X..200 and ISO 7498:

- a) *application-entity*;
- b) *Application Layer*;
- c) *application-process*;
- d) *application protocol data unit*;
- f) application service element.
- e) Network Service Access Point.

3.2 Directory model definitions

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

- a) *access control*;
- b) *Administration Directory Management Domain*;
- c) *alias*;
- d) *attribute*;
- e) *attribute type*;
- f) *attribute value*;
- g) *authentication*;
- h) *Directory Information Tree (DIT)*;
- i) *Directory Management Domain (DMD)*;
- j) *Directory System Agent (DSA)*;
- k) *Directory User Agent (DUA)*;
- l) *distinguished name*;
- m) *entry*;
- n) *name*;
- o) *object (of interest)*;
- p) *Private Directory Management Domain*;
- q) *relative distinguished name*;
- r) *root*;
- s) *schema*;
- t) *security policy*;

- u) *subordinate object*;
- v) *superior entry*;
- w) *superior object*;
- x) *tree*.

3.3 Distributed Operation definitions

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

- a) *uni-chaining*;
- b) *multi-chaining*;
- c) *referral*.

3.4 Replication definitions

The following terms are defined in ITU-T Rec. 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:

- a) *the Directory*: a collection of open systems cooperating to provide directory services;
- b) *Directory Information Base (DIB)*: the set of information managed by the Directory;
- c) *(directory) user*: the end user of the Directory, i.e., the entity or person which accesses the Directory.

4 Abbreviations

ACI	Access Control Information
ADDMD	Administration Directory Management Domain
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
NSAP	Network Service Access Point
OSI	Open Systems Interconnection
PRDMD	Private Directory Management Domain
RDN	Relative Distinguished Name

5 Conventions

With minor exceptions this Directory Specification has been prepared according to the "Presentation of ITU-TS/ISO/IEC common text" guidelines in the Guide for ITU-TS and ISO/IEC JTC 1 Cooperation, March 1993.

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

This Directory Specification uses the term “1988 edition systems” to refer to systems conforming to the previous (1988) edition of the Directory Specifications, i.e., the 1988 edition of the series of CCITT X.500 Recommendations and the ISO/IEC 9594:1990 edition. Systems conforming to the current Directory Specifications are referred to as “1993 edition systems”.

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), 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.




Figure 1 — Access to the Directory

The information held in the Directory is collectively known as the *Directory Information Base* (DIB). Clause 7 of this Directory Specification overviews 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 overviewed in clause 8 of this Directory Specification 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.

It is likely that the Directory will be distributed, perhaps widely distributed, both along functional and organizational lines. Clause 9 overviews 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. Access control is overviewed in clause 10.

When the Directory is distributed, it may be desirable to replicate information to improve performance and availability. The Directory replication mechanism is overviewed in clause 11.

The provision and consumption of the directory services requires that the users (actually the DUAs) 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, overviewed 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 will govern which objects are listed in the Directory, which users will access the information, and which kinds of access they will 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 will be useful across a range of applications. These definitions are contained in ITU-T Rec. X.520 | ISO/IEC 9594-6 and ITU-T Rec. X.521 | ISO/IEC 9594-7;
- certain patterns of use of the Directory will be common across a range of applications: this area is overviewed further in annex A.

7 The Directory Information Base (DIB)

Note — The DIB, and its structure, are defined in ITU-T Rec. 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. 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.

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 — 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 all other entries are object entries. Alias entries point to object entries, and provide the basis for alternative names for the corresponding objects.

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.




Figure 3 — A Hypothetical Directory Information Tree

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.

8 The Directory Service

Note - The definition of the abstract service of the Directory can be found in ITU-T Rec. X.511 | ISO/IEC 9594-3.

8.1 Introduction

This clause provides an overview of the service provided to users, as represented by their DUAs, by the Directory. All services are provided by the Directory in response to requests from DUAs. There are requests which allow interrogation of the Directory, as described in 8.3, and those for modification, as described in 8.4. In addition, requests for service can be qualified, as described in 8.2. The Directory always reports the outcome of each request that is made of it. The form of the normal outcome is specific to the request, and is evident from the description of the request. Most abnormal outcomes are common to several requests. The possibilities are described in 8.5.

