

INTERNATIONAL
STANDARD

ISO/IEC
9594-3

Second edition
1995-09-15

Information technology — Open Systems Interconnection — The Directory: Abstract service definition

*Technologies de l'information — Interconnexion de systèmes ouverts
(OSI) — L'Annuaire: Définition du service abstrait*



Reference number
ISO/IEC 9594-3:1995(E)



Documents
Document Sharing Hub

Contents

	Page
1 Scope.....	1
2 Normative references	1
2.1 Identical Recommendations International Standards	1
2.2 Paired Recommendations International Standards equivalent in technical content	2
3 Definitions.....	2
3.1 Basic Directory definitions	2
3.2 Directory model definitions	2
3.3 Directory Information Base definitions	2
3.4 Directory entry definitions.....	2
3.5 Name definitions.....	3
3.6 Distributed operations definitions.....	3
3.7 Abstract Service definitions	3
4 Abbreviations	3
5 Conventions.....	3
6 Overview of the Directory Service.....	4
7 Information types and common procedures	4
7.1 Introduction.....	4
7.2 Information types defined elsewhere	5
7.3 Common arguments	5
7.4 Common results	6
7.5 Service controls.....	7
7.6 Entry information selection.....	8
7.7 Entry information.....	9
7.8 Filter.....	9
7.9 Paged results	11
7.10 Security parameters.....	12
7.11 Common elements of procedure for basic-access-control	13
7.12 Optionally-signed parameters	14
8 Bind and Unbind operations.....	14
8.1 Directory Bind	14
8.2 Directory Unbind	16

9	Directory Read operations.....	16
9.1	Read	16
9.2	Compare.....	18
9.3	Abandon.....	19
10	Directory Search operations	19
10.1	List	19
10.2	Search.....	21
11	Directory Modify operations.....	24
11.1	Add Entry.....	24
11.2	Remove Entry	25
11.3	Modify Entry.....	26
11.4	Modify DN.....	28
12	Errors.....	30
12.1	Error precedence	30
12.2	Abandoned	30
12.3	Abandon Failed.....	31
12.4	Attribute Error.....	31
12.5	Name Error.....	32
12.6	Referral	32
12.7	Security Error.....	32
12.8	Service Error	33
12.9	Update Error	34
	Annex A – Abstract Service in ASN.1	35
	Annex B – Operational semantics for Basic Access Control	42
	Annex C – Amendments and corrigenda.....	58

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 9594-3 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 21, *Open systems interconnection, data management and open distributed processing*, in collaboration with ITU-T. The identical text is published as ITU-T Recommendation X.511.

Implementors should note that a defect resolution process exists and that corrections may be applied to this part of ISO/IEC 9594 in the form of technical corrigenda. A list of approved technical corrigenda for this part of ISO/IEC 9594 can be obtained from the subcommittee secretariat. Published technical corrigenda are available from your national standards organization.

This second edition technically revises and enhances ISO/IEC 9594-3:1990. It also incorporates technical corrigendum 1:1991, technical corrigendum 2:1992, technical corrigendum 3:1992 and technical corrigendum 4:1993. Implementations may still claim conformance to the first edition of this part of ISO/IEC 9594. However, at some point, the first edition will no longer be supported (i.e. reported defects will no longer be resolved). It is recommended that implementations conform to this second edition as soon as possible.

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*

Annexes A and B form an integral part of this part of ISO/IEC 9594. Annex C is for information only.

Introduction

This Recommendation | International Standard, together with the 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 defines the capabilities provided by the Directory to its users.

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, provides the ASN.1 module for the directory abstract service.

Annex B, which is an integral part of this Recommendation | International Standard, provides charts that describe the semantics associated with Basic Access Control as it applies to the processing of a Directory operation.

Annex C, 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.

This page intentionally left blank

INTERNATIONAL STANDARD

ITU-T RECOMMENDATION

INFORMATION TECHNOLOGY – OPEN SYSTEMS INTERCONNECTION – THE DIRECTORY: ABSTRACT SERVICE DEFINITION

1 Scope

This Recommendation | International Standard defines in an abstract way the externally visible service provided by the Directory.

This Recommendation | International Standard does not specify individual implementations or products.

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.500 (1993) | ISO/IEC 9594-1:1995, *Information technology – Open Systems Interconnection – The Directory: Overview of concepts, models and services.*
- ITU-T Recommendation X.501 (1993) | ISO/IEC 9594-2:1995, *Information technology – Open Systems Interconnection – The Directory: Models.*
- ITU-T Recommendation X.518 (1993) | ISO/IEC 9594-4:1995, *Information technology – Open Systems Interconnection – The Directory: Procedures for distributed operation.*
- ITU-T Recommendation X.519 (1993) | ISO/IEC 9594-5:1995, *Information technology – Open Systems Interconnection – The Directory: Protocol specifications.*
- ITU-T Recommendation X.520 (1993) | ISO/IEC 9594-6:1995, *Information technology – Open Systems Interconnection – The Directory: Selected attribute types.*
- ITU-T Recommendation X.521 (1993) | ISO/IEC 9594-7:1995, *Information technology – Open Systems Interconnection – The Directory: Selected object classes.*
- ITU-T Recommendation X.509 (1993) | ISO/IEC 9594-8:1995, *Information technology – Open Systems Interconnection – The Directory: Authentication framework.*
- ITU-T Recommendation X.525 (1993) | ISO/IEC 9594-9:1995, *Information technology – Open Systems Interconnection – The Directory: Replication*
- ITU-T Recommendation X.680 (1994) | ISO/IEC 8824-1:1995, *Information technology – Abstract Syntax Notation One (ASN.1): Specification of basic notation.*
- ITU-T Recommendation X.681 (1994) | ISO/IEC 8824-2:1995, *Information technology – Abstract Syntax Notation One (ASN.1): Information object specification.*
- ITU-T Recommendation X.682 (1994) | ISO/IEC 8824-3:1995, *Information technology – Abstract Syntax Notation One (ASN.1): Constraint specification.*
- ITU-T Recommendation X.683 (1994) | ISO/IEC 8824-4:1995, *Information technology – Abstract Syntax Notation One (ASN.1): Parameterization of ASN.1 specifications.*

