Chapter 1

Introduction

1.1 Distributed Security Framework - Snapshot

1.1.1 Purpose and Scope

This document, the X/Open Distributed Security Framework Snapshot, captures and presents part of the ongoing work of the X/Open Security Working Group on the Distributed Security Framework. It is the first stage of a longer term programme to produce an X/Open Distributed Security Framework - Guide (see Section 1.2 on page 3).

The principal objective in producing a Snapshot is to provide initial guidance to those responsible for the development of X/Open specifications. It focuses on the information system security concepts, requirements and models. The intention is to provided early guidance on the information security needs and obligations that specifiers of major functional areas should take into account.

1.1.2 Approach and Structure

Background

Chapter 4 describes a wide range of security countermeasures that need to be considered when addressing the information security system needs of different markets. However, comprehensive and effective solutions employing many of the measures identified are either beyond currently available *commercial off the shelf (COTS)* technology or, if available, are not easily integrated to provide distributed system wide solutions. (The customised development of integrated security solutions can not be cost justified by all but a few.)

Offset against this is the fact that, for most market sectors, awareness of information system security needs is evolving. The developing security awareness and concerns generally accompany the slow emergence of enterprise wide information systems based on distributed system technology.

The reality is that, for large parts of the commercial IT market, there is not yet the demand for distributed information system security solutions that deploy all of the measures identified. What is being asked for are solutions that incorporate those measures that counter the most obvious threats associated with distributed systems.

Approach

Any framework of distributed information system security must accommodate the evolutionary nature the IT security needs and solutions. It must cover both immediate and longer term requirements.

In this framework the longer term or target needs and solutions are addressed through the use of abstract models of information system security and general descriptions of threats, vulnerabilities and countermeasures. In this way a framework is created for marshalling requirements and organising how those requirements may be best addressed. At the target level the framework needs to avoid being overly prescriptive and detailed in describing solutions.

It is envisaged that the development of X/Open distributed system solution specifications will occur in a staged manner. Each stage reflecting the evolving commercial market security requirements and the availability of cost effective security solutions, security capable products, and suitable security standards.

Each stage will offer a *protection profile* to the market. The term protection profile comes from the continuing work in the IT security criteria and evaluation field. Simply, a protection profile is the specification of an integrated set of IT security measures and the degree of confidence (assurance) to be placed in these measures.

As to what each stage addresses is determined by:

- The work of the X/Open Requirement Task Group (RTG) in collecting and prioritising information system security requirements.
- The matching of these requirements to a suitable protection profile either formulated by X/Open or adopted by X/Open from elsewhere.
- The need for consistency with the target framework objectives and models, particularly with respect to support for interworking between the X/Open protection profiles.

This, the XDSF Snapshot, version of the framework proposes that the first stage be known as the X/Open "Baseline" protection profile. It avoids anticipating the full collation and resolution of all of the security requirements to be met by the baseline. It does, however, describe generic distributed security services expected to be in the baseline and target.

Structure

This document is structured as follows:

Chapter 2 — Security Concepts
: Introduces the top level information system security concepts, terms and models used within the framework. These are security policy, security authority, security domain, and security domain interactions.

Chapter 3 — Threats
: The security threats to information systems and their resources are first introduced by generic type. They are then examined from the point of view of general vulnerabilities and methods of attack. The analysis of the security threats, vulnerabilities and methods of attack is scoped to include all of the areas that will eventually have to be addressed by the framework.

Chapter 4 — Countermeasures
: The description of countermeasures provides a general introduction the range of measures used to counter system security threats. The countermeasures are grouped under the headings of; authentication, authorisation and access control, accountability and audit, availability, administration, and assurance. They are described at the level of general measures, strategies and techniques.

Chapter 5 — Security Architectures
: The concept of system security architectures within system architectures is introduced. The objectives to be met by such architectures are identified, particularly with regard to the specification deployment and integration of system security services. The chapter also describes the characteristics of the interfaces to security services. These are particularly important to the specifiers of APIs to security services.

Chapter 6 — Framework Baseline
: This chapter provides a first pass over the baseline distributed system security

requirements. It identifies a set of generic security services designed to meet these requirements. A top level analysis is presented of the security obligations to be placed on the major functional areas into which IT is typically divided. These security obligations are derived from the proposed baseline and the preceding analysis of threats and vulnerabilities.

Chapter 7 — Generic Security Services

More detailed descriptions of some of the generic security services are provided. These are currently:

- Authentication
- Access Control
- Audit
- Cryptographic Services
- User Authentication
- Secure Association.

1.2 Distributed Security Framework - Guide

1.2.1 Purpose and Scope

The X/Open Distributed Security Framework identifies, at a generic level, the information system security services required to meet a range of distributed information system security needs. It provides guidance on the specification, deployment and integration of these services.

The principles and objectives governing the development of the Framework are to:

- support X/Open's mission to produce open systems specifications that promote portability and interoperability
- satisfy the current requirements, identified by the X/Open Requirement Task Group, for security in distributed systems, and accommodate the anticipated future requirements
- provide a framework that allows the implementation of different, user organisation specified information system security policies and security architectures. The system security architectures shall be those appropriate for single processor through to large, heterogeneous distributed system configurations
- identify and specify the common system security services that need to deployed, and describe how these services are to be integrated to provide effective system security architectures
- guide the specifiers of X/Open components and profiles on how to contribute to the development of integrated security solutions, and how to satisfy the security requirements and obligations that are best met in their specific areas
- allow for the evolutionary nature of information system security needs, standards and solutions
- use as input to the Framework:
  - — ISO 7498/2 (Basic Reference Model, Security Architecture)
  - — ISO/IEC CD 10181-1 (Security Frameworks in Open Systems - Part 1 Overview)

- ISO 9594-8, CCITT X.509 (The Directory: Authentication Framework)
- Work ongoing in IEEE POSIX 1003.22.[1]
- ECMA TR/46 (Security in Open Systems. A Security Framework)
- OSF DCE
- SESAME
- GSS-API
- XDCS, X/Open Distributed Computing Services Framework.

1. **Note to Reviewers:** This document requires an acknowledgement of joint working between IEEE POSIX 1003.22 and the X/Open Security Working Group.

Chapter 2

Security Concepts

2.1 Introduction

Enterprises are increasingly dependent on the use of IT systems to support their business operations. As the technology of IT systems improves then the architectures of the IT systems themselves are changing to reflect the actual organisation and operations of the enterprises they support. The IT system resources are becoming distributed throughout an enterprise and the demands for interoperability between, and the interconnectivity of, IT systems both within and between enterprises is increasing.

In order to survive an enterprise is responsible for the safe custody and maintenance of its assets. These assets comprise both tangible and intangible assets. The greatest intangible asset is generally the information on which the enterprise relies to manages and conduct its business, in fact its business may be the provision of information.

The responsibility for the protection of the assets of an enterprise rests with the board of directors, or equivalent body or individual, who are therefore the ultimate security authority for the enterprise. The security policy for an enterprise identifies the security relevant assets of the enterprise, who is responsible for them, and how they should be accessed and managed. The scope of applicability of the security policy is the management domain of the security authority.

The board of directors may delegate to heads of departments the authority and responsibility for enforcing the specific aspects of the enterprise security policy that apply to the assets under the control of the department. When those assets are information held on IT systems, then aspects of security policy are further delegated to the IT systems themselves. The departments and IT system can themselves be considered as security domains enforcing a security policy specific to the activities and assets within their own domains. Such security domains are subdomains of the overall enterprise domain.

In order for an IT system to take responsibility for enforcing aspects of the enterprise security policy, the information, policy rules and procedures of the enterprise security policy must be translated or mapped to IT system elements, rules and activities.

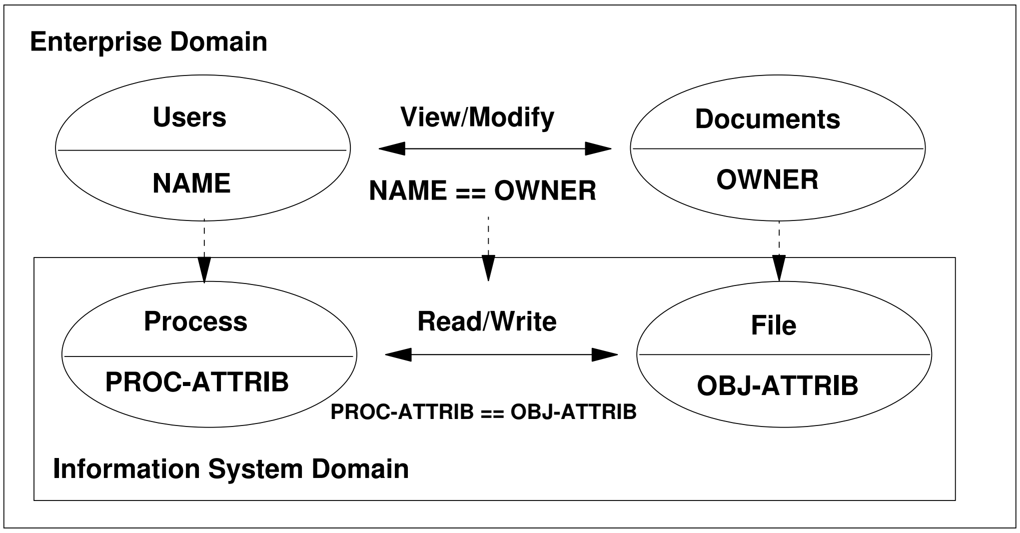

Figure 2-1 The IT System as a Subdomain of the Enterprise Domain.

A simplified example is illustrated in Figure 2.1. Within the enterprise security domain, the policy states that only the owner of a document is permitted to view it or modify its contents. Each individual has an attribute, his or her name; each document has an attribute, its owner. Within the IT system domain, processes represent users; files represent documents. Each system element is assigned an attribute corresponding to that of the entity it represents in the enterprise domain. Thus, the process attribute identifies the user; the file attribute identifies the owner. The policy rule is mapped to controlling read and write access by the process to the file, and permits access if and only if the process attribute equals the file attribute.

The operations of an enterprise involve interactions between security domains, both between its own subdomains and between its own security domain and the security domains of other enterprises. Similarly, an IT system may be decomposed into a set of security subdomains and its operations modelled as a set of interactions between security domains both between its own subdomains and with other IT system security domains.

The purpose of this security framework is to support the modelling of such a decomposition and to identify the security services, and thence the APIs, necessary to support the enforcement of security within an IT system.

The security concepts introduced above are now more formally described based upon the ISO definitions.

2.2 Security Policy

"A security policy is a set of rules that constrain one or more sets of security-relevant activities of one or more sets of elements. A security policy need not apply to all activities or all elements in a distributed system. This means that instances of its use must include a definition to which activities and to which elements the security policy applies." [ISO 10181-1]

A security policy must include a definition of the scope of the security domain to which it applies. In particular, a security policy may be specific to a particular service within an IT system.

2.3 Security Authority

"A security authority:

- is responsible for the implementation of a security policy (which may be null in some cases);
- may be a composite entity; such an entity must be identifiable and responsibility rests with it and not its components;
- may, depending on any security policy that the security authority may be subject to, delegate the enforcement of the security policy to one or more entities;
- may, depending on any security policy the security authority may be subject to, delegate responsibility over some or all of its activities and some or all of its elements within its security domain to one or more entities.

Two security authorities are said to be linked if they are constrained to coordinate their security policies." [ISO 10181-1]

Within an IT system the security authority is represented by the security services supported by the system. The activities of the security services are controlled by the security information configured by the security authority of the enterprise domain of which the IT system is a subdomain. That is, authority within an IT system is derived, by delegation, from the security authority of the enterprise the IT system is managed by. In addition a security authority is typically represented by more than one entity within an IT system, for example an operating system together with a set of security relevant applications each of which is responsible for some aspect of the overall security policy and/or some subset of the elements and activities of the IT system.

2.4 Security Domain

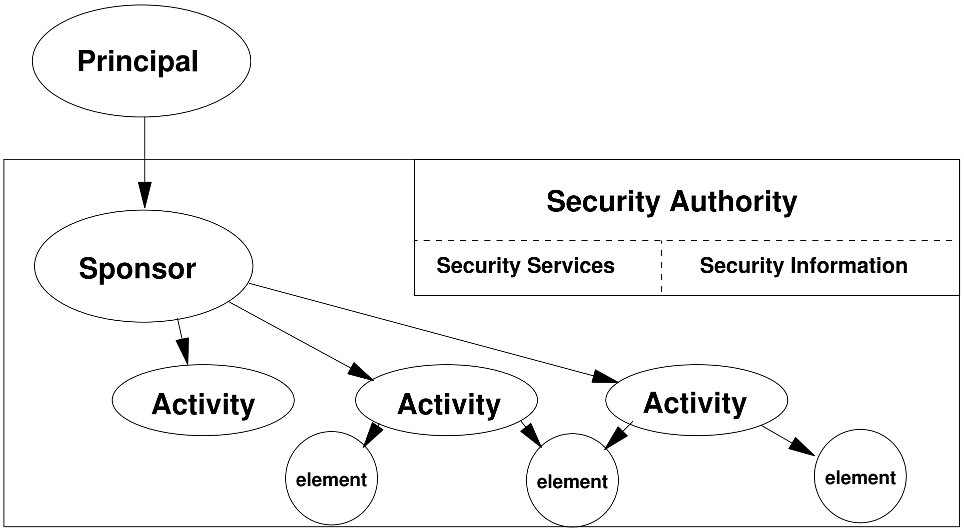

Figure 2-2 Security Domains

"A security domain is a set of elements under a given security policy administered by a single security authority for some specific security relevant activities. The activities of a security domain involve one or more elements from that domain, however at least one of the elements must be in that domain." [ISO 10181-1]

Within this framework a security domain provides services to principals that are external to the domain. See Figure 2.2. A principal may be an end-user (that is, an individual, or a service within another security domain acting on behalf of an end-user or service. Principals are required to interact only with a sponsor located within the security domain. The sponsor promotes the principal within the security domain both facilitating and constraining the principal's interactions with the activities and elements of the domain. The sponsor is responsible for the enforcement of the domain security policy through the invocation of the domain security services.

The security authority is represented within the security domain by a set of security information together with a set of security services that utilise the security information maintained by the security domain to enforce the security policy. The security information is associated with the elements and activities of the domain and with elements and services external to the security domain such as principals or external security authorities with which an inter-relationship exists or an interaction has been established.

A subset of the services supported by the security domain will be management services for the maintenance of the security information on which the security services are based.

In general a security domain is required to:

- protect the integrity, and maybe confidentiality, of its elements and activities,

- ensure the availability of, and account for the use of, the elements and activities under its protection.

To meet the generic requirements identified above the sponsor and the domain security services must:

- verify the principal's authorisation to use the domain's services, requiring authentication of the principal's identity or security information presented as proof of authorisation,
- mediate the principal's access to the activities and elements of the security domain on the basis of the access rights associated with the principal within the security domain. This includes mediation of access to the security services and security information representing the security authority of the domain,
- record a principal's use of the security domains services and access to the security domains elements to provide accountability for their use.
- protect the integrity and confidentiality of the activities and elements of the domain including providing for the segregation of the domain's elements and activities from any other domain's activities and elements.

It is a responsibility of a security policy of a security domain to identify the specific threats and detail the security measures to be taken to address those threats within that security domain.

A security domain does not necessarily provide all the security measures detailed above but may place reliance upon security domains comprising its environment to provide some of the protective measures required. A security policy is required to specify assertions on the security relevant behaviour of the environment of the security domain necessary to ensure the successful enforcement of the security policy within the domain.

This includes assertions on the responsibilities of principals authorised to use the security domain services, for example, not to disclose passwords, and especially upon those principals with administrative responsibilities over the security domain.

2.5 Interactions and Inter-relationships Between Security Domains

The security relevant operations of enterprises and IT systems generally require interactions between security domains. These interactions take the form of activities that cross domain boundaries and elements that are shared by the domains. To establish security management and control of these interactions it is necessary to formulate a secure interaction policy.

A secure interaction policy defines which subset of elements of each associating security domain comprise the common security domain and the set of permitted activities.

A secure interaction policy is established by the exchange of security information between the security domains. This occurs within two contexts, an administrative context and an operational context.

Administrative Context

In the administrative context the exchange of security information is persistent. It defines the elements, permitted activities and the policy that must operate between the associating security domains. An example is the exchange and installation of security information pertaining to the authentication of principals, this may be cryptographic keys, passwords, or address of an authentication service.

If the associating security domains are independent security domains then negotiation is required between the security authorities of the associating security domains.

If the associating security domains are both subdomains of a common superdomain then the secure interaction policy may be defined, and the associated security information established, by the security authority of the common superdomain.

Operational Context

In the operational context the exchange of security information is transient and related to the formation of an instance of an association to service a particular set of transactions between the associating security domains. The interaction occurs within the context of secure interaction policy and the security information exchanged is, or is derived from, that established previously by the administrative action creating or maintaining the common security domain.

Chapter 3

Threats

3.1 Introduction

This section describes the security threats, abbreviated here to threats, that commonly apply to the information system domain. First, these threats are identified generically. Then, they are dealt with in greater detail in the form of vulnerabilities and modes of attack.

The approach adopted is to describe threats in the round in order to acquaint the reader with both the areas of concern and the interdependent nature of the threats. This approach also helps to get across the message that the effective counters to these threats are to be found in the integration of measures deployed in the non-system (physical, personnel) and system domains.

From the framework point of view, the threats and vulnerabilities identified here will all eventually have to addressed within X/Open Security Guides and, where appropriate, by X/Open conformant systems.

3.2 Generic Threats

From the perspective of users and owners, the threats to information and information processing systems are usually categorised under the following main headings:

- Unauthorised Modification
- Unauthorised Disclosure
- Unauthorised Use of Resources
- Denial of Service
- Repudiation.

Even at this level of definition it is impossible to clearly separate each area of threat. The realisation of a threat in one area often increases the risks of other types of threat occurring.

3.2.1 Unauthorised Modification

The unauthorised modification of information and information processing resources and services that can prejudice the business of nations, organisations and individuals.

The concern here is with the maintenance of the integrity of information. The threats being to the procedures and processes involved in ensuring that business sensitive and critical information is introduced, changed and removed in an authorised manner. Examples are: the modification of sales and purchase order, funds transfer, electronic mail and military signals information with the intention of defrauding or confusing.

The threats include the unauthorised modification of the software and hardware base of the information system, which may then accidentally or intentionally facilitate the realisation of other threats.

3.2.2 Unauthorised Disclosure

The unauthorised disclosure of information that may damage the interests of a nation, organisation, or individual.

Concerns over the confidentiality and privacy of information vary. At one end there are government and defence concerns over the disclosure of national and alliance classified information. For most enterprises the need is protect business confidential information. From the point of view of individuals the primary need is for privacy. This last need is variously addressed in data privacy legislation and similar obligations on the holders of personal data.

It is important to note that certain classes of system management and configuration information, especially that which is security relevant, constitutes sensitive information that should not be disclosed without authority.

3.2.3 Unauthorised Use of Resources

The unauthorised use of system resources includes the theft of hardware and software components, as well as the unauthorised use of information systems and the services they provide.

Apart from physical theft, the most common concern here is masquerade, the unauthorised *impersonation* of an authorised user or other entity by discovering and using their authentication credentials.

The threat of masquerade is not limited to users, malfeasant systems and services may attempt to impersonate authorised systems and services.

The unauthorised use of information system resources and services may result in embarrassment and cost penalties to the owners, as well as leading to unauthorised disclosure, modification, and denial of service.

3.2.4 Denial of Service

Denial of service covers actions and events that prevent information processing systems from providing agreed levels of service to authorised users.

Denial of service threats range over the environmental acts of nature and man, such as fire, flood, bombs, through to more targeted threats, for example the accidental or intentional swamping of an electronic mail service with a large volume of unwanted messages.

The different types of denial of service threats can result in catastrophic loss through to no more than a temporary suspension of a service. But even the latter could have a significant impact on the business of an enterprise. For example in the stocks and shares market, the suspension, for a few active trading hours, of some part of an organisation's IT based trading services.

3.2.5 Repudiation

The subsequent denial by the initiator of having performed some business or security sensitive information processing operation.

This category of threat covers many forms of denial of accountability or responsibility. It includes the denial that a message was sent or received by the sender or receiver. Where such messages are being relied on for the conduct of business or some other endeavour, then such a denial may damage the interests of one or both parties involved.

3.3 Vulnerabilities and Methods of Attack

The term security vulnerabilities is used to cover the weaknesses in the construction, functionality, and operation of information systems that expose them to the accidental or intentional realisation of security threats.

The intentional exploitation of information system security vulnerabilities is known as a method of attack. The methods of attack used may vary according to the nature and functionality of the system and its components as well as the objectives and skill of the attacker.

The following analysis mainly concentrates on the vulnerabilities and methods of attack that generally fall within the scope of IT to address.

3.4 Unauthorised Modification

With respect to unauthorised modification there are two areas of concern and vulnerability:

- Modification of software and hardware
- Modification of user, application and system data.

3.4.1 System Software and Hardware Modification

Absence of or poor system configuration management facilitates the introduction of unauthorised software and hardware components. Such components may accidentally or intentionally contain malfeasant functionality that damage the integrity of the system. Principal examples of such attacks are:

Introduction of Malfeasant Software

The common forms of which are:

- Virus software with the capability of propagating itself and often carrying a payload designed to perform some malicious act.
- *Trojan Horse* components that partially or fully emulate security relevant software components in order to capture and relay security sensitive information.
- Security and business relevant software components containing *trapdoor* functionality that permits those who know about it to bypass some part of its normal functionality. This category includes the wide range of flaws and undocumented *features* common to much software.
- Software components designed to perpetrate some malicious or fraudulent act either at some specific time or when signalled to do so. These are sometimes referred to as *logic bombs*. They may be inserted as the part of a virus, or be specifically engineered, usually by someone with detailed knowledge of the target system.

Introduction of Unapproved Hardware

Unapproved or unchecked hardware containing hardware and firmware features and functionality that allow existing security mechanisms to be bypassed.

3.4.2 Data Modification

The preservation of the integrity of enterprise sensitive and critical information is usually based on the concepts of well formed transactions. Such well formed transaction models specify the stages in the life cycle of each type of enterprise information. For each stage, whether it be information creation, modification, transmission, or destruction, the well formed transaction model should capture and define the authorisations and cross checks required in order to perform the stage.