The Directory ensures that changes to the DIB, whether the result of a Directory service request, or by some other (local) means, result in a DIB which continues to obey the rules of the Directory schema.

A user and the Directory are bound together for a period of time at an access point to the Directory. At the time of binding, the user and the Directory optionally verify each other's identity.

8.2 Service Qualification

8.2.1 Service Controls

A number of controls can be applied to the various service requests, primarily to allow the user to impose limits on the use of resources which the Directory shall not surpass. Controls are provided on, among other things: the amount of time, the size of results, the scope of search, the interaction modes, and on the priority of the request.

8.2.2 Security Parameters

Each request may be accompanied by information in support of security mechanisms for protecting the Directory information. Such information may include the user's request for various kinds of protection; a digital signature of the request, together with information to assist the correct party to verify the signature.

8.2.3 Filters

A number of requests whose outcome involves information from or concerning a number of entries, may carry with them a filter. A filter expresses one or more conditions that an entry shall satisfy in order to be returned as part of the outcome. This allows the set of entries returned to be reduced to only those relevant.

8.3 Directory Interrogation

8.3.1 Read

A read request is aimed at a particular entry, and causes the values of some or all of the attributes of that entry to be returned. Where only some attributes are to be returned, the DUA supplies the list of attribute types of interest.

8.3.2 Compare

A compare request is aimed at a particular attribute of a particular entry, and causes the Directory to check whether a supplied value matches a value of that attribute.

Note - For example, this can be used to carry out password checking, where the password, held in the Directory, might be inaccessible for read, but accessible for compare.

8.3.3 List

A list request causes the Directory to return the list of immediate subordinates of a particular named entry in the DIT.

8.3.4 Search

A search request causes the Directory to return information from all of the entries within a certain portion of the DIT which satisfy some filter. The information returned from each entry consists of some or all of the attributes of that entry, as with read.

8.3.5 Abandon

An abandon request, as applied to an outstanding interrogation request, informs the Directory that the originator of the request is no longer interested in the request being carried out. The Directory may, for example, cease processing the request, and may discard any results so far achieved.

8.4 Directory Modification

8.4.1 Add Entry

An add entry request causes a new leaf entry (either an object entry, or an alias entry) to be added to the DIT.

8.4.2 Remove Entry

A remove entry request causes a leaf entry to be removed from the DIT.

Note - As with add entry, this service is presently intended for operation on "true leaf" entries, and will be enhanced in the future for the general case.

8.4.3 Modify Entry

A modify entry request causes the Directory to execute a sequence of changes to a particular entry. Either all of the changes are made, or none of them, and the DIB is always left in a state consistent with the schema. The changes allowed include the addition, removal, or replacement of attributes or attribute values.

8.4.4 Modify Distinguished Name

A modify distinguished name (DN) request is used to change the relative distinguished name of an entry (either an object entry or an alias) or to move an entry to a new superior in the DIT. If an entry has subordinates, then all subordinates are renamed or moved accordingly.

8.5 Other Outcomes

8.5.1 Errors

Any service may fail, for example because of problems with the user supplied parameters, in which case an error is reported. Information is returned with the error, where possible, to assist in correcting the problem. However, in general only the first error encountered by the Directory is reported. Besides the above-mentioned example of problems with the parameters supplied by the user (particularly invalid names for entries or invalid attribute types) errors may arise from violations of security policy, schema rules, and service controls.

8.5.2 Referrals

A service may fail because the particular access point to which the DUA is bound is not the most suitable for carrying out the request, e.g., because the information affected by the request is (logically) far away from the access point. In this case the Directory may return a referral, which suggests an alternative access point at which the DUA can make its request.

Note - The Directory and the DUA may each have a preference as to whether referrals are used, or whether the requests are *chained* (see 8.3.3.2). The DUA can express its preference by means of service controls. The Directory makes the final decision as to which approach is used.

9 The Distributed Directory

Note - The models of the Directory are defined in ITU-T Rec. X.501 | ISO/IEC 9594-2, while the procedures for the operation of the distributed Directory are specified in ITU-T Rec. X.518 | ISO/IEC 9594-4.