- ITU-T Recommendation X.880 (1994) | ISO/IEC 13712-1:1995, *Information technology – Remote Operations: Concepts, model and notation.*
- ITU-T Recommendation X.881 (1994) | ISO/IEC 13712-2:1995, *Information technology – Remote Operations: OSI realizations – Remote Operations Service Element (ROSE) service definition.*

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 Basic Directory definitions

The following terms are defined in ITU-T Rec. X.500 | ISO/IEC 9594-1:

- a) *Directory;*
- b) *Directory Information Base;*
- c) *(Directory) User.*

3.2 Directory model definitions

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

- a) *Directory System Agent;*
- b) *Directory User Agent.*

3.3 Directory Information Base definitions

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

- a) *alias entry;*
- b) *Directory Information Tree;*
- c) *(Directory) entry;*
- d) *immediate superior;*
- e) *immediately superior entry/object;*
- f) *object;*
- g) *object class;*
- h) *object entry;*
- i) *subordinate;*
- j) *superior.*

3.4 Directory entry definitions

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

- a) *attribute;*
- b) *attribute type;*
- c) *attribute value;*
- d) *attribute value assertion;*

- e) *operational attribute*;
- f) *user attribute*;
- g) *matching rule*.

3.5 Name definitions

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

- a) *alias, alias name*;
- b) *distinguished name*;
- c) *(directory) name*;
- d) *purported name*;
- e) *relative distinguished name*.

3.6 Distributed operations definitions

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

- a) *chaining*;
- b) *referral*.

3.7 Abstract Service definitions

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

3.7.1 filter: An assertion about the presence or value of certain attributes of an entry in order to limit the scope of a search.

3.7.2 originator: The user that originated an operation.

3.7.3 service controls: Parameters conveyed as part of an operation which constrain various aspects of its performance.

4 Abbreviations

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

AVA	Attribute Value Assertion
DIB	Directory Information Base
DIT	Directory Information Tree
DSA	Directory System Agent
DUA	Directory User Agent
DMD	Directory Management Domain
RDN	Relative Distinguished Name

5 Conventions

With minor exceptions this Directory Specification has been prepared according to the "Presentation of ITU-T | 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.511 | ISO/IEC 9594-3. The term "Directory Specifications" shall be taken to mean the X.500-Series 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 ITU-T X.500 Recommendations and the ISO/IEC 9594:1990 edition. Systems conforming to the current Directory Specifications are referred to as "1993 edition systems".

If the items in a list are numbered (as opposed to using “–” or letters), then the items shall be considered steps in a procedure.

This Directory Specification defines directory operations using the Remote Operation notation defined in ITU-T Rec. X.880 | ISO/IEC 9072-1.

6 Overview of the Directory Service

As described in ITU-T Rec. X.501 | ISO/IEC 9594-2, the services of the Directory are provided through access points to DUAs, each acting on behalf of a user. These concepts are depicted in Figure 1. Through an access point the Directory provides service to its users by means of a number of Directory operations.




Figure 1 – Access to the Directory

The Directory operations are of three different kinds:

- a) Directory Read operations, which interrogate a single Directory entry;
- b) Directory Search operations, which interrogate potentially several Directory entries; and
- c) Directory Modify operations.

The Directory Read operations, the Directory Search operations and the Directory Modify operations are specified in clauses 9, 10, and 11, respectively. Conformance to Directory operations is specified in ITU-T Rec. X.519 | ISO/IEC 9594-5.

7 Information types and common procedures

7.1 Introduction

This clause identifies, and in some cases defines, a number of information types which are subsequently used in the definition of Directory operations. The information types concerned are those which are common to more than one operation, are likely to be in the future, or which are sufficiently complex or self-contained as to merit being defined separately from the operation which uses them.

Several of the information types used in the definition of the Directory service are actually defined elsewhere. Subclause 7.2 identifies these types and indicates the source of their definition. Each of the remaining subclauses (7.3 through 7.11) identifies and defines an information type.

This clause also specifies some common elements of procedure that apply to most or all of the Directory operations.

7.2 Information types defined elsewhere

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

- a) **Attribute**;
- b) **AttributeType**;
- c) **AttributeValue**;
- d) **AttributeValueAssertion**;
- e) **DistinguishedName**;
- f) **Name**;
- g) **RelativeDistinguishedName**.

The following information type is defined in ITU-T Rec. X.520 | ISO/IEC 9594-6:

- **PresentationAddress**.

The following information types are defined in ITU-T Rec. X.509 | ISO/IEC 9594-8:

- a) **Certificate**;
- b) **SIGNED**;
- c) **CertificationPath**.

The following information types are defined in ITU-T Rec. X.880 | ISO/IEC 9072-1:

- **InvokeID**.