Vulnerabilities and methods of attack on the integrity of system, user and application data depend on the completeness and correctness of the appropriate transaction models as well as their IT implementation.

What follows is a generalisation of the weakness of and attacks on the integrity of system, user and application data.

3.4.3 Transmission

Interception and Modification

Information communicated over networks can be intercepted and modified as it passes through the weakly defended or uncontrolled links, routers and nodes. The classic example given is the modification, upwards, of the financial amount field of an authorised message requesting a credit to the attacker's account. (This is a some what naive example given the level of secondary checking usually accompanying such transactions, but more realistic examples can be envisaged for other areas of integrity sensitive communications.)

Insertion and Replay

Messages, possibly previously intercepted and modified, are introduced into the communication traffic either as part of a masquerade or fraud attack, or just to damage the integrity and availability of the communication services.

3.4.4 Storage

The integrity of stored data can be accidentally or intentionally compromised in a number of different ways. Many of these result from poor access control regimes governing not just who has access to what stored information but what tools and utilities can be used to access the information. (The more physical ways in which data storage media can be corrupted lies in the province of disaster planning and recovery to address.)

Specifically Targeted Data Modification

As for transmitted data, specific stored data item values can be changed in order to support a fraud attack or to facilitate the operations of the attacker in some other endeavour.

Data Contamination and Corruption

Localised or general corruption of stored data that may be either immediately obvious or only discovered after some time. The latter is more insidious as it could be copied into the entire backup archive.

Deletion

The unauthorised removal of sensitive stored data.

Creation and Replacement

The unauthorised introduction or replacement of sensitive information.

3.4.5 Processing

Malfeasant Software

The concern here is in the trustworthiness of the software components used to process sensitive information. Attackers may introduce their own versions of regularly used tools and application components, see malfeasant software above. These can be designed to modify specific instances of data in order to achieve any one of range of objectives identified above.

3.4.6 Modification of System Data

The system and security administration data, and the services used to set this system data, determine the current operational configuration of the system. If not adequately protected, system and security management data is susceptible to all of the attacks already described, however, as it is critical to the general preservation of system security, it is explored further here.

There are a wide range of options available to the attacker as to how to exploit vulnerabilities in this area. Some of which are reasonably direct, obvious and detectable, others are more subtle, insidious and less easily detected. Examples follow:

User Account Modification

Unauthorised access to user account or profile data and services can be used to:

- Increase or decrease the privileges and authorisations assigned.
- Discover and use redundant user accounts.
- Remove active user accounts.

Communication Configuration Modification

Unauthorised access here can be used to create or modify communication addressing and routing parameters at different levels of communication service in order to open up new channels and even redirect communications.

Changing Clocks

There are an increasing number of security mechanisms that are dependent on correct time synchronisation. An attacker who manages to falsify the system clock(s) may well succeed in defeating these mechanisms.

System Service Modification

More generally, inadequately controlled access to system administration and management services can be exploited by an attacker to perform disclosure, modification and denial of service acts.

3.5 Unauthorised Disclosure

There are a variety of ways in which the confidentiality and privacy of information within information systems can be compromised. The vulnerabilities to unauthorised disclosure threats occur in the following areas.

3.5.1 Transmission

Eavesdropping

Sensitive information transmitted over local and particularly wide area networks is liable to eavesdropping at diverse points in the layered communication services and networks. The term *sniffer* is used to cover a software component designed for eavesdropping. Typically a sniffer is placed at some strategic point in the communications services from where it can observe and report sensitive transmitted information, for example user ids and passwords.

Traffic Analysis

In certain cases the frequency, volume, origin and destination of messages is considered to be confidential information. An attacker may discover such information through the analysis of communication traffic flow.

3.5.2 Storage

Copying

Information, as variously stored, can be subject to unauthorised copying using anyone of the many IT techniques and tools for reading and producing copies of stored information.

Garbage Searching

A sometimes overlooked vulnerability arises in the tendency within all levels of information system services to make transient copies of information as it is processed and transmitted. These copies often take the form of large internal buffers or temporary files occupying reusable *garbage space*. Thus an attacker may use the electronic equivalent to garbage collection and analysis.

3.5.3 Processing

Malfeasant Software

As for unauthorised modification, attackers may introduce their own versions of regularly used tools and application components. These can be designed to capture and relay the sensitive information they process.

Electromagnetic Radiation

Computers, terminals and electronic communication equipment if not adequately shielded, radiate electromagnetic signals which can be picked up at distances varying according to the strength of the signal. With the appropriate equipment the displayed, processed or transmitted information can be reconstructed.

3.5.4 Indirect Disclosure Vulnerabilities and Attacks

There are more derived classes of confidential and private information that are subject to indirect and subtle forms of attack. Examples are given here to ensure that they are included in the future development of X/Open security guidelines.

Aggregation
Sensitive information may be arrived at through the assembly of less sensitive information. For example, the unauthorised construction of a detailed description of an individual by aggregating the parts of the description from different sources.

Deduction
The very fact that certain data is being held on a system or database may be considered to be sensitive information. Although unauthorised direct access to this data may be effectively prohibited its presence can sometimes deduced from the responses to carefully constructed searches.

Listening
Microphones and voice recording services are features recently introduced on some workstations. It appears that in certain cases the default configuration of these services permits other users, possibly remotely, to turn voice recording on without there being any visible indication to the owner of the workstation. This disclosure vulnerability is included here to indicate the ever widening range of disclosure weaknesses in modern information systems.

3.6 Unauthorised Use of Resources

Unauthorised use in its widest context includes the theft of hardware and software resources. Although schemes and mechanisms for the electronic tagging of physical computer assets are emerging, addressing these is not currently within the scope of the Framework. Similarly software licence management is outside the current scope.

Within scope are the vulnerabilities to the unauthorised use of systems and the services they offer. Weaknesses in the control of access to these resources can then lead to disclosure, modification, repudiation and denial of service attacks. The converse is also true, in particular disclosure and modification vulnerabilities can facilitate unauthorised use attacks.

3.6.1 Discovery and Use of Authentication Credentials

Identification, but principally, authentication credentials are used to establish the authenticity of authorised users and other entities, such as systems and their component services, and even hardware. In many cases the need is for mutual authentication in order to counter masquerade by either party.

With respect to entity authentication credentials, the main area of vulnerability lies in poor measures governing use, selection, renewal and safe transmission and keeping of these credentials.

Where entity authentication credentials are used, the receipt, passing (some times over dubious long haul networks) and storage of authentication credentials may be liable to the attacks described below. Discovered authentication credentials may then be used to masquerade as the entity by reusing the credentials or replaying them over the network.

Trojan Horse
Trojanisation of some part of the identification and authentication functions.

Eavesdropping
Interception when passed over local and wide area networks.

Accessing the Credential Store
Discovery of authentication credentials by access to their stored representation. Here, even though the credentials may be stored encrypted, poorly selected credentials may be liable to brute force decryption attacks.

3.6.2 Acquisition of Authorisations

As already identified under unauthorised modification of system data, an attacker can exploit weaknesses in the protection of authorisation information to acquire further authorisations (privileges) and thus access services and information beyond that normally allowed to them. For example, in UNIX like systems there are a variety of ways in which an attacker can acquire *root* or *superuser* privileges.

3.7 Denial of Service

Information system vulnerabilities to denial of service and availability threats abound. They include areas of obvious concern such as:

- Dependency on continuous power supply and communications links.
- System hardware component reliability.
- System software component reliability.

These are generally well recognised areas of system availability management and planning. The establishment of disaster strategies and plans, and the regular taking and securing of backup copies of sensitive information, is appreciated by most and practised by many. Apart from physical events and acts the accidental or intentional denial of service can be caused by:

Unreliable Software
The most common cause.

Malfeasant Software
Most of the types of malfeasant software attacks can be used to destroy or cripple a service.

Swamping
The availability of a system or service to authorised users can be severely reduced by flooding it with requests. Swamping may also overwhelmed the system or service by exhausting the storage associated with it.

Of concern to information system security are attacks that aim to disable some part of the existing security mechanisms. By taking advantage of the fact that, either the system or some security relevant service, on failing does not fully recover to a secure state, an attacker may bypass the normal defences.

3.8 Repudiation

Repudiation attacks are predicated on the existence of some measures for making users and other entities accountable for their business or security sensitive actions. There are two such general mechanisms:

Audit logging
Recording each action and the identity of the responsible entity in an audit trail.

Sealing and Signing
Binding the identity of the responsible entity to information produced by the action.

The often quoted example of the latter case is the binding of digital signatures to electronic mail messages. In both cases additional information such as location, context and particularly time is usually recorded.

Attacks on such measures and mechanisms are levered off more basic attacks.

Masquerade
By using the identity of another entity hide the true responsibility.

Audit Trail Modification
Breaching the access controls to the audit trail and then change or delete all or only relevant audit records.

Malfeasant Software
Use of unauthorised software to perform the operation and to omit or falsify the record or signature.

3.9 Lack of Awareness and Utility

To put things in perspective, currently the most common vulnerabilities arise from the lack of awareness of information system security concerns, needs and practices by the users and owners.

There is often failure to clearly assign to users and administrators responsibility for information security. Addressing this falls outside the scope of X/Open specifiers. However, the concept of security domains, as described in Chapter 2, does provide a useful way of establishing boundaries of security responsibility and jurisdiction.

There are system security awareness and utility issues that do fall within the scope the framework. They are summarised below:

Awareness

Not knowing in sufficient time that a system has been attacked or that its security has been compromised. The vulnerabilities arise from inadequate measures for detecting and reporting attempted or successful attacks. Damage limitation and control is obviously dependent on timely detection and warning.

Utility

Even though strong security measures may have been deployed, if these are difficult to use and managed, then their effectiveness is significantly diminished.

From the users' point of view, should the security behaviour of the system appear to impede normal operational efficiency, then the users will be encouraged to find ways round the system's security.

The effectiveness and utility of system and security management are current one of the greatest areas of vulnerability within distributed systems. Lack of facilities that allow large, heterogeneous distributed system to be easily managed almost guarantees that they will be riddled with vulnerabilities.

3.10 Assessment of Threats, Vulnerabilities and Risks

Security threat and risk assessment is typically an iterative process applied at each level; enterprise, information system, and information sub-system and IT services. Simply, the assessment steps are:

- Identify what has to be protected against which threats.
- Determine the vulnerabilities to these threats.
- Estimate the risks of these vulnerabilities being exploited.
- Where the risks are unacceptable, identify and specify a set of countermeasures to the threats with the aim of reducing the vulnerabilities and associated risks to acceptable levels.

The information system security concerns (risk weighted threats and vulnerabilities) may vary from one enterprise to another. The proposal here is that protection profiles are used to generalise these concerns.

However, in the specification of standards for each of the major functional areas of IT and their component services, list of threats, vulnerabilities and methods of attack, given above, should be interpreted and adapted to the context and functionality of the each area and component.

Chapter 4

Countermeasures

4.1 Introduction

This chapter provides an introduction to system security countermeasures. The countermeasures are grouped under:

- Authentication
- Authorisation and Access Control
- Accountability and Audit
- Availability
- Administration and Management
- Assurance

4.2 General Strategy

To meet an enterprise's information security objectives an integrated set of security countermeasures need to be deployed, operated and maintained across the enterprise (non-system) and system domains.

Security Countermeasures are use to prevent, deter, detect, limit, and contain the realisation of identified threats.

Based on some security threat, vulnerability and risk assessment, security objectives are defined and a set of security countermeasures devised. The type, strength and deployment of these countermeasures is that which is considered to be adequate to counter the identified range of security threats and vulnerabilities in each and all of the domains.

The term security countermeasure, is abbreviated here to measures. Measures encompass security strategies, procedures and mechanisms. Procedures, sometimes known as security operating procedures, define the human duties and actions to be followed in order to prevent, limit and handle security threatening events and activities. In the case of information systems, security mechanisms are automated procedures established and maintained within the information system.

The effective counter to each threat and vulnerability is usually achieved through a combination of measures. Thus a desired level of security is achieved through the use of integrated and interdependent procedures and mechanisms.

For each information system security domain the principal information concerns, objectives and rules are captured and specified in the domain's security policy. The security procedures associated with the domain and the mechanisms implemented within the domain are the vehicles for enforcing the security policy.

4.3 Authentication

To counter unauthorised use through masquerade, identification and authentication procedures and mechanisms are required at the points of access to the information systems and the services they provide. The purpose is to prove that the identity of entity requesting service is as claimed. In many cases these measures need to support mutual authentication in order to counter masquerade by either party.

The basic authentication strategy is to use either secret or unforgeable and unique information shared by the requesting entity and the identification and authentication mechanisms, possibly on both sides. For example passwords, secret information held on tokens or smart cards, or biometric information (finger prints, retinal scans, etc) unique to each user.

The I&A procedures and mechanisms need to effectively counter the range of unauthorised disclosure attacks which can result in the discovery and the use of an entity's credentials. For this they are mainly dependent on cryptographic mechanisms and the procedures to be used in the protection of credentials in the non-system domain.

The identity of each entity is used by other security procedures and mechanisms when making security relevant decisions and performing operations for the entity. Consequently the deployment and use of strong I&A procedures and mechanisms is fundamental to the success of the overall security solution.

4.4 Authorisation and Access Control

Authorisation and access control strategies and techniques are used to counter unauthorised modification and disclosure threats.

A range of authorisation and access control strategies and measures can be combined to form an access control regime. The general purpose being to ensure that only those entities (known as initiators in access control parlance) with the necessary authorisations are allowed to access sensitive information and information system resources and services (known as targets). These access control strategies and measures can individually or in combination variously meet organisation and individual needs to preserve the confidentiality, integrity and availability of sensitive information and information processing resources.

The level of application (granularity) of authorisation and access control measures varies considerably. The access controlled targets range over systems, services, programs, devices, files, documents, memory segments, communication channels of various forms, records, data fields and even data field by content. Initiators cover systems, users, devices, processes and other instances of active executables.

Access control strategies and measures used to protect the confidentiality and integrity of information and information processing resource use four protection and separation techniques.

Physical
: This is the simplest in that it is based on the physical separation of information processing resources — essentially standalone systems. This is a simple and effective technique, but it is inconsistent with distributed system objectives.

Temporal
: Temporal separation techniques operate at different levels. At the coarsest level of granularity, an information processing system is only used to process information of one sensitivity at any one time and is totally purged before being used again. This technique when applied at a finer level of granularity to targets, such as transient files, buffers, and memory address space, is used to counter disclosure through garbage collection and similar

attacks. As such it is known as the object reuse measure. Another dimension to temporal separation is object locking as used by logical integrity mechanisms. (Single threading and atomic transaction techniques are important to the preservation of the integrity of security sensitive information.)

Cryptographic

Cryptographic separation and protection techniques are used extensively in information system security solutions. Cryptographic mechanisms range in type, application, performance, and most importantly strength of protection they offer. Principal types are based on symmetric and asymmetric algorithms; in the former the same key is used to encrypt and decrypt information, in the latter case different keys are used for these operations. The terms *private key* and *public/private key* encryption are often used to denote each type. Cryptographic mechanisms are used within authentication, accountability and audit, and to protect the integrity and confidentiality of information in storage and transmission.

Logical

Protection and separation techniques of this type are based on the combination of hardware and software functionality to create separation environments and properties at different layers within an information system. At the lower levels, the operating system, usually through the use of specific hardware features, separates and protects itself from the vagaries of higher level functionality. At varying levels of strength it provides separation of the different services and resources it offers, for example devices, directories and files. It also separates and protects the higher functional components from each other by establishing a multiprocessing environment, controlling the mapping of memory to each process, and mediating the different ways processes can communicate with each other. These operating system protection and separation features are fundamental to the implementation of any reasonably robust system, let alone a *secure* one. Based on their sound implementation at the operating system level, logical separation and protection techniques are used, where required within higher level functional components.

Authorisation and access control strategies include:

Capability based strategies

(Here **capability** is being used in the more restrictive POSIX sense.)

Capabilities (Authorisations) are assigned to each initiator, these determine what functionality is available to that initiator through each service invoke by the initiator. Capabilities are often used to control which initiator may override or bypass the normal functional behaviour of a service or system.

Discretionary Access Control (DAC) strategies

(Alternatively known as the Access Control List (ACL) strategy)

These are typically, though not exclusively, based on the UNIX file, device, etc., access control mechanisms. Authorisations take the form of the individual entity and group identifiers a initiator are allowed to adopt. The initiator which owns an target, and any other initiator given the permission to do so, assigns to other initiators, identified as individual entities and groups, permissions to perform specific operations, such as read, write and execute, on the target. The term *discretionary* reflects the fact that it is originally the owner of the target who decides who else can access the target. DAC is the most commonly used strategy.

Mandatory Access Control (MAC) strategies

(Alternatively know as Label Based or Information Flow strategies)

Mandatory implies that system mechanisms are assigned much greater responsibility for making access control decisions. MAC strategies originate in the government and military worlds. They have been devised to protect the confidentiality of Nationally Classified information, though they are adaptable to similar requirements in other areas. Here authorisations assigned to initiators take the form of clearances. These in government and military application include components for; hierarchically organised attributes (Top-Secret, Secret, Confidential, Restricted, Unclassified), and categories which are usually used to identify specific projects and topics. Sometimes a third component, caveats, is added. Caveats are often used to distinguish clearances on the basis of nationality. Targets are assigned classification labels based on the same component structure and attribute values used for clearances. Essentially, MAC strategies require the creation and maintenance of classification labelled information processing, communication and storage compartments. The strategies specify the rules governing access to and information flow between compartments based on comparisons between the relevant clearances and classification labels. Although MAC strategies have mainly been specified to protect information confidentiality others have been devised to meet integrity and confidentiality objectives.

To provide solutions to confidentiality and integrity protection requirements authorisation and access control regimes and mechanisms may be constructed from a combination of the techniques and strategies outlined above. Thus a particular access control policy may be enforced by mechanisms based on the DAC strategy using both cryptographic and logical separation techniques.

Typically, capability and DAC strategies are combined. In government and military applications the MAC strategy is often added to these two.

The integrity focused measures based on and implementing the well formed transaction modelling approach (introduced in chapter 3) are usually engineered within the appropriate services. However, they are dependent on lower level implementation of capability and DAC mechanisms to ensure protection of the services themselves.

Finally, to avoid the risks of assigning too many authorisations and responsibilities to anyone entity (user), the separation of duties strategy can be used. The range of security sensitive responsibilities and duties are separated into distinct roles (entity types). The authorisations allowed to anyone role are then limited to only those required to perform the duties assigned to that role. Closely associated with this approach is the application of "n-person" rules. Here two or more initiators have to collaborate in order to combine their authorisations to perform some sensitive operation.

4.5 Accountability and Audit

Audit and Accountability addresses repudiation and awareness concerns. Procedures and mechanisms are established to record, and where appropriate, report the occurrence of all actions and events which are considered to be business and security sensitive or threatening.

For each action the identity of the initiating entity and/or performing entity, together with other pertinent descriptors of the action, eg date and time, are recorded in a manner such that they cannot be subsequently changed or contradicted. Procedures, mechanisms and tools are required to ensure that records of security sensitive actions and usage are regularly analysed. Similarly, procedures and mechanisms are required to ensure that business sensitive actions are subsequently available to support the resolution of legal and contractual disputes.

For each security threatening event, which may be associated with an action, an alarm is raised and reported to a designated authority. The speed with which the alarm is raised is determined by the severity of the event. Ideally, the reporting of the alarm should give as much time as possible to those responsible for limiting the potential damage.

Audit and accountability strategies include:

Audit Logging
Procedures and mechanisms for detecting security relevant actions and events, wherever they occur, and recording them in a protected audit trail. What qualifies as a security relevant event is determined by tuneable audit selection criteria. Examples of these selection criteria are; type of entity performing the event, sensitivity of the action and/or information being processed, location, and time of day (allowing for higher level of recording out of normal working hours).

Audit Analysis and Archiving
Measures for analysing the audit trails produced and generating useful reports on security relevant usage and activity. The audit trails need to be securely archived in order to allow for subsequent analysis, possibly some distance into the future.

Alarm Detection and Reporting
This is often closely coupled with audit action and event logging. However, not all the security threatening events may be detectable by audit detection mechanisms. As for audit logging, tuneable alarm selection criteria are required. The alarms are reported over one or more robust communication services to designated user interfaces.

Signing and Sealing
Audit logging has severe limitations in countering repudiation of business sensitive actions performed by entities in different security domains, particularly where the domains reside in different enterprises. The problem being the level of trust and mutual recognition that would have to be established by each domain in the audit procedures and mechanisms used by the other. Here Signing and sealing mechanisms based on cryptographic techniques provide a better solution. For each business sensitive transaction the relevant information is cryptographic signed and sealed using a key (digital signature). The key used is one which in some way can be proved to bind the initiating entity to the transaction. To overcome the remaining mutual trust issues, it is proposed that trusted third parties are used to establish and maintain an adequate level of trust in the digital signatures used. As for audit, non-repudiation through signing and sealing needs to be accompanied by procedures and mechanisms which allow a proof of signature to be produce sometime after the event, in some cases this may be several years later.

For their secure working, audit and accountability measures place several dependencies on authentication, authorisation and access control measures.

4.6 Availability

Information system security availability concerns and solutions merge with those of high availability and safety critical system engineering. Availability concerns are generally addressed by strategies and techniques for:

- Prevention
- Detection
- Duplication and Redundancy
- Recovery and Fall Back

Addressing many of the threats to information system and service availability, in the context of security, are based on these strategies.

Of specific security interest are the principles of:

Fail Secure
Devising and implementing measures which, when they fail, do not leave the system or service in an insecure state.

Recover Secure
Establishing measures which ensure that, following a failure, the recovery procedures return the system or service to a secure operational state, before it is made available for use.

4.7 Administration and Management

In general, good system and security administration is necessary to counter the lack of responsibility and control threats. Weakness in these areas currently present some of the greatest threats to information system security.

There is a fuzzy boundary between system and security management. Like any other IT service, system security services and mechanisms have management and control requirements for:

- operational control
- reporting of events and states
- collection and analysis of utilisation and performance statistics
- maintenance and consistent propagation of configuration parameters and equivalent information.

The first objective is, the provision of administration facilities that support these service management and control operations within distributed systems. Close behind this is the need to counter insecurity through complexity and confusion. The management facilities need to be easy to use, and consistent in their user interfaces and in the views they present of the managed system.