9.1 Functional Model

The functional model of the Directory is shown in Figure 4.




Figure 4 — Functional Model of the Directory

A *Directory System Agent (DSA)* is an OSI application process which is part of the Directory and whose role is to provide access to the DIB to DUAs and/or other DSAs. A DSA may use information stored in its local database or interact with other DSAs to carry out requests. Alternatively, the DSA may direct a requester to another DSA which can help carry out the request. Local databases are entirely implementation dependent.

9.2 Organizational Model

A set of one or more DSAs and zero or more DUAs managed by a single organization may form a *Directory Management Domain (DMD)*. The organization concerned may or may not elect to make use of the Directory Specifications to govern the communications among the functional components within the DMD.

The other Directory Specifications specify certain aspects of the behavior of DSAs. For this purpose, a group of DSAs within one DMD may, at the option of the organization which manages the DMD, behave as a single DSA .

A DMD may be an *Administration DMD (ADDMD)*, or a *Private DMD (PRDMD)*, depending on whether or not it is being operated by a public telecommunications organization.

Note: It should be recognized that the provision of support for private directory systems by ITU-T members falls within the framework of national regulations. Thus, the technical possibilities described may or may not be offered by an Administration which provides directory services. The internal operation and configuration of private DMDs is not within the scope of envisaged Directory Specifications.

9.3 Operation of the model

The DUA interacts with the Directory by communicating with one or more DSAs. A DUA need not be bound to any particular DSA. It may interact directly with various DSAs to make requests. For some administrative reasons, it may not always be possible to interact directly with the DSA which needs to carry out the request, e.g., to return some directory information. It is also possible that the DUA can access the Directory through a single DSA. For this purpose, DSAs will need to interact with each other.

The DSA is concerned with carrying out the requests of DUAs, and with obtaining the information where it does not have the necessary information. It may take the responsibility to obtain the information by interacting with other DSAs on behalf of the DUA.

A number of cases of request handling have been identified, as illustrated in Figures 5 through 7, and described below.

In Figure 5a, DSA C receives a referral from DSA A and is responsible for either conveying the request to the DSA B (named in the referral from DSA A), or conveying the referral back to the originating DUA.




Figure 5a — Referrals

Note — If DSA C returns the referral to the DUA, the "request (to B)" will not occur. Similarly, if DSA C conveys the request to DSA B, it will not return a referral to the DUA.

In Figure 5b, the DUA receives the referral from DSA C, and is responsible for reissuing the request directly to DSA A (named in the referral from DSA C).




Figure 5b — Referrals

Figure 6 shows DSA uni-chaining, whereby the request can be passed through several DSAs before the response is returned.




Figure 6 — Uni-chaining

Figure 7 shows multi-chaining, where the DSA associated with the DUA carries out the request by forwarding it to two or more other DSAs, the request to each DSA being identical.




Figure 7 — Multi-chaining

All of the approaches have their merits. For example, the approach in Figure 5 may be used where it is desirable to offload the burden from the local DSA. In other circumstances a hybrid approach that combines a more elaborate set of functional interactions may be needed to satisfy the initiator's request, as illustrated in Figure 8.




Figure 8 — Mixed modes hybrid approach

10 Access control in the Directory

Note — The directory access control model is defined in ITU-T Rec. X.501 | ISO/IEC 9594-2

Access to Directory information is determined by some administratively controlled security policy. Two aspects of the security policy which affect access to the Directory are the authentication procedures and the access control scheme.

Authentication procedures and mechanisms to support the Directory include methods to verify and propagate, where necessary, the identity of DSAs, Directory users, and the origin of information received at an access point. General authentication procedures are defined in ITU-T Rec. X.509 | ISO/IEC 9594-8.

The definition of an access control scheme to support the Directory includes methods to specify access control information, enforce access rights defined by that access control information, and to maintain access control information. The enforcement of access rights encompasses controlling access to Directory information related to DIT structure, Directory user information, and Directory operational information including access control information.

ITU-T Rec. X.501 | ISO/IEC 9594-2 defines one specific access control scheme (of potentially many), referred to as “basic access control” for the Directory. Administrative authorities may make use of all or parts of this scheme in implementing their security policies, or may freely define their own schemes at their discretion.. The basic access control scheme provides a means of controlling access to the Directory information within the DIB (potentially including structure and access control information). Control of access to information enables the prevention of unauthorized detection, disclosure, or modification of that information.