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

- a) **OperationProgress**;
- b) **ContinuationReference**.

7.3 Common arguments

The **CommonArguments** information may be present to qualify the invocation of each operation that the Directory can perform.

CommonArguments	::=	SET {
serviceControls	[30]	ServiceControls DEFAULT {},
securityParameters	[29]	SecurityParameters OPTIONAL,
requestor	[28]	DistinguishedName OPTIONAL,
operationProgress	[27]	OperationProgress
		DEFAULT { nameResolutionPhase notStarted },
aliasedRDNs	[26]	INTEGER OPTIONAL,
criticalExtensions	[25]	BIT STRING OPTIONAL,
referenceType	[24]	ReferenceType OPTIONAL,
entryOnly	[23]	BOOLEAN DEFAULT TRUE,
exclusions	[22]	Exclusions OPTIONAL,
nameResolveOnMaster	[21]	BOOLEAN DEFAULT FALSE }

The **ServiceControls** component is specified in 7.5. Its absence is deemed equivalent to there being an empty set of controls.

The **SecurityParameters** component is specified in 7.9. Its absence is deemed equivalent to there being an empty set of security parameters.

The **requestor** Distinguished Name identifies the originator of a particular operation. It holds the name of the user as identified at the time of binding to the Directory. It may be required when the request is to be signed (see 7.10), and shall hold the name of the user who initiated the request.

The **operationProgress**, **referenceType**, **entryOnly**, **exclusions**, and **nameResolveOnMaster** components are defined in ITU-T Rec. X.518 | ISO/IEC 9594-4. They are supplied by a DUA only when acting on a continuation reference returned by a DSA in response to an earlier operation, and their values are copied by the DUA from the continuation reference.

The **aliasedRDNs** component indicates to the DSA that the **object** component of the operation was created by the dereferencing of an alias on an earlier operation attempt. The integer value indicates the number of RDNs in the name that came from dereferencing the alias. (The value would have been set in the referral response of the previous operation.)

NOTE – This component is provided for compatibility with 1988 edition implementations of the Directory. DUAs (and DSAs) implemented according to later editions of the Directory specifications shall always omit this parameter from the **CommonArguments** of a subsequent request. In this way, the Directory will not signal an error if aliases dereference to further aliases.

7.3.1 Critical extensions

The **criticalExtensions** component provides a mechanism to list a set of extensions which are critical to the performance of a Directory operation. If the originator of the extended operation wishes to indicate that the operation shall be performed with one or more extensions (i.e. that performing the operation without these extensions is not acceptable), it does so by setting the **criticalExtensions** bit(s) which corresponds to the extension(s). If the Directory, or some part of it, is unable to perform a critical extension it returns an indication of **unavailableCriticalExtension** (as a **ServiceError** or **PartialOutcomeQualifier**). If the Directory is unable to perform an extension which is not critical, it ignores the presence of the extension.

This document defines a number of extensions which are available to 1993 edition implementations of the Directory. The extensions take such forms as additional numbered bits in a BIT STRING, or additional components of a SET or SEQUENCE, and are ignored by 1988 edition systems. Each such extension is assigned an integer *identifier*, which is the number of the bit which may be set in **criticalExtensions**. If the criticality of an extension is defined to be critical, the DUA shall set the corresponding bit in **criticalExtensions**. If the defined criticality is non-critical, the DUA may or may not set the corresponding bit in **criticalExtensions**.

The extensions, their identifiers, the operations in which they are permitted, the recommended criticality, and the clauses in which they are defined are shown in Table 1.

Table 1 – Extensions

Extension	Identifier	Operations	Criticality	Defined (subclauses)
subentries	1	all	non-critical	7.5
copyShallDo	2	Read, Compare, List, Search	non-critical	7.5
attribute size limit	3	Read, Search	non-critical	7.5
extraAttributes	4	Read, Search	non-critical	7.6
modifyRightsRequest	5	Read	non-critical	9.1
pagedResultsRequest	6	List, Search	non-critical	10.1
matchedValuesOnly	7	Search	non-critical	10.2
extendedFilter	8	Search	non-critical	10.2
targetSystem	9	AddEntry	critical	11.1
useAliasOnUpdate	10	AddEntry, RemoveEntry, ModifyEntry	critical	11.1
newSuperior	11	ModifyDN	critical	11.4

7.4 Common results

The **CommonResults** information should be present to qualify the result of each retrieval operation that the Directory can perform.

```

CommonResults ::= SET {
    securityParameters [30] SecurityParameters OPTIONAL,
    performer [29] DistinguishedName OPTIONAL,
    aliasDereferenced [28] BOOLEAN DEFAULT FALSE }

```

The **SecurityParameters** component is specified in 7.9. Its absence is deemed equivalent to there being an empty set of security parameters.

The **performer** Distinguished Name identifies the performer of a particular operation. It may be required when the result is to be signed (see 7.10) and shall hold the name of the DSA which signed the result.

The **aliasDereferenced** component is set to TRUE when the purported name of an object or base object which is the target of the operation included any aliases which were dereferenced.

