
**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

IECNORM.COM : Click to view the full PDF of ISO/IEC 9594-1:2014



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2014

All rights reserved. Unless otherwise specified, 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
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

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.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of the joint technical committee is to prepare International Standards. 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.

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.

ISO/IEC 9594-1 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 6, *Telecommunications and information exchange between systems*, in collaboration with ITU-T. The identical text is published as Rec. ITU-T X.500 (10/2012).

This seventh edition cancels and replaces the sixth edition (ISO/IEC 9594-1:2008), which has been technically revised.

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: Public-key and attribute certificate frameworks*
- *Part 9: Replication*

IECNORM.COM : Click to view the full PDF of ISO/IEC 9594-1:2014

CONTENTS

	<i>Page</i>
1 Scope	1
2 Normative references	1
2.1 Identical Recommendations International Standards	1
3 Definitions	2
3.1 Communication model definitions	2
3.2 Directory model definitions	2
3.3 Distributed Operation definitions	3
3.4 Replication definitions	3
3.5 Basic directory definitions	3
4 Abbreviations	3
5 Conventions	3
6 Overview of the Directory	4
7 The Directory Information Base (DIB)	5
8 The Directory service	7
8.1 Introduction	7
8.2 Service qualification	7
8.3 Directory interrogation	8
8.4 Directory modification	8
8.5 Other outcomes	9
9 The distributed Directory	9
9.1 Functional model	9
9.2 Organizational model	10
9.3 Operation of the model	10
10 Access control in the Directory	13
11 Service administration	14
12 Replication in the Directory	15
12.1 Introduction	15
12.2 Forms of Directory replication	15
12.3 Replication and consistency of Directory information	16
12.4 Views of replication	16
12.5 Replication and Access Control	17
13 Directory protocols	17
14 Systems management of the Directory	17
14.1 Introduction	17
14.2 Management of the DIT domain	18
14.3 Management of Directory components	18
Annex A – Applying the Directory	19
A.1 The Directory environment	19
A.2 Directory service characteristics	19
A.3 Patterns of use of the Directory	19
Annex B – Amendments and corrigenda	23

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 seventh edition technically revises and enhances, the sixth edition of this Recommendation | International Standard.

This seventh 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 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. 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 (2012) | ISO/IEC 9594-2:2014, *Information technology – Open Systems Interconnection – The Directory: Models.*
- Recommendation ITU-T X.509 (2012) | ISO/IEC 9594-8:2014, *Information technology – Open Systems Interconnection – The Directory: Public-key and attribute certificate frameworks.*
- Recommendation ITU-T X.511 (2012) | ISO/IEC 9594-3:2014, *Information technology – Open Systems Interconnection – The Directory: Abstract service definition.*
- Recommendation ITU-T X.518 (2012) | ISO/IEC 9594-4:2014, *Information technology – Open Systems Interconnection – The Directory: Procedures for distributed operation.*
- Recommendation ITU-T X.519 (2012) | ISO/IEC 9594-5:2014, *Information technology – Open Systems Interconnection – The Directory: Protocol specifications.*
- Recommendation ITU-T X.520 (2012) | ISO/IEC 9594-6:2014, *Information technology – Open Systems Interconnection – The Directory: Selected attribute types.*
- Recommendation ITU-T X.521 (2012) | ISO/IEC 9594-7:2014, *Information technology – Open Systems Interconnection – The Directory: Selected object classes.*
- Recommendation ITU-T X.525 (2012) | ISO/IEC 9594-9:2014, *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 February 2010 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 and all parts of ISO/IEC 9594.

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 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.




Figure 3 – A hypothetical Directory Information Tree

The hierarchical group function allows an alternative hierarchical relationship to be established among entries independent of the hierarchical relationship reflected by the DIT structure. The Directory Search operation (see 8.3.4) can return not only information from matched entries, but also other members of the hierarchical group to which the matched entry might belong. The hierarchical group function also has the advantage that it allows hierarchical relationships to be changed without changing the DIT structure and thereby the distinguished names of the entries.