The basic access control model for the directory defines, for every operation, one or more points at which access control decisions may take place. Each access control decision involves:

- that component within the Directory being accessed;
- the user requesting the operation;
- a specific right necessary to complete a portion of the operation; and,
- the security policy governing access to that item.

11 Replication in the Directory

Note — Directory replication is defined in ITU-T Rec. X.525 | ISO/IEC 9594-9

11.1 Introduction

Replication in the Directory refers to the existence of copies of directory entry information and operational information held by DSAs other than the DSA responsible for the creation and update of the information. This DSA, containing the original information, is called the master DSA.

It is possible to construct directory systems that make no use of replicated information.

Replication of directory information serves to satisfy two general sorts of requirements, one related to the general quality of the service provided by the Directory and the other related to the management of directory systems.

The deployment of additional copies of directory entry information may be of use in the improvement of the service provided by the Directory by:

- a) improving the performance of directory systems by moving directory information "closer" to particular directory users;
- b) improving the availability of the directory service by introducing redundant directory information and directory components so that an individual component failure does not prevent all access to the information in some portion of the DIT.

The deployment of additional copies of directory entry information may be of use in the management of directory systems:

- a) by facilitating the distribution of certain operational information (e.g., knowledge); and
- b) by providing an opportunity to recover from severe system failures through the reconstruction of the information to be held in a component of the Directory from a copy of that information held in another component of the Directory.

11.2 Forms of directory replication

There are two forms of replicated entry information that may be held by the components of the Directory, cache copies and shadowed information.

Cache copies are copies of entry information that a component of the Directory obtains and uses in ways not specified in these Directory Specifications.

Shadowed copies are copies of directory information that a component of the Directory obtains and uses in ways specified in ITU-T Rec. X.525 | ISO/IEC 9594-9.

Directory System Agents may retain information obtained from another DSA only if permitted in the policy and agreement under which the information was originally supplied. A DSA retaining such information may only supply it to DUAs in accordance with the access control policy pertaining to the information. If it is known that there are no read access controls on the information, it may be supplied as if read permission were granted.

A DSA holding copied information forwards all requests that would modify the copy information, and all requests that indicate that copy information shall not be used, to the master DSA holding the information.

When responding to an interrogation with copy information, a DSA holding copy information indicates that a copy was used to satisfy the request.

The administrative authorities responsible for two DSAs may establish a shadowing agreement whereby one DSA, a shadow supplier, contracts to provide another DSA, a shadow consumer, with shadowed information from an agreed portion of the DIT. If permitted by the shadowing agreement under which shadowed information is obtained, a shadow consumer may enter into agreements with other DSAs to be a shadow supplier for that information.

In addition to the provision of updates to copies of entry information held in the shadow consumer, operational information (e.g., knowledge) may also be provided to the shadow consumer by the shadow supplier.

In any one shadowing agreement, the information to be replicated will typically comprise three elements:

- Replicated entry information from within a subtree of the DIT
- Relevant operational information, including access control information, required to give full read access to the replicated information
- Optionally, subordinate knowledge information.

The replicated information may form a subset of the complete information within the subtree, in that:

- A selection of the entries may be made by specifying only those that meet certain criteria on their object classes
- Within each entry, a selection of the attributes may be made in accordance with a specification of attributes.

11.3 Replication and consistency of directory information

Consistency in the Directory is achieved when all copies of a specific attribute are the same. At times consistency may be subject to compromise because transient inconsistencies can exist within the Directory for shadowed information and permanent inconsistencies can exist for cached information.

Cached entry information may become and indefinitely remain inconsistent with entry information maintained by that component of the Directory to which updates are directed. In contrast, shadowed information held by a shadow consumer is brought into agreement with the corresponding information held by a shadow supplier according to a schedule contracted to as part of the shadowing agreement.

It is essential that the information contained within an instance of an individual object entry be internally consistent. Any mechanism for replication shall be accompanied by mechanisms to maintain the internal consistency of replicated information and the reliability of the service. The Directory defines schema procedures to ensure the internal consistency of an entry.