7.5 Service controls

A **ServiceControls** parameter contains the controls, if any, that are to direct or constrain the provision of the service.

```

ServiceControls ::= SET {
    options [0] BIT STRING {
        preferChaining (0),
        chainingProhibited (1),
        localScope (2),
        dontUseCopy (3),
        dontDereferenceAliases (4),
        subentries (5),
        copyShallDo (6) } DEFAULT {},
    priority [1] INTEGER { low (0), medium (1), high (2) } DEFAULT medium,
    timeLimit [2] INTEGER OPTIONAL,
    sizeLimit [3] INTEGER OPTIONAL,
    scopeOfReferral [4] INTEGER { dmd(0), country(1) } OPTIONAL,
    attributeSizeLimit [5] INTEGER OPTIONAL }

```

The **options** component contains a number of indications, each of which, if set, asserts the condition suggested. Thus:

- a) **preferChaining** indicates that the preference is that chaining, rather than referrals, be used to provide the service. The Directory is not obliged to follow this preference.
- b) **chainingProhibited** indicates that chaining, and other methods of distributing the request around the Directory, are prohibited.
- c) **localScope** indicates that the operation is to be limited to a local scope. The definition of this option is itself a local matter, for example, within a single DSA or a single DMD.
- d) **dontUseCopy** indicates that copied information (as defined in ITU-T Rec. X.518 | ISO/IEC 9594-4 shall not be used to provide the service.
- e) **dontDereferenceAliases** indicates that any alias used to identify the entry affected by an operation is not to be dereferenced.

NOTE 1 – This is necessary to allow reference to an alias entry itself rather than the aliased entry, e.g., in order to read the alias entry.

- f) **subentries** indicates that a **Search** or **List** operation is to access subentries only; normal entries become inaccessible – i.e. the Directory behaves as though normal entries do not exist. If this service control is not set, then the operation accesses normal entries only and subentries become inaccessible. The service control is ignored for operations other than **Search** or **List**.

NOTE 2 – The effects of subentries on access control, schema, and collective attributes are still observed even if subentries are inaccessible.

NOTE 3 – If this service control is set, normal entries may still be specified as the base object of an operation.

- g) **copyShallDo** indicates that if the Directory is able to partly but not fully satisfy a query at a copy of an entry, it shall not chain the query. It is meaningful only if **dontUseCopy** is not set. If **copyShallDo** is not set, the Directory will use shadow data only if it is sufficiently complete to allow the operation to be fully satisfied at the copy. A query may be only partly satisfied because some of the requested attributes are missing in the shadow copy, or because the DSA holding the shadowed data does not support the requested matching rules on that data. If **copyShallDo** is set and the Directory is not able to fully satisfy a query, it shall set **incompleteEntry** in the the returned entry information.

If this component is omitted, the following are assumed: no preference for chaining but chaining not prohibited, no limit on the scope of the operation, use of copy permitted, aliases shall be dereferenced (except for modify operations for which alias dereferencing is not supported), subentries are not accessible, and operations not fully satisfiable with shadowed data are subject to further chaining.

The **priority** (low, medium, or high) at which the service is to be provided. Note that this is not a guaranteed service in that the Directory, as a whole, does not implement queuing. There is no relationship implied with the use of priorities in underlying layers.

The **timeLimit** indicates the maximum elapsed time, in seconds, within which the service shall be provided. If the constraint cannot be met, an error is reported. If this component is omitted, no time limit is implied. In the case of time limit exceeded on a List or Search, the result is an arbitrary selection of the accumulated results.

NOTE 4 – This component does not imply the length of time spent processing the request during the elapsed time: any number of DSAs may be involved in processing the request during the elapsed time.

The **sizeLimit** is only applicable to List and Search operations. It indicates the maximum number of objects to be returned. In the case of size limit exceeded, the results of List and Search may be an arbitrary selection of the accumulated results, equal in number to the size limit. Any further results shall be discarded.

The **scopeOfReferral** indicates the scope to which a referral returned by a DSA should be relevant. Depending on whether the values **dmd** or **country** are selected, only referrals to other DSAs within the selected scope shall be returned. This applies to the referrals in both a **Referral** error and the **unexplored** parameter of List and Search results.

The **attributeSizeLimit** indicates the largest size of any attribute (i.e. the type and all its values) that is included in returned entry information. If an attribute exceeds this limit, all of its values are omitted from the returned entry information and **incompleteEntry** is set in the returned entry information. The size of an attribute is taken to be its size in octets in the local concrete syntax of the DSA holding the data. Because of different ways applications store the data, the limit is imprecise. If this parameter is not specified, no limit is implied.

NOTE 5 – Attribute values returned as part of an entry's Distinguished Name are exempt from this limit.

Certain combinations of **priority**, **timeLimit**, and **sizeLimit** may result in conflicts. For example, a short time limit could conflict with low priority; a high size limit could conflict with a low time limit, etc.

7.6 Entry information selection

An **EntryInformationSelection** parameter indicates what information is being requested from an entry in a retrieval service.

```
EntryInformationSelection ::= SET {
  attributes CHOICE {
    allUserAttributes [0] NULL,
    select [1] SET OF AttributeType
    -- empty set implies no attributes are requested -- } DEFAULT allUserAttributes : NULL,
  infoTypes [2] INTEGER {
    attributeTypesOnly (0),
    attributeTypesAndValues (1) } DEFAULT attributeTypesAndValues,
  extraAttributes CHOICE {
    allOperationalAttributes [3] NULL,
    select [4] SET OF AttributeType } OPTIONAL }
```

The **attributes** component specifies the user and operational attributes about which information is requested:

- a) If the **select** option is chosen, then the attributes involved are listed. If the list is empty, then no attributes shall be returned. Information about a selected attribute shall be returned if the attribute is present. An **AttributeError** with the **noSuchAttributeOrValue** problem shall only be returned if none of the attributes selected is present.
- b) If the **allUserAttributes** option is selected, then information is requested about all user attributes in the entry.

