
**Information technology — Open Systems
Interconnection — Security frameworks for
open systems: Access control framework**

*Technologies de l'information — Interconnexion de systèmes
ouverts (OSI) — Cadres généraux pour la sécurité des systèmes ouverts:
Cadre général de contrôle d'accès*

Contents

	<i>Page</i>
1 Scope	1
2 Normative references	2
2.1 Identical Recommendations International Standards	2
2.2 Paired Recommendations International Standards equivalent in technical content	2
3 Definitions	2
4 Abbreviations	4
5 General discussion of access control	4
5.1 Goal of access control	4
5.2 Basic aspects of access control	5
5.2.1 Performing access control functions	5
5.2.2 Other access control activities	7
5.2.3 ACI forwarding	8
5.3 Distribution of access control components	9
5.3.1 Incoming access control	10
5.3.2 Outgoing access control	10
5.3.3 Interposed access control	10
5.4 Distribution of access control components across multiple security domains	10
5.5 Threats to access control	10
6 Access control policies	11
6.1 Access control policy expression	11
6.1.1 Access control policy categories	11
6.1.2 Groups and roles	11
6.1.3 Security labels	11
6.1.4 Multiple initiator access control policies	12
6.2 Policy management	12
6.2.1 Fixed policies	12
6.2.2 Administratively-imposed policies	12
6.2.3 User-selected policies	12
6.3 Granularity and containment	12
6.4 Inheritance rules	12
6.5 Precedence among access control policy rules	13
6.6 Default access control policy rules	13
6.7 Policy mapping through cooperating security domains	13
7 Access control information and facilities	13
7.1 ACI	13
7.1.1 Initiator ACI	14

© ISO/IEC 1996

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the publisher.

ISO/IEC Copyright Office • Case postale 56 • CH-1211 Genève 20 • Switzerland

Printed in Switzerland

7.1.2	Target ACI	14
7.1.3	Access request ACI	14
7.1.4	Operand ACI	14
7.1.5	Contextual information	14
7.1.6	Initiator-bound ACI	15
7.1.7	Target-bound ACI	15
7.1.8	Access request-bound ACI	15
7.2	Protection of ACI	15
7.2.1	Access control certificates	15
7.2.2	Access control tokens	16
7.3	Access control facilities	16
7.3.1	Management related facilities	16
7.3.2	Operation related facilities	17
8	Classification of access control mechanisms	19
8.1	Introduction	19
8.2	ACL scheme	20
8.2.1	Basic features	20
8.2.2	ACI	20
8.2.3	Supporting mechanisms	20
8.2.4	Variations of this scheme	21
8.3	Capability scheme	22
8.3.1	Basic features	22
8.3.2	ACI	22
8.3.3	Supporting mechanisms	22
8.3.4	Variation of this scheme – Capabilities without specific operations	22
8.4	Label based scheme	23
8.4.1	Basic features	23
8.4.2	ACI	23
8.4.3	Supporting mechanisms	23
8.4.4	Labeled channels as targets	24
8.5	Context based scheme	24
8.5.1	Basic features	24
8.5.2	ACI	25
8.5.3	Supporting mechanisms	25
8.5.4	Variations of this scheme	25
9	Interaction with other security services and mechanisms	25
9.1	Authentication	25
9.2	Data integrity	25
9.3	Data confidentiality	26
9.4	Audit	26
9.5	Other access-related services	26
Annex A	Exchange of access control certificates among components	27
A.1	Introduction	27
A.2	Forwarding access control certificates	27
A.3	Forwarding multiple access control certificates	27
A.3.1	Example	27
A.3.2	Generalization	28
A.3.3	Simplifications	28
Annex B	Access control in the OSI reference model	29
B.1	General	29
B.2	Use of access control within the OSI layers	29
B.2.1	Use of access control at the network layer	29
B.2.2	Use of access control at the transport layer	29
B.2.3	Use of access control at the application layer	29
Annex C	Non-uniqueness of access control identities	30

Annex D – Distribution of access control components	31
D.1 Aspects considered.....	31
D.2 AEC and ADC locations	31
D.3 Interactions among access control components	32
Annex E – Rule-based versus identity-based policies.....	34
Annex F – A mechanism to support ACI forwarding through an initiator.....	35
Annex G – Access control security service outline.....	36

IECNORM.COM : Click to view the full PDF of ISO/IEC 10181-3:1996

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 10181-3 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.812.

ISO/IEC 10181 consists of the following parts, under the general title *Information technology — Open Systems Interconnection — Security frameworks for open systems*:

- *Part 1: Overview*
- *Part 2: Authentication framework*
- *Part 3: Access control framework*
- *Part 4: Non-repudiation framework*
- *Part 5: Confidentiality framework*
- *Part 6: Integrity framework*
- *Part 7: Security audit framework*

Annexes A to G of this part of ISO/IEC 10181 are for information only.

Introduction

This Recommendation | International Standard defines a general framework for the provision of access control. The primary goal of access control is to counter the threat of unauthorized operations involving a computer or communications system; these threats are frequently subdivided into classes known as unauthorized use, disclosure, modification, destruction and denial of service.

IECNORM.COM : Click to view the full PDF of ISO/IEC 10181-3:1996

INTERNATIONAL STANDARD

ITU-T RECOMMENDATION

**INFORMATION TECHNOLOGY –
OPEN SYSTEMS INTERCONNECTION –
SECURITY FRAMEWORKS FOR OPEN SYSTEMS:
ACCESS CONTROL FRAMEWORK**

1 Scope

The Security Frameworks are intended to address the application of security services in an Open Systems environment, where the term *Open Systems* is taken to include areas such as Database, Distributed Applications, ODP and OSI. The Security Frameworks are concerned with defining the means of providing protection for systems and objects within systems, and with the interactions between systems. The Security Frameworks are not concerned with the methodology for constructing systems or mechanisms.

The Security Frameworks address both data elements and sequences of operations (but not protocol elements) that are used to obtain specific security services. These security services may apply to the communicating entities of systems as well as to data exchanged between systems, and to data managed by systems.

In the case of Access Control, accesses may either be *to* a system (i.e. to an entity that is the communicating part of a system) or *within* a system. The information items that need to be presented to obtain the access, as well as the sequence of operations to request the access and for notification of the results of the access, are considered to be within the scope of the Security Frameworks. However, any information items and operations that are dependent solely on a particular application and that are strictly concerned with local access within a system are considered to be outside the scope of the Security Frameworks.

Many applications have requirements for security to protect against threats to resources, including information, resulting from the interconnection of Open Systems. Some commonly known threats, together with the security services and mechanisms that can be used to protect against them, in an OSI environment, are described in CCITT Rec. X.800 | ISO 7498-2.

The process of determining which uses of resources within an Open System environment are permitted and, where appropriate, preventing unauthorized access is called access control. This Recommendation | International Standard defines a general framework for the provision of access control services.

This Security Framework:

- a) defines the basic concepts for access control;
- b) demonstrates the manner in which the basic concepts of access control can be specialized to support some commonly recognized access control services and mechanisms;
- c) defines these services and corresponding access control mechanisms;
- d) identifies functional requirements for protocols to support these access control services and mechanisms;
- e) identifies management requirements to support these access control services and mechanisms;
- f) addresses the interaction of access control services and mechanisms with other security services and mechanisms.

As with other security services, access control can be provided only within the context of a defined security policy for a particular application. The definition of access control policies is outside the scope of this Recommendation | International Standard, however, some characteristics of access control policies are discussed.

It is not a matter for this Recommendation | International Standard to specify details of the protocol exchanges which may need to be performed in order to provide access control services.

This Recommendation | International Standard does not specify particular mechanisms to support these access control services nor the details of security management services and protocols.

A number of different types of standard can use this framework including:

- a) standards that incorporate the concept of access control;
- b) standards that specify abstract services that include access control;
- c) standards that specify uses of an access control service;
- d) standards that specify the means of providing access control within an Open System environment; and
- e) standards that specify access control mechanisms.

Such standards can use this framework as follows:

- standard types a, b, c, d, and e can use the terminology of this framework;
- standard types b, c, d, and e can use the facilities defined in clause 7 of this framework; and
- standard type e can be based upon the classes of mechanism defined in clause 8.

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

- ITU-T Recommendation X.200 (1994) | ISO/IEC 7498-1:1994, *Information technology – Open Systems Interconnection – Basic Reference Model: The Basic Model*.
- ITU-T Recommendation X.810 (1995) | ISO/IEC 10181-1:1996, *Information technology – Open Systems Interconnection – Security frameworks for open systems: Overview*.
- ITU-T Recommendation X.811(1995) | ISO/IEC 10181-2:1996, *Information technology – Open Systems Interconnection – Security frameworks for open systems: Authentication framework*.
- ITU-T Recommendation X.880 (1994) | ISO/IEC 13712-1: 1995, *Information technology – Remote Operations: Concepts model and notation*.