The more rigorous security management strategies and measures are to:

- separate system and security management duties into clearly defined management roles.
- assign only those capabilities necessary to each role.
- apply n-person controls the use of specified security critical management operations.

Generally, system and security management facilities need to be protected and separated through the use of the access control mechanisms.

4.8 Assurance

Any delegation of responsibility for the enforcement of security to the information system must be based on some measure of confidence in the correct and continuing working of the security relevant measures and functionality.

There much debate on how to rate, measure and achieve confidence (assurance). The government and military worlds have spent considerable effort over the years in devising first national and now international schemes for specifying and evaluating system security assurance. However, it is the opinion of many, that the costs associated with achieving the higher levels of assurance specified by these schemes is beyond the budgets of the major commercial markets.

Whatever position is taken on the suitability and acceptability of current and emerging IT security evaluation schemes, the basic assurance principles and objectives need to addressed in the development of system security solutions and architectures.

In summary, these are to establish confidence in the correctness and effectiveness of security solutions and their components at each stage in the system life cycle:

- Architectural design
- Component and product specification
- Development
- Supply
- Deployment and integration
- Maintenance and management
- Disposal.

Chapter 5

Security Architectures

5.1 Introduction

An information system security architecture is a top level specification, that is included within the overall architectural specifications for a system. It covers the high level identification, specification and deployment of a set of interdependent and co-operating security relevant elements (IT services and components of IT services). These elements are collectively responsible for enforcing the system security policy in the system domain.

The security functional elements are variously and collectively responsible for providing; identification and authentication, access control, accounting, audit and security administration services.

The security relevant and enforcing elements are sometimes known as the *trusted* elements. Here *trust* implies that the required degree of confidence (assurance) in the correct and continued working of these elements is established and maintained.

In the first instance this confidence is established by testing the correctness of the specification and implementation of each element. Where an element is provided as a distinct product, this confidence may be initially achieved through the evaluation and testing of the product under some recognised security evaluation scheme.

However, confidence in the overall system security solution is established by designing, implementing and proving a system security architecture against these general objectives and criteria:

- The trusted elements are adequately separated and protected from the *untrusted* elements, so that the opportunities for the latter to interfere with the correct working of the former are minimal.
- The trusted elements are deployed such; that they can appropriately mediate the security relevant operations of the untrusted elements, and that this mediation is not easily bypassed.
- The trusted elements are closely integrated and can thus be relied on to collectively enforce the system security policy.

Confidence in the continued correct working of the trusted elements is addressed by the system security architecture design and through the application of well defined security operating procedures and strong operational configuration management and administration to the system.

5.2 The Framework and System Security Architectures

One of the principal objectives of this framework is to establish a generic target system security architecture that supports a range of instantiations. These security architectures will vary in:

System Security Policies
The types of system security policies reflect the different types of security functional and assurance requirements demanded by different market sectors. The main example being the differences between military and commercial market requirements. Although variations in the security functional requirements are disappearing, differences in assurance requirements still remain.

System Type and Scale
The system architectures range from standalone systems to large scale distributed systems required to operate, within enterprises, and increasingly across enterprises. The distributed systems are likely to be heterogeneous, using different vendor platforms.

It is envisaged that the range of system security architectures covered is primarily determined by the set of protection profiles X/Open decides to support.

In focusing on distributed system security needs the framework views isolated or standalone systems as no more than restrictions of the distributed case. The framework does not cover the *optimised* security solutions available only in the isolated system case.

5.3 Distributed System Security Architectures - Concepts and Models

5.3.1 Basic Definitions and Concepts

Services

A service is a combination of functional and data elements that are exercised and accessed through a well defined interface. The range of operations performed by the service functional elements on the service data elements and/or the passed data elements are specified by the interface.

The layering of services is a common characteristic of information system functional architectures. For services located at all but the lowest levels within the system, both service functional and data elements are often formed out of the more primitive elements of underlying services through the use of the interfaces to those underlying services. The term substantiate is used to cover this service formation through functional and data element abstraction.

Distributed Services

A distributed service is one in which one or more of its functional elements do not belong to the same process hierarchy. Typically, one or more the functional elements are located on a different processing platform to the others.

Chain of Execution

The chain of execution is the execution path taken in performing any operation on an information system. Depending on the operation a chain of execution may traverse several services. Where these services are distributed, then the chain of execution has to be carried between the different process hierarchies.

5.3.2 Distributed System Frameworks and Architectures

There are several dimensions to (distributed) information systems. At the architectural level an information system can be described from the following points of view:

Major Functional Areas

An information system can be separated into:

- Operating system services
- Communication services
- Distributed processing support services
- User interface (HCI) services
- Management services
- Application services.

This separation reflects the type of services provided, their general position within a hierarchically layered functional architecture, the supply of these services, and the way in which the standards community tends to group them. However, they are far from being distinct areas. There are many obvious and necessary places in which one area overlaps with another, and they build both vertically and horizontally with each other to provide an integrated system.

Qualities

Information systems qualities are those general and pervasive properties required of systems. Prime examples are:

- Interoperability
- Availability
- Performance
- Reliability
- Utility
- Security.

These are requirements that place obligations on most elements and interfaces within the system. For the typical system, requirements for all of the qualities are defined either explicitly or implicitly, and they have to be addressed in combination. Conflicts between quality requirements often arise and compromise and balance between them has to be found. Qualities requirements are meet through the application of standards, principles and functionality. They often require the adaptation of existing elements and the inclusion of new elements.

Physical and Temporal Dimensions

Physical and temporal dimensions characterise specific systems, but they also need to be addressed generically at the level of frameworks. Generic physical dimensions cover distribution, size and scale objectives and constraints. Generic temporal dimensions are concerned with issues such as evolution and migration.

Perspectives

There are various types and levels of view of a system that system frameworks and architectures need to accommodate. Principal examples are:

- Inter Enterprise
- Enterprise
- User
- System and Security Management
- Auditor
- Supplier
- Designer
- Implementor
- Evaluator.

Perspectives can be expressed in terms of needs translated into requirements and as responsibilities and duties.

5.3.3 Security Domains

Security, typical of most information system architecture qualities, has an impact on all other the architectural dimensions described above.

To help bring the various architectural aspects of information system security down to manageable proportions this framework employs the concept of security domains, introduced in Chapter 1. The concept of nested and interacting security domains and their accompanying security policies is used to:

- Capture many of the perspectives in terms of requirements and responsibilities.
- Handle the scaling of distributed systems, particularly those which cross enterprise boundaries.
- Define boundaries for security responsibility and jurisdiction within systems as well as between them.

Functional Areas - Platform and Service Domains

At the highest level, the framework separates the major architectural functional areas into the following types of security domain:

- *Platform domains* provide the services that control and structure the physical, processing and communication resources of an IT system into elements that may be used by the service domains.
- *Service domains* provide information processing services to principals in the context of enterprise operations (applications). These correspond to the application services. The

service domains also comprise a set of *middleware* services that together with the platform services constitute the system infrastructure services.

The infrastructure services collectively provide the system environment that facilitates, communication, distributed processing, interactions with users, and system management and control.

Platform domains structure the physical resources of an IT system, processor memory, data storage and communication media into identifiable elements with assigned security characteristics and which mediate the flow of data between these elements. In order to successfully enforce data flow mediation the services of a platform domain are presented at a protected, non-bypassable interface.

A platform domain is typically composed of operating systems and communication systems elements. The actual interfaces and services presented by a platform domain are dependent upon the implementation. They may be represented by an interface equivalent to POSIX.1 system service interface, such as traditional UNIX kernel implementations, or by a microkernel interface with POSIX.1 system services implemented as service domains.

Service domains are formed from the elements of platform domains, eg., processes, files, and communication channels. Thus inter and intra service domain operations involving interactions and data flow between platform domain entities are subject to the security policy enforced by the platform domains.

The inter-relationship between service domain components and a platform domain is that of subdomain to superdomain. The platform domain enforces a basic policy on the activities of the service domain but may delegate some aspects of policy enforcement to a particular service domain component assigned that responsibility, ie., a process which possesses a particular authorisation (capability, etc).

A service domain will rely upon the platform domain to enforce policies that provide for its protection, and may rely upon a platform to control its interactions with other service domains.

It is common for a service domain to provide services in support of other service domains, in particular some service domains may provide specific security services such as authentication and/or security attribute services. The system security architecture needs to identify such dependencies between service domains.

The point of separation between platform and service domains is based on the current realities of most open system implementations. Typically, the primary, hardware supported, separation and protection boundary is to be found at the interface to the operating system kernel. As to what parts of the infrastructure services lie below that protected interface the framework is not prescriptive. Nor does the framework prohibit the use of a hardware supported multi-ring protection domains. In this case the higher layers of the infrastructure and application service domains can be separated and structured using a stronger protection mechanism.

Where the operating system services of several processors are closely coupled and they are all enforcing the same security policy, then they can be considered to constitute a single platform domain.

Distributed Systems and Security Domains

The following figure provides an example of the structuring of security domains within a simple distributed system. Here the relationship between service and platform security domains, that of subdomains to superdomain, is illustrated.

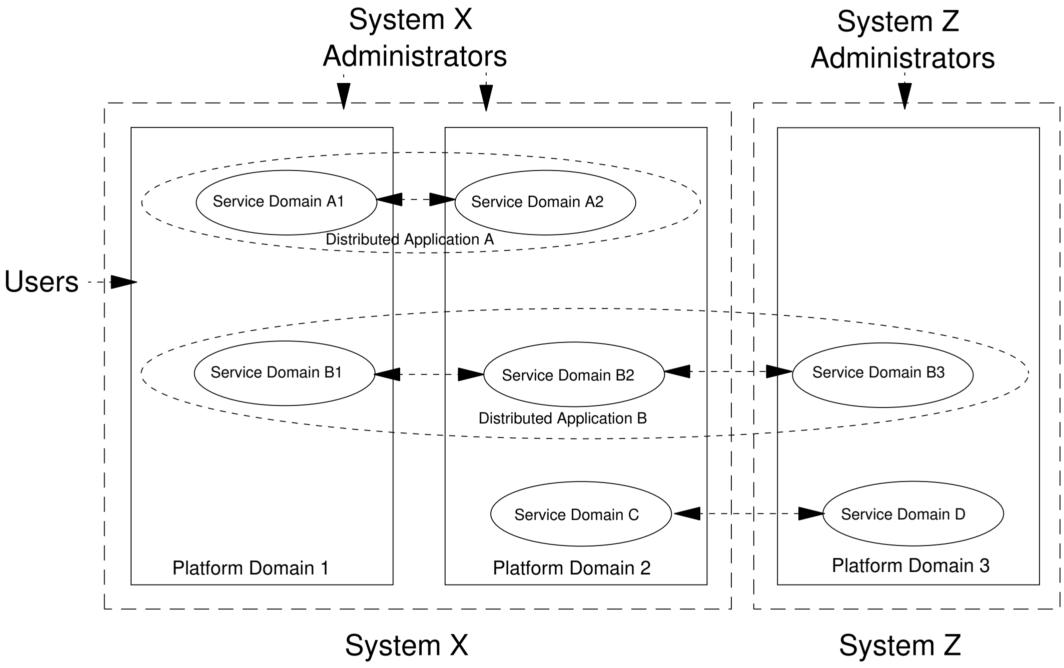

Figure 5-1 Security Domains in Distributed Systems

A distributed application may span more than a single platform with elements of the application represented by service domains on more than one platform. Each service domain will enforce a security policy common to the distributed application. However, the functions of each service domain will be subject to the security policy enforced by the platform on the platform domain elements that comprise it. The interactions between the service domain elements of the distributed application will be subject to any interconnection policy enforced between the platforms.

A distributed system may therefore be constructed so that some security services are implemented either within the platform domains, or within the service domains, or within both.

5.3.4 Trust

The concept of trust was introduced as a fundamental security architecture objective. It is now expanded.

Element x trusts element y for some classes of activity in the context of a security policy if and only if element x has confidence that element y will behave in a well-defined way that does not violate the security policy. [ISO 10181-1]

The administrative actions that establish a secure interaction policy between two security domains form a trust relationship between them. Each trusts the other to behave in a manner that does not violate the secure interaction policy. An example of this form of relationship will be a secure interaction policy permitting the exchange of information whose privacy is to be protected. Regardless of any security measures implemented within or over the communication services transferring the information the originator must trust the recipient to uphold the privacy security policy once the information is delivered into the recipient's custody.

Similarly a trust relationship exists between a subdomain and a superdomain and is established by the administrative actions on integrating the components of an IT system into a whole. For example the assignment of security attributes, including capabilities, to service domain components, ie., program and data files.

Particular trust dependencies arise between a service domain and the platform domains that substantiate it in order to maintain the integrity of the security services of the service domain. These are, derived from ECMA-138:

- Protect the security services and objects against external corruption. That is provide protection for the elements substantiating the service domain against unauthorised access by elements not constituting part of the service domain. For example, privacy of process address space.
- Provide the OSI communications security services (for example peer entity authentication) requested when forming associations, or indicate that it is unable to do so.
- Provide the necessary initial security information to security services and objects. For example, cryptographic master keys and identities for service domains.
- Enforce use of an appropriate authorisation service for each access request. That is the mediation of access to platform entities cannot be bypassed.
- Direct all initial human user contact to a subject sponsor and secure this link according to the local security policy. This is may include the provision of a trusted path. Within this framework this requirement extends to any initial contact by a principal with a security domain must be directed to a component acting as sponsor for a security domain.

5.4 Information System Services and Security

5.4.1 Introduction

A system comprises many different categories of service that in combination provide facilities to perform information processing. As described above, these services are generally organised into major functional areas and interact with each other within a layered functional architecture.

The security requirements to be met by any particular service within a given system are derived from the system's security policy and are mapped to the system's services by the system's security architecture.

With respect to these security requirements, services can be classed as:

Security Enforcing

The service is required to meet some specified security requirements in the way in which it operates, ie enforce some aspect or aspects of the system security policy. For example, a database application may be required to ensure that the system access control policy is applied with respect to the access to the database data elements (tables, records and fields). Operating system services are usually required to enforce those parts of the system security policy appropriate to operating system services.

Security Enforced

Here the service is not required to enforce some aspect of the system security policy. Instead, that aspect of the policy is enforced on the operations of the service by the information processing environment in which it runs, ie the relevant security enforcing services that provide that environment, typically the operating system services within the platform domain.

Within the layered functional architecture, the higher level security enforcing services or elements of those services, often are dependent on particular security functions of the underlying services to help them meet their security requirements. For example, an EDI application service, required to protect the confidentiality of EDI messages, will look to use the message confidentiality services provided by the communication services.

The behaviours and properties of security enforcement performed by a service are defined by parts of its interface. The term *security service* is used here to cover all services that offer security enforcing functionality as part of their interface.

The callers of these security services can be categorised as:

Security Aware

Here the caller is aware of some or all of the security functionality provided by a security service. Security aware callers typically are engineered to make use of the security services available to them. They may also be required, in certain circumstances, to take responsibility for some part of the security policy normally enforced by the security services. For example, some elements in a service domain may be assigned the capability, as specified by the operating system interface, to *override* the normal access control policy enforced by the operating system.

Security Unaware

Here the caller is unaware of some or all of the security functionality provided by a security service. In these cases the caller is unaware of either the security functionality available or the way that functionality is controlling the security of the environment in which it is operating.

5.4.2 Portability and Interoperability

The dimensions of security services and their interfaces, outlined above, have to be addressed in the specification of APIs. In particularly, they have to be taken into account in the way the APIs are specified to support the engineering of portable and interoperable system security solutions.

For any one system, which services are to be security enforcing and which services are to be unaware of aspects of security enforcement, is determined by the system security architecture and security policy.

This framework endeavours to identify and place generic security service APIs. However, from the perspective of callers of these security service APIs the impact on their security behaviour may be explicit or implicit. That is, the caller may be unaware of any aspect of security as it affects the behaviour of a security service, or the caller may be aware of some or all aspects of security and be required to manipulate the security services. It is therefore insufficient to just define security service APIs for portability and security information exchange formats to support interoperability. The interfaces need to be structured and extended to meet the different types of usage.

5.4.3 Security Service Model

In order to support the description of the characteristics of the different security services, their placement within a system, and the potential impact on principal security service APIs a security service model is defined. This model of security services and their APIs is applicable to services at all levels within a system.

A security service is based on the provision and/or use of security information. The security information is of two types; firstly security management information established by administrative action, secondly security operational information established by operational actions. Each security service provides two sets of APIs, a set of security management APIs which handle the security management information and a set of security operational APIs via which the actual security service is provided. This model is illustrated below.

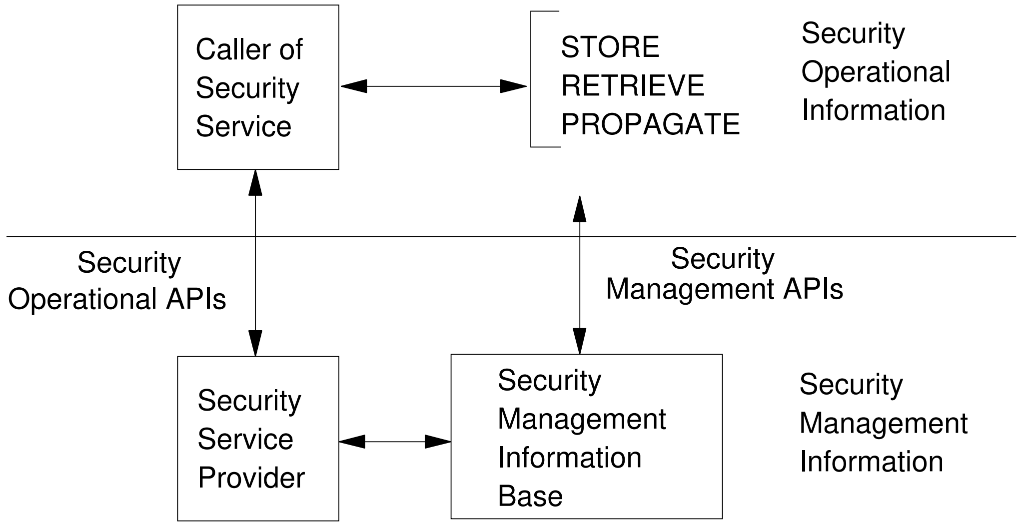

Figure 5-2 Security Service Model

Security Management APIs

The security management information is manipulated by a set of security management services and accessed via security management APIs. The security management information can be considered to be held in a Security Management Information Base (SMIB) which is accessed by the security service provider as required. The SMIB requires protection from access by any functions other than the security service provider and the security management functions. That is, the integrity and confidentiality of the SMIB must be maintained.

The SMIB is a set of security information associated with the elements of the security domain. The Management Services relate to the installation, modification, listing etc of the security information and the enabling and disabling of a particular security service. That is, the security management APIs are used to control and configure a particular security service.

Each security service will have a specific set of security management APIs and a specific set of security information associated with each element within a security domain. That is, each element may have several sets of security information associated with it, potentially one set for each of the individual security services.

Security Operational APIs

The security operational information is handled by the entities that call the security services. The security operational information is passed between the invoking functions and the security service provider via the operational interfaces to the security services. The Operational Related Services are the means of delivery of the actual security service.

The security operational information is associated with a sponsor/principal pairing and represents the security context established to service the principal. The security context has to be generated initially as a result of principal authentication. It then has to be stored and propagated with the chain of execution (the sponsor of the principal) within the system. That is, it has to be associated with every process or thread created and executed as part of the session servicing the principal and it must be propagated to any independent service that provides services to that principal. The integrity and confidentiality of the operational security information must also be maintained, including its binding with the sponsor/principal pairing.

The current security state of a security domain at a point in time is represented by the combination of the SMIB and the current set of security operational information maintained by that security domain. The security state of a system will be represented by the combination of the security states of every security domain within the system.

Security and Primary APIs

For any service, which is required to provide some security functionality, it is necessary to specify a service interface that accommodates security unaware and security aware callers. The term primary interface is used for the service interface visible to both types of callers. The term security operational interface applies to a secondary interface which may be accessible to security aware callers.

Security Unaware Caller

In this example the callers of the primary service are unaware of security and the API to the primary service makes no provision for the control of the security context of the primary service. The implementation of the primary service itself has to include functions to invoke the appropriate security services, eg., authorisation, and to acquire and handle any security operational information required by the security service.

The primary service API operates within whatever default security context it is invoked. This default security context will have to have been previously established by security aware functions eg., a login program. In this case the Primary Service API includes no provision for security information within its input and output parameters. However, its behaviour such as success or failure will have to reflect the implicit security policy and its return values or error status codes may include security related failures, eg. authorisation failure.

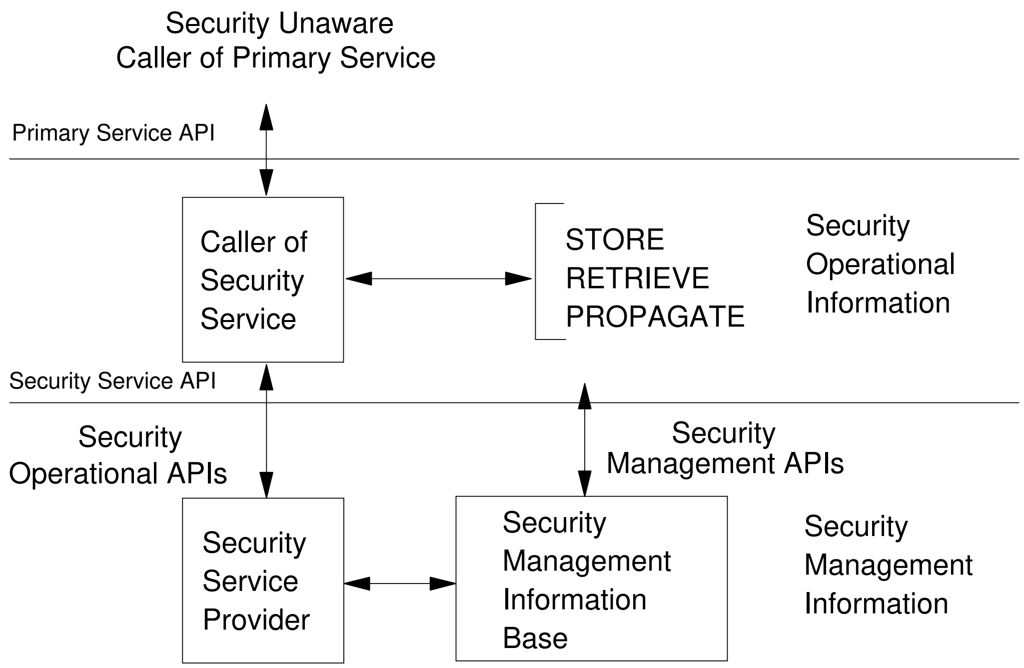

Figure 5-3 Security Unaware Caller

Security Aware Caller

In this case the Caller of the primary service is aware of one or more security services and will be responsible for handling security information related to those security services.