Attribute information is only returned if access rights are sufficient. A **SecurityError** (with an **insufficientAccessRights** problem) shall only be returned in the case where access rights preclude the reading of all attribute values requested.

The **infoTypes** component specifies whether both attribute type and attribute value information (the default) or attribute type information only is requested. If the **attributes** component is such as to request no attributes, then this component is not meaningful.

The **extraAttributes** component specifies a set of additional user and operational attributes about which information is requested. If the **allOperationalAttributes** option is chosen, then information is requested about all directory operational attributes in the entry. If the **select** option is chosen, then information about the listed attributes is requested.

NOTE – This component may be used to request information about, for example, specific operational attributes when **attributes** is set to **allUserAttributes**, or about all operational attributes. If the same attribute is listed or implied in both **attributes** and **extraAttributes**, it is treated as though it has been requested only once.

A request for a particular attribute is always treated as a request for the attribute and all *subtypes* of that attribute (except for requests processed by 1988-edition systems).

In responding to a request for attribute information, the Directory treats all *collective attributes* of an entry as if they were actual user attributes of the entry, i.e. they are selected like other user attributes and are merged into the returned entry information. A request for **allUserAttributes** requests all collective attributes of the entry as well as ordinary attributes of the entry. An attribute is a collective attribute of an entry if all of the following are true:

- a) it is located in a subentry whose subtree specification includes the entry;
- b) it is not excluded by the presence in the entry of a **collectiveExclusions** attribute value equal to the collective attribute type; and
- c) it is permitted by the content rule for the structural object class for the entry.

7.7 Entry information

An **EntryInformation** parameter conveys selected information from an entry.

```
EntryInformation ::= SEQUENCE {
    name                Name,
    fromEntry           BOOLEAN DEFAULT TRUE,
    information          SET OF CHOICE {
        attributeType   AttributeType,
        attribute        Attribute } OPTIONAL,
    incompleteEntry     [3] BOOLEAN DEFAULT FALSE -- not in 1988-edition systems -- }
```

The **Name** parameter indicates the Distinguished Name of the entry or the name of an alias to the entry. The Distinguished Name of the entry is returned whenever permitted by the access control policy. If access is allowed to the attributes of the entry but not to its Distinguished Name, the Directory may return either an error or the name of a valid alias to the entry.

NOTES

1 If the entry was located using an alias, then that alias is known to be a valid alias. Otherwise, how it is ensured that the alias is valid is outside the scope of these Directory Specifications.

2 Where a particular component of the Directory has a choice of alias names available to it for return, it is recommended that where possible it choose the same alias name for repeated requests by the same requestor, in order to provide a consistent service.

The **fromEntry** parameter indicates whether the information was obtained from the entry (**TRUE**) or a copy of the entry (**FALSE**).

The **information** parameter is included if any attribute information from the entry is being returned, and contains a set of **attributeTypes** and **attributes**, as appropriate.

The **incompleteEntry** parameter is included and set to **TRUE** whenever the returned entry information is incomplete in relation to the user's request, e.g. because attributes or attribute values are omitted for reasons of access control (and their existence is permitted to be disclosed), the presence of incomplete shadow information together with **copyShallDo**, or because the **attributeSizeLimit** has been exceeded. It is not set to **TRUE** because an alias name has been returned instead of the Distinguished Name.

7.8 Filter

7.8.1 Filter

A **Filter** parameter applies a test that is either satisfied or not by a particular entry. The filter is expressed in terms of assertions about the presence or value of certain attributes of the entry, and is satisfied if and only if it evaluates to **TRUE**.

NOTE – A Filter may be **TRUE**, **FALSE**, or undefined.

Filter ::= CHOICE {		
item	[0]	FilterItem,
and	[1]	SET OF Filter,
or	[2]	SET OF Filter,
not	[3]	Filter }
FilterItem ::= CHOICE {		
equality	[0]	AttributeValueAssertion,
substrings	[1]	SEQUENCE {
type		ATTRIBUTE.&id(({SupportedAttributes})),
strings		SEQUENCE OF CHOICE {
initial	[0]	ATTRIBUTE.&Type
		((SupportedAttributes){@substrings.type}),
any	[1]	ATTRIBUTE.&Type
		((SupportedAttributes){@substrings.type}),
final	[2]	ATTRIBUTE.&Type
		((SupportedAttributes){@substrings.type}}),
greaterOrEqual	[2]	AttributeValueAssertion,
lessOrEqual	[3]	AttributeValueAssertion,
present	[4]	AttributeType,
approximateMatch	[5]	AttributeValueAssertion,
extensibleMatch	[6]	MatchingRuleAssertion }
MatchingRuleAssertion ::= SEQUENCE {		
matchingRule	[1]	SET SIZE (1..MAX) OF MATCHING-RULE.&id,
type	[2]	AttributeType OPTIONAL,
matchValue	[3]	MATCHING-RULE.&AssertionType (CONSTRAINED BY {
– <i>matchValue must be a value of type specified by the &AssertionType field of</i>		
– <i>one of the MATCHING-RULE information objects identified by matchingRule -- }),</i>		
dnAttributes	[4]	BOOLEAN DEFAULT FALSE }

A **Filter** is either a **FilterItem** (see 7.8.2), or an expression involving simpler filters composed together with the logical operators **and**, **or**, and **not**.