2.2 Paired Recommendations | International Standards equivalent in technical content

- CCITT Recommendation X.800 (1991), *Security Architecture for Open Systems Interconnection for CCITT applications*.

ISO 7498-2:1989, *Information processing systems – Open Systems Interconnection – Basic Reference Model – Part 2: Security Architecture*.

3 Definitions

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

3.1 This Recommendation | International Standard makes use of the following terms defined in CCITT Rec. X.800 | ISO 7498-2:

- a) access control;
- b) access control list;
- c) accountability;
- d) authentication;
- e) authentication information;
- f) authorization;

- g) capability;
- h) identity-based security policy;
- i) rule-based security policy;
- j) security audit;
- k) security label;
- l) security policy;
- m) security service;
- n) sensitivity.

3.2 This Recommendation | International Standard makes use of the following terms defined in ITU-T Rec. X.810 | ISO/IEC 10181-1:

- a) secure interaction policy;
- b) security certificate;
- c) security domain;
- d) security domain authority;
- e) security information;
- f) security policy rules;
- g) security token;
- h) trust.

3.3 This Recommendation | International Standard makes use of the following term defined in ITU-T Rec. X.200 | ISO/IEC 7498-1:

- real system.

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

3.4.1 **access control certificate:** A security certificate that contains ACI.

3.4.2 **Access Control Decision Information (ADI):** The portion (possibly all) of the ACI made available to the ADF in making a particular access control decision.

3.4.3 **Access Control Decision Function (ADF):** A specialized function that makes access control decisions by applying access control policy rules to an access request, ADI (of initiators, targets, access requests, or that retained from prior decisions), and the context in which the access request is made.

3.4.4 **Access Control Enforcement Function (AEF):** A specialized function that is part of the access path between an initiator and a target on each access request and enforces the decision made by the ADF.

3.4.5 **Access Control Information (ACI):** Any information used for access control purposes, including contextual information.

3.4.6 **access control policy:** The set of rules that define the conditions under which an access may take place.

3.4.7 **access control policy rules:** Security policy rules concerning the provision of the access control service.

3.4.8 **access control token:** A security token that contains ACI.

3.4.9 **access request:** The operations and operands that form part of an attempted access.

3.4.10 **access request access control decision information; access request ADI:** ADI derived from access request-bound ACI.

3.4.11 **access request access control information; access request ACI:** ACI about an access request.

3.4.12 **access request-bound access control information; access request-bound ACI:** ACI bound to an access request.

3.4.13 **clearance:** Initiator-bound ACI that can be compared with security labels of targets.

- 3.4.14 contextual information:** Information about or derived from the context in which an access request is made (e.g. time of day).
- 3.4.15 initiator:** An entity (e.g. human user or computer-based entity) that attempts to access other entities.
- 3.4.16 initiator access control decision information; initiator ADI:** ADI derived from initiator-bound ACI.
- 3.4.17 initiator access control information; initiator ACI:** ACI about an initiator.
- 3.4.18 initiator-bound access control information; initiator-bound ACI:** ACI bound to an initiator.
- 3.4.19 operand access control decision information; operand ADI:** ADI derived from operand-bound ACI.
- 3.4.20 operand access control information; operand ACI:** ACI about the operands of an access request.
- 3.4.21 operand-bound access control information; operand-bound ACI:** ACI bound to the operands of an access request.
- 3.4.22 retained ADI:** ADI which has been retained by an ADF from earlier access control decisions for use in future access control decisions.
- 3.4.23 target:** An entity to which access may be attempted.
- 3.4.24 target access control decision information; target ADI:** ADI derived from target-bound ACI.
- 3.4.25 target access control information; target ACI:** ACI about a target.
- 3.4.26 target-bound access control information; target-bound ACI:** ACI bound to a target.

4 Abbreviations

ACI	Access Control Information
ADI	Access Control Decision Information
ADF	Access Control Decision Function
AEF	Access Control Enforcement Function
SI	Security Information
SDA	Security Domain Authority

5 General discussion of access control

5.1 Goal of access control

For the purpose of this Security Framework, the primary goal of access control is to counter the threat of unauthorized operations involving a computer or communications system; these threats are frequently subdivided into classes known as:

- unauthorized use;
- disclosure;
- modification;
- destruction; and
- denial of service.

The subgoals of this Security Framework are:

- control of access by processes (which may be acting on behalf of humans or other processes) to data, different processes or other computing resources;
- control of access within a security domain or across one or more security domains;
- control of access according to its context; for example, dependent on such factors as time of attempted access, location of the accessor or route of access;
- control of access that is reactive to changes in authorization during access.

5.2 Basic aspects of access control

The following subclauses describe abstract access control functions largely independent of access control policies and system designs. Access control in real systems is concerned with many types of entities, such as:

- physical entities (e.g. real systems);
- logical entities (e.g. OSI layer entities, files, organizations, and enterprises);
- human users.

Access control in real systems can require a complex set of activities. The activities are:

- establishing access control policy representation;
- establishing ACI representations;
- allocating ACI to elements (initiators, targets or access requests);
- binding ACI to elements;
- making ADI available to the ADF;
- performing access control functions;
- modification of ACI (any time after allocating ACI values; includes revocation);
- revocation of ADI.

These activities can be divided into two groups:

- the operational activities (making ADI available to the ADF and performing access control functions);
- and
- the management activities (all the remaining activities).

Some of the activities above may be grouped as a single identifiable activity within a real system. Although some access control activities necessarily precede others, there is often overlap among them and some activities may be performed repeatedly.

A detailed discussion of the notions involved in performing access control functions will be presented first since all the other activities support this one.

5.2.1 Performing access control functions

For the purpose of this subclause, the functions which are fundamental to access control are illustrated in Figures 5-1 and 5-2. Other functions may be necessary for the overall operation of access control. In later discussions, a variety of ways in which these functions may be implemented are presented, including different ways of distributing the access control functions and ACI, and different styles of communication among access control functions in the same or cooperating security domains.




Figure 5-1 – Illustration of fundamental access control functions




Figure 5-2 – Illustration of ADF

The basic entities and functions involved in access control are the initiator, the Access Control Enforcement function (AEF), the Access Control Decision Function (ADF), and the target.

Initiators represent both the human beings and computer-based entities that access or attempt to access targets. Within a real system, an initiator is represented by a computer-based entity, although the access requests of the computer-based entity on behalf of the initiator may be further limited by the ACI of the computer based-entity.

Targets represent computer-based or communications entities to which access is attempted or that are accessed by initiators. A target may be, for example, an OSI layer entity, a file, or a real system.

An access request represents the operations and operands that form part of an attempted access.

The AEF ensures that only allowable accesses, as determined by the ADF, are performed by the initiator on the target. When the initiator makes a request to perform a particular access on the target, the AEF informs the ADF that a decision is required so that a determination can be made.

In order to perform this decision, the ADF is provided with the access request (as part of the decision request) and the following types of Access Control Decision Information (ADI):

- initiator ADI (ADI derived from the ACI bound to the initiator);
- target ADI (ADI derived from the ACI bound to the target);
- access request ADI (ADI derived from the ACI bound to the access request).

The other inputs to the ADF are the access control policy rules (from the ADF's security domain authority), and any contextual information needed to interpret the ADI or policy. Examples of contextual information include the location of the initiator, the time of access, or the particular communications path in use.

Based on these inputs, and possibly from ADI retained from prior decisions, the ADF arrives at a decision to allow or deny the initiator's attempted access to the target. The decision is conveyed to the AEF which then either allows the access request to pass to the target or takes other appropriate actions.

In many situations, successive access requests by an initiator on a target are related. A typical example is in an application that opens a connection to a peer target application process and then attempts to perform several accesses using the same (retained) ADI. For some succeeding access requests communicated over the connection, additional ADI may need to be provided to the ADF for it to allow the access request. In other situations, a security policy may demand that certain related access requests between one or more initiators and one or more targets are subject to restrictions. In such cases, the ADF may use retained ADI from prior decisions involving multiple initiators and targets to make the decision on a particular access request.

For the purposes of this subclause, an access request involves a single interaction by an initiator with a target, if permitted by the AEF. Although some access requests between an initiator and a target are purely independent of others, it is often the case that two entities enter into a set of related access requests such as a query-response paradigm. In such cases, the initiator and target roles are assumed by the entities as necessary, either simultaneously or alternately, and the access control functions are carried out for each access request, possibly by separate AEF components, ADF components and access control policies.

5.2.2 Other access control activities

5.2.2.1 Establishing access control policy representations

Access control policies commonly are stated in natural languages as broad principles; for example: only managers of at least a certain rank are allowed to examine salary information of employees. The conversion of these principles into rules is an engineering design activity that necessarily precedes the other access control activities and is not within the scope of this Security Framework. An overview of access control policy concepts is given in clause 6.

5.2.2.2 Establishing ACI representations

In this activity, choices are made for the representation of ACI within real systems (data structures) and for exchange between real systems (syntaxes). A broad range of possible representations is discussed in this Security Framework. ACI representations must be able to support the requirements of specific access control policies. Some ACI representations may be appropriate both for use within and between real systems. Different ACI representations may be used for different purposes and among particular elements.