It is also essential that the knowledge information which allows the DIT to be distributed across DSAs be accurate. Any mechanism for replication shall be accompanied by mechanisms to maintain the accuracy of knowledge information and the reliability of the service. The Directory defines procedures for manipulating the minimum knowledge information needed by a DSA to ensure the coherency of the DIT.

In an environment where directory information is replicated, the Directory has no specific time constraints to achieve consistency. A user of shadowed information will have a high level of confidence in it because:

- the shadowed information is internally consistent;
- the knowledge relating it to the DIT is accurate; and
- the shadowed entry will ultimately become consistent with the entry in the master DSA.

11.4 Views of replication

This section describes the distinct ways in which the existence of replication of directory information manifests itself to:

- a) directory users;
- b) administrative users; and
- c) the operational components of the Directory (DSAs).

11.4.1 Directory user view

Because of the nature of the operation of the Directory, replicated information will be generally consistent with information held by the master DSA for that information. Therefore, in the general case, requested information, returned to the end user, will be of an acceptable nature and the fact that it is from a copy will not be important.

The directory user is always notified if a request has been satisfied from entry copy information. In the case when the user has a critical need, or can detect an inconsistency, he has the option of requesting access to information held by the master DSA.

The user of the Directory is therefore offered the choice between increased levels of performance and availability at the cost of occasionally receiving information that is out of date and a maximum level of information timeliness at the cost of potentially reduced levels of performance and availability.

11.4.2 Administrative user view

An administrative user is charged with the management of the information held in and the service provided by a DSA. To perform this management function the administrative user requires tools to monitor, control and optimize the DSA's service.

The standardized (and local) capability of a DSA to support replication is one of the principal tools available to the administrative user to optimize the service provided by a DSA.

11.4.3 DSA view

Although a DSA can detect the difference between replicated information and information which is held by a master, it generally uses both in the same way, i.e., it satisfies user interrogation requests with either, depending on which is most conveniently available to it.

There are two exceptions to this equivalence of master and replicated information. A DSA only uses entry information to satisfy requests to modify the DIB and interrogation requests that signal that replicated information is not acceptable.

In addition, since the information held locally may be known to be partial (see 11.2), a DSA may pass an inquiry to another DSA better able to provide the information required.

11.5 Replication and Access Control

The Access Control Model allows access control information to be specified for an area of the DIT. That area may span DSA boundaries. If multiple DSAs are involved, each will hold the appropriate access control information.

Any time entries are replicated to another DSA, the access control information shall also be replicated.

12 Directory protocols

Note- the OSI application layer protocols defined to allow DUAs and DSAs in different open system to cooperate are specified in ITU-T Rec. X.519 | ISO/IEC 9594-5.

There are four Directory protocols:

- the Directory Access Protocol (DAP), which defines the exchange of requests and outcomes between a DUA and a DSA;
- the Directory System Protocol (DSP), which defines the exchange of requests and outcomes between two DSAs;

- The Directory Information Shadowing Protocol (DISP), which defines the exchange of replication information between two DSAs that have established shadowing agreements;
- The Directory Operational Binding Management Protocol (DOP), which defines the exchange of administrative information between two DSAs to administer operational bindings between them.

Each protocol is defined by one or more application contexts, each containing a set of protocol elements. For example, the DAP contains protocol elements associated with interrogating and modifying the Directory.

Each application context is made up of application service elements. These application service elements are defined to use the Remote Operations Service Element (ROSE) of ITU-T Rec. X.880 | ISO/IEC 9072-1 to structure and support their interactions. Thus the DAP, DSP, DISP, and DOP are defined as sets of remote operations and errors using the ROS notation.