A **Filter** which is a **FilterItem** has the value of the **FilterItem** (i.e. TRUE, FALSE, or undefined).

A **Filter** which is the **and** of a set of filters is TRUE if the set is empty or if each filter is TRUE; it is FALSE if at least one filter is FALSE; otherwise it is undefined (i.e. if at least one filter is undefined and no filters are FALSE).

A **Filter** which is the **or** of a set of filters is FALSE if the set is empty or if each filter is FALSE; it is TRUE if at least one filter is TRUE; otherwise it is undefined (i.e. if at least one filter is undefined and no filters are TRUE).

A **Filter** which is the **not** of a filter is TRUE if the filter is FALSE; FALSE if it is TRUE; and undefined if it is undefined.

7.8.2 Filter item

A **FilterItem** is an assertion about the presence or value(s) of attributes in the entry under test. An assertion about a particular attribute type is also satisfied if the entry contains a subtype of the attribute and the assertion is TRUE for the subtype, or if there is a collective attribute of the entry (see 7.6) for which the assertion is TRUE. Each assertion is TRUE, FALSE, or undefined.

Every **FilterItem** includes or implies one or more **AttributeTypes** which identifies the particular attribute(s) concerned.

Any assertion about the values of such an attribute is only defined if the **AttributeType** is known by the evaluating mechanism, the purported **AttributeValue(s)** conforms to the attribute syntax defined for that attribute type, the implied or indicated matching rule is applicable to that attribute type, and (when used) a presented **matchValue** conforms to the syntax defined for the indicated matching rules.

NOTE 1 – Where these conditions are not met the **FilterItem** is undefined.

NOTE 2 – Access control restrictions may affect the evaluation of the **FilterItem**.

Attribute value assertions in filter items are evaluated using the matching rules defined for that attribute type. Matching rule assertions are evaluated as specified in their definition. A matching rule defined for a particular syntax can only be used to make assertions about attributes of that syntax or subtypes of that syntax.

A **FilterItem** may be undefined (as described above). Otherwise, where the **FilterItem** asserts:

- a) **equality** – It is TRUE if and only if there is a value of the attribute or one of its subtypes for which the **equality** matching rule applied to that value and the presented value returns TRUE.

- b) **substrings** – It is TRUE if and only if there is a value of the attribute or one of its subtypes for which the **substring** matching rule applied to that value and the presented value in **strings** returns TRUE. See ITU-T Rec. X.520 | ISO/IEC 9594-6 for a description of the semantics of the presented value.
- c) **greaterOrEqual** – It is TRUE if and only if there is a value of the attribute or one of its subtypes for which the **ordering** matching rule applied to that value and the presented value returns FALSE. I.E., there is a value of the attribute which is *greater than or equal to* the presented value.
- d) **lessOrEqual** – It is TRUE if and only if there is a value of the attribute or one of its subtypes for which either the **equality** matching rule *or* the **ordering** matching rule applied to that value and the presented value returns TRUE, i.e. there is a value of the attribute which is *less than or equal to* the presented value.
- e) **present** – It is TRUE if and only if the attribute or one of its subtypes is present in the entry.
- f) **approximateMatch** – It is TRUE if and only if there is a value of the attribute or one of its subtypes for which some locally-defined approximate matching algorithm (e.g. spelling variations, phonetic match, etc.) returns TRUE. There are no specific guidelines for approximate matching in this edition of this Directory Specification. If approximate matching is not supported, this **FilterItem** should be treated as a match for **equality**.
- g) **extensibleMatch** – It is TRUE if and only if there is a value of the attribute with the indicated **type** or one of its subtypes for which the matching rule specified in **matchingRule** applied to that value and the presented value **matchValue** returns TRUE.

If several matching rules are given, the way in which these rules are combined into a new rule is unspecified (it is a locally-defined algorithm, which reflects the semantics of the constituent matching rules – e.g. **phonetic + keyword** match).

If **type** is omitted, the match is made against all attribute types which are compatible with that matching rule. If **dnAttributes** is TRUE, the attributes of the Distinguished Name of the entry are used in addition to those of the entry in evaluating the match.

If an **extensibleMatch** is requested in a **filter** (rather than an **extendedFilter**), the **extendedFilter** bit in the **criticalExtensions** parameter in **CommonArguments** shall be set, indicating that the extension is critical.

NOTE 3 – An **extensibleMatch** is not permitted for 1988-edition systems.

7.9 Paged results

A **PagedResultsRequest** parameter is used by the DUA to request that the results of a list or search operation be returned to it “page-by-page”: It requests the DSA to return only a subset – a *page* – of the results of the operation, in particular the next **pageSize** subordinates or entries, and to return a **queryReference** which can be used to request the next set of results on a follow-up query. It shall not be used if results are to be signed, and is not supported by 1988-edition systems. Although a DUA may request **pagedResults**, a DSA is permitted to ignore the request and return its results in the normal manner.

```

PagedResultsRequest ::= CHOICE {
    newRequest          SEQUENCE {
        pageSize          INTEGER,
        sortKeys          SEQUENCE OF SortKey OPTIONAL,
        reverse            [1] BOOLEAN DEFAULT FALSE,
        unmerged          [2] BOOLEAN DEFAULT FALSE },
    queryReference      OCTET STRING }

SortKey ::= SEQUENCE {
    type                AttributeType,
    orderingRule        MATCHING-RULE.&id OPTIONAL }

```

For a new list or search operation, the **PagedResultsRequest** is set to **newRequest**, which consists of the following parameters:

- a) The **pageSize** parameter specifies the maximum number of subordinates or entries to return in the results. The DSA shall return up to but not more than the requested number of entries. The **sizeLimit**, if any, is ignored.
- b) The **sortKeys** parameter specifies a sequence of attribute types with optional ordering matching rules to use as sort keys for sorting the returned entries prior to return to the DUA. The entries are sorted according to their values of the **type** attribute of the first **SortKey** in the sequence, and in the event of multiple entries having the same sort position, of the next **SortKey** in the sequence, and so on.