The chosen ACI representations can be considered as templates for the assignment of particular ACI values for elements in a security domain (as discussed in the next subclause). One aspect of establishing ACI representations is the determination of the types and ranges of the ACI values that may be assigned to elements in a security domain (but not which types can be assigned to specific elements).

Representation of ACI exchanged between real systems for access control management purposes or for ACI exchanges among entities and access control functions are candidates for OSI standardization. How ACI is represented within real systems or for presentation to a local ADF are not matters for OSI standardization. Protecting the exchange of ACI is discussed in 7.2. For OSI applications (and, possibly, others), it is appropriate to consider ACI representations as attributes which consist of attribute type-attribute value pairs.

5.2.2.3 Allocating ACI to initiators and targets

In this activity, the specific attribute types and attribute values of ACI assigned to an element are designated by an SDA, its agents, or other entities (e.g. resource owners). These entities may specify or modify ACI allocations in accordance with the security domain policy. ACI allocated by an entity may be limited by the ACI that has been bound to it by another entity. The allocation of ACI to elements is a continuing activity as new elements are added to a security domain.

NOTE – The administrative act of granting “access rights” is sometimes referred to as authorization. This meaning is included in the allocation of ACI to initiators or targets.

ACI can be either information about a single entity or information about a relationship among entities. The ACI allocated to an initiator may be purely about that initiator, or it may be about relationships between that initiator and particular targets, or about relationships between that initiator and possible contexts. Thus, the ACI allocated to an initiator may include initiator ACI, target ACI, or contextual information. Similarly, ACI allocated to a target may include target ACI, initiator ACI (about one or more initiators), or contextual information.

In actual operation, ACI must be bound to an element (see 5.2.2.4) so that an ADF that uses ADI derived from the bound ACI has confidence in that information. Thus, although the allocation of ACI to elements is a prerequisite to constructing bound ACI, only ACI that is bound to an element is actually present in real open systems.

5.2.2.4 Binding ACI to initiators, targets and access requests

Binding ACI to an element (i.e. an initiator, target or access request) creates a secure linkage between an element and the ACI allocated to that element. The binding provides assurance to access control functions and other elements both that the ACI is indeed assigned to the particular element and that no modification has occurred since the binding was made. Binding is achieved by using an integrity service. Several binding mechanisms are possible, including some that are dependent on the location of the element and the ACI, while others may depend on some cryptographic signing or sealing process. The integrity of the binding of ACI to elements needs to be protected within initiator and target systems (e.g. by relying upon operating system functions such as file protection and process separation) and, also, in the exchange of ACI. Since there may be several possible representations of an element's ACI (both within systems and between systems), different binding mechanisms may be used for the same ACI. Under some security policies, the confidentiality of ACI also needs to be maintained.

The binding of ACI to elements is a continuing activity as new elements are added to a security domain. An SDA, its agents, or other allowed entities, may delete or add ACI bindings at will in accordance with the applicable security policy. An SDA may modify the ACI bound to an element as needed to express changing security policy or attributes. The bound ACI may include validity period indicators, thereby minimizing the ACI that may later need to be revoked.

The time at which ACI is bound to an element and the entity that causes the binding mechanism to be invoked depends on the type of element. Initiators will have ACI bound to them by an SDA or its agents by the time they are able to make accesses.

All targets will have ACI bound to them by an SDA or its agents by the time they become accessible. Targets that are created by an application on behalf of a user or another application will have their ACI bound to them at the time of creation or after their creation. The ACI bound to such targets may be restricted by limitations in the ACI bound to the user or the application.

ACI is bound to an access request by a user or application, or by an SDA or its agent on behalf of the user or application, before the access is attempted. Again, the ACI bound to the access request may be restricted by limitations in the ACI bound to the user or application. It is often the case that an access request causes a new target entity to be created (e.g. when a file is transferred between systems). Such a target's ACI may be specified in (or derived from) ACI bound to the access request.

5.2.2.5 Making ADI available to the ADF

If allowed by the access control policy and if the binding mechanism in use permits, a subset of the ACI bound to an initiator or target may be selected by the initiator or target for use at the ADF in making particular access control decisions. The ACI bound to one element may be temporarily bound to another element, for example, when one entity acts on behalf of another entity.

In order to perform its functions, the various ADI of Figure 5-2 must be made available to the ADF. Note that no assumptions are made in this subclause about the physical distribution of entities, functions, or ADI, nor how any of the inputs are made available to the ADF. Some possible relationships among entities and distributed access control components are discussed in 5.3, 5.4 and Annex D.

Three possibilities exist for initiator ADI, target ADI or access request ADI:

- a) ADI may be pre-placed at one or more ADF components after allocation of ACI values;
- b) ADI may be derived from bound ACI delivered to the ADF components during the access control process (possibly in conjunction with the attempted access);
- c) ADI may be derived from bound ACI obtained from other sources (e.g. a Directory Service Agent). Either the initiator or target obtains the bound ACI [which for the ADF is indistinguishable from b)] or the ADF obtains the bound ACI as needed [which for the initiator or target is indistinguishable from a)].

The means through which the ADF obtains the bound ACI and derives this ADI is not specified. Initiator-bound ACI is not necessarily delivered by the initiator, target-bound ACI is not necessarily delivered by the target, nor is access request-bound ACI necessarily delivered with an access request.

The ADF must be able to determine unequivocally that the ADI has been derived from ACI bound to elements by an appropriate SDA. Means to provide this assurance are discussed in 7.2.

5.2.2.6 Modification of ACI

The SDA may modify the ACI allocated and bound to an element as needed to express changing security attributes. ACI may be modified at any time following its allocation to elements. If the modification reduces the accesses that are permissible by an initiator on targets, then this change may require the revocation of the ACI and of ADI derived from it that may be retained by ADFs.

5.2.2.7 Revocation of ADI

After ACI is revoked, any attempted use of ADI derived from that ACI must result in an access being disallowed. Further use of ADI derived from that ACI before it was revoked should be prevented, or its attempted use must result in an access being denied. If an access based on such previously derived ADI is continuing when the ACI is revoked, the access control policy in effect may require that the access be terminated.

5.2.3 ACI forwarding

It is a common requirement in distributed systems for entities to request other entities to perform accesses for them on their behalf. Initiators and targets are roles assumed by entities, although not all entities may assume both roles. An entity may simultaneously assume the initiator role relative to one entity while being a target itself relative to another entity acting as initiator.

Figure 5-3 demonstrates the basic notion of an entity, A, requesting another, B, to perform an access upon yet another entity, C. The several access control components that might be involved in such a chained access are not shown in Figure 5-3.




Figure 5-3 – ACI forwarding

There are any number of variations of this basic notion. The variations are visibly different in the combinations of ACI required by policy that must be present to allow such chained accesses to proceed and how that ACI is made available to the appropriate access control components. Under some policies, B may need no ACI beyond that already bound to it to carry out the access for A; under some other policies B will use only ACI obtained from A that are relevant to the access, while in the general case, ACI bound to A and B must be used.

A few examples will serve to indicate some of the range of variation possible:

- a) Among the most simple possibilities, A might request B to carry out an access for which B's ACI is sufficient to carry out the access request for A.
- b) A may need to provide some or all of the ACI needed for the access request to be approved by the appropriate access control components:
 - 1) A might provide this ACI by passing it to B with the access request.
 - 2) A might request prior authorization from C before requesting B to carry out the access. In this case, A would provide ACI to C which would in turn provide A a token. This token would be sent by A with the requested access to B and then C would recognize the token as a record of the previous authorization. (See Annex F for additional details on this case.)

Figure 5-3 generalizes to any number of intermediate entities, with the AEF of the final target entity obtaining an access decision mainly based on the ACI obtained from one or more of the entities in the sequence. Annex B provides additional detail on the interactions between initiators and targets in complex chains of indirect access.

NOTE – The designer of an access control policy should be aware that without care such transitive accesses may allow accesses that would not be permitted directly.

5.3 Distribution of access control components

An AEF or an ADF may be composed of one or more access control components. The access control functions may be distributed among these components as allowed by the access control policy. The basic access control functions presented above are independent of considerations of location of the components, communications among them, or their possible distribution.

An AEF is placed between each initiator-target instance so that the initiator can act on a target only through the AEF. There are several possible physical instantiations of AEF and ADF components. An ADF component may or may not be collocated (tightly coupled) with an AEF component. An ADF component may serve one or more AEF components. Likewise an AEF component may use one or more ADF components.

Collocation (tight coupling) of an AEF component and an ADF component may have advantages with respect to efficiency and timeliness (reducing delay) and may also avoid the need to protect the communications between the AEF and the ADF. ADF components that serve several AEF components may be advantageous in reducing the need to distribute ACI and in making certain associated security functions, such as audit, less complex.

A discussion of AEF and ADF components, location and examples of relationships that apply to a single initiator and a single target are found in Annex D. Component locations may be based on one or more of the following considerations.