It may interface to these security services in two ways:

- *Directly* — by invoking the security service API to configure default values applicable to the primary service API. In this case the primary service API may be identical to the security unaware case, or
- *Indirectly* — by passing the security information as parameters to the primary service APIs. In this case the primary service API specification must include the security information, which may be optional parameters.

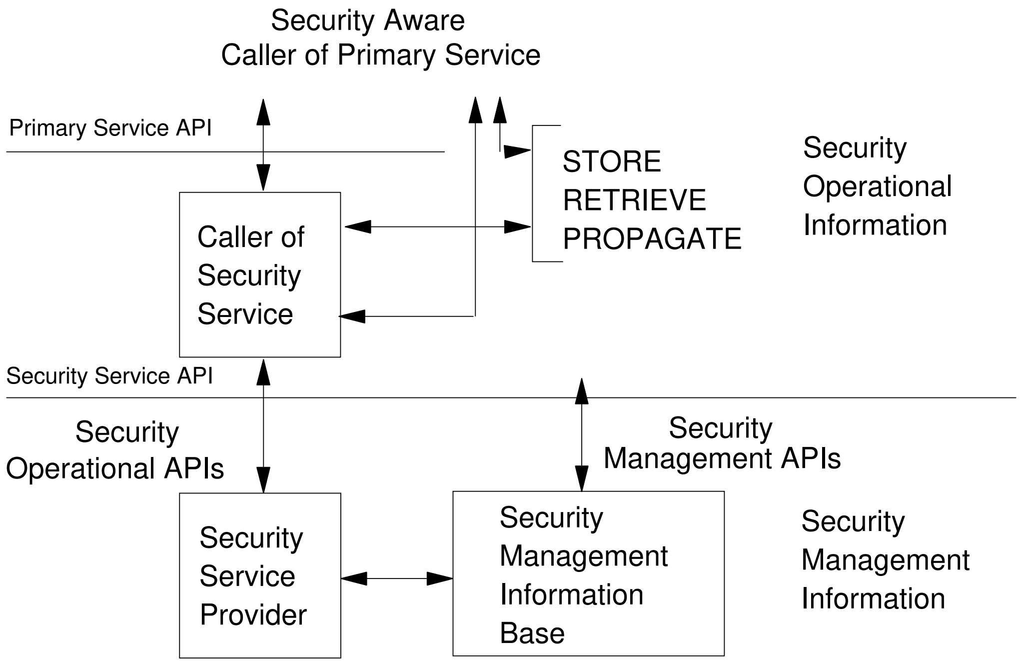

Figure 5-4 Security Aware Caller

Chapter 6

Distributed Security Framework Baseline

6.1 Introduction

This chapter outlines the generic system security services necessary to satisfy the more immediate market requirements for distributed system security. In doing so it is describing the proposed X/Open Distributed System Security Framework Baseline. This baseline forms a fundamental part of the longer term, target framework specification.

The chapter concludes with a initial mapping of the security requirements and obligations placed on the major functional areas by the baseline and more generally system security requirements.

6.2 Types of System Security Architectures

Three basic types of system architectures are to be supported:

- Standalone
- Interconnected
- Distributed

The security requirements for the distributed case are the most challenging. The stance is taken here that in identifying generic solutions to these requirements solutions are provided for the two other types.

The three types are outlined below using the security domain approach adopted by the framework. Each description builds on the preceding description, thus demonstrating how each type can be viewed as a subset of the following type. In addition special emphasis is given to the way in which the security attributes associated with a principal need to be propagated with the chain of execution. The chain of execution being the sequence or processing operations invoked by each operation imitated by a principal.

6.2.1 Standalone System

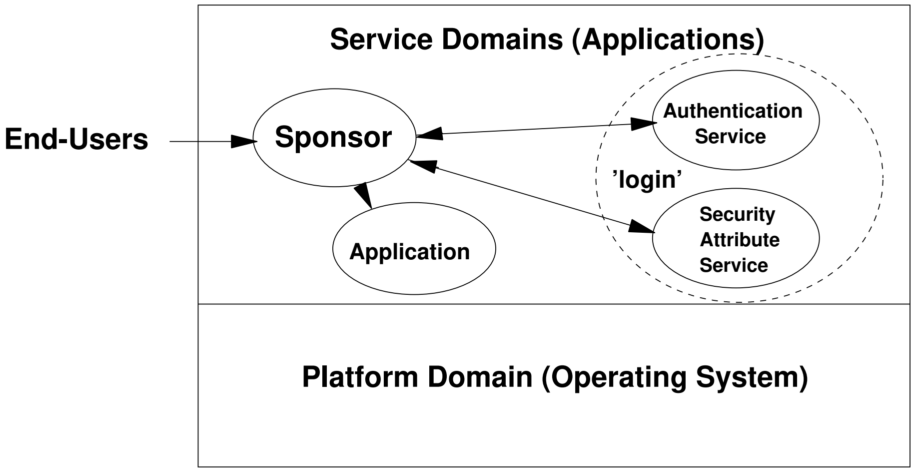

Figure 6-1 Standalone System

A standalone system comprises a single platform domain which supports a set of applications, or service domains as illustrated in Figure 6.1. The platform domain and its applications (service domains) are under the control of a single administrative authority. This authority establishes the trust relationships within the system by configuring security attributes on software installation. Specific capabilities, for example, may be assigned to an application or security enforcing service elements within an application. The attributes on the elements of the platform, ie., on its processes, files, intercommunications channels, may be pre-set, or are set using trusted utilities, viz., trusted management services.

The platform domain then enforces security policy with respect to those elements, and the security attributes assigned to them. All applications trust the platform domain to maintain the association of the principal's security attributes (security context) with the chain of execution. Thus the security attributes are made available, through platform services, to the relevant security enforcing services in order that they can ensure that the security policy is enforced on all operations performed on behalf of the principal.

The platform maintains the segregation of service domains, their data resources; it maintains the integrity and confidentiality of their contents.

6.2.2 Interconnected Systems

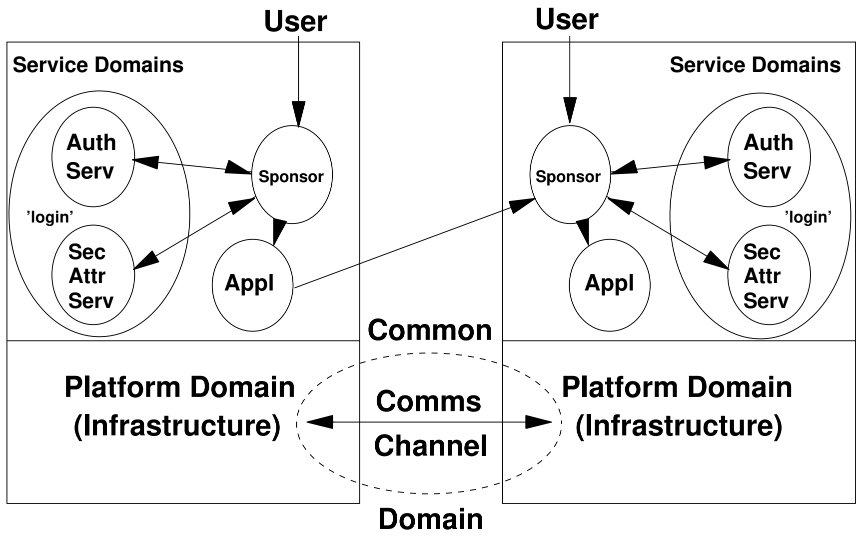

Figure 6-2 Interconnected Systems

When two systems are connected as illustrated in Figure 6.2, each system remains responsible for its own security services. An interconnection policy needs to be established, defining how, and when the systems can communicate, and the nature of those interactions. Each system then needs to be configured such that the security information they contain supports the formation of operational inter- relationships between the two systems. A communications channel provides the means by which elements in one system may be accessible to the other.

Each system remains responsible for the authentication of users of its services. Hence, a principal is required to authenticate itself separately to each system it wishes to use. This involves providing a set of security information meaningful to the system accessed.

The communications channel needs to protect the data transmitted through it. Thus, it must provide segregation of data resources and protection of the integrity and confidentiality of their contents. This can rarely be achieved by solely physical means and often cryptographic measures are used to provide the basic protection required.

On a large scale the interconnected system architecture is representative of many inter-enterprise system cases.

6.2.3 Distributed Systems

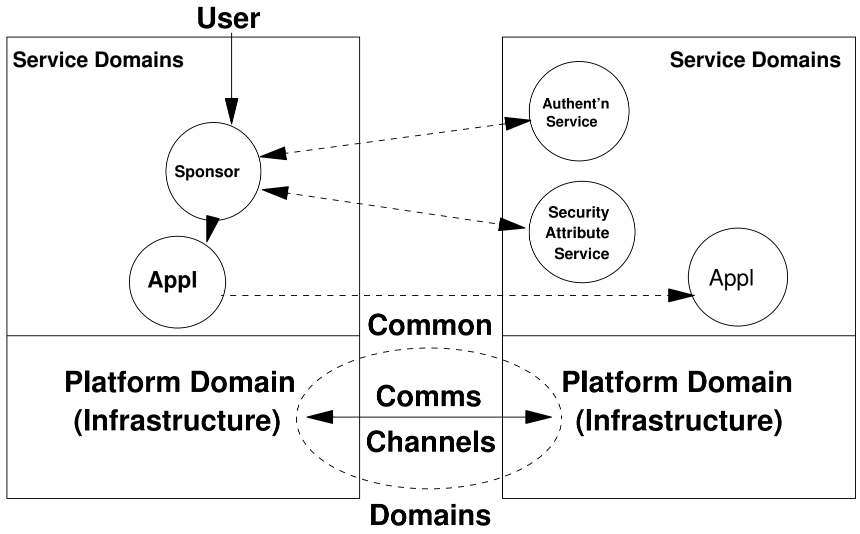

Figure 6-3 Distributed Systems

A distributed system is a set of interconnected systems which place mutual trust in certain security services which are shared by or are common to the systems, such as authentication and security attribute services. In order to do so, security information needs to be propagated throughout the distributed system in a reliable way. In a standalone system, the platform is trusted to maintain the association of security attributes with a chain of execution. In a distributed system, the association of security attributes is more complex. The security information needs to be propagated in such a way that the attributes are protected from threats within the communications system. In addition, security services must be deployed within the distributed processing support functional area. These are assigned the responsibility of propagating the security context associated with a chain of execution as it passes from one platform domain to another.

6.3 Baseline Distributed System Security Requirements

The baseline distributed system security requirements are derived from the available X/Open RTG analyses and the preceding descriptions of system security architectures and needs. They are collated and summarised below.

6.3.1 Authentication

The requirements are to support and provide:

- The authentication of principals (users and other entity types) within distributed systems. The authentication services are required to support different mechanisms and devices (challenge and response, passwords, smart cards, biometric devices).
- Services and functions for securely; storing authentication credentials, binding the authentication credentials to each chain of execution, and, where necessary, transmitting authentication credentials, over distributed systems. These services are required to support different mechanisms.
- Support for the authentication of objects such as messages and the origin of messages.

Note that support for non-repudiation requirements is addressed under accountability.

6.3.2 Authorisation and Access Control

The requirements are to support and provide:

- Access control services to be used in the protection of the confidentiality and integrity of information in storage and transmission. These services shall support mechanisms based on cryptographic and logical separation and protection techniques.
- Authorisation services which support the assignment and modification of access control information for initiators and targets.
- Authorisation state services which maintain the current state of operational (in use) access control information being used by initiators.
- Services and functions for securely; storing access control information; attaching access control information to initiators and targets, binding the access control information to each chain of execution, and, where necessary, securely transmitting access control information, over distributed systems. These services are required to support different mechanisms.
- Access Control Decision Functions (ADF) which make decisions as to whether a requested access action should be allowed. The decision is based on the access control information associated with the action, target and initiator in accordance with the access rules derived from a system security policy.
- Access Control Enforcement Functions (AEF) which mediating all system security policy defined access actions on targets and enforce the decisions made by the ADF.

6.3.3 Accountability and Audit

Accountability includes non-repudiation requirements.

The requirements are to support and provide:

- Services and functions for the detection, recording, analysis and archiving of security relevant audit events. All security services are required to detect security relevant events which occur within their boundaries.
- Services and functions for the detection and reporting of security alarm events.
- Services and functions for setting audit and alarm detection, analysis and reporting criteria.
- Portable audit record formats.
- Services and functions for binding the initiator's id or credentials with some information through the use of signing and sealing mechanisms. The services shall support both automatic and voluntary application of such a signature. In the case of automatic application this shall be applied the information resulting from an action performed the initiator. The automatic application of signature shall be determined by customised selection criteria based on the identity of the initiator, the type of action, and the sensitivity of the information produced.
- Services for checking the signature

Note: The requirements for trusted third party services in support of accountability have not yet been addressed.

6.3.4 Availability

The general requirement is to ensure that, following failure, systems and their security services recover to a secure state.

6.3.5 Administration

Security relevant administration requirements derive from the need to manage each of the security services.

6.4 Generic System Security Architecture

To meet the baseline requirements the framework proposes the adoption of the generic system security architecture outlined in Figure 6.4. This architecture is mainly based on the ECMA Security Framework. However, some of the security services within the proposed architecture are modifications of those identified by the ECMA Security Framework.

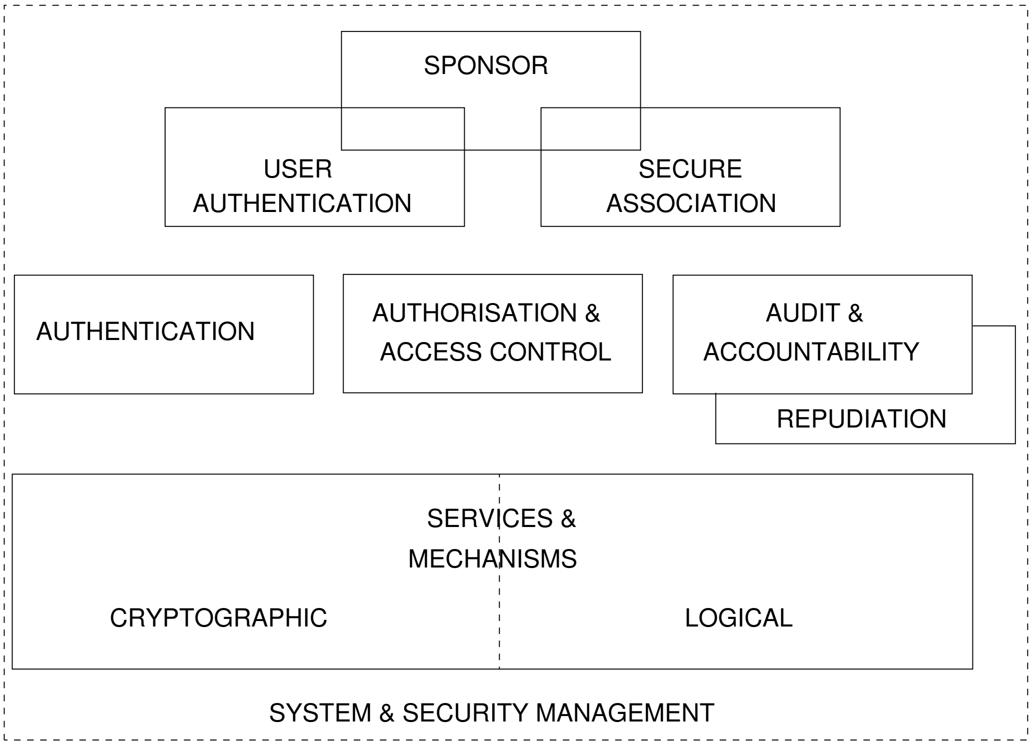

Figure 6-4 Generic Security Services

The proposed generic security architecture is partly realised by SESAME, DCE and some of the interfaces required are provided by the GSS-API.

The component security services are introduced below and then described in more detail in following chapters.

6.4.1 Sponsor Services

A Sponsor service promotes a principal within a security domain. Sponsor services are deployed within each security domain. Thus there is a Sponsor service promoting users within a system domain, these are known as User Sponsors. In the many of the service domains a sponsor can take the form the top level control functionality exercising control and selection of use the services within the domain. For example the top level control functionality in a database server within some distributed application. Here, the sponsor services have as principals the client components within that application.

They are responsible for:

- The initial identification and authentication of the principal. For this purpose the Sponsor uses the appropriate Authentication service. In the case of the User Sponsor this would be the User Authentication service.

- Establishing one or more security contexts for the principal. The security contexts contain security attributes to be used in subsequent security operations, typically these attributes are; principal identification, audit identification and access control information. Here the Sponsor services uses the Access Control and Audit services to acquire the relevant access control information passing to these services a request based on the principal's id. Alternatively the security context information may be obtained from functionally rich version of the User Authentication service, as in SESAME.

- Facilitating interactions between the principal and the other services within the security domain. In certain cases the Sponsor may be assigned responsibility to mediate access to these services, and in this respect it will operate as an Access Control Enforcement Function.

The Sponsor is often responsible for the switching of security contexts on behalf of the principal in order to support different operations. Where these operations result in a request to pass the chain of execution to another security domain, then the Secure Association service is invoke to form the necessary secure association and pass the security context to that domain. Note that in the formation of secure associations between domains the Sponsor service in the target domain is the equivalent to the ECMA target association management service.

6.4.2 Authentication Services

Authentication services are responsible for:

- Proving that the identity of an entity is as claimed.

- Authenticating messages and their origins.

The framework distinguishes between two types of Authentication service; the User Authentication service outlined below and this the internal Authentication service.

The main users of the authentication services are used by:

- Secure Association services for mutual authentication of either party when forming a secure association, and data origin authentication as part of the process of establishing the association.

- Communication services for peer and data origin authentication at some levels in the OSI stack.

They may be used by the Non-repudiation services when checking the identity of a principal, and by application services in support of application specific security protocols.

The Authentication services are required to support a number of different authentication mechanisms. In that most of these are based on crytographic techniques, the Authentication service uses the Cryptographic services.

6.4.3 User Authentication Services

User Authentication services differ from Authentication services in that they are required to support a range of procedures and mechanisms suitable the identification and authentication of users. Again most of these mechanisms are based on cryptographic techniques.

The User Authentication services are used by User Sponsor services. As already identified, the User Authentication service may be required to provide a full security context for the user. In these case the Access Control and Audit services are also used.

Note that the User Authentication services may be required to provide support for the mutual exchange of authentication credentials, ie the user requires the system to authenticate itself to him or her. This can be considered to be a very user visible implementation of a trusted path.

6.4.4 Secure Association Services

The Secure Association services are responsible for forming associations between security domains in accordance with the appropriate interdomain policy. They are also responsible for propagating the current security context from initiating domain to the target domain. They may also be responsible for checking the association and establishing a communication channel for the association as required by the interdomain policy.

The Secure Association services may use:

- Authentication services to authenticate each party and authenticate the security context passed.
- Communication security services to establish an a communication channel for the association which meets the requirements of the interdomain policy with respect to communication confidentiality and integrity protection.
- The ADF within the Access Control services to confirm that the requested association is permitted under the interdomain policy.

The Secure Association services are used by any service either in aware or unaware mode.

Note: There some debate as to the extent the secure association services are used by the other generic security services to form associations amongst themselves.

6.4.5 Access Control Services

There three types of access control services and functions:

- Access control information assignment
- Access control decision functions (ADF)
- Access control enforcement functions (AEF)

Collectively they support the application of the access control aspects of a security policy to each domain. The may be realised as functions and services which are exclusive to a service domain. But, more typically and usefully, they realised as common services deployed within the platform or infrastructure service domains.

They are responsible for:

- The assignment of access control information to entities (initiators and targets). This covers to related types of functions; those that are used to apply the access control information, and those that bind the access control information to the entity.
- Making decisions as to whether a request by an initiator to access a target is permitted under the security policy based on the current access control information assigned to both the initiator and target.
- Mediating all initiator requests to access targets and enforcing the decisions made by the ADF.

All three services and functions make use of the cryptographic and logical separation and protection mechanisms and functions in meeting their responsibilities.

Extensive use is made of these access control services and function by the other security enforcing services as well as the other generic security services. In particular, explicit and implicit use is made of them to meet basic service and service element separation and protection requirements.

6.4.6 Audit and Accountability Services

There several types of accountability through audit services and functions:

- Event detection
- Event discrimination
- Event recording
- Analysis and archiving
- Alarm detection and reporting
- Event and alarm configuration
- Audit identification assignment

These services and functions are variously deployed. The for each type of event, the event detection functions need to be located in the services where the auditable event is apparent. There is a general requirement for all security services to detect and report auditable events which are particular to their functionality. The rest of the audit services and functions are usually established as either platform domain or infrastructure service domain services and functions.

The accountability and audit services and functions use the access control functions to protect themselves and the audit trails.

The accountability and audit functions are used by all other security services to record auditable events. The audit assignment service is used to allocated and propagate audit ids for principals. It is therefore used by sponsor and possibly user authentication services to acquire the relevant audit ids.

6.4.7 Repudiation Services

[Note repudiation services are still to be outlined - they will be partly modelled on the accountability and audit services.]

6.4.8 Security Administration Services

The security administration services are an intersection of system management and security management services. They encompass a set of protected and trusted services which are responsible for:

- Controlling the starting and stopping of security services.
- Detecting and handling security service failure and ensuring recovery to a secure state.
- Facilitating and managing the creation, maintenance and consistency in distribution of security management information for each security service. This may include the distribution cryptovariables (keys) used by the crypto services.
- Maintaining historic (archived) records of system security configuration. This information is usually required to support historic audit trail analysis, and may therefore, be included as part of the function of the accountability and audit services.
- Collection and management of the current security operational state information from the system. Here there is an overlap in responsibility between access control and management services. But as this security operational state information is not exclusive to access control services, it has been allocated as a security management responsibility. (In the SESAME solution, the maintenance of this information is assigned to the authentication and privilege

attribute services.)