For a particular **SortKey**, the DSA uses the **orderingRule** matching rule if it is present, otherwise the **ordering** matching rule of the attribute if one is defined; it ignores the sort key if none is defined. If the attribute type is multivalued, the “least” value is used; if the attribute type is missing from the returned results, it is regarded as “greater” than all other matched values. A DSA is permitted to support only certain sort key sequences (thus, a DSA that holds and returns its data in the internal order “alphabetic by surname” will be able to comply with only one sort key sequence). If it cannot support the requested sequence, it shall use a default sort sequence.

- c) If the **reverse** parameter is TRUE then the DSA will return the sorted results in reverse order (i.e. from “greatest” to “least” – if the attribute type is multivalued, the “greatest” is used; if the attribute type is missing from the returned results, it is regarded as “less” than all other matched values). If it is false, the DSA returns them in forward order. If no **sortKeys** parameter is specified, this parameter is ignored.
- d) If the **unmerged** parameter is TRUE and the DSA must merge results from a number of other DSAs, it shall return all the data from one DSA (in sort order) before returning data from the next DSA. If the parameter is false, the DSA shall collect the results from all other DSAs and sort the merged data before returning any of it. If no **sortKeys** parameter is specified, this parameter is ignored.

For a followup request, i.e. to request the next set of paged results, the DUA makes the same list or search request as before, but sets **PagedResultsRequest** to **queryReference**, with the value of this parameter the same as that returned in the **PartialOutcomeQualifier** of the previous results. The DUA has no understanding of the **queryReference**, which is available to a DSA to use as it wishes to record context information for the query. The DSA uses this information to determine which results to return next.

NOTES

- 1 If the DIB changes between search requests, the DUA may not see the effects of these changes. This is implementation dependent.
- 2 A query-reference may remain valid even if a DUA begins a new list or search operation. A DUA may request paged results with several queries and then return to an earlier query and request the next page of results using the query-reference supplied for it. The number of “active” query references to which a DUA can return is a local DSA implementation option, as is the lifetime of those query-references.
- 3 Paged results are not supported in the Directory System Protocol. Paged results are provided entirely by the DSA to which the DUA has connected.

7.10 Security parameters

The **SecurityParameters** govern the operation of various security features associated with a Directory operation.

NOTE – These parameters are conveyed from sender to recipient. Where the parameters appear in the argument of an operation the requestor is the sender, and the performer is the recipient. In a result, the roles are reversed.

SecurityParameters ::= SET {
 certification-path [0] **CertificationPath** OPTIONAL,
 name [1] **DistinguishedName** OPTIONAL,
 time [2] **UTCTime** OPTIONAL,
 random [3] **BIT STRING** OPTIONAL,
 target [4] **ProtectionRequest** OPTIONAL }

ProtectionRequest ::= INTEGER { none(0), signed (1) }

The **CertificationPath** component consists of the sender's certificate, and, optionally, a sequence of certificate pairs. The certificate is used to associate the sender's public key and distinguished name, and may be used to verify the signature on the argument or result. This parameter shall be present if the argument or result is signed. The sequence of certification pairs consists of certification authority cross certificates. It is used to enable the sender's certificate to be validated. It is not required if the recipient shares the same certification authority as the sender. If the recipient requires a valid set of certificate pairs, and this parameter is not present, whether the recipient rejects the signature on the argument or result, or attempts to generate the certification path, is a local matter.

The **name** is the distinguished name of the first intended recipient of the argument or result. For example, if a DUA generates a signed argument, the name is the distinguished name of the DSA to which the operation is submitted.

The **time** is the intended expiry time for the validity of the signature, when signed arguments are used. It is used in conjunction with the random number to enable the detection of replay attacks.

The **random** number is a number which should be different for each unexpired token. It is used in conjunction with the time parameter to enable the detection of replay attacks when the argument or result has been signed.

The **target ProtectionRequest** may appear only in the request for an operation to be carried out, and indicates the requestor's preference regarding the degree of protection to be provided to the result. Two levels are provided: **none** (no protection requested, the default), and **signed** (the Directory is requested to sign the result). The degree of protection actually provided to the result is indicated by the form of result and may be equal to or lower than that requested, based on the limitations of the Directory.

7.11 Common elements of procedure for basic-access-control

This subclause defines the elements of procedure that are common to all abstract service operations when **basic-access-control** is in effect.

7.11.1 Alias dereferencing

If, in the process of locating a target object entry (identified in the argument of an abstract service operation), alias dereferencing is required, no specific permissions are necessary for alias dereferencing to take place. However, if alias dereferencing would result in a **ContinuationReference** being returned (i.e. in a **Referral**), the following sequence of access controls applies. These access controls shall also be applied to a referral that is received in a response from another DSA. That is, the DSA shall police all referrals whether they were generated locally or remotely.