8 The Directory service

NOTE – The definition of the abstract service of the Directory can be found in Rec. ITU-T 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 and/or LDAP clients, by the Directory. All services are provided by the Directory in response to requests from DUAs and/or LDAP clients. 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 that the Directory shall not surpass, but also to control the progress of the Directory operation. Controls are provided on, among other things: the amount of time, the size of results, the scope of search, the interaction modes, and 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 one or more filters. A filter expresses one or more conditions that an entry or a compound 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, or a compound entry and causes the values of some or all of the attributes of that entry to be returned. In the case of compound entries, family member information is contained in a package (of syntax similar to that of an attribute) comprising the selected family information. Where only some attributes are to be returned, the DUA supplies the list of attribute types of interest as part of the request. A DUA may also supply one or more contexts for one or more attribute types of interest, in order to select only those values that apply in the specified contexts.

NOTE – LDAP clients do not support the Read operation.

8.3.2 Compare

A compare request is aimed at a particular attribute of a particular entry or a compound entry, and causes the Directory to check whether a supplied value matches a value of that attribute. A DUA may also supply one or more contexts for the attribute value of interest to constrain the comparison operation.

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.

NOTE – LDAP clients do not support the List operation.

8.3.4 Search

A search request causes the Directory to return information from all of the entries or compound entries within one or more portions of the DIT that satisfy some filters. The information returned from each entry consists of some or all of the attributes of that entry as with read. Information returned from related entries may be combined according to some join criteria.

It is possible to put restrictions on the types of searches that can be performed by use of search-rules. It is also possible, as a facility of search-rules, to progressively relax or tighten up searches within a single Directory operation, if too few or too many items of entry information would otherwise be returned.

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 to be added to the DIT. Contexts may be included with the attribute values for the new entry.

8.4.2 Remove entry

A remove entry request causes a leaf entry or, if required, the entries comprising a compound entry to be removed from the DIT.

8.4.3 Modify entry

A modify entry request causes the Directory to execute a sequence of changes to a particular entry or family member. 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. Contexts may be included with attribute values that are added to the entry. This operation can only be used on a single family member, but cannot manipulate a compound entry as a whole.

A Modify Entry operation can supply in the result, if required, the information contained in the entry or compound entry after a successful modification has taken place.

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, an alias, or a family member) or to move an entry, if not a family member, to a new superior in the DIT. If an entry has subordinates, then all subordinates are renamed or moved accordingly. In the case of family members, these can be moved to have new superiors, provided that they remain within the same compound entry.

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 or LDAP client 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 or LDAP client 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 9.3). 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 Rec. ITU-T X.501 | ISO/IEC 9594-2, while the procedures for the operation of the distributed Directory are specified in Rec. ITU-T X.518 | ISO/IEC 9594-4.

9.1 Functional model

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

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

An *LDAP server* is an application process which is part of the Directory, that responds to requests via the LDAP protocol, and whose role is to provide access to the DIB to LDAP clients and/or LDAP requesters. An LDAP server may use information stored in its local database or may direct a requester to another LDAP responder or LDAP server which can help carry out the request. As with DSAs, local databases are entirely implementation dependent.




Figure 4 – Functional model of the Directory

9.2 Organizational model

A set of one or more DSAs and/or LDAP servers and zero or more DUAs and/or LDAP clients 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 behaviour 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.

9.3 Operation of the model

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

The DSA is concerned with carrying out the requests of DUAs and LDAP clients, 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 and/or LDAP servers on behalf of the DUA or LDAP client.

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 1 – 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, 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 LDAP client.




Figure 5b – Referrals

NOTE 2 – In addition to any other ITU-T X.500 protocols it may support, DSA C in Figure 5b shall also be an LDAP responder.

NOTE 3 – If DSA C returns the referral to the LDAP client, 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 LDAP client.