6.5 Mapping of Security to Major Functional Areas

One of the objective of the framework is to provide guidance to the specifiers of component and profile interfaces as to what security requirements and obligations need to be met. This section offers a first pass in providing this guidance.

6.5.1 General Approach

The general approach to the specification and design of security mechanisms and measures within a functional area is summarised here. It is based on the principles, objectives and strategies introduced in earlier chapters. The specification of security measures and mechanisms within functional areas needs to be conducted within the general constraints and objectives of the framework.

- Adopt the concepts of Chapter 2. View the Functional area as one or more security domains with dependencies and interactions amongst themselves and the security domains representing other functional areas.
- Perform a threat and vulnerability analysis for the functional area using the list of threats and vulnerabilities outlined in Chapter 3.
- Identify security measures and map these the services within the functional area, and establish which services are security enforcing with respect to each security strategy and measure.
- Define interfaces to the security relevant services which support security aware and unaware callers as defined in Chapter 5.

6.5.2 General Functional Requirements

The following sections identify security requirements which need to be considered by the specifiers of major functional areas. They are derived from the outline of the baseline and the more general analysis of security requirements described in chapters 2 to 5.

The security requirements common to most functional areas are:

Access Control

Provision of service elements which meet object reuse requirements by purging redundant communication service objects of all information.

Availability

General compliance with the principles of fail and recovery secure.

Accounting and Audit

Provision of audit event detection functions

6.5.3 Operating System Services

The main security requirements to be met at the operating system level, have already been introduced in the preceding descriptions of platform domains. They are summarised below:

- Provide low level elements from which the higher level services can be substantiated and on which the protected and separated service domains can formed.
- Mediate access to all its elements according any one or combination of the access control strategies (privilege, DAC, MAC) appropriate and required.

Based on these fundamental security services, the operating system may also be required to provide specific service elements to perform or support:

Authentication

Secure storage of and access to authentication credentials.

Support for a trusted path between user interface devices controlled by the operating system and higher level authentication services.

Access Control

In addition to the basic operating system access control services the following may also be required of the operating system:

Secure storage of and access to cryptovariables (keys) and cryptoservices.

Secure storage of and access to the security context information (authentication and authorisation information) associated with each chain of execution passing through the environment controlled by the platform domain.

Secure storage and access to security relevant object information on behalf of itself or higher level services.

Support for the combination of authorisations in order to perform "n-person" operations.

Accounting and Audit

Provision of the low level service functional elements for; security audit event detection, recording, alarm reporting, and the setting of audit and alarm selection criteria.

Secure storage and management of audit trails.

Administration

Provision of the low level service elements which are used to maintain security relevant operating system management information.

Miscellaneous

Translation services which map security relevant information from one representational or semantic form to another.

6.5.4 Communication Services

This mapping of security requirements and obligations on communications services does not, as yet, take the placement down to specific ISO OSI layers.

Authentication

Provision of service elements and protocols to support authentication between peer entities.

Access Control

Provision of confidentiality and integrity protection services to higher level service elements which counter the transmission threats and vulnerabilities.

Secure storage of and access to cryptovariables (keys) and cryptoservices, when not provided by some specific service.

Accounting and Audit

Provision of the functional elements responsible for countering communication repudiation threats.

Availability

General compliance with the principles of fail and recovery secure.

Administration

Provision of service elements which are used to maintain security relevant communication service management information.

Miscellaneous

Translation services which map security relevant communication information from one representational or semantic form to another.

6.5.5 Distributed Processing Support Services

Authentication

Provision of service elements and protocols to support authentication between instances of distributed processing support services.

Provision of authentication services which securely pass authentication credentials between the elements of higher level services.

Access Control

Provision of secure association services which securely pass the security context between elements of higher level services.

Provision of confidentiality and integrity protection service elements which counter the transmission threats and vulnerabilities.

Secure storage of and access to cryptovariables (keys) and cryptoservices, when not provided by some specific service.

Administration

Provision of service elements which are used to maintain security relevant distributed processing support service management information.

Miscellaneous

Translation services which map security relevant communication information from one representational or semantic form to another.

6.5.6 HCI Services

Authentication

Provision of service elements and protocols to securely exchange authentication credentials between the user and the system authentication services.

Support for trusted path mechanisms

6.5.7 Management Services

Note: This text still to be completed.

Chapter 7

Security Services

7.1 Introduction

This chapter presents a description of each security service. The section on each service is structured as:

- Introduction — briefly describing the purpose of the service and place it into context.
- Basic Model of Service — describing the elements of the service and their inter-relationship.
- Management Services — describing the management aspects of the service.
- Operational Services — describing the operational aspects of the services.
- Impact of service on primary service APIs — describing the different ways in which a primary service provider may interact with the security service and the relative responsibility of the primary service provider and its caller depending upon their relative awareness of the particular security service.

7.2 Authentication

7.2.1 Introduction

Authentication is the process of providing assurance of the claimed identity of an entity. There are two forms of authentication: entity authentication providing corroboration of the identity of a principal, and data origin authentication providing corroboration of the identity of the principal responsible for a specific data item.

7.2.2 Basic Service Model

The following figure illustrates the basic ISO model of authentication [ISO 10181-2]:

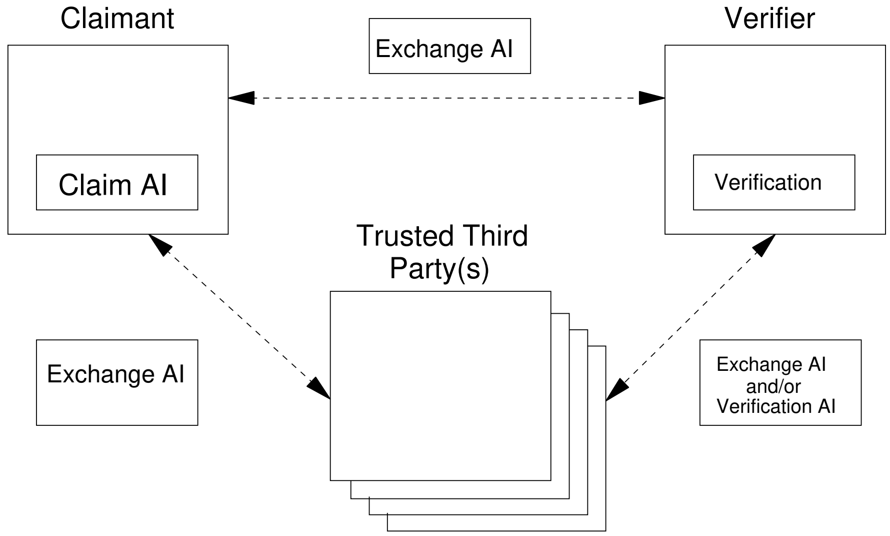

Figure 7-1 Basic Authentication Service Model

The claimant is the entity which is, or represents a principal, for the purposes of authentication.

The verifier is the entity which is or represents the entity requiring an authenticated identity.

The trusted third party is a security authority or its agent, trusted by the other entities with respect to security-related activities.

Claim AI is information used to generate exchange AI needed to authenticate the principal. eg., password, secret key, private key.

Verification AI is used to verify an identity claimed through exchange AI, for example *password* related to identity of principal, *secret key* related to identity of principal or authority, *public key* related to identity of principal or authority.

Exchange AI is information exchanged between a claimant and a verifier during the process of authenticating a principal. Examples are: *claimed distinguishing identifier*, password, challenge, response to challenge, test number, verifier identifier, result of transformation function, on-line certificate, off-line certificate.

Variations of this model are identified in the ISO work, namely No Third Party, In-Line Authentication, On-Line Authentication, and Off-Line Authentication.

7.2.3 Management Services

The Management Related services for authentication are:

Install

Installs claim AI and verification AI into the authentication service SMIB. The install operation may be a compound operation comprising:

enrollment
: The recording of authentication information associated with a principal

testify
: The provision of additional information by a witness entity

validate
: Introduces a principal into a security domain (activates a principal record) and is performed on behalf of a security authority

confirm
: Invoked after a validate service, returns specific information such as acknowledgement or rejection to principal. (Note when is this used?)

De-install

De-install causes a principal to be removed from a population of authenticatable principals. De-install may be refined with:

invalidate
: Causes a principal to be disabled from authenticating, but leaves installation record present in SMIB

notify
: May be invoked by security authority after invalidate service to notify principal of its invalidation and possibly information on how to re-enroll.

unenroll
: Causes a principal to be suppressed from a security domain.

Change AI

Change AI is invoked on behalf of a principal or security authority to cause authentication information to change, for example change password

Distribute

Distribute enables any entity to acquire sufficient verification AI upon which to verify exchange AI. An example will be the distribution of cryptographic keys used to verify cryptographic seals applied to authentication certificates.

Disable and Re-enable

Disable causes a state to be established whereby a principal is temporarily unable to authenticate. Re-enable causes the disabled state to be terminated.

7.2.4 Operational Services

Acquire

Acquire allows a claimant or verifier to obtain information required to generate exchange AI for an instance of authentication. This may require interacting with a trusted third party (for example, authentication service).

Generate

Generate is invoked by a claimant to generate exchange AI, and/or to process received exchange AI.

Verify

Verify is invoked by a verifier to verify received exchange AI from the claimant, and/or to generate exchange AI for transfer to the claimant.

Within an IT System the necessity to authenticate a principal's identity arises in two specific circumstances: the identification and authentication of a user, physically distinct to the IT system, and as part of the creation of a secure association between security domains within an IT system, or between IT systems.

7.2.5 Impact of Authentication Service on Primary Service APIs

The authentication service is a service that supports the formation of associations between entities. As such the impact with primary services is covered under the section on Secure Association Service.

7.3 Access Control

7.3.1 Introduction

It has been stated earlier that it is a responsibility of the sponsor and other security services of a security domain to mediate access by a principal to the activities and elements of the security domain. The overall provision of this mediation is termed Access Control. Access Control is applied at many places within a system. Within a single security domain it may be applied:

- as part of the initial authentication of a principal at the granularity of access to the domain as a whole
- at the granularity of access to an activity, or group of activities, supported by the domain
- at the granularity of an element, or group of elements, of the domain
- or a combination of the above.

When a service is provided by a combination of security domains within a system, in particular a combination of service domains and platform domains, then the overall system access control policy will be enforced by the combination of access control services within all the constituent security domains.

This section describes the ISO model of access control. This is based upon the concept of making an access authorisation decision and then enforcing that decision. However, although mediation of access requests is the central concept of an access control service, mediation of access is only effective if it cannot be bypassed. That is access to the entities of a security domain can only be made via activities mediated by the domain under the jurisdiction of which the entities and activities fall. In support of mediation there is therefore a requirement for the maintenance of the segregation of the resources used to substantiate the elements of the security domain from other security domains and the maintenance of the integrity of the content of, and binding to elements, of access control information associated with those elements. This is primarily the responsibility of each security domain. Where one domain utilises the elements of one or more other domains to store or transmit its elements then it must:

- either trust those other domains to enforce access control services such that access to those elements is still mediated by the security services of the owning domain (for example the use of platform domain storage services as part of a protected subsystem).
- or utilise security services to encapsulate the data using cryptographic services to provide such protection itself (for example in the use of communication services when the communication services and physical environment do not provide adequate protective measures).

The establishment of such trusts or the identification of the need for specific measures is decided by the security authority of the enterprise as part of the system design process.

7.3.2 Basic Service Model

The following figure illustrates the fundamental model of access control defined by ISO [ISO 10181-3]:

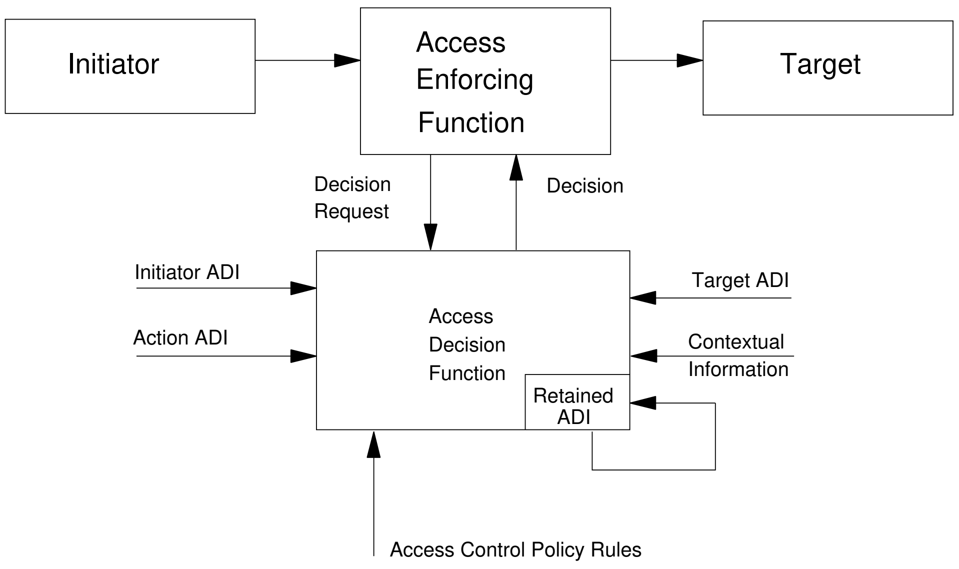

Figure 7-2 Basic Access Control Service Model

This model describes access control in terms of Access Control Information and Access Control Functions.

Access Control Information (ACI) is information used for access control purposes. ACI may be associated with entities as initiators or targets, may be associated with actions, and may include contextual information.

Access Control Decision Information (ADI) is the portion (possibly all) of the ACI associated with an entity or action that is made available for use in making a particular access control decision.

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.

The AEF is responsible for ensuring that any actions by the initiator on the target are authorised (by the ADF). When an initiator makes a request to perform an action on the target, the AEF invokes the services of the ADF so that a decision can be made. In order to make the decision, the ADF is provided with, or will acquire Access Control Decision Information (ADI) associated with the initiator, the target, and the action. Other inputs to the ADF are the access control policy rules and contextual information needed to interpret the ADI or the policy. Contextual information may include location, time of access, or communication path.

When the AEF requests a decision of the ADF the ADF requires access to the appropriate ADI and contextual information. This may be:

- preplaced at one or more ADF components after allocation of the ACI values
- ADI may be delivered to the ADF as part of the decision request,
- ADI may be obtained from another functional element (eg., a Directory Service Agent). The ADI may be obtained by the initiator or target, or by the ADF itself.

As described in Chapter 1, the enterprise security authority as part of the system design process will map enterprise security policy rules to access control mechanisms supported by the system and map enterprise security attributes to a representation as ACI within the system.

The ISO framework includes the following examples of Access Control Information:

- Initiator ACI:
  - — individual access control identities
  - — identifier of hierarchical group in which membership is asserted
  - — identifier of functional group in which membership is asserted
  - — role that may be taken
  - — sensitivity markings to which access is allowed
  - — integrity markings to which access is allowed
  - — a target access control identity and the actions allowed on the target — that is a *capability*
- Target ACI:
  - — target access control identities
  - — individual initiator access control identities and the actions on the target allowed or denied to them
  - — hierarchical group membership access control identities and the actions on the target allowed or denied to them

  - functional group membership access control identities and the actions on the target allowed or denied to them
  - role access control identities and the actions on the target allowed or denied to them
  - authorities and the actions authorised for them
  - sensitivity markings
  - integrity markings
- Contextual Information:
  - Time periods
  - route (an access may be granted only if the route being used has specific characteristics)
  - location (an access may be granted only to initiators at specific end-systems, workstations or terminals, or only to initiators at a specific physical location).
  - system status (an access may be granted only for particular ADI when the system has a particular status, for example during a disaster recovery)
  - strength of authentication (an access may only be granted when authentication mechanisms of at least a given strength are used)
  - other access currently active for this or other initiators

ACI have to be allocated to elements by the Security Authority. This is part of the process of installing or creating elements within the domain. ACI must be allocated to initiators before they may request any actions. ACI must be allocated to a target before it may be used. The allocation of ACI may require specific administrative action on the part of the Security Authority, for example configuration of ACI for a principal to be permitted to utilise the services of a domain, or may be performed automatically as part of the element creation process on the basis of the security context of the action creating the element, for example creation of a file.

ACI and ADI require binding to the elements of a domain, including a sponsor/principal pairing as an initiator, to provide assurance to the access control functions that the ACI or ADI is correctly associated with the element. The ACI and its binding to an element requires propagation and the integrity of the ACI and its binding to an element requires protection when the ACI and/or the entity is stored, processed, or exchanged. The methods by which ACI may be bound to elements include:

- on the basis of storage location, for example storage of security attributes as part of the inode of a file,
- on the basis of entity identity combined with a security attribute security service, this may require the support of an authentication service to verify an entity's identity (see User Identification and Authentication)
- on the basis of cryptographic processes for the signing or sealing of sets of information, particularly relevant in communication services (see Secure Association Service).

The functions that support the binding of ACI also need to support the selection of subsets of an element's ACI by the element and the binding of a subset of the ACI to another element to support the delegation of access rights to another element which is to operate as proxy for the initiating element.

If the ACI, or part of the ACI, associated with a principal are revoked, any ADI generated on the basis of that ACI is also required to be revoked. This may be achieved by specific action or by relying on the time expiry of such ADI when in the form of a security certificate or security token.

7.3.3 Management Services

Install ACI
Install ACI establishes an initial set of ACI (for example capabilities for use by initiators, security labels for use by initiator/targets and ACLs for targets) associated with an element. This service may be integral with the creation of an element based upon the security context of the creating activity.

Modify ACI
Modify ACI modifies (for example adds, deletes, or changes) the ACI associated with an element

Revoke ACI
Revoke ACI revokes use of ACI relating to an element. This differs from Modify ACI in that any ADI relating to this ACI is revoked.

List ACI
List ACI lists the ACI of a given element

Disable Component
Disable component disables the use of an access control function component. In the case of an AEF component the service inhibits all access through the AEF.

Re-enable Component
Re-enable component re-enables use of an access control function component.

Within a system there will be several sets of such services for each of the distinct security domains comprising the system. As an example there will be a set of such services for the platform domain, a set of services for particular applications, for example a RDBMS, and a set of services to support a common access control service shared by distributed applications, for example DCE.

7.3.4 Operational Services

The following services relate specifically to the exchange and generation of ADI in the formation of associations between a principal and the sponsor of a security domain. These are detailed further in the sections dealing with User identification and Authentication and the Secure Association services.

Acquire Initiator ADI
This service obtains initiator ADI or an access control certificate containing initiator ADI prior to an action.

Acquire Target ADI
This service acquires protected target ADI.

Generate Data ADI
This service binds data ADI to data at the time it becomes available.

Generate Action ADI
This service binds initiator ADI and data ADI to an action necessary for an access control decision to be made.

Revoke ADI
This service revokes the validity of ADI.

Verify Action ADI
Verifies the validity of action ADI.

Get Contextual Information
This service gets the contextual information required for an access control decision to be made.

Bind Initiator ADI
This service provides for the secure binding of Initiator ADI with a chain of execution within a service domain on an end-system such that the Initiator ADI are securely propagated with the chain of execution.

Retrieve Initiator ADI
This service retrieves the initiator ADI bound to a chain of execution within a security domain on an end-system.

The following service is related to the actual access control mediation process.

Decide Access
This service determines if an access is allowed based on ADI.

7.4 Impact of Access Control Services on Primary Service APIs

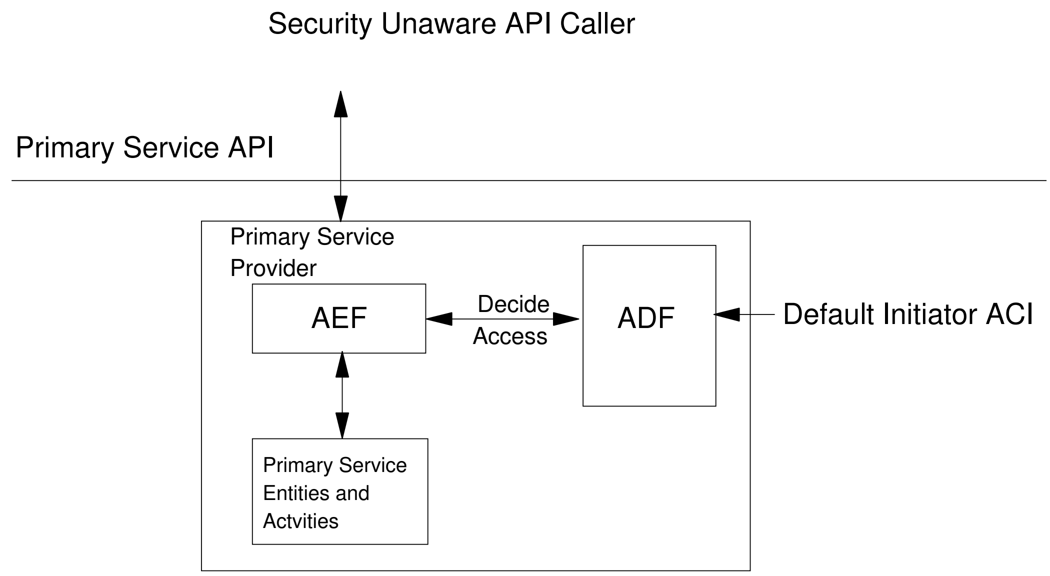

Figure 7-3 Access Control Enforcing API / Security Unaware Caller

In this case the primary service APIs provide no support for the caller of the interface to assume responsibility for access control policy when invoking the API. The access control policy is enforced entirely by the primary service provider. The ADI and contextual information required by the ADF as the basis for the authorisation decision is acquired by the primary service provider (or the ADF). The default Initiator ADI bound to the chain of execution of which invocation of the API forms a part is used as input to the ADF. The primary service provider must implement the AEF but may or may not implement the ADF.