5.3.1 Incoming access control

An SDA may consider that incoming access control at a target is sufficient. In this case, a target AEF component enforces an incoming access control policy and a target cannot receive a request that is not in conformance with the access control policy for the target. This means that the access requests sent by the initiator will reach the target AEF and will be subject to examination by the target AEF to verify that they satisfy the access control policy enforced by the ADF component.

5.3.2 Outgoing access control

An SDA may consider that it is important to prevent unauthorized accesses to targets by employing access control components local to the initiator (e.g. when the target access control system implementation is not of high quality, or if the available network resources should not be expended without first checking that the requested access is authorized), in which case outgoing access control by an initiator AEF is necessary. In this case, an initiator cannot perform an access that is not in conformance with the access control policy of the initiator's security domain.

5.3.3 Interposed access control

An SDA may conclude that it is important to filter the accesses between initiators and targets, in which case, an AEF is interposed between the initiator and the target. The interposed AEF may then enforce both outgoing and incoming access control policies. These access control policies may be independent of the initiator and target security domain access control policies.

5.4 Distribution of access control components across multiple security domains

It is possible for security domains to enter into relationships such that resources in one security domain can be accessed from another security domain. Multiple security domains may be involved, but in many instances not all will be distinct. Some of these security domains contribute ACI, some exercise control over an access, and some do both. These security domains may include:

- the security domain where ACI is bound to the initiator;
- the security domain in which the initiator resides;
- the security domain where ACI is bound to the access request;
- the security domain where ACI is bound to the target;
- the security domain in which the target resides;
- the security domains where the access control decisions are made;
- the security domains where the access control decisions are enforced.

The access control process is then similar to the case of all the AEF and ADF components being under the same SDA, as described in 5.3, with the additional complications of inter-SDA and inter-domain relationships and inter-domain communications.

Inter-domain communications include:

- notifications between SDAs or their agents of new bindings of ACI or modifications of ACI;
- requests, at the time of attempted access, for verification and translation of representations of ACI and access control policies, and responses to those requests;
- requests for access and responses to those requests.

5.5 Threats to access control

ACI and access control functions may be distributed over several real systems and security domains. ACI may be communicated over insecure communication facilities and it may be handled by components operating under different SDAs. When different SDAs are involved, a trust relationship among the SDAs is needed. Among the threats that should be considered are:

- masquerade by an entity appearing to be a proper AEF or ADF;
- bypass of an AEF;
- interception, replay and modification of ACI or other communications related to access control;
- use of ACI by other than the intended initiator;

- use of ACI against other than the intended target;
- use of ACI for other than the intended access requests;
- use of ACI at the wrong ADF;
- use of ACI outside intended constraints.

Possible means for achieving protection against threats to access control are given in 7.2.

6 Access control policies

Access control policies express certain security requirements in a security domain. An access control policy is a set of rules acted upon by ADFs. There are several considerations that may be included in access control policies and in their expression as rules. One or more of these considerations may be applicable to a particular security policy. Some access control mechanisms more easily accommodate particular considerations than others (see clause 8).

NOTE – Security policies that could be satisfied by access control mechanisms but which are concerned with other security services (e.g. confidentiality, integrity) are not considered here.

Two important and distinct aspects of an access control policy are the way in which it is expressed and the way in which it is managed (6.1 and 6.2). Commonly, administratively-imposed access control policies are expressed and implemented using security labels, while user-selected access control policies are expressed and implemented in alternative ways. Nonetheless, the expression of an access control policy, its management, and the mechanisms used to support it are logically independent of one another.

6.1 Access control policy expression

6.1.1 Access control policy categories

Two categories of security policy, rule-based and identity-based, are identified in CCITT Rec. X.800 | ISO 7498-2. Rule-based access control policies are intended to apply to all access requests by any initiator on any target in a security domain. Identity-based access control policies are based on rules specific to an individual initiator, a group of initiators, entities acting on behalf of initiators, or originators acting in a specific role. Context can modify rule-based or identity-based access control policies. Context rules may define the entire policy in effect. Real systems will usually employ a combination of these policy types; if a rule-based policy is used, then an identity-based policy is usually in effect also.

6.1.2 Groups and roles

Access control policies stated in terms of groups of initiators or in terms of initiators acting in specific roles are particular types of identity-based policies.

A group is a set of initiators whose members are considered equivalent when a particular access control policy is enforced. Groups allow access to particular targets by a set of initiators without the necessity of including the identity of individual initiators in a target's ACI, and without explicitly allocating the same ACI to each initiator. The composition of a group is determined by a management action; the ability to create or modify groups must be subject to access control. Audit of access requests by the group without distinguishing the members may or may not be required.

A role characterizes the functions a user is allowed to perform within an organization. A given role may apply to a single individual (e.g. director of a department) or to several individuals (e.g. teller, loan officer, member of a board).

Groups and roles may be used hierarchically to combine initiator identities, groups, and roles.

6.1.3 Security labels

Access control policies stated in terms of security labels are particular types of rule-based security policies. Initiators and targets are separately associated with named security labels. Access decisions are based on a comparison of the initiator and target security labels. These policies are expressed by rules describing which accesses may take place between initiators and targets with specified security labels.

Expression of access control policies in terms of security labels are particularly useful when used to provide a form of integrity or confidentiality.

6.1.4 Multiple initiator access control policies

There are many access control policies that are stated in terms of multiple initiators. These policies might identify individual initiators, or initiators that are members of the same or different groups, or initiators that are exercising different roles, or some combination of these. Examples of such multiple party access control policies include:

- Specifically identified individuals must agree to an access for it to be performed. More often initiators assuming particular roles must agree to an access, such as a company president and treasurer.
- Two members of different groups must agree to an access, such as any company officer and any member of the board of directors. In this example, the policy would probably require that the same individual cannot act for both groups, so the individual identities and group memberships would be part of the ADI used by the ADF.
- A specified number of members of a group (perhaps a majority) must agree on an access.

6.2 Policy management

This subclause identifies three aspects on a spectrum of policy management.

6.2.1 Fixed policies

Fixed policies are those that always apply and cannot be changed, for example because they are built into the system.

6.2.2 Administratively-imposed policies

Administratively-imposed policies are those that are always applied and may only be changed by duly authorized persons.

6.2.3 User-selected policies

User-selected policies are those that are available at the request of an initiator or target and are applied only for access requests involving that initiator or target or resources of that initiator or target.

6.3 Granularity and containment

Access control policies may define targets at varying levels of granularity. Each level of granularity may have its own logically separate policy and may entail the use of different AEF and ADF components (although they might use the same ADI). For example, access to a database server might be controlled only to the server as a whole; that is, either an initiator is denied access entirely or is allowed access to anything in the server. Alternatively, access might be controlled to individual files, records within files, or even data items within records. A particular database might be a Directory Information Tree, access to which might be controlled at the granularity of the entire tree, or sub-trees within the tree, or entries in the tree, or even attribute values in entries. Another example of granularity is a computer system and applications within the system.

Containment may be used to control access to a set of targets by specifying a policy that allows access to these targets only if access is allowed to a target that encompasses them. Containment might also be applied to subgroups of initiators contained in a larger group. Often the notion of containment is applied to targets that are related to one another, such as files in a database or data items in records. In the case of an element being contained within another, it is necessary for the initiator to be given the access right required to “pass through” the enclosing element before attempting to access the enclosed element. Unless designers of these security policies exercise care, access denied by one policy may effectively be allowed by another when this is not the intention.

6.4 Inheritance rules

A new element may be created by copying an existing element, or by changing an existing element, or by combining existing elements, or by construction. The ACI of the new element may be dependent on factors such as the ACI of its creator or the ACI of the copied, modified or merged elements. Inheritance rules specify these ACI dependencies, although the creator of the element may be allowed to further restrict its ACI.

Inheritance rules are the parts of an access control policy that determine the creation and modification of ACI, or the indirect application of ACI to an element based on its membership in a security domain or by the containment of one target in another.

Inheritance rules may or may not themselves be inherited by copied, modified, or merged elements. An initiator may be permitted to copy a target for its own use, but be forbidden to make further copies or to allow other initiators to copy or use it. Alternatively, once the copy is made, there may be no control of future uses.

When an element is contained in another, some (or all) of its ACI may be implied from the ACI of the containing element according to inheritance rules. Such inheritance rules can simplify the administration of uniform policies applied to a large number of elements.

6.5 Precedence among access control policy rules

It is possible for access control policy rules to conflict with one another. Precedence rules specify the order that access control policy rules are applied and which rules have precedence over others. For example, if rule A and rule B of an access control policy would individually cause an ADF to reach a different decision about a requested access, a precedence rule could give priority to rule A, in which case rule in B would not be considered, or the precedence rule might require that both rules allow the access for the request to be allowed.