NOTE 4 – If DSA C returns the referral to the LDAP client, the referral shall be in the form of an LDAP referral. If the referral returned by DSA A is in the form of an LDAP referral, DSA C may return that referral directly to the LDAP client; otherwise,

DSA C shall either convey the request to DSA B or translate the referral into an LDAP referral. If DSA C returns the referral to the LDAP client, the client will bind directly to DSA B, which shall also be an LDAP responder. It will also be necessary for DSA B to be an LDAP responder if DSA A returns an LDAP referral and DSA C conveys the request directly to DSA B.

In Figure 5c, 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 5c – Referrals

In Figure 5d, the LDAP client 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 5d – Referrals

NOTE 5 – Both DSA A and DSA C in Figure 5d shall be LDAP responders. Alternatively, either of these two DSAs could be an LDAP server.

NOTE 6 – The referral that is returned to the LDAP client shall be in the form of an LDAP referral.

Figures 6a through 6c show DSA uni-chaining, whereby the request can be passed through several DSAs before the response is returned.




Figure 6a – Uni-chaining




Figure 6b – Uni-chaining

NOTE 7 – In addition to any other ITU-T X.500 protocols it may support, the DSA on the left in Figure 6b shall also be an LDAP responder.




Figure 6c – Uni-chaining

NOTE 8 – In addition to any other ITU-T X.500 protocols it may support, the DSA in Figure 6c shall also be both an LDAP responder and an LDAP requester.

Figures 7a through 7c show multi-chaining, where the DSA associated with the DUA or LDAP client carries out the request by forwarding it to two or more other DSAs and/or LDAP servers, the request to each DSA or LDAP server being identical.




Figure 7a – Multi-chaining




Figure 7b – Multi-chaining

NOTE 9 – In addition to any other ITU-T X.500 protocols it may support, the DSA on the left in Figure 7b shall also be an LDAP responder.




Figure 7c – Multi-chaining

NOTE 10 – In addition to any other ITU-T X.500 protocols it may support, the DSA on the left in Figure 7c shall also be both an LDAP responder and an LDAP requester.

All of the approaches have their merits. For example, the approaches in Figures 5b and 5d 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 Figures 8a and 8b.




Figure 8a – Mixed modes hybrid approach




Figure 8b – Mixed modes hybrid approach

10 Access control in the Directory

NOTE – The Directory access control model is defined in Rec. ITU-T 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 Rec. ITU-T 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.

Rec. ITU-T 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.

Control of access to information enables the prevention of unauthorized detection, disclosure, or modification of that information. Rec. ITU-T X.501 | ISO/IEC 9594-2 defines three specific access control schemes for the Directory, referred to as "basic access control", "simplified access control", and "rule-based access control". Administrative authorities may make use of all or parts of these schemes 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). The simplified access control scheme provides a subset of the functionality of the basic access control scheme. The rule-based access control scheme provides additional means of controlling access to the Directory information within the DIB (potentially including structure and access control information), based on clearances and labels. The rule-based access control scheme can be used in conjunction with either the simplified or basic access control, or the rule based access control scheme can be used alone.

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, possibly a complete compound entry;
- 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.

The rule-based 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:

- a clearance associated with the user requesting an access;
- security label(s) associated with the information being accessed;
- the security policy rules governing that access which defines whether access should be denied, given a relationship between the clearance and security label.

11 Service administration

The Directory Abstract Service as defined in Rec. ITU-T X.511 | ISO/IEC 9594-3 provides to the user powerful and efficient means for browsing and reading Directory information.