7.4.1 Access Control Enforcing API / Security Selecting Caller

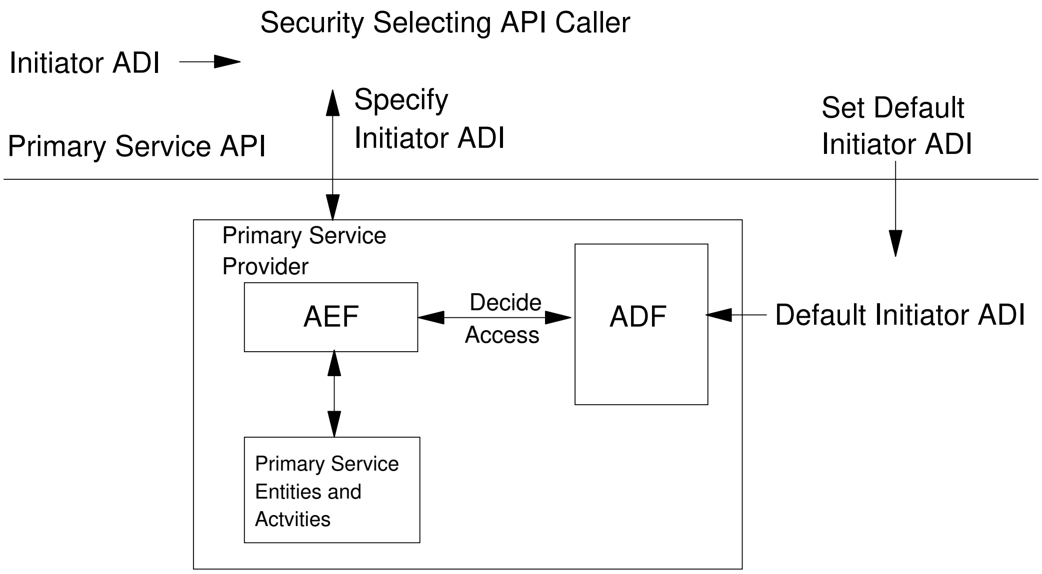

Figure 7-4 Access Control Enforcing API / Security Selecting Caller

In this case the caller of the API is responsible for selecting the Initiator ADI to be used in the access control decision. This initiator ADI may be passed directly to the primary service provider as parameters to the primary service API or may be provided via a separate security service to bind a particular set of initiator ADI with subsequent primary service invocations as a default. The primary service caller is responsible for managing the initiator ADI and their binding to a particular chain of service invocations.

7.4.2 Access Control Non-Enforcing API / Access Control Enforcing Caller

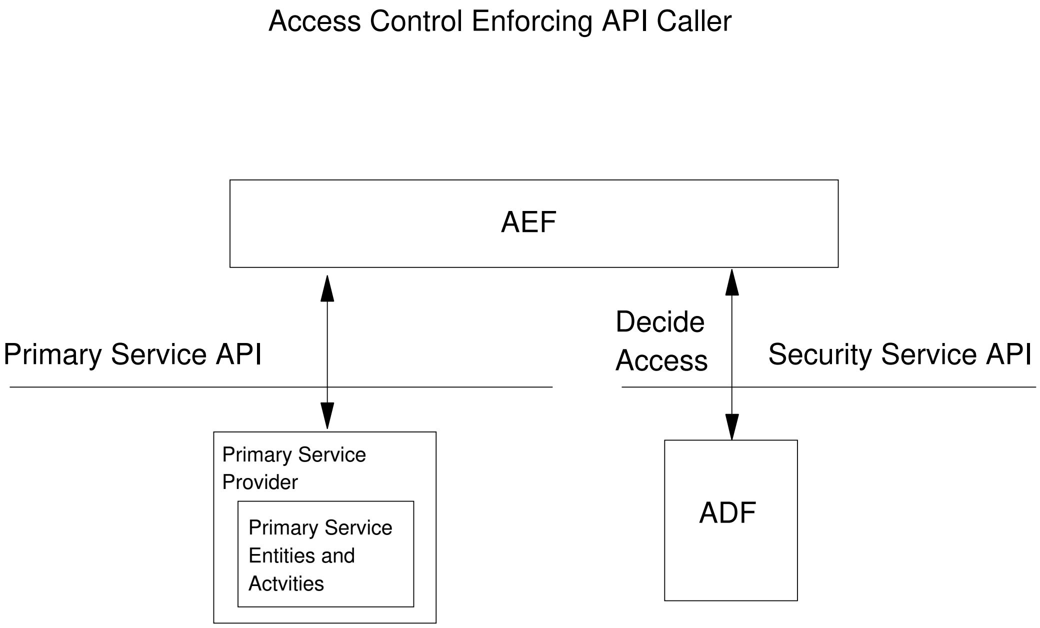

Figure 7-5 Access Control Non-Enforcing API / Access Control Enforcing Caller

In this case the primary service provider does not enforce access control on use of its services. The caller of the primary service API is responsible for enforcement of access control policy and therefore contains the AEF. The caller may also contain the ADF or may call a separate service that provides the ADF functionality.

In practice the primary service provider may be required to enforce an access control policy but on different entities to the caller of the primary service. This is to ensure that only callers that are trusted with the responsibility to enforce the appropriate access control policy may invoke the primary service APIs. An example would be a platform service interface that requires the calling process to possess an appropriate capability.

Note: This model may also apply if the caller is permitted to assume responsibility for access control policy, in which case the Primary Service Provider permits the caller to override the primary service provider access controls.

7.5 Security Accounting and Audit Service

7.5.1 Introduction

A security accounting and audit service provides for the detection, recording and analysis of security relevant events within an IT system. The events that require recording within a security domain will be defined by the security policy applicable to the domain.

This service description does not encompass a non-repudiation service except to the extent that the recording of security audit information supports such a service.

7.5.2 Basic Service Model

The security audit service comprises the following services:

Event Detection
The detection of events must occur within the software that provides the services that comprise the event. That is event detection is part of the primary service provision and must be included within the design of that service provision. This means that an API specification for the primary service shall include a requirement for the generation of an audit event.

Event Discrimination
Event detection has to be included as a fundamental part of a service provision. However, the security policy may permit the security authority, as represented by the Security Auditor, to configure the actions to be taken on the detection of an event based on local or current context (for example actions may be varied by time of day). The actions may be, for example: no action, generate alarm, record event, record event and generate alarm. A detected event therefore has to be screened against the actions configured for the event.

Audit Event Recording
An event that requires recording has to be appended to an audit trail. The information that requires recording includes the operational security context of the event together with such security domain management information that permits the record to be interpreted (for example, mappings of internal security attribute values to text strings). This may include the recording of cryptographic keys used for enciphering or sealing data.

For efficiency purposes the security domain management information may be recorded periodically (for example, once per day) and the security management state applicable to any individual event regenerated by tracing all audited changes to that information since the last record of the state.

Audit Archiving and Analysis
Audit trails require archiving securely. Analysis of audit trails may require the merging of several audit trails and the tracing of chains events.

Audit Alarm Event Collection and Alarm Generation
An event that is configured as an alarm event must either immediately initiate an alarm or be added to a collection of such events to subsequently trigger an alarm when a configured threshold of such events has occurred.

Audit Alarm Examination
Audit alarms that have been generated have to be examined and responded to by a representative of the security authority of the security domain, usually the Security Auditor.

Distributed Audit Services
Within this Security Framework an IT system comprises many security domains. An audit service will generally be supported as one or more specialist security domains providing all

or parts of the total audit service to all, or a subset of, the other security domains.

The individual events detected within each security domain will vary according to the function of the security domain. Additionally if activities span security domains under the ultimate control of different security authorities, then some mapping of security attribute syntax and semantics may be required.

In order to support the centralised archiving and analysis of audit trails in a distributed heterogeneous environment it is necessary to establish a common audit trail interchange format, including a common interpretation of audit events types.

A common interpretation of audit event types will be best achieved by considering security relevant operations at the abstract level. Examples of such abstract events may comprise:

- Association Creation and Termination
- Security Relevant Entity Creation and Deletion
- Security Management Actions (for example install, deinstall, disable, re-enable and so on).

An audit trail export function is then required to convert any local representation of an audit trail to such a common interchange and analysis format.

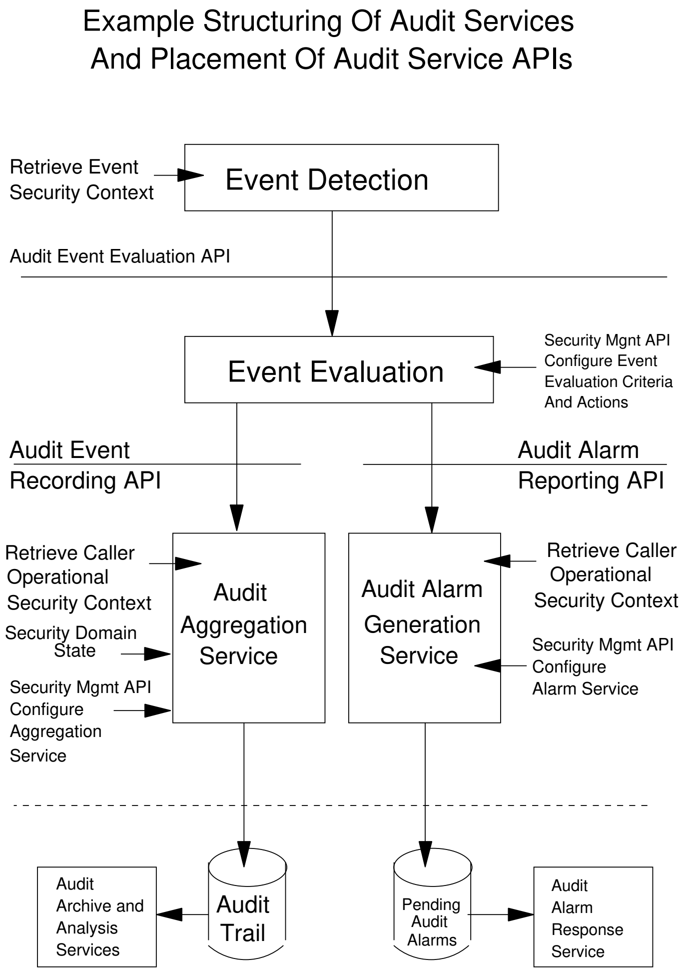

Figure 7-6 Basic Security Audit Service Model

7.5.3 Management Services

These are an initial list of possible services. More careful consideration needs to be given to their definition.

- Audit Event Discrimination Services
  - — Install Audit Event Discrimination Information
  - — Audit event discrimination information may be bound to different elements and functions that in combination determine the audit event discrimination configuration.
  - — Modify Audit Event Discrimination Information
  - — Deinstall Audit Event Discrimination Information
- Audit Event Alarm Management

- — Install Audit Event Alarm
- — This will require the specification of such information as threshold values, destination(s) of alarm message, contents of alarm message, action(s) to be taken as well as sending alarm (eg., disable user account and terminal).
- — Modify Audit Event Alarm
- — Deinstall Audit Event Alarm
- — Disable/Re-enable Audit Event Alarm

- Audit Trail Management
  - — Install Audit Trail
  - — Deinstall Audit Trail
  - — Disable/Re-enable Audit Trail
  - — Archive Audit Trail
  - — Export Audit Trail
- Audit Trail Analysis
  - — Import Audit Trail Archive
  - — Merge Audit Trail
  - — Produce Reports
- Operational Services
  - — Audit Event Discrimination Service
  - — Audit Event Recording Service
  - — Audit Alarm Reporting Service
  - — Suspend Primary Service Auditing.

7.5.4 Impact of Audit Service on Primary Service APIs

Security Audit Detecting API / Security Unaware Caller

If a primary service provider supports services that may be interpreted as security auditable events by the security policy then the primary service provider is required to incorporate the audit event detection functionality and to invoke the other audit services if an auditable event is detected. The specification of the primary service APIs should include the requirement to audit the API. The caller of the security event may be unaware of any auditing requirements. The primary service implementation, or the audit services it invokes, must derive all necessary security context information for inclusion in the audit record.

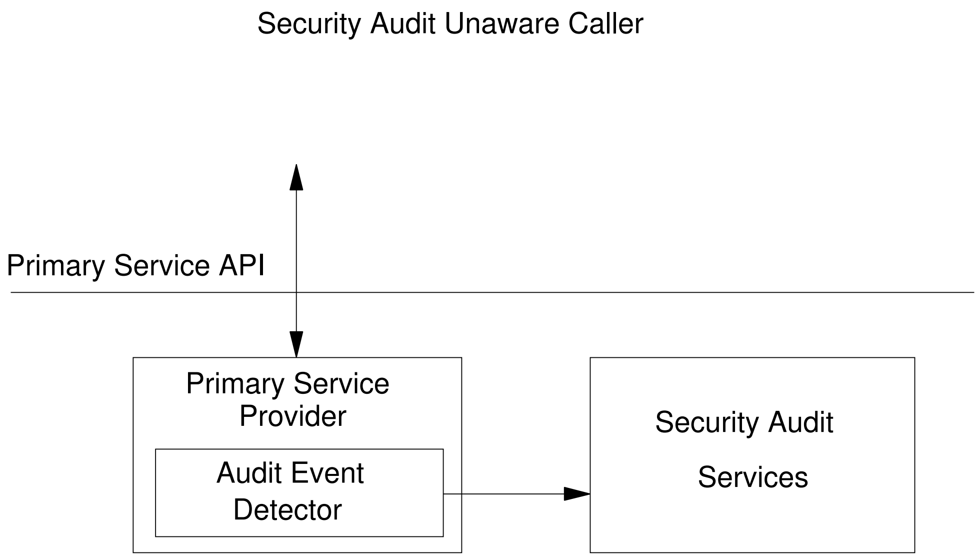

Figure 7-7 Security Audit Detecting API / Security Unaware Caller

Security Audit Detecting Caller

If the caller of a primary service API also provides services that are interpreted as auditable events then it is also required to incorporate audit event detection functionality and to invoke the other audit services. In this case the security policy may permit the caller of the primary service API to modify the event discrimination configuration for the primary service API to suppress the generation of any audit records or alarms from its use of the primary service.

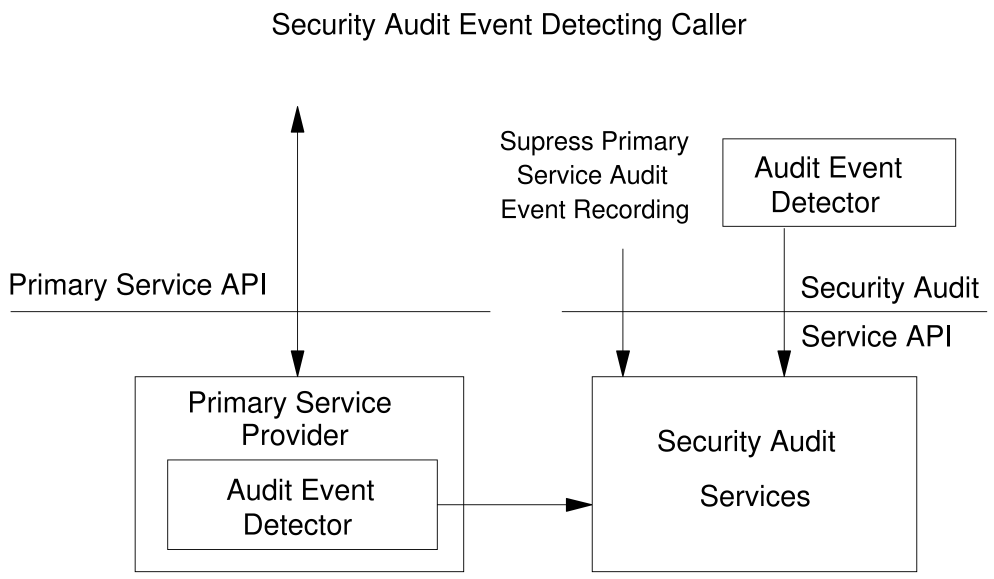

Figure 7-8 Security Audit Detecting Caller

7.6 Cryptographic Services

7.6.1 Introduction

Cryptography provides a basic mechanism that may be used to support higher level security services such as authentication of identities and data-origin, non-repudiation, and data separation and protection for the purposes of confidentiality and integrity.

7.6.2 Basic Model

To be added, including {Figure C07F09 Basic cryptographic Service Model}.

7.6.3 Management Services

Install Mechanism

This service installs an encryption mechanism for use within a system. This may include specification of the uses to which it may be applied and its assignment to support a specified communications security Quality of Service.

Deinstall Mechanism

This service removes a mechanism from a system.

Disable Mechanism

This service disables a mechanism and thus inhibits its use.

Re-enable Mechanism

This service re-enables a mechanism previously disabled.

Generate Key(s)

This service generates keys applicable to a particular mechanism. This may be a single key for a symmetric mechanism, or a key pair for an asymmetric mechanism.

Delete Key

This service deletes cryptographic keys.

Export Key

This service converts a cryptographic key, or set of keys, into a form suitable for distribution.

Import Key

This service imports a cryptographic key, or set of keys, previously exported and places them in secure storage.

Activate Key

This service makes a key a currently available key.

De-activate Key

This service removes a key from the currently available keys.

Archive Keys

This service archives a set of keys for retrospective use, for instance as part of audit analysis or if required for non-repudiation purposes.

7.6.4 Operational Services

Encipher Data

This service permits an entity to encipher data prior to, or as part of data transfer or storage for purposes of confidentiality.

Decipher Data

This service deciphers previously enciphered data. This may be performed as part of data transfer or retrieval or as a separate operation.

Seal Data

This service applies a signature to data prior to, or as part of data transfer for purposes of integrity and/or data-origin authentication.

Verify Data Seal

This service validates the integrity of the contents or authenticate the origin of a message against seal applied by the seal data operation.

7.6.5 Impact of Cryptographic Service on Primary Service APIs

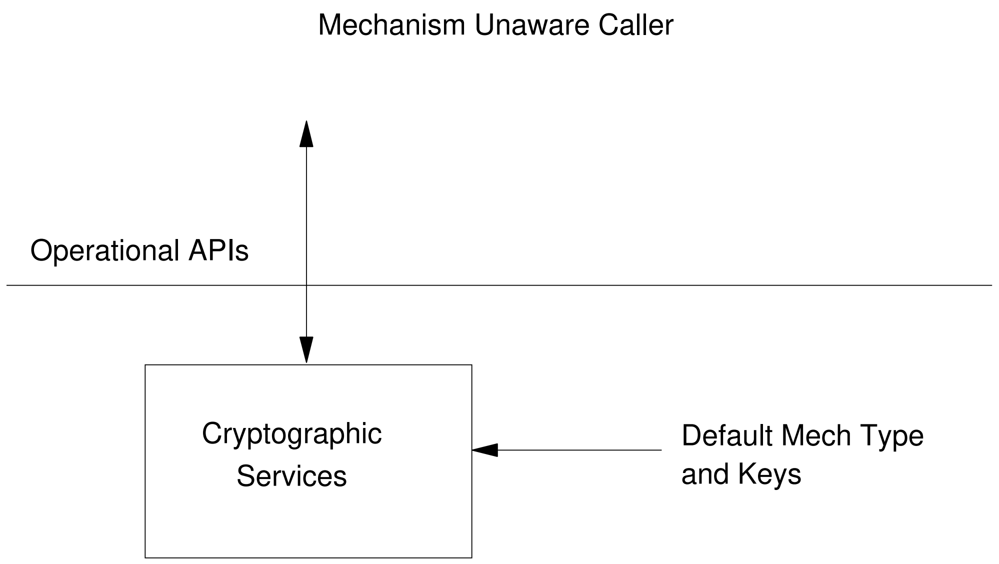

Figure 7-9 Cryptographic Mechanism Unware Caller

Note: Text to support this diagram has still to be added.

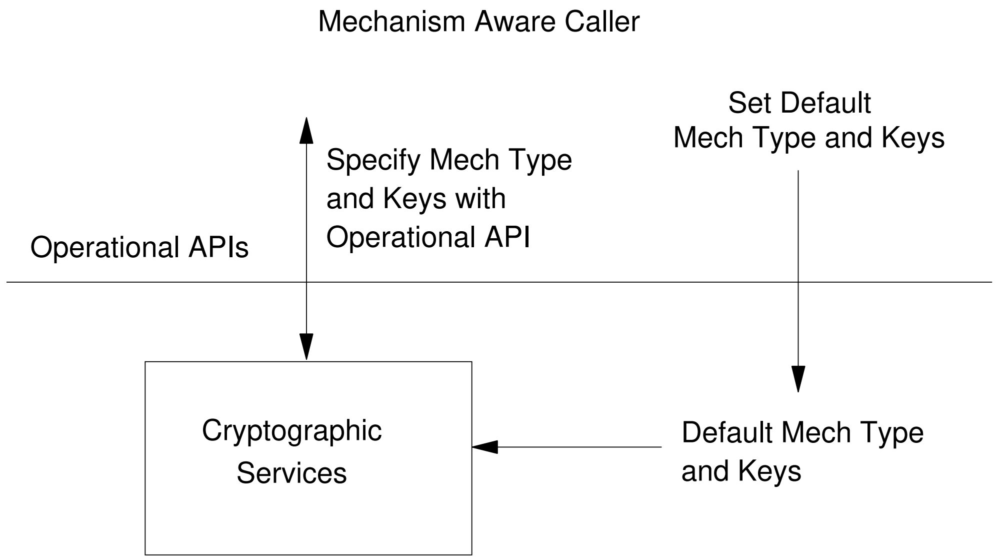

Figure 7-10 Cryptographic Mechanism Aware Caller

Note: Text to support this diagram has still to be added.

7.7 User Authentication

7.7.1 Introduction

Users are physically distinct from an IT System. In order to utilise the services of an IT system a user must create an association between himself as a principal and the IT system. In order to create such an association a user must identify himself to the system and have that identity authenticated. The process of user authentication may include the application of access control in that the success of the authentication process may depend upon the context of user interaction such as time of day, location of terminal or workstation etc.