- 1) *Read* permission is required to the alias entry. If permission is not granted, the operation fails in accordance to the procedure described in 7.11.3.
- 2) *Read* permission is required to the **AliasedObjectName** attribute and to the single value that it contains. If permission is not granted the operation fails and the error **NameError** with problem **aliasDereferencingProblem** shall be returned. The **matched** element shall contain the name of the alias entry.

NOTE – In addition to the access controls described above, security policy may prevent the disclosure of knowledge information which would otherwise be conveyed as a **ContinuationReference** in **Referral**. If such a policy is in effect and if a DUA constrains the service by specifying **chainingProhibited** the Directory may return a **ServiceError** with problem **chainingRequired**. Otherwise, a **SecurityError** with problem **insufficientAccessRights** or **noInformation** shall be returned.

7.11.2 Return of NameError

If, while performing an abstract service operation, the specified target object (alias or entry) – e.g. the Name of an entry to be read or the **baseObject** in a **Search** – could not be found, a **NameError** with problem **noSuchObject** shall be returned. The **matched** element shall either contain the name of the next superior entry to which *DiscloseOnError* permission is granted, or the name of the DIT root (i.e. an empty **RDNSequence**).

NOTE – The second alternative may be taken by a DSA which does not have access to all superior entries.

7.11.3 Non-disclosure of the existence of an entry

If, while performing an abstract service operation, the necessary entry level permission is not granted to the specified target object entry – e.g. the entry to be read – the operation fails and the error returned is one of: if *DiscloseOnError* permission is granted to the target entry, a **SecurityError** with problem **insufficientAccessRights** or **noInformation** shall be returned; otherwise, a **NameError** with problem **noSuchObject** shall be returned. The **matched** element shall either contain the name of the next superior entry to which *DiscloseOnError* permission is granted, or the name of the DIT root (i.e. an empty **RDNSequence**).

NOTE – The second alternative may be taken by a DSA which does not have access to all superior entries.

Additionally, whenever the Directory detects an operational error (including a **Referral**) it shall ensure that in returning that error it does not compromise the existence of the named target entry and any of its superiors. For example, before returning a **ServiceError** with problem **timeLimitExceeded** or an **UpdateError** with problem **notAllowedOnNonLeaf**, the Directory verifies that *discloseOnError* permission is granted to the target entry. If it is not, the procedure described in the paragraph above shall be followed.

7.11.4 Return of Distinguished Name

In a Compare, List, or Search operation, *ReturnDN* permission is required to the **object** (or **baseObject**) entry if as a result of dereferencing an alias, the object's distinguished name is to be returned in the **name** parameter of the operation result (see 9.2.3). If this permission is not granted, the Directory shall return an alias name for the entry instead, as described in 7.7, or shall omit the name parameter altogether.

In a Read or Search operation, *ReturnDN* permission is required to an entry in order to return its distinguished name in **EntryInformation**. If this permission is not granted, the Directory shall return the name of an alias instead, as described in 7.7, or if no alias name is available shall fail the operation with a **NameError** (in the case of Read) or omit the entry from the results (in the case of Search).

If the user supplied alias name is returned in the result, then the **aliasDeferenced** flag of **CommonResults** shall not be set to **TRUE**.

7.12 Optionally-signed parameters

An **OPTIONALLY-SIGNED** information type is one whose values may, at the option of the generator, be accompanied by their digital signature. This capability is specified by means of the following type:

```

OPTIONALLY-SIGNED {Type} ::= CHOICE {
    unsigned      Type,
    signed        SIGNED {Type}}

```

The **SIGNED** type, which describes the form of the signed form of the information, is specified in ITU-T Rec. X.509 | ISO/IEC 9594-8.

8 Bind and Unbind operations

The **DirectoryBind** and **DirectoryUnbind** operations, defined in 8.1 and 8.2 respectively, are used by the DUA at the beginning and end of a particular period of accessing the Directory.

8.1 Directory Bind

8.1.1 Directory Bind syntax

A **DirectoryBind** operation is used at the beginning of a period of accessing the Directory.

directoryBind	OPERATION ::=	{
ARGUMENT		DirectoryBindArgument
RESULT		DirectoryBindResult
ERRORS		{directoryBindError }
DirectoryBindArgument	::=	SET {
credentials	[0]	Credentials OPTIONAL,
versions	[1]	Versions DEFAULT {v1}}
Credentials	::=	CHOICE {
simple	[0]	SimpleCredentials,
strong	[1]	StrongCredentials,
externalProcedure	[2]	EXTERNAL }
SimpleCredentials	::=	SEQUENCE {
name	[0]	DistinguishedName,
validity	[1]	SET {
time1	[0]	UTCTime OPTIONAL,
time2	[1]	UTCTime OPTIONAL,
random1	[2]	BIT STRING OPTIONAL,
random2	[3]	BIT STRING OPTIONAL} OPTIONAL,
password	[2]	CHOICE {
unprotected		OCTET STRING,
protected		SIGNATURE {OCTET STRING} } OPTIONAL}
StrongCredentials	::=	SET {
certification-path	[0]	CertificationPath OPTIONAL,
bind-token	[1]	Token,
name	[2]	DistinguishedName OPTIONAL }
Token	::=	SIGNED { SEQUENCE {
algorithm	[0]	AlgorithmIdentifier,
name	[1]	DistinguishedName,
time	[2]	UTCTime,
random	[3]	BIT STRING }
Versions	::=	BIT STRING {v1(0)}
DirectoryBindResult	::=	DirectoryBindArgument