Annex A — Applying the Directory

(This Annex forms an integral part of this Recommendation | Standard)

A.1 The Directory environment

Note - In this subclause, the term *network* is used with its general meaning to denote the set of interlinked systems and processes relevant to any telecommunications service, not only one which relates to the OSI network layer

The Directory exists in and provides services in the following environment:

- a) Many telecommunications networks will be on a large scale, and will constantly undergo change:
 - 1) Objects of various kinds will enter and leave the network without warning and may do so either singly or in groups;
 - 2) The connectivity of the objects (particularly network nodes) will change, owing to the addition or removal of paths between them;
 - 3) Various characteristics of the objects, such as their addresses, availability, and physical locations, may change at any time;
- b) Although the overall rate of changes is high, the useful lifetime of any particular object is not short. An object will typically be involved in communications much more frequently than it will change its address, availability, physical location, etc.;
- c) The objects involved in current telecommunications services are typically identified by numbers or other strings of symbols, selected for their ease of allocation or processing but not for ease of use by human beings.

A.2 Directory service characteristics

The need for directory capabilities arises from:

- a) the desire to isolate (as far as possible) the user of the network from the frequent changes to it. This can be accomplished by placing a 'level of indirection' between the users and the objects with which they deal. This involves the users referring to objects by name, rather than by, for example, address. The Directory provides the necessary mapping service;
- b) the desire to provide a more "user-friendly" view of the network. For example, the use of aliases, the provision of *yellow-pages* (see A.3.5) etc., helps to relieve the burden of finding and using network information.

The Directory allows users to obtain a variety of information about the network, and provides for the maintenance, distribution and security of that information.

A.3 Patterns of use of the Directory

Note — this clause is concerned only with Directory retrieval: it is assumed that the Directory modification services are used solely to maintain the DIB in the form necessary for the application over time.

A.3.1 Introduction

The Directory service is defined in the Directory Specifications in terms of particular requests that a DUA can make and the parameters of them. An application designer is likely, however, to think in more goal-oriented terms when considering the information retrieval requirements of the Directory in that application. Accordingly, this clause describes a number of high-level patterns of use of the Directory service that are likely to be relevant to many applications.

A.3.2 Look-up

The straight Directory look-up, which is likely to be the most frequent type of query of the Directory, involves the DUA supplying the distinguished name of an object, together with an attribute type. The Directory will return any value(s) corresponding to that attribute type. This is a generalization of the classic directory function, which is obtained when the attribute type requested corresponds to a particular type of address. Attribute types for various kinds of address are standardized, including OSI PSAP address, Message Handling O/R address, and telephone and telex numbers.

Look-up is supported by the read service, which also provides the following further generalizations:

- look-up can be based upon names other than the distinguished name of the object, e.g., aliases;
- the values from a number of attribute types can be requested with a single request: the extreme case being that the values of all attributes in the entry are to be returned.

A.3.3 User-friendly naming

Names can be given to objects in such a way as to maximize the chances that these names can be predicted (or perhaps remembered) by human users. Names which have this property would typically be made up of attributes which are somehow inherent to the object, rather than being fabricated for the purpose. The name of an object will be common among all of the applications which refer to it.

A.3.4 Browsing

In many human-oriented uses of the Directory, it may not be possible for the user (or DUA) to directly quote a name, user-friendly or otherwise, for the object about which information is sought. However, perhaps the user will 'know it when he sees it'. The browsing capability will allow a human user to wander about the DIB, looking for the appropriate entries.

Browsing is accomplished by combinations of the list and search services, possibly in conjunction with read (although the search service includes the capability of read).

A.3.5 Yellow Pages

There are a variety of ways to provide a *Yellow Pages* type capability. The simplest is based upon filtering, using assertions about particular attributes whose values are the categories (e.g., the 'Business Category' attribute type defined in ITU-T Rec. X. 520 | ISO/IEC 9594-6). This approach does not require any special information being set-up in the DIT, except to ensure that the requisite attributes are present. However, in the general case it may be expensive to search where there is a large population, because filtering requires the generation of the universal set which is to be filtered.

An alternative approach is possible, based upon the setting up of special subtrees, whose naming structures are designed especially for *Yellow Pages* type searching. Shown in Figure A.1 is an example of a *Yellow Pages* subtree populated by alias entries only. In reality, the entries within the *Yellow Pages* subtrees may be a mixture of object and alias entries, so long as there exists only one object entry for each object stored in the Directory.




Figure A.1 — An Approach to *Yellow Pages*

A.3.6 Groups

A group is a set whose membership can change over time by explicit addition and removal of members. The group is an object, as are its members. The Directory can be requested to:

- indicate whether or not a particular object is a member of a group;
- list the membership of a group.