7.7.2 Basic Service Model

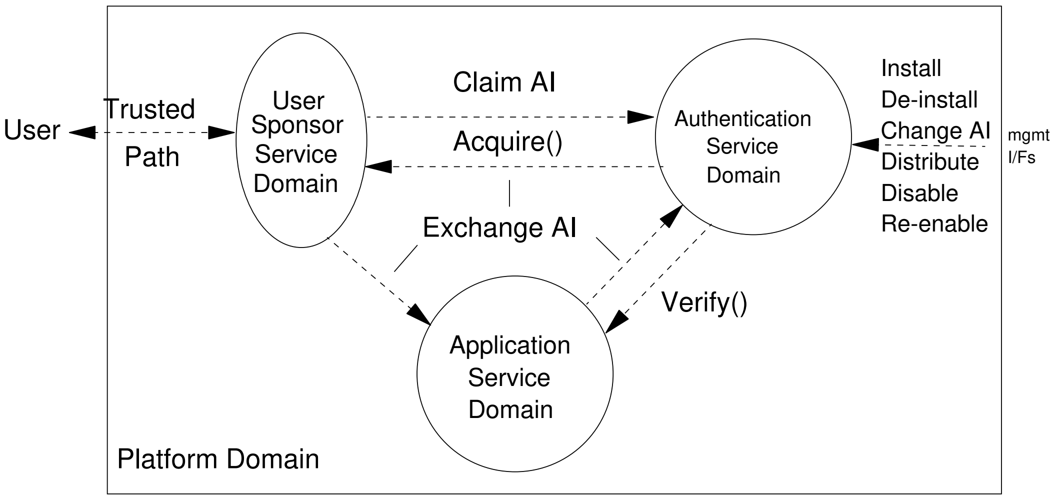

Figure 7-11 Basic User Authentication Service Model

Platform domains, through which users interact with with an IT system, are required to direct all initial user contact with the system to a user sponsor service domain. The system security policy may require that this interconnection is via a trusted path.

An authentication trusted path is a platform supported mechanism that provides the user with assurance that communication is only with the user sponsor security domain, and provides the user sponsor security domain with assurance that it is communicating directly with a user I/O device. The purpose is to provide assurance of the integrity and confidentiality of the communication.

The user sponsor service is responsible for promoting the user to the authentication service and, if successfully authenticated, to promote the authenticated user to other services of the system.

The user sponsor service solicits claim authentication information from the user and submits this, together with contextual information such as location if required by the system security policy, to the authentication service domain as part of an Acquire AI operation. If user authentication is successful, Exchange Authentication Information, possibly in the form of an authentication certificate, is returned by the authentication service to the user sponsor. Additional security information such as initial access control information and accountability information may also be returned.

The claim AI that may be used to support human user authentication include passwords, magnetic stripe card or smartcard, and physical characteristics such as fingerprint, hand written signature.

The exchange AI, maybe represented by an authentication certificate, can then be submitted by the user sponsor to an application service domain in assertion of the user principal's identity. The application service domain may invoke a verify operation to verify the exchange AI submitted by the user sponsor. See the section on Secure Association Service.

The exchange AI requires binding to the chain of execution initiated on behalf of the authenticated user and requires protection by the user sponsor, and any applications to which it is submitted as it provides assertion of the user principal's identity.

7.7.3 Extended Service Model

As well as having an identity authenticated, a user generally needs to obtain a set of initiator ACI to provide authorisation to access services within the system. The security policy may also require a user to be assigned an identity for accountability purposes and initial audit control information. These services may be combined with the authentication process, as in traditional local platform user authentication or the SESAME implementation of distributed security services, or accessed separately, as in DCE.

The authentication service and access control information services may be collocated and under a single security authority or may be under different security authorities. Note that an IT system may comprise many security domains each of which may have its own access control policy and access control services.

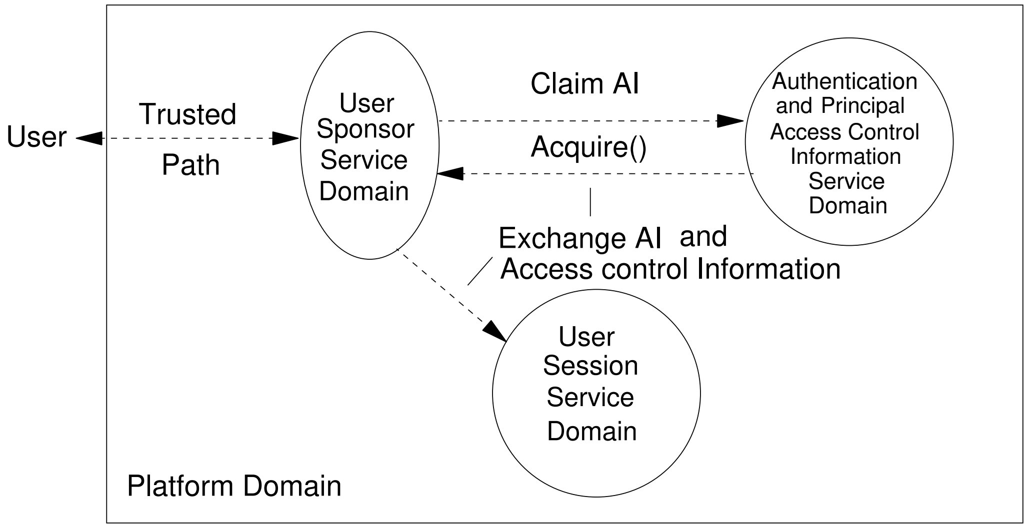

Figure 7-12 Extended User Authentication Service Model

In the traditional local platform user authentication model the exchange AI is returned as an authenticated identity together with a set of initiator ACI and audit control information, if the security policy requires, as a set of process attribute values. The user sponsor invokes a user session service domain to which the returned process attributes are applied to the user session process group leader and subsequently inherited by all processes invoked within the user session service domain. That is the authenticated control information are bound to the entity that represents the user within a user session service domain, the chain of execution within the session, and propagated by inheritance as process attributes.

7.7.4 Management Services

The management related services may include:

- Install User, which may comprise enrollment, testify, and validate
- De-install User, which may comprise invalidate, notify, unenroll
- Modify User AI
- Distribute User AI
- Disable and Re-enable User
- List Users, installed, disabled and so on.

These services will require both APIs and HCIs to be defined.

Other related services may provide facilities for attribute to name translation.

7.7.5 Operational Services

The operational related services may include:

- Solicit User Claim AI (HCI)
- Generate Exchange AI
- Verify Exchange AI.

7.8 Secure Association Service

7.8.1 Introduction

An essential part of the operation of an IT system is the formation of associations. The section on User I&A has described the secure formation of associations between Users as principals and the IT system. This section deals with the formation of secure associations between peer entities within the IT system, these entities being the sponsors within security domains.

The creation of a secure association is subject to the secure interaction policy configured for associations between the specific security domains by previous management action via the management related secure association services. The operational related secure association services actually create the association between the security domains. The establishment of a secure association depends upon other security services namely authentication, access control, audit and cryptographic services.

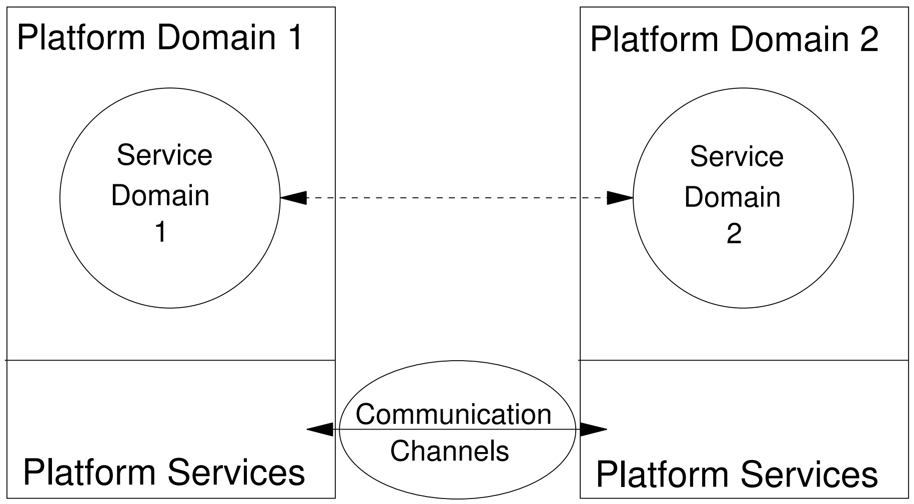

Figure 7-13 C7F14 Layered Nature of Associations Between Security Domains

The creation of an association between one set of entities within an IT system will of necessity require the formation of an association between entities that contain the primary associating entities. Thus an association between service domains also requires the formation of an association between platform domains (unless collocated within a single platform domain). The creation of the association will be subject to the secure interaction policies applicable to each of the constituent associations required. This equates to associations being formed, and secure interaction policy requirements applying to different layers (for example layer 7 and layer 4) of the OSI model.

7.8.2 Basic Service Model

The stages in the creation of a secure association are illustrated in the following figure. The numbers in the figure refer to the section numbers below.

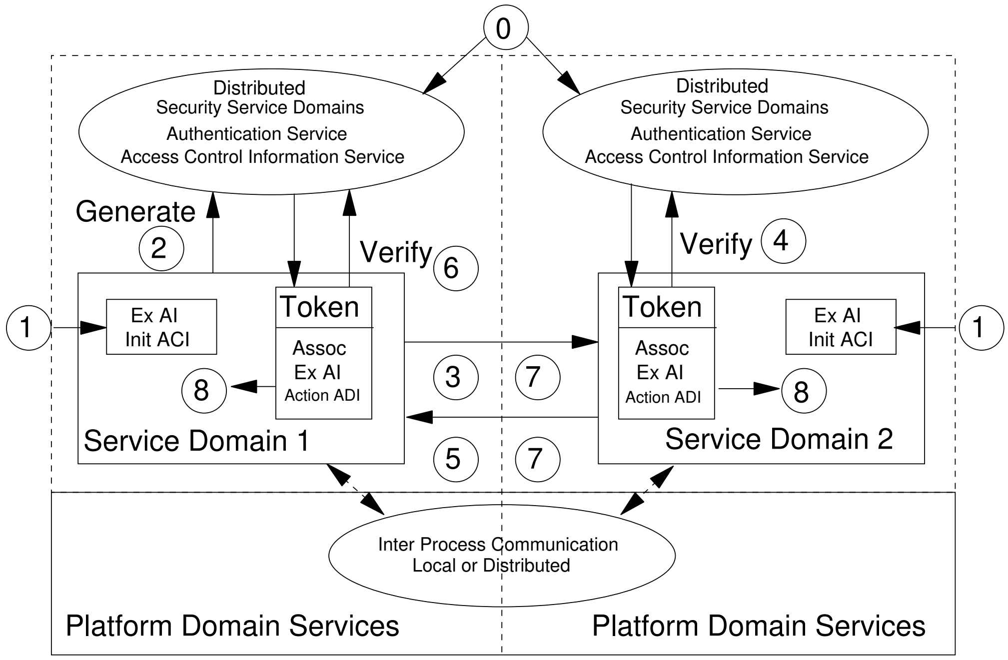

Figure 7-14 Basic Secure Association Service Model

Establish Secure Interaction Policy (0)

An association may only be created within the bounds of a secure interaction policy established for the formation of a particular association (or set of associations). The detailed rules of the secure interaction policy must be interpreted into a set of criteria for authentication, access control, audit, and cryptography services which need to be configured within the SMIBs of the interacting security domains by management action. This will include any requirements for an interdomain service to map the syntax and semantics of security attributes between their representations within each security domain, or to a common interchange format.

Initial Security Information (1)

A service domain which requires to form an association is required to possess some initiator security information. This will be either:

- security information such as exchange AI and/or initiator ACI appropriate to the principal to whom the service domain is providing service, or
- will be similar initiator security information that is appropriate to the service domain itself as a principal, or
- may be both.

Security information appropriate to a user as principal will have been derived by the user sponsor service as part of the user authentication process and/or acquisition of access control security information and propagated securely to the service domain now operating on behalf of that user principal. Security information appropriate to the service domain itself as a principal will have been established when the service domain was created from initial information associated with that service domain by the platform domain (for example, service domain identity, cryptographic key, or security certificate).

A service domain, providing services on behalf of many principals may possess more than one set of initiator security information and may select which of the sets of initiator security information is to be used for the creation of a particular association.

Generate Association Security Context (2)

The initiator invokes the generate service of the secure association service which in turn invokes other security services to establish:

- the authorisation of the initiator to form an association with the requested target with a specified or default security context under the secure interaction policy. Authorisation will be mediated upon the basis of the initiator security information specified, or taken as default.
- the security context of the association. The initiator may specify particular aspects of the security context, or may accept, or have imposed, specified defaults under the secure interaction policy. This will include specification of the Security Quality of Service (confidentiality and Integrity Services) and the strength of authentication mechanisms to be used.

The security context for the association (Action ADI) is returned to the initiator as a security certificate (maybe referenced by a handle).

Initial Security Information Exchange (3)

If the generation of an association security context is successful, the initiator sends the security token representing the requested security context for the association to the target service domain.

Verify Association Security Context (4)

The target service domain invokes the verify security context service of the secure association service, with the security token received from the initiator as a parameter, to verify the security context requested. The verify security context service will:

- verify the authenticity of the security token
- verify the identity of the initiator
- verify the authorisation for the target to form an association with the initiator with the requested security context. This will be based upon security information included within the security token and security information possessed by the target. This may include verification of the chain of delegation used to request this association.

If the association is authorised a further security token is returned by the secure association service to the target.

Responding Security Information Exchange (5)

The target sends the security token generated by the target verify security context operation to the initiator. This will convey security information regarding the acceptance or otherwise of the originally specified security context and levels of service and also include cryptographic keys to be used within the secure association.

Verify Target Security Context (6)

The initiator invokes the generate security context service of the secure association service with the security token returned by the target service domain. The secure association service will:

- verify the authenticity of the target security token
- verify the authenticated ID of the target if mutual peer-entity authentication was required.
- modify the security context if necessary.

Association Security Context Negotiation (7)

In some circumstances it may be necessary for the initiator and target to negotiate over some aspects of the security context establishment. In this case further calls to the operational related services of the secure association service and exchange of security tokens will occur.

Extract Security Information From Association Context (8)

In order to enforce security policy within their own security domains the target and/or the initiator may require to extract security information such as access control information, audit identity and control information from the security context associated with the association.

In addition the target and/or initiator may need to extract the cryptographic keys to be used as the session key and bind this to the association.

7.8.3 Management Services

Install Association

This service will install a permitted initiator/target pairing together with associated permitted security context as defined by the secure interaction policy. This will include the specification of aspects such as the quality of security service to be used as part of the association and the types of cryptographic mechanisms to be employed. It may also include the configuration of an interdomain service to map the semantics (and maybe the syntax) of a set of security attributes within another security domain into the semantics (and syntax) of the the security authority's own security domain.

Modify Association

This service will modify the details of a permitted association.

Uninstall Association

This service will delete a permitted association.

Disable

This service will set an association state such that it is unavailable for use.

Re-enable

This service will terminate the disabled state of an association.

7.8.4 Operational Services

Generate Association Security Context

An initiator uses this service to request a security context to be generated for use in creating an association with a specified target, or to verify a security context returned by a target.

Verify Association Security Context

A target uses this service to verify an association security context provided by an initiator.

Release Association Security Context

On termination of an association this service releases and deletes a security context created by a generate or verify service.

Extract Security Information

This service permits an initiator or target to extract security information from an association security context (for example security token) for use by other security services, for example AI, ACI, Audit Information.

7.8.5 Services Used Within An Association

Bind ACI to Message (Data)

An association security context may permit the transfer of data of varying access control attributes. This service permits an entity to specify the ADI to be bound to data to be transferred via an association.

Message Protection Services

These services invoke confidentiality and integrity services for specific messages and data to provide protection during transmission via an association. The specific cryptographic mechanisms used to provide these services will have been determined as part of the security context established for the association. (See Cryptographic Operations Services.)

The concept of Quality of Service applied to operations via associations represents a service to support a trade off of performance against strength of security mechanism, and also support for protection of higher strength mechanisms through their reduced use and hence reduced exposure of encrypted data.

7.8.6 Impact of Secure Association Service on Primary Service APIs

Security Enforcing API / Security Unaware Caller

Security enforcing APIs provide no support for the caller of the interface to assume responsibility for any aspect of security association policy enforcement when calling that primary service API. The policy will be enforced entirely by the primary service provider. It is therefore suitable for supporting Security Unaware Callers.

A Security Enforcing API requires that the default security context within which the primary service is invoked has been previously established, that is the initiator security information has been generated and made available to the primary service provider either by inheritance from a previous process or from previous actions of the current process. In addition, the Security Quality of Service has been previously defined, either by management action as a default or via an operational API.

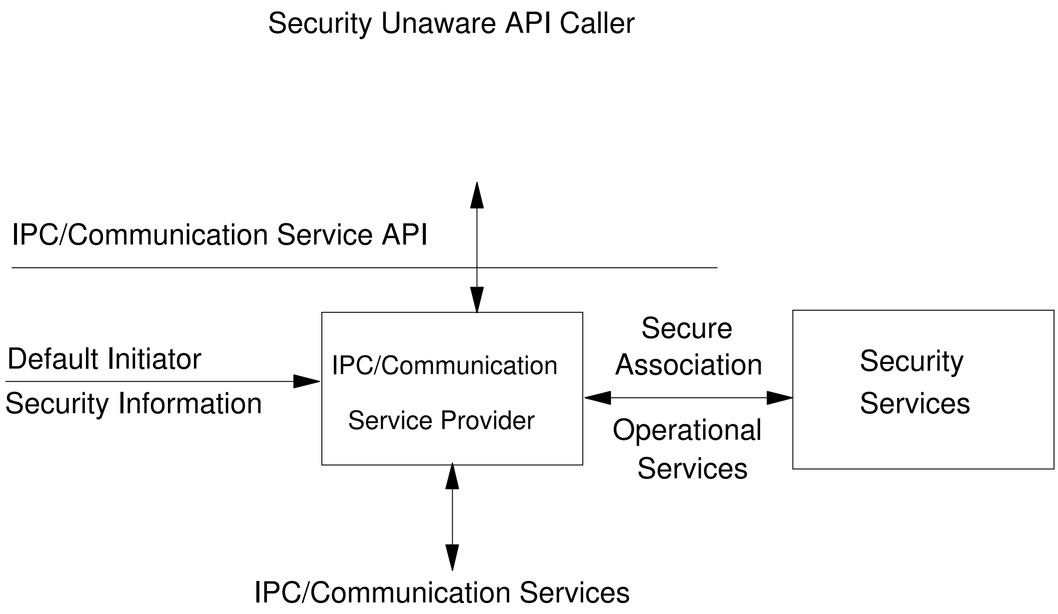

Figure 7-15 Security Enforcing API / Security Unaware Caller

Security Quality of Service Supporting API /Simple Security Selecting Caller

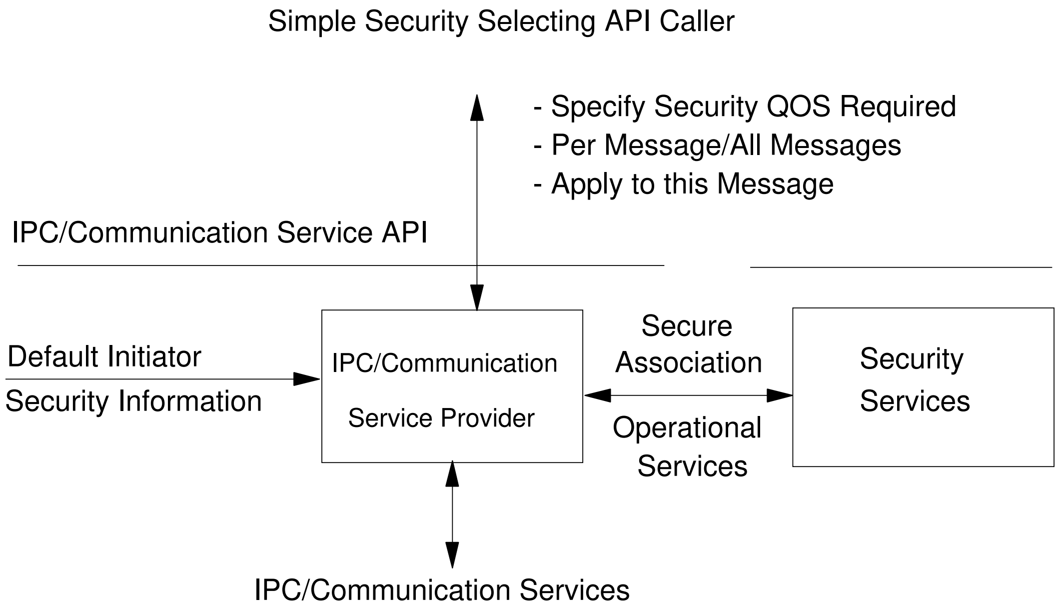

Figure 7-16 Security QOS Supporting API / Simple Security Selecting Caller

Security Quality of Service Supporting APIs provides support for the caller of the primary service interface to select the strength of mechanism (Security Quality of Service) to be applied to information exchanged via an association.

The specification of the Security Quality of Service requires provision within the primary service APIs that create associations. This may include an indication of whether this applies to all messages or only messages for which the service is specified.

If Security Quality of Service is to be specified on a per message basis then a parameter specifying this needs to be included in each API that creates a message.

Security Context Supporting API / Security Selecting Caller

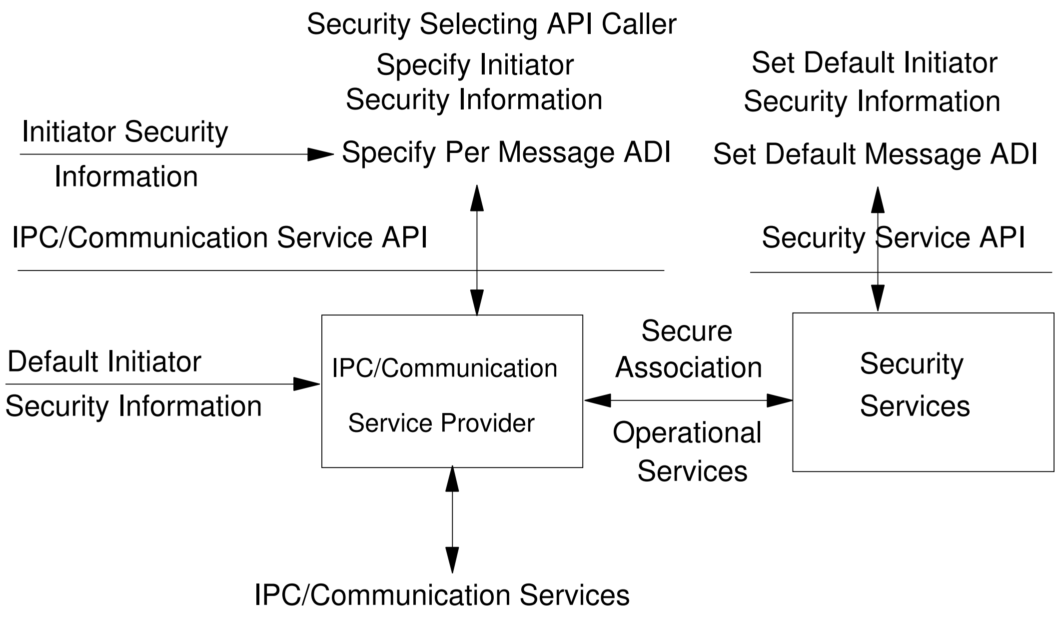

Figure 7-17 Figure Title

Security context supporting APIs support the specification of a security context for the creation of an association and/or the binding of security information to information exchanged via an association. The caller of Security Context Supporting APIs is responsible for manipulating security information. That is, it is responsible for manipulating the initiator security information to be used in the establishment of the security context of a secure association or for the specification of the security information to be bound with information exchanged via the association.