Precedence rules may need to be applied to the use of initiator-bound ACI when an initiator is acting as a group member or in a particular role. The precedence rule might allow the initiator's own ACI to be combined with the ACI of the group or role being assumed, in which case it must also specify how conflicting ACI are to be combined. Alternatively, the precedence rule might require that only the group or role ACI be applied to a particular access request.

In cases when an access request involves multiple security domains, the principles described in Rec. X.810 | ISO/IEC 10181-1 regarding secure interaction policies must be observed.

6.6 Default access control policy rules

An access control policy may include default access control policy rules. These might be used when one or more initiators have not been explicitly granted or denied access to a specific target. For example, a default access control policy rule could allow access to a target if the access is not explicitly forbidden by other access control policy rules applied to relevant ADI.

6.7 Policy mapping through cooperating security domains

When providing access control for access requests between cooperating security domains, there will sometimes be a need to map or translate the ACI which is bound to the access request. This may result from the cooperating security domains having different representations for ACI, or from different security policy interpretations of the same ACI. Examples of information that might be mapped between cooperating security domains include:

- individual, group, or role identifiers (e.g. individual JSmith in security domain X may be recognized as individual XJSmith in security domain Y);
- roles and their attributes (e.g. *security administrator* in a private network attached to a public carrier may be recognized as *subscriber security administrator* in the public carrier network);
- individual identifiers to role or group (e.g. all individuals in a private network may be mapped to the role of subscriber individual in a public carrier network).

7 Access control information and facilities

7.1 ACI

The types of ACI include initiator, target, access request, operation, operand and contextual information, as described in this clause. ACI may need to be exchanged between real systems as part of the access control function. When such exchanges occur, it is essential that the cooperating entities have an agreed understanding of the abstract syntax. The discussion of ACI presented in this clause provides the basis for the detailed description of particular access control schemes described in clause 8.

NOTE – In order to maximize interoperability between real systems, there is a need to standardize representations of ACI. Other ACI, which is not considered necessary to be standardized (e.g. retained ADI), is not covered in this clause.

Depending upon the chosen security policy it will be necessary to define which ACI is required.

7.1.1 Initiator ACI

Initiator ACI is ACI about an initiator.

Example contents of initiator ACI include:

- a) the access control identity of an individual;
- b) identifier of the hierarchical group in which membership is asserted;
- c) identifier of the functional group in which membership is asserted;
- d) identifiers of roles that may be taken;
- e) sensitivity markings;
- f) integrity markings.

NOTE – An individual access control identity is not necessarily the same as that used for authentication, audit, or charging. The individual access control identity is unique within the name space of the SDA (see Annex C).

7.1.2 Target ACI

Target ACI is ACI about a target.

Examples of target ACI include:

- a) target access control identities;
- b) sensitivity markings;
- c) integrity markings;
- d) identifier of the container that encompasses a target.

7.1.3 Access request ACI

Access request ACI is ACI about an access request.

Examples of access request ACI include:

- a) allowed class of operation (e.g. read, write);
- b) integrity level required for use of operation;
- c) data type of the operation.

7.1.4 Operand ACI

Operand ACI is ACI about an access request operand.

Examples of operand ACI include:

- a) sensitivity markings;
- b) integrity markings;

7.1.5 Contextual information

Examples of contextual information include:

- a) time periods: an access may be granted only within precise periods specified by day, week, month, year, etc.;
- b) route: an access may be granted only if the route being used has specific characteristics;
- c) location: an access may be granted only to initiators at specific systems, workstations or terminals, or only to initiators at a specific physical location;
- d) system status: an access may be granted only for particular ADI when the system has a particular status (e.g. during a disaster recovery period);
- e) strength of authentication: an access may only be granted when authentication mechanisms of at least a given strength are used;
- f) other accesses currently active for this or other initiators.

7.1.6 Initiator-bound ACI

Initiator-bound ACI may contain initiator ACI, some target ACI and selected contextual information. Forms of initiator-bound ACI are discussed in clause 8, such as security labels, capabilities, and access control certificates. Examples include:

- a) initiator ACI;
- b) a target access control identity and the accesses allowed on the target (i.e. a capability);
- c) initiator location.

7.1.7 Target-bound ACI

Target-bound ACI may contain some initiator ACI, target ACI and selected contextual information. Forms of target-bound ACI are discussed in clause 8, such as labels and access control lists. Examples include:

- a) individual initiator access control identities and the accesses on the target allowed or denied to them;
- b) hierarchical group membership access control identities and the accesses on the target allowed or denied to them;
- c) functional group membership access control identities and the accesses on the target allowed or denied to them;
- d) role access control identities and the accesses on the target allowed or denied to them;
- e) authorities and the accesses authorized for them.

7.1.8 Access request-bound ACI

Access request-bound ACI may contain initiator ACI, target ACI and contextual information. Examples include:

- a) initiator/target pairs allowed to take part in an access;
- b) targets allowed to take part in an access;
- c) initiators allowed to take part in an access.

7.2 Protection of ACI

7.2.1 Access control certificates

ACI exchanged between real systems requires protection against a variety of threats to access control as described in 5.5. The authority under which ACI was issued must be verifiable by the ADF that uses ADI derived from it. One means to provide this verification is to package the ACI in a security certificate signed or sealed by the issuing authority. Such a package is called an access control certificate.

An access control certificate may contain information of various forms. Many of these are common to protecting security certificates generally and are discussed in ITU-T Rec. X.810 | ISO/IEC 10181-1.

The following information items specific to the initiator may be included:

- initiator ACI;
- a means to validate the binding of the access control certificate to a specific initiator so that it cannot be used by another initiator;
- an identifier for an account to which the access can be charged;
- identifiers for the entities accountable (i.e. responsible) for the access for accountability or audit purposes;
- the number of times the access control certificate may be used by a particular initiator.

The following information items specific to the target may be included:

- target ACI;
- a means to validate the binding of the access control certificate to a specific target so that it cannot be used to access another target;
- the number of times the access control certificate may be used by a particular initiator.

The following information items specific to the access request may be included:

- a means to validate the binding of the access control certificate to a specific access request so that it cannot be used with another access request;
- a means to validate the binding of the access control certificate to one or more access requests so that it can be used with other access requests (e.g. for access control forwarding);
- the number of times the access control certificate may be used to access a particular target;
- access request ACI.

7.2.2 Access control tokens

Another general means of protecting ACI is to place it within a security token. A security token, unlike an access control certificate which is signed or sealed by an authority, can be produced by the initiator. In the case of access control a security token is particularly relevant to access request-bound ACI.

An access control certificate may be obtained from an SDA for use in several access requests. However, the initiator may generate a security token to bind the access control certificate to a specific access request.

A security token may contain information of various forms. Many of these are common to protecting security tokens generally and are discussed in ITU-T Rec. X.810 | ISO/IEC 10181-1.

The same information items specific to the initiator, the target, and the access request that may be included in an access control certificate may be included in an access control token.

7.3 Access control facilities

This subclause identifies a number of access control facilities that may be used to provide access control in real systems. Generic descriptions of access control facilities are provided that are independent of specific mechanisms. Specific interface primitives for use with particular real systems are not prescribed.

NOTE – Although the access control facilities are described generically, they tend to demonstrate one general approach to providing access control service of the many possible. Other approaches are not deprecated by their absence in this subclause.

The access control facilities are separated into those related to management which may be invoked, for example, by a security administrator, and those which relate to the operation of access control. In particular, management related facilities support the activities “binding ACI to elements”, as described in 5.2.2.4, “modification of ACI”, as described in 5.2.2.6, and “revocation of ACI”, as described in 5.2.2.7. Operation related facilities support the activities “making ADI available to the ADF”, as described in 5.2.2.5, and “performing access control functions”, as described in 5.2.1. When different real systems or security domains use differing ACI representations, additional facilities are required to map ACI representations among them.

7.3.1 Management related facilities

Of the activities in 5.2.2, the establishment of policy and ACI representations and allocation of ACI to elements are not treated here. The Install ACI facility relates to the binding of ACI to elements. The Change ACI and Revoke ACI facilities relate to the modification and revocation of ACI. Facilities to enable and disable access control components and to list ACI of an element are in addition to the activities identified in 5.2.1.

- Install ACI – This facility binds an initial set of ACI (e.g. capabilities for use by initiators, security labels for use by initiators and targets, and ACLs for targets) to an element.
- Change ACI – This facility modifies (e.g. adds to or deletes from) the ACI bound to an element.
- Revoke ACI – This facility revokes use of ACI bound to an element so that the ACI is no longer relevant to that element. This differs from Change ACI in that any ADI relating to this ACI is also revoked.
- Revoke retained ADI – This facility revokes the validity of retained ADI.
- List ACI – This facility lists the specified ACI bound to a given element.
- Disable component – This facility disables the use of an access control function component. In the case of an AEF component, the facility inhibits all accesses through that AEF component (this prevents any access to targets exclusively served by this AEF component).
- Re-enable component – This facility re-enables the use of an access control function component.