Groups are supported by having the entry for the group contain a multiple valued 'Member' attribute (such an attribute type is defined in ITU-T Rec. X.520 | ISO/IEC 9594-6). The two capabilities mentioned can then be carried out by means of compare and read respectively.

A member of a group could itself be a group, if this is meaningful for the application. However, the necessary recursive verification and expansion services would have to be created by the DUA out of the non-recursive versions provided.

A.3.7 Authentication

Many applications require the objects taking part to offer some proof of their identity before they are permitted to carry out some action. The Directory provides support for this authentication process. (As a separate matter, the Directory itself requires its users to authenticate themselves, so as to support access control).

The more straightforward approach to authentication, called 'simple authentication', is based upon the Directory holding a 'User Password' attribute in the entry for any user that wishes to be able to authenticate itself to a Service. At the request of the Service, the Directory will confirm or deny that a particular value supplied is actually the user's password. This avoids the user needing a different password for every Service. In cases where the exchange of passwords in a local environment that uses simple authentication is considered to be inappropriate, the Directory optionally provides means to protect those passwords against replay or misuse by a one way function.

The more complex approach, called 'strong authentication' is based upon public key cryptography, where the Directory acts as a repository of users' public encryption keys, suitably protected against tampering. The steps that users can take to obtain each others' public keys from the Directory, and then to authenticate with each other using them, are described in detail in ITU-T Rec. X.509 | ISO/IEC 9594-8.

A.4 Generic applications

A.4.1 Introduction

There are a number of generic applications which can be imagined as implicitly supported by the Directory: applications which are not specific to any particular telecommunications service. Two such applications are described herein: the inter-personal communications directory and the inter-system communications directory (for OSI).

Note - Authentication, described in A.3.7 as an *access pattern*, could alternatively be thought of as a generic Directory application.

A.4.2 Inter-personal communications

The intent of this application is to provide humans or their agents with information on how to communicate with other humans, or groups thereof.

The following classes of object are certainly involved: person, organizational role, and group. Many other classes are involved too, perhaps in a less direct way, including: country, organization, organizational unit.

The attribute types concerned, other than those used in naming, are generally the addressing attributes. Typically the entry for a particular person will have the addresses corresponding to each of the communication methods by which that person can be reached, selected from an open-ended list which includes at least the following: telephony, electronic mail, telex, ISDN, physical delivery (e.g., the postal system), facsimile. In some cases, such as electronic mail, the entry will have some additional information such as the types of information which the user's equipment can handle. If authentication is to be supported, then User Password and/or Credentials will be needed.

The naming schemes used for the various object classes should be user-friendly, with aliases being set up as appropriate to provide alternative names, provide continuity after a name change, etc.

The following access patterns will be manifested in this application: look-up, user-friendly naming, browsing, *Yellow Pages*, and groups. To varying degrees, authentication will also be used.

A.4.3 Inter-system communications (for OSI)

According to the OSI Reference Model, two directory functions are required in OSI, one, operating in the Application Layer, which maps application-entity titles onto presentation-addresses, and one, in the Network Layer, which maps NSAP-addresses onto SNPA-addresses (SNPA = Subnetwork Point of Attachment).

Note - For the remainder of this subclause, only the Application Layer case is dealt with.

This function is carried out by consulting the Directory if the information required to accomplish the mapping is not conveniently available by other means.

The users are application-entities and the object classes of interest are also application-entities, or subclasses thereof.

The main attribute type concerned, other than those used for naming, is the presentation-address. Other attribute types, not viewed as necessary for the directory function itself, could support verifying or finding out the application-entity type, or the lists of application_contexts, abstract syntaxes, etc. supported. The authentication-related attribute types could also be relevant.

The main access pattern to be manifested will be look-up.

Annex B — Amendments & corrigenda

(This annex does not form an integral part of this Recommendation | International Standard)

This edition of this Directory Specification includes the following amendments:

- Amendment 1 for Replication, Schema, and Access Control

This edition of this Directory Specification includes the following technical corrigenda correcting the defects reported in the following defect reports:

- There were no defect reports against the previous edition of this Directory Specification.