These Directory Specifications provide extensive service administration capabilities that allow administrative authorities to administer and constrain the service to a user. There may be several reasons for an administrative authority to constrain and adjust the service given to a user:

- An administrative authority has knowledge about the quality of information it holds. To improve the rate of successful Directory searches and to ensure that only quality information is returned, an administrative authority can constrain what attribute types that are allowed in a search filter and what information can be returned.
- To protect investment in verified and cleaned-up information, an administrative authority may have quite stringent restriction on what information can be returned as adopted to the type of user and the particular type of service provided.
- An administrative authority may want to prevent misuse of information, for example for mass marketing purposes, e.g., by picking up all people on a particular street, by picking-up all people with a certain profession, etc.
- Protection of personal data beyond what is possible using access control. This includes returning fake postal addresses, not allowing searches based on very short character strings, not allowing searches using certain combinations of attributes, or requiring certain combinations, etc.
- What restriction and what adaptation that should be made on the provided service may depend on the user group.

12 Replication in the Directory

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

12.1 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.

12.2 Forms of Directory replication

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

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 Rec. ITU-T X.525 | ISO/IEC 9594-9.

Multiple master implementations maintain more than one writeable instance of each entry within a given set of directory entries. Each writeable copy of a directory entry is complete (i.e., it holds all user attributes and DSA-shared operational attributes). Exactly one of the instances is identified in a manner that permits the Directory to identify that instance as the primary master in order to support deployment scenarios in which it is necessary to perform updates against a single DSA (e.g., when incrementing an attribute value used as a counter). The manner in which a Directory component obtains writeable copies of an entry, and the manner in which the writeable copies are brought into a state of consistency following a modification is outside the scope of this Recommendation | International Standard.

DSAs 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 and/or LDAP clients 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 cached or shadowed information forwards all requests that would modify the copy information to a master DSA holding the information. A DSA holding copied information forwards all requests that indicate that copy information shall not be used, to the master DSA holding the information.

When responding to an interrogation with cached or shadowed copy information, a DSA holding that 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 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;
- within each attribute, a selection of the attribute values may be made based on their contexts.

12.3 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 each view 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 its view of the DIT is accurate; and
- the shadowed entry will ultimately become consistent with the entry in the master DSA.

12.4 Views of replication

This subclause 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).

12.4.1 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.

12.4.2 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.

12.4.3 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 12.2), a DSA may pass an inquiry to another DSA better able to provide the information required.

NOTE – A DSA may contain replicated information from several sources, and this information may overlap. If this is the case, the DSA shall separately maintain each such view of the information, as provided by replication.

12.5 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.

13 Directory protocols

NOTE – The Directory protocols defined to allow DUs and DSAs in different open system to cooperate are specified in Rec. ITU-T 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 as a set of protocol elements. For example, the DAP contains protocol elements associated with interrogating and modifying the Directory.

14 Systems management of the Directory

NOTE – Directory systems management is defined in Rec. ITU-T X.530 | ISO/IEC 9594-10.

14.1 Introduction

The purpose of Directory management is to assure that needed, accurate Directory information is available to users as scheduled with the expected response time, integrity, security, and level of consistency. Furthermore, systems management should be accomplished with the minimum burden on processing time and memory on platforms and the communications system.

Management of the Directory is divided into four major segments:

- a) management of the DIT Domain: Management of Directory information;
- b) management of the operation of a single DSA;
- c) management of a single DUA; and
- d) management of the DMD – Integrated management of the functional components of the Directory.

The systems management specification addresses the first three segments. Management of the Directory Management Domain is for further study.

14.2 Management of the DIT domain

The user attributes in the Directory are managed by the Directory Access Protocol. Operational attributes may also be managed using DAP. This includes attributes in the information framework, subschema attributes, access control attributes, and the attributes in the DSA information tree, including knowledge. Knowledge may also be managed using the Directory Operational Binding Management Protocol, the Directory Information Shadowing Protocol, and the Directory System Protocol.

14.3 Management of Directory components

The System Management specification defines the OSI Systems Management managed objects used to manage the Directory components (DUAs and DSAs) within a Directory Domain. Management of these Directory components can be accomplished using the Common Management Information Services and protocol.

Some management requirements are not fulfilled by the Directory or Management Services but are by locally defined services.