7.3.2 Operation related facilities

The operation related facilities are expected to be used as follows, but not every access control interaction will require the use of all these steps:

- a) The initiator of the first access request of an activity determines the SDAs for elements involved in the activity by using the Identify Trusted Security Authorities facility (see ITU-T Rec. X.810 | ISO/IEC 10181-1).
- b) A secure interaction policy is established for use in the activity (see ITU-T Rec. X.810 | ISO/IEC 10181-1).
- c) ACI is bound to elements as described in 5.2.2.4 using the Acquire and Generate ACI facilities.
- d) ADI is made available to the ADF through the use of the Verify bound ACI and derive ADI facility.
- e) Contextual Information, as needed under the secure interaction policy, is obtained using the Get Contextual Information facility.
- f) The access control decision is obtained through the Decide Access facility.

Many of the facilities described below make use of protected ACI (to ensure integrity or confidentiality as required by the security policy) as discussed in 7.2.

7.3.2.1 Acquire initiator-bound ACI

This facility obtains initiator-bound ACI, or an access control certificate or an access control token containing initiator-bound ACI prior to an access request.

Invoked by initiator or ADF.

Candidate inputs are:

- authenticated initiator identity (as obtained from the Verify Facility as defined in ITU-T X.811 | ISO/IEC 10181-2);
- initiator-bound ACI selection criteria;
- validity period;
- identity of a target or group of targets;
- secure interaction policy.

Candidate outputs are:

- status (success or failure of acquire initiator-bound facility);
- initiator-bound ACI or access control certificate or access control token containing initiator-bound ACI.

7.3.2.2 Acquire target-bound ACI

This facility obtains target-bound ACI.

Invoked by ADF.

Candidate inputs are:

- target identity;
- target-bound ACI selection criteria;
- validity period;
- secure interaction policy.

Candidate outputs are:

- status;
- target-bound ACI.

7.3.2.3 Generate access request-bound ACI

This facility binds the initiator-bound ACI, access request ACI and operand-bound ACI to an access request that is necessary for an access control decision to be made.

Invoked by initiator.

Candidate inputs are:

- initiator-bound ACI (an access control certificate containing initiator-bound ACI or retained ADI);
- operand-bound ACI;
- target identity;
- operations and operands;
- validity period;
- secure interaction policy.

Candidate outputs are:

- status;
- access request-bound ACI;
- access control token;
- access control certificate (generated by an SDA on behalf of the initiator);
- retained ADI.

NOTE – The first of a sequence of access requests may return retained ADI that may be used instead of initiator-bound ACI.

7.3.2.4 Verify bound ACI and derive ADI

This facility verifies the validity of bound ACI and derives ADI from it. In cases where some or all ADI is pre-stored at the ADF, this service would be augmented or replaced by a retrieval of the pre-stored ADI.

Invoked by ADF.

Candidate inputs are:

- bound ACI (initiator, target, access request or operand);
- access control token;
- access control certificate;
- operations and operands;
- validity period;
- secure interaction policy.

Candidate outputs are:

- status;
- operation and operands;
- ADI (initiator, target, access request, or operand).

7.3.2.5 Get contextual information

This facility obtains contextual information required for an access control decision to be made.

Invoked by initiator or ADF.

Candidate inputs are:

- operations and operands;
- context information required;
- secure interaction policy.

Candidate outputs are:

- status;
- contextual information.

7.3.2.6 Decide access

This facility determines if an access is allowed.

Invoked by ADF.

Candidate inputs are:

- operations and operands;
- initiator ADI;
- operand ADI;
- target ADI;
- contextual information;
- retained ADI;
- secure interaction policy.

Candidate outputs are:

- access control decision;
- validity period of decision;
- sequence of access requests authorized;
- retained ADI.

8 Classification of access control mechanisms

8.1 Introduction

An access control mechanism is composed of an access control scheme (e.g. based on access control lists, capabilities, labels, and context) and supporting mechanisms to provide ADI to the ADF for that scheme. This clause describes a range of access control schemes which are defined in terms of the ACI that needs to be kept at different locations (principally at the initiator or at the target) and the common supporting mechanisms which are used in the Decide Access facility of 7.3.2.6. Both a basic scheme and common or likely variations on these schemes are described.

This clause discusses major categories of access control schemes and mechanisms; its objective is to show that different schemes, each having its own advantages and disadvantages, can fit into a unifying framework. The typical access control schemes can be defined, in terms of initiator-bound and target-bound ACI, as follows:

- a) If one considers a set of (target identity, operation type) pairs as initiator-bound ACI, and target identities as target-bound ACI, under an appropriate access control policy, one obtains what is essentially a capability scheme.
- b) If one considers what are commonly called “clearance” and “classification” as initiator-bound ACI and target-bound ACI, respectively, under an appropriate access control policy, one obtains what is essentially a label-based scheme.
- c) If one considers an initiator identity as initiator-bound ACI, and a set of (initiator identity, operation type) pairs as target-bound ACI, under an appropriate access control policy, one obtains what is essentially an access control list scheme.
- d) Rules concerning contextual information are most often used in conjunction with other access control schemes, but they may be used alone to create a context-based access control scheme. The contextual information may be part of initiator-bound ACI, access request-bound ACI or target-bound ACI, or it may be made available to the ADF independently of other ACI.

It is easy to devise more sophisticated variants of a) above in which the target identity becomes a target type with more than one target possessing a given “type” attribute, giving the capability a wider applicability. It is a small step further to consider this “type” attribute as a “clearance” which is matched against security label, and so arrive at b) above. Similarly, each of the first three schemes can be considered as cases of their neighbouring scheme. Each scheme can be thought of as different parts of a continuum in which schemes overlap and are not completely distinct.

When initiator names are held at targets for use as target-bound ACI (e.g. in ACL entries), day-to-day management of the target-bound ACI is difficult for systems with a dynamic population of initiators. Conversely, when target names are held as initiator-bound ACI (e.g. in capabilities) day-to-day management of initiator-bound ACI is difficult for systems with dynamic target populations.

Therefore management clearly is a factor which should influence choice of expression of policy, and to define a standard for all systems based on one or another approach is inappropriate. A practical system is likely to require a number of access control schemes from different points in the spectrum.

8.2 ACL scheme

8.2.1 Basic features

The basic features of the access control list scheme are:

- a) access control is managed as a list of (initiator qualifier, operation qualifier) pairs as target-bound ACI and individual, group or role identifiers as initiator-bound ACI;
- b) this class of access control scheme is convenient when a very fine granularity of access control is required;
- c) this class of access control scheme is convenient when there are few initiators or groupings of initiators;
- d) this class of access control scheme is convenient for revoking access to a target or group of targets;
- e) this class of access control scheme is convenient where access control management is performed on a per-target rather than per-initiator basis;
- f) this class of access control scheme is not convenient when the population of individual or groups of initiators changes frequently, but is convenient when target populations are dynamic.

8.2.2 ACI

8.2.2.1 Initiator-bound ACI

An individual, group or role identifier is the primary initiator-bound ACI in the ACL scheme.

8.2.2.2 Target-bound ACI

An ACL is the primary target-bound ACI in the ACL scheme. An ACL is a set or sequence of entries. Each entry has two fields:

a) *Initiator Qualifier*

In a simple ACL, the qualifier is the distinguishing identifier of an initiator to which an "operation qualifier" (see below) is applied. The initiator qualifier can, however, be less specific, representing more general initiator ACI such as its role or group membership.

b) *Operation Qualifier*

This describes the operations, or classes of operations (in an access request), allowed or denied for the associated initiator qualifier.

NOTE – In addition to the operations or classes of operations, limitations on the values of operands may be imposed to refine the intended access conditions.

8.2.3 Supporting mechanisms

Two mechanisms may be used to obtain the initiator-bound ACI from which the ADI required in the Decide Access facility is derived:

a) *Using authentication*

If access control is based on an individual initiator's identity, then this identity can be validated using authentication, either directly or indirectly.

If access control is based on a group or role identity, the authenticated identity is a parameter to an Acquire initiator-bound ACI facility used to obtain a validated group or role.

b) *Using access control certificates or access control tokens*

The initiator obtains an access control certificate or access control token (or both) using the Acquire initiator-bound ACI facility. This access control certificate or token is then bound to an access request by the initiator using the Generate Access request-bound ACI facility and finally is verified by the ADF using the Verify Bound ACI and Derive ADI facility.

The acceptability of the certificate authority identified in an access control certificate, or the initiator in the case of an access control token, is determined as part of the Verify Bound ACI and Derive ADI facility.

The initiator ADI (i.e. the individual, group or role identifiers), the access request and the target ADI (i.e. the access request qualifier) are parameters to the Decide Access facility. Using a suitable matching algorithm, the initiator ADI and the operation derived from the access request are compared with each (initiator qualifier, access request qualifier) entry of the access control list. The access control decision is made on the basis of whether or not a match is made. The decision returned will indicate that access should be denied if there is a match against a list of exclusions or if no match is made against a list of inclusions. Otherwise, the decision returned will indicate that access should be granted.

8.2.4 Variations of this scheme

This subclause describes common variations of the basic access control list scheme described above.

8.2.4.1 Ordered ACLs

In some ACLs employing sequences of entries, the search rule is such that the first qualifying entry terminates the search. The ordering of such ACLs is therefore significant, permitting the expression of policies in which individual initiators can be specifically denied access even though for another more general match, e.g. for initiators in a group, the initiators are given right of access.

8.2.4.2 ACLs with grouped initiators

ACL information can be structured to reflect grouping of similar access rights for a set of initiators. In addition, when targets themselves are grouped, ACL may be associated with groups of targets. A hierarchy of ACLs may be used with top level ACLs providing a coarse grain access control information over a large group of targets which may be overridden by ACLs for subgroups of targets.

8.2.4.3 ACLs with target qualifier

This extension is particularly relevant where an access control list is not collocated with a specific target. A target must be specified in each entry of the ACL. ACL entries are structured as a triple:

- initiator qualifier;
- access request qualifier; and
- target qualifier.

The matching algorithm compares the initiator ACI, requested access, and target ACI with each initiator qualifier, action qualifier, target qualifier entry of the access control list.

8.2.4.4 ACLs with grouped targets

This extension involves sharing a single ACL amongst many targets so that the decisions determined by one ACL refer to many targets. When a single target is subject to the decision criteria of more than one ACL, the ACL mechanism access control policy must specify the rule required to combine the resulting decisions.

8.2.4.5 ACLs with context qualifier

This extension involves the use of contextual information. ACL entries are structured as a triple:

- initiator qualifier;
- access request qualifier; and
- context qualifier.

The context qualifier is an additional qualifier which describes the contextual restrictions for this entry. The matching algorithm compares the initiator ACI, requested access and contextual information with each initiator qualifier, access request qualifier, context qualifier entry of the access control list.

8.2.4.6 ACLs with partial matching

In some implementations partial matching qualifiers, where parts of the identity or other initiator ACI are to be matched against the initiator qualifier, are supported. For example, if an initiator has a name which is constructed of hierarchical sequence of component names (such as country, organization, organizational unit, personal-name), the ACL can be constructed to recognize one or more components which may be viewed as group identities.

8.2.4.7 ACLs without access request qualifier

In this variation of the ACL scheme, the sets or sequences of entries in an ACL do not contain access request qualifiers. No access request qualifier is involved in the Decide Access facility. If an access by an initiator is allowed, it is allowed for all access requests.

8.3 Capability scheme

8.3.1 Basic features

The basic features of the capability scheme are:

- a) access control is managed in terms of initiator-bound ACI (a capability) that defines a set of allowed operations on an identified set of targets;
- b) this access control scheme is convenient when there are few targets;
- c) this access control scheme is not convenient for revoking access to a target at the target unless it is possible to individually identify the capabilities once granted to the initiator, but is convenient for an initiator's SDA to revoke that initiator's access rights;
- d) this access control scheme is convenient where access control management is performed at initiators;
- e) capabilities are convenient when there are "many" users or "many" groups of users accessing "few" targets and the target and the users are in different security domains.

NOTE – Use of passwords for access control is similar to, but distinct from capabilities. The basic features of passwords are:

- access control is based on ACI shared between an initiator and target;
- access control depends on the confidentiality of the ACI being maintained by the initiator and target as well as in transfer (it is often difficult to maintain the confidentiality of passwords);
- password changes can be difficult if several initiators share the same password.

8.3.2 ACI

8.3.2.1 Initiator-bound ACI

The initiator-bound ACI is a set of capabilities.

A capability has two main components:

- a) the name of the target or set of targets;
- b) the list of operations authorized on the target.

Capabilities may be conveyed by an access control certificate signed or sealed under the authority of the SDA.

8.3.2.2 Target-bound ACI

The target-bound ACI is a set of entries. Each entry has two components:

- a) the identity of the SDA;
- b) the operations which the SDA may authorize.

8.3.3 Supporting mechanisms

The initiator obtains an access control certificate or access control token using the Acquire initiator-bound ACI facility which is then bound to an access request by the initiator using the Generate Access request-bound ACI facility and, finally, is verified by the ADF using the Verify Bound ACI and Derive ADI facility.

The initiator ADI (i.e. the capability contents), the operation name and the target ADI are parameters to the Decide Access facility. The target ADI is checked to verify that it is one of the target names in the capability and the operation is checked to verify that it is one of those named in the capability. The access is allowed if both checks succeed.

The Decide Access facility will indicate that access should be denied whenever:

- a) the capability presented is not recognized as a valid capability; or
- b) access to the target is predicated on operations the SDA allowed improperly (i.e. the SDA is not allowed to enable these operations); or
- c) the operation derived from the access request does not match the capability.

8.3.4 Variation of this scheme – Capabilities without specific operations

In this variation of the capability scheme, no set of allowed operations is contained in the capability and no operation name is supplied to the Decide Access facility. If an access by an initiator is allowed, it is allowed for all operations.

8.4 Label based scheme

8.4.1 Basic features

The basic features of the label based scheme are:

- a) This scheme makes use of security labels which can be assigned to initiators and targets, and data passed between systems.
- b) This scheme is most convenient when there are many initiators accessing many targets and only a coarse granularity of access control is required.
- c) This scheme, given certain policy restrictions, can be used to control the flow of data within a security domain. Security labels also may be convenient for providing access control between security domains.
- d) The allowed operations are not explicitly included in the initiator-bound or target-bound ACI, but are defined as part of the security policy.

NOTES

- 1 Labels are not necessarily simple structures.
- 2 When an initiator is a human user (or an initiator process represents a human user), the label bound to the initiator often is called a clearance. In these cases, the label bound to the target is called a classification.

8.4.2 ACI

8.4.2.1 Initiator-bound ACI

The initiator-bound ACI is a security label.

8.4.2.2 Target-bound ACI

The target-bound ACI is a security label.

NOTE – The representations of the initiator-bound ACI and the target-bound ACI are usually structured in such a way as to facilitate their comparison, however, the same representation need not be used for both. The translation of security information representations is discussed in ITU-T Rec.X.810 | ISO/IEC 10181-1.

8.4.2.3 Operand-bound ACI

Operands of an access request may have labels bound to them. Labeled operands are a particular case of labeled data.

Two security properties of labeled data must be assured: the integrity of the binding of the label to the data, and the right of the initiator to create data with that label.

Security labeling, given certain policy restrictions, can be used to provide general access control to data within a security domain or between security domains.

Examples of labeled data include:

- documents;
- messages;
- connectionless data units;
- files in transfer.

8.4.3 Supporting mechanisms

Four mechanisms may be used to obtain the initiator-bound or operand-bound ACI used in the Decide Access facility.

- a) *Using access control certificates or access control tokens*

See 8.2.3.

- b) *Using authentication and look up*

The ADF obtains an authenticated initiator identity and uses it to look up its clearance.

c) *Using a labeled channel*

The clearance of the initiator or the label of data may be implied from the label of the channel used to convey the access request. The integrity of the binding of a label to a channel can be assured through use of an integrity service. Assurance that the channel has been assigned "correctly" can be achieved by trusting the provider of the communications service to verify it. Similarly, assurance can be obtained that a target entity is authorized to accept a channel by trusting the provider of the communications service to verify the authorization before the channel is established.

d) *Using labeled data*

The clearance of the initiator may be implied from the labels of the operands of the access request. The integrity of the binding of a label to data can be provided either by the integrity of the underlying channel or through the use of an integrity check code or digital signature over the data and the security label produced by an SDA.

A security label can be used as target ACI to protect a target. Access rules define the access permissions (operations) granted given the security label of the initiator and the security label assigned to a target.

If the security policy requires that the ACI held in the security label are used for target ACI, then overall flow of data in and out of that target can be controlled. Hence, the overall flow of data in and out of targets may be analyzed for security domains applying the same security policy.

Targets can be created within other targets. The security label of the containing target limits the security labels that may be assigned to the contained target under the rules for the appropriate security policy.

Examples of targets to which labels may be applied include:

- OSI n-entities;
- Directory Service entries;
- files held in a file store;
- database entries.

8.4.4 Labeled channels as targets

The creator of the channel (e.g. an SDA) assigns a security label to a channel. To make use of the channel, an initiator's ACI and the security label assigned to the channel are input to the Decide Access facility. That is, the channel is treated as a target. Labels for data carried within a channel must be consistent with the label of the channel.

The label assigned to the channel can also be used to control the route of the channel. In OSI terms, the N-layer entities and relay systems access the N-1 layer connections or connectionless data units, thus the N-layer entities must meet the access rules for the N-1 layer connections or connectionless data units.

Examples of labeled channels include:

- A-Associations;
- OSI N-layer connections;
- interprocess channels.

8.5 Context based scheme

8.5.1 Basic features

In certain cases, the ADF may require contextual information to interpret ADI or security policy rules. The basic features of the context based scheme are:

- a) access control is managed in terms of initiator-bound or target-bound ACI or independently as information obtained by the ADF;
- b) this scheme is convenient to enforce rules applicable to all initiators.

8.5.2 ACI

8.5.2.1 Context control lists

Context control lists are sets or sequences of entries. Each entry has two fields:

a) *Context qualifier*

The context qualifier is a sequence of contextual conditions (e.g. time, route, location) to which an operation qualifier is applied. Each contextual condition is individually associated with a true or false statement.

b) *Operation qualifier*

This describes the operations allowed for the associated context qualifier.

8.5.2.2 Contextual information

This information is obtained from the context where the access request is performed.

Context is dependent upon the environment where the access request is received by the ADF. There are various ways to obtain contextual information, such as from an underlying layered service interface or a local management interface.

8.5.3 Supporting mechanisms

The ADF uses the Get Contextual Information facility to obtain the contextual information. The contextual information and the access request are inputs to the Decide Access facility. The requested operation derived from the access request and the contextual information provided are matched against the operation qualifier and context qualifier, respectively, to determine whether access is allowed or denied.

8.5.4 Variations of this scheme

In some context control lists employing sequences of entries, the search rule is such that the first qualifying entry terminates the search. For each entry the rule is that access is denied if the contextual information does not conform to all the contextual conditions. For example, this would allow policies such as those allowing a particular operation only from some location during some period of time, but not using a particular route.

9 Interaction with other security services and mechanisms

This clause describes how other security services and mechanisms can be used to support access control. The use of access control to support other security services is not described here.

9.1 Authentication

The nature of the Access Control and Authentication services is sometimes misunderstood. Although there are some commonalities and interrelationships, the services are not the same. Some access control schemes (e.g. ACLs) rely upon identities and, thus, require authentication for assurance of identity. Successful authentication can lead to the initiator obtaining some ACI. Note that in some systems the Verify Facility for authentication and the ADF are collocated. In these cases, an authentication exchange is the single visible protocol. In distributed systems, these functions are not necessarily collocated and separate initiator ACI may be used. An identity is then simply considered as part of the initiator-bound ACI.

The relationship between authentication and access control can be specified by the access control policy. For example, if the initiator is authenticated by a less secure mechanism, the access control policy may dictate that certain operations (e.g. modify) could not be performed on the target. On the other hand, if the initiator is authenticated by a more secure mechanism, these operations would be permitted.

9.2 Data integrity

The data integrity service can be used to ensure the integrity of inputs and outputs within and between access control components, e.g. to prevent modification of capabilities, ACLs and contextual information that is stored or transferred.

9.3 Data confidentiality

The data confidentiality service may be required under some security policies to establish the confidentiality of certain inputs and outputs in and between access control components, e.g. to protect against aggregation of sensitive information.

9.4 Audit

ACI may be used to audit the access requests of a particular initiator. It may be necessary to gather several audit trails to be able to identify exactly which access requests have been performed by which initiator.

An audit policy may require that some or all access attempts are recorded. Therefore, a reliable recording mechanism may be required to be available to an access control mechanism. An audit policy may also require information about the operation of the access control mechanism to be recorded (e.g. the circumstances under which accesses have been denied). An access control policy may require that no accesses take place that are not audited, in which case the access control mechanism will be functionally dependent on the reliable recording service.

In cases where accountability of the initiator is required, the initiator is always subject to authentication prior to an access. It is important to understand that authentication and access control, while often strongly related, are not always performed by functions under the control of the same authorities, nor need the functions be collocated. The information used for authentication may be needed to obtain initiator-bound ACI (see 8.5 and 9.1 for further discussion).

Anonymous access with accountability can be provided in the following way:

- The initiator obtains from an SDA ACI that includes an associated audit identifier. The acquisition of the ACI is recorded: the identity of the initiator and the audit identifier are kept in an audit trail of the ACI issuing security domain.
- The initiator uses its initiator-bound ACI to access the target. The target security domain ADF that receives the initiator-bound ACI stores the audit identifier and the access request in its audit trail.
- An SDA with access to the audit information from both the target security domain and the initiator-bound ACI issuing security domain may, using the audit identifier, identify the initiator. By this means the initiator may be held accountable for its accesses.

If there is a conflict between the initiator's desire for anonymity and the target security domain's requirement for knowledge of the initiator identity, the access may be rejected; the decision depends on the access control policy of the target security domain.

9.5 Other access-related services

Access control is not the only service whose implementation is realized at the time when an access request is made. Audit (above), accountability, and charging are other security-related services that operate at the time of an access request:

- an audit service records arbitrary information about the access request;
- an accountability service specifically audits the name or names of entities formally responsible for invoking the access request;
- a charging service ensures that an account is debited by an amount appropriate for the use of the accessed resource.

The information required at the time of an access request to support each of these services is logically different. The ADI provided for access control, the account name for charging, and the responsible entity identification for accountability may all be different. However, in some implementations, it is required to use the same information (e.g. the access control identity) for each of these. This may result in confusion, especially in the presence of forwarded access requests. It is preferable to keep the different types of information separate.

Annex A

Exchange of access control certificates among components

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

A.1 Introduction

The purpose of this annex is to give a practical example of how access control certificates may be forwarded between components where some components act at the same time as targets and initiators, and to establish the general requirements for the passing of multiple access control certificates between different components in a chained access.

A.2 Forwarding access control certificates

The Security Frameworks Overview describes a number of mechanisms which enable an entity which has the right to use a security certificate to transfer this right to other entities. In situations where one entity, B, makes access requests on behalf of another entity, A, these mechanisms may be used to transfer the right to use an access control certificate from A to B.

A.3 Forwarding multiple access control certificates

In some instances it may be necessary to use several access control certificates to accomplish a complex interaction. This requirement is first illustrated by means of an example. This gives a view of the sources and uses of the various access control certificates that might be required. Three classes of access control certificates are identified, each of which has different characteristics. The example is then extended to give a more general view.

A.3.1 Example

Assume that application A2 is being accessed by the application A1 used by user U. Each access request is by A1 on A2. However, A2 may use the services of another application, A3, to satisfy the access request which, in turn, might need the services of application A4, as illustrated in Figure A.1.




Figure A.1 – Forwarding multiple access control certificates

First, consider the relationship between A1 and A2 and the access control certificates that may need to be associated with a requested access by A1 on A2. Two access control certificates may be required: for the user U and for the application A1.

There are three classes of access control certificates that may be required both for the user U and for the application A1:

- access control certificates needed to access A2 valid for all operations;
- access control certificates needed to access A2 valid for a specified set of operations;
- access control certificates needed to access A2 valid for a single operation.

In principle, each access control certificate may be obtained from a different SDA.

The access control certificates valid for all operations are conveyed at the beginning of a connection or an association.

When access control certificates define a valid set of operations, they remain unchanged until other access control certificates of this class are conveyed.

The access control certificates valid for a single operation are bound to the single operation.

A.3.2 Generalization

Next, consider the relationship between A2 and A3 and the access control certificates that may need to be associated with a requested access by A2 on A3. Three access control certificates may be required: for the user U, for the application A1, and for the application A2.

The access control certificates for the user U and application A1 intended for use at A2 may or may not be acceptable for use at A3. If they are acceptable, each of these certificates can be of any of the three classes described above. If they are not acceptable, then the user U or the application A1 (or both), when making an access request on A2, must provide an additional access control certificate intended for use at A3 which, again, may be any of the three classes above.

This scheme may be generalized for the relationship between A3 and A4, with possible additional certificates being needed from U, A1, or A2.

A.3.3 Simplifications

Usually only the access control certificate of the user U or the access control certificate of the application A1 is needed. Access control certificates valid only for a single operation are only seldom used. An access control certificate from U intended for use at A2 may be forwarded by A1 even if it is not used at A1.

IECNORM.COM : Click to view the full PDF of ISO/IEC 10181-3:1996

Annex B

Access control in the OSI reference model

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

NOTE – This text is based on CCITT Rec. X.800 | ISO 7498-2.

B.1 General

Access control may be used at the establishment of the data transfer phase of a connection or at times during the connection. This service is available in both connection oriented and connectionless protocols.

B.2 Use of access control within the OSI layers

Access control is only relevant for the following OSI layers:

- network layer (layer 3);
- transport layer (layer 4);
- application layer (layer 7).

B.2.1 Use of access control at the network layer

Access control when used at the network layer allows control of access to and/or from network nodes, subnetwork nodes or relays. Access control in the network layer can serve many purposes. For example, it allows an end system to control establishment of network connections and to reject unwanted calls. It also allows one or more subnetworks to control usage of network layer resources. In some cases, this latter purpose is related to charging for network usage.

The access control mechanisms used by the network layer are within the same layer.

B.2.2 Use of access control at the transport layer

Access control when used at the transport layer allows control of access to and/or from session entities. Different applications supported by the same end system cannot be separately controlled if they share a transport connection.

The mechanisms used by the transport layer are within the same layer.

B.2.3 Use of access control at the application layer

See the Open Systems Interconnection – Upper Layers Security Model (ISO 10745).