The specification of the security information may be supported in band, that is parameters to the primary service API, or out of band, that is as separate interfaces that set default values to be applied to subsequent primary service API invocations.

Security Non-Enforcing API / Security Integrating Caller

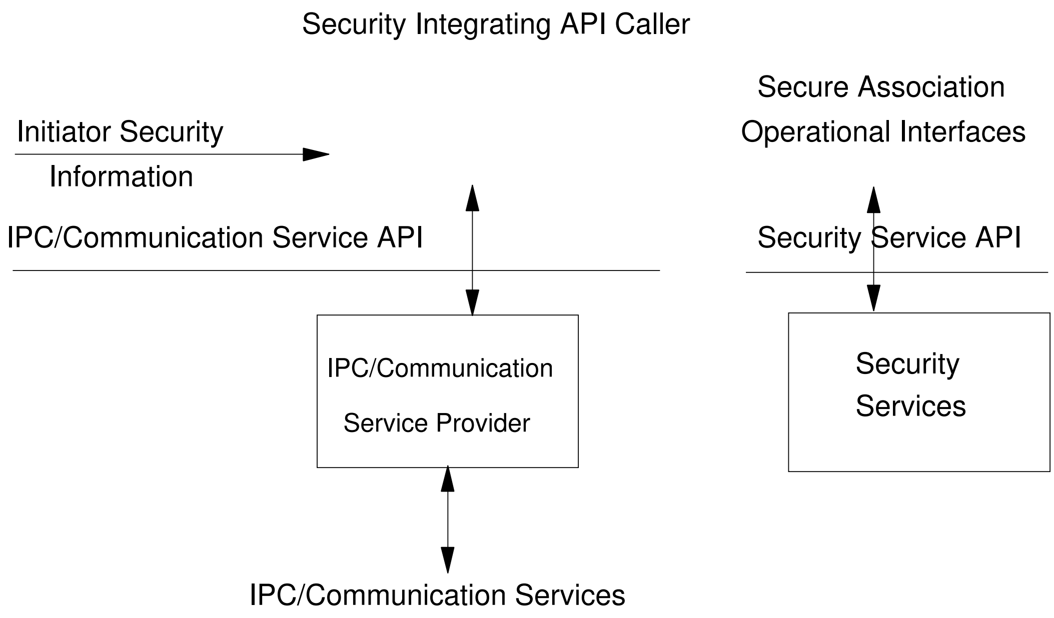

Figure 7-18 Security Non-Enforcing API / Security Integrating Caller

A primary service provider presenting Security Non-Enforcing APIs provides no support for secure association security services. It is therefore the responsibility of the caller of such APIs to implement the Secure Association Services. The implementation of the primary service provider would be required to restrict the use of such APIs to callers that were authorised as responsible for the implementation of the Secure Association Service. That is possessed a capability assigned by the security authority to invoke these APIs.

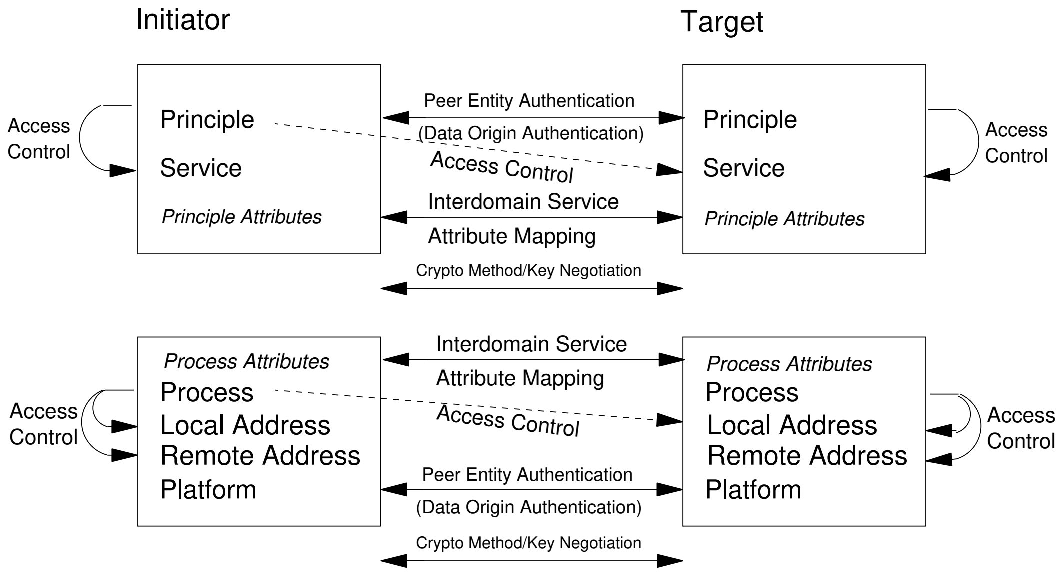

Figure 7-19 Example Disposition of Security Services in Association Establishment

The figure above illustrates how various security services within different security domains may be deployed in supporting a secure association service.

Note: Text to support this diagram has still to be added.

Glossary

The glossary of terms presented below is an initial list. Some of the terms are included to prompt consideration of further detail in the main body of the document. Some terms may need to ba added.

access control
The prevention of unauthorized use of a resource including the prevention of use of a resource in an unauthorized manner.[ISO 7498-2:1989]

access control certificate
ADI in the form of a security certificate[ISO/IEC CD 10181-3 Oct 1991]

access control decision function (ADF)
A specialized function that makes access control decisions by applying access control policy rules to a requested action, ACI (of initiators, targets, actions, or that retained from prior actions), and the context in which the request is made.[ISO/IEC CD 10181-3 Oct 1991]

access control decision information (ADI)
The portion (possibly all) of the ACI made available to the ADF in making a particular access control decision.[ISO/IEC CD 10181-3 Oct 1991]

access control enforcement function (AEF)
A specialized function that is part of the access path between an initiator and a target on each access and enforces the decisions made by the ADF.[ISO/IEC CD 10181-3 Oct 1991]

access control information (ACI)
Any information used for access control purposes, including contextual information.[ISO/IEC CD 10181-3 Oct 1991]

access control list
A list of entities, together with their access rights which are authorized to have access to a resource.[ISO 7498-2:1989]

access control policy
The set of rules that define the conditions under which an access may take place.[IS0/IEC CD 10181-3 Oct 1991]

accountability
The property that ensures that the actions of an entity may be traced to that entity.[ISO 7498-2:1989]

action
The operations and operands that form part of an attempted access.[ISO/IEC CD 10181-3 Oct 1991]

action access control decision information (action ADI)
Action decision information associated with the action.[ISO/IEC CD 10181-3 Oct 1991]

active threat
The threat of a deliberate unauthorized change to the state of the system.

Note: Examples of security-relevant active threats may be: modification of messages, replay of messages, insertion of spurious messages, masquerading as an authorized entity and denial of service.[ISO 7498- 2:1989]

alarm collector function

A function that collects the security alarm messages, translates them into security alarm records, and writes them to the security alarm log.[ISO/IEC CD 10181-7 Aug 1992]

alarm examiner function

A function that interfaces with a security alarm administrator.[ISO/IEC CD 10181-7 Aug 1992]

application program interface

The interface between the application software and the application platform, across which all services are provided.

The application program interface is primarily in support of application portability, but system and application interoperability are also supported by a communication API.[POSIX.0/D15 Jun 1992]

assertion

Explicit statement in a System Security policy that security measures in one security domain constitute an adequate basis for security measures (or lack of them) in another.[CESG Memorandum No.1 Issue 1.2 Oct 1992]

association-security-state

The collection of information which is relevant to the control of communications security for a particular application-association.[ISO/IEC DIS 10745 May 1992]

asymmetric authentication method

Method for demonstrating knowledge of a secret, in which not all authentication information is shared by both entities.[ISO/IEC DIS 10181-2 Jul 1991]

audit

See Security Audit[ISO 7498-2:1989]

audit authority

The manager responsible for defining those aspects of a security policy applicable to maintaining a security audit.[ISO/IEC CD 10181-7 Aug 1992]

audit recorder function

A function that records the security-relevant messages in a security audit trail.[ISO/IEC CD 10181-7 Aug 1992]

audit trail

See Security Audit Trail[ISO 7498-2:1989]

audit trail analyzer function

A function that checks a security audit trail in order to produce, if appropriate, security alarm messages.[ISO/IEC CD 10181-7 Aug 1992]

audit trail archiver function

A function that archives a part of the security audit trail.[ISO/IEC CD 10181- 7 Aug 1992]

audit trail collector function

A function that collects individual audit trail records into a security audit trail.[ISO/IEC CD 10181-7 Aug 1992]

audit trail examiner function

A function that builds security reports out of one or more security audit trails.[ISO/IEC CD 10181-7 Aug 1992]

audit trail provider function
A function that provides security audit trails according to some criteria.[ISO/IEC CD 10181-7 Aug 1992]

authenticated identity
An identity of a principal that has been assured through authentication.[ISO/IEC DIS 10181-2 Jul 1991]

authentication
See data origin authentication, and peer entity authentication. [ISO 7498-2:1989]

authentication certificate
Authentication information in the form of a security certificate which may be used to assure the identity of an entity guaranteed by an authentication authority.[ISO/IEC DIS 10181-2 Jul 1991]

authentication exchange
A sequence of one or more transfers of exchange authentication information (AI) for the purposes of performing an authentication.[ISO/IEC DIS 10181-2 Jul 1991]

authentication information
Information used to establish the validity of a claimed identity.[ISO 7498- 2:1989]

authentication initiator
The entity which starts an authentication exchange.[ISO/IEC DIS 10181-2 Jul 1991]

authorization
The granting of rights, which includes the granting of access based on access rights.[ISO 7498-2:1989]

authorization policy
A set of rules, part of an access control policy, by which access by security subjects to security objects is granted or denied. An authorization policy may be defined in terms of access control lists, capabilities or attributes assigned to security subjects, security objects or both.[ECMA TR/46 Jul 1988]

availability
The property of being accessible and useable upon demand by an authorized entity.[ISO 7498-2:1989]

capability
A token used as an identifier for a resource such that possession of the token confers access rights for the resource.[ISO 7498-2:1989]

ciphertext
Data produced through the use of encipherment. The semantic content of the resulting data is not available.

Note: Ciphertext may itself be input to encipherment, such that super- enciphered output is produced.[ISO 7498-2:1989]

claim authentication information (claim AI)
Information used by a claimant to generate exchange AI needed to authenticate a principal.[ISO/IEC DIS 10181-2 Jul 1991]

claimant
An entity which is or represents a principal for the purposes of authentication. A claimant includes the functions necessary for engaging in authentication exchanges on behalf of a principal.[ISO/IEC DIS 10181-2 Jul 1991]

cleartext
Intelligible data, the semantic content of which is available.[ISO 7498- 2:1989]

confidentiality
The property that information is not made available or disclosed to unauthorized individuals, entities, or processes.[ISO 7498-2:1989]

contextual information
Information derived from the context in which an access is made (e.g., time of day).[ISO/IEC CD 10181-3 Oct 1991]

corporate security policy
The set of laws, rules and practices that regulate how assets including the sensitive information are managed, protected and distributed within a user organisation.[ITSEC Ver 1.2 1991]

countermeasure
See Security Measure.[]

credentials
Data that is transferred to establish the claimed identity of an entity.[ISO 7498-2:1989]

cryptanalysis
The analysis of a cryptographic system and/or its inputs and outputs to derive confidential variables and/or sensitive data including cleartext.[ISO 7498- 2:1989]

cryptographic checkvalue
Information which is derived by performing a cryptographic transformation (see cryptography) on a data unit.

Note: The derivation of the checkvalue may be performed in one or more steps and is a result of a mathematical function of the key and data unit. It is usually used to check the integrity of a data unit.[ISO 7498-2:1989]

cryptography
The discipline which embodies principles, means, and the methods for the transformation of data in order to hide its information content, prevent its undetected modification and/or prevent its unauthorized use.

Note: Cryptography determines the methods used in encipherment and decipherment. An attack on a cryptographic principle, means, or methods is cryptanalysis.[ISO 7498-2:1989]

data integrity
The property that data has not been altered or destroyed in an unauthorized manner.[ISO 7498-2:1989]

decipherment
The reversal of a corresponding reversible encipherment.[ISO 7498-2:1989]

decryption
See Decipherment.[ISO 7498-2:1989]

denial of service
The prevention of authorized access to resources or the delaying of time- critical operations.[ISO 7498-2:1989]

digital fingerprint
A characteristic of a data item, such as a cryptographic checkvalue or the result of performing a one-way hash function on the data, that is sufficiently peculiar to the data

item that it is computationally infeasible to find another data item that will possess the same characteristics.[ISO/IEC CD 10181-1:Dec 1992]

digital signature
Data appended to, or a cryptographic transformation (see cryptography) of, a data unit that allows a recipient of the data unit to prove the source and integrity of the data unit and protect against forgery e.g. by the recipient.[ISO 7498-2:1989]

discretionary access control
Methods of restricting access to objects or other resources based primarily on the instructions of arbitrary unprivileged users.[FC Ver 1.0 Dec 1992]

distributed application
A set of information processing resources distributed over one or more open systems which provides a well defined set of functionality to (human) users, to assist a given (office) task.[ECMA TR/46 Jul 1988]

encapsulated subsystem
A collection of procedures and data objects that is protected in a domain of its own so that the internal structure of a data object is accessible only to the procedures of the encapsulated subsystem that the procedures may be called only at designated domain entry points. Encapsulated subsystem, protected subsystem, and protected mechanisms of the TCB are terms that may be used interchangeably.[FC Ver 1.0 Dec 1992]

encipherment
The cryptographic transformation of data (see cryptography) to produce ciphertext.

Note: Encipherment may be irreversible, in which case the corresponding decipherment process cannot feasibly be performed.[ISO 7498-2:1989]

encryption
See encipherment.[ISO 7498-2:1989]

end-to-end encipherment
Encipherment of data within or at the source end system, with the corresponding decipherment occurring only within or at the destination end system. (See also link-by-link encipherment.)[ISO 7498-2:1989]

exchange authentication information (exchange AI)
Information exchanged between a claimant and a verifier during the process of authenticating a principal.[ISO/IEC DIS 10181-2 Jul 1991]

identity-based security policy
A security policy based on the identities and/or attributes of users, a group of users, or entities acting on behalf of the users and the resources/objects being accessed.[ISO 7498-2:1989]

distinguishing identifier
Data which unambiguously distinguishes an entity in the authentication process. This [Recommendation | Part of this International Standard] requires that such an identifier be unambiguous at least within a security domain.[ISO/IEC DIS 10182-2]

initiator
An entity (e.g., human user or computer based entity) that attempts to access other entities.[ISO/IEC CD 10181-3 Oct 1991]

initiator access control decision information (initiator ADI)
ADI associated with the initiator.[ISO/IEC CD 10181-3 Oct 1991]

initiator access control information (initiator ACI)
Access control information relating to the initiator.[ISO/IEC CD 10181-3 Oct 1991]

integrity
See Data Integrity.[ISO 7498-2:1989]

key
A sequence of symbols that controls the operations of encipherment and decipherment.[ISO 7498-2:1989]

key management
The generation, storage, distribution, deletion, archiving and application of keys in accordance with a security policy.[ISO 7498-2:1989]

masquerade
The pretence by an entity to be a different entity.[ISO 7498-2:1989]

non-discretionary access control
Means of restricting access to objects based largely on administrative actions.[FC Ver 1.0 Dec 1992]

off-line authentication certificate
A particular form of authentication information binding an entity to a cryptographic key, certified by a trusted authority, which may be used for authentication without directly interacting with the authority.[ISO/IEC DIS 10181-2 Jul 1991]

on-line authentication certificate
A particular form of authentication information,certified by a trusted authority, which may be used for authentication following direct interaction with the authority[ISO/IEC DIS 10181-2 Jul 1991]

organizational security policy
Set of laws, rules, and practices that regulate how an organization manages, protects, and distributes sensitive information.[FC Ver 1.0 Dec 1992]

password
Confidential authentication information, usually composed of a string of characters.[ISO 7498-2:1989]

peer-entity authentication
The corroboration that a peer entity in an association is the one claimed.[ISO 7498-2:1989]

physical security
The measures used to provide physical protection of resources against deliberate and accidental threats.[ISO 7498-2:1989]

policy
See: Security Policy.[ISO 7498-2:1989]

principal
An entity whose identity can be authenticated.[ISO/IEC DIS 10181-2 Jul 1991]

privacy
The right of individuals to control or influence what information related to them may be collected and stored and by whom and to whom that information may be disclosed.

Note: because this term relates to the right of individuals, it cannot be very precise and its use should be avoided except as a motivation for requiring security.[ISO 7498-2:1989]

private key
A key used in an asymmetric algorithm. Possession of this key is restricted, usually to only one entity.[ISO/IEC CD 10181-1:Dec 1992]

public key
The key, used in an asymmetric algorithm, that is publicly available.[ISO/IEC CD 10181-1:Dec 1992]

repudiation
Denial by one of the entities involved in a communication of having participated in all or part of the communication.[ISO 7498-2:1989]

rule-based security policy
A security policy based on global rules imposed for all users. These rules usually rely on a comparison of the sensitivity of the resources being accessed and the possession of corresponding attributes of users, a group of users, or entities acting on behalf of users.[ISO 7498-2:1989]

seal
A cryptographic checkvalue that supports integrity but does not protect against forgery by the recipient (i.e., it does not support non-repudiation). When a seal is associated with a data element, that data element is

secret key
In a symmetric cryptographic algorithm the key shared between two entities.[ISO/IEC CD 10181-1:Dec 1992]

secure interaction policy
The common aspects of the security policies in effect at each of the communicating application-processes.[ISO/IEC DIS 10745 May 1992]

security architecture
A high level description of the structure of a system, with security functions assigned to components within this structure.[CESG Memorandum No.1 Issue 1.2 Oct 1992]

security attribute
A security attribute is a piece of security information which is associated with an entity in a distributed system.[ECMA-138 Dec 1989]

security audit
An independent review and examination of system records and activities in order to test for adequacy of system controls, to ensure compliance with established policy and operational procedures, to detect breaches in security, and to recommend any indicated changes in control, policy and procedures.[ISO 7498-2:1989]

security audit message
A message generated following the occurrence of an auditable security-related event.[ISO/IEC CD 10181-7 Aug 1992]

security audit record
A single record in a security audit trail corresponding to a single security- related event.[ISO/IEC CD 10181-7 aug 1992]

security audit trail
Data collected and potentially used to facilitate a security audit.[ISO 7498- 2:1989]

security auditor
An individual or a process allowed to have access to the security audit trail and to build audit reports.[ISO/IEC CD 10181-7 Aug 1992]

security certificate

A set of security relevant data which is protected by integrity and data origin authentication from an issuing security authority, and includes an indication of a time period of validity.

Note: All certificates are deemed to be security certificates (see the relevant definitions in 7498-2). The term "security certificate" is adopted in order to avoid terminology conflicts with [X.509 | ISO 9594- 8] (i.e. the directory authentication standard).[ISO/IEC CD 10181-1:Dec 1992]

security domain

A set of elements, a security policy, a security authority and a set of security relevant activities in which the set of elements are subject to the security policy, administered by the security authority, for the specified activities.[ISO/IEC CD 10181-1:Dec 1992]

security event manager

An individual or process allowed to specify and manage the events which may generate a security message and to establish the action(s) to be taken for each security message type.[ISO/IEC CD 10181-7 Aug 1992]

security label

The marking bound to a resource (which may be a data unit) that names or designates the security attributes of that resource.

Note: The marking may be explicit or implicit.[ISO 7498-2:1989]

security object

An entity in a passive role to which a security policy applies.[ECMA TR/46 Jul 1988]

security policy

The set of criteria for the provision of security services (see also identity- based and rule-based security policy.)

Note: A complete security policy will necessarily address many concerns which are outside the scope of OSI.[ISO 7498-2:1989]

security service

A service, provided by a layer of communicating open systems, which ensures adequate security of the systems or of data transfers.[ISO 7498-2:1989]

security state

State information that is held in an open system and which is required for the provision of OSI security services.[ISO/IEC DIS 10745 May 1992]

security token

A set of security relevant data which is protected by integrity and data origin authentication from a source which is not considered a security authority.[ISO/IEC CD 10181-1:Dec 1992]

sensitivity

The characteristic of a resource which implies its value or importance, and may include its vulnerability.[ISO 7498-2:1989]

separation

The concept of keeping information of different security classes apart in a system.

Note: Separation may be implemented by temporal, physical, logical or cryptographic techniques.[CESG Memorandum No.1 Issue 1.2 Oct 1992]

signature

See Digital Signature.[ISO 7498-2:1989]

strength of mechanism
An aspect of the assessment of the effectiveness of a Target of Evaluation, namely the ability of its security mechanisms to withstand direct attack against deficiencies in their underlying algorithms, principles and properties.[ITSEC Ver 1.2 1991]

subject
Abbreviation of security subject.[ECMA TR/46 Jul 1988]

symmetric authentication method
Method for demonstrating knowledge of a secret, in which both entities share a common authentication information.[ISO/IEC DIS 10181-2 Jul 1991]

system security function
A capability of an open system to perform security-related processing.[ISO/IEC DIS 10745 May 1992]

target
An entity to which access may be attempted.[ISO/IEC CD 10181-3 Oct 1991]

target access control decision information (target ADI)
ADI associated with the target.[ISO/IEC CD 10181-3 Oct 1991]

target access control information (target ACI)
Access control information relating to the target.[ISO/IEC CD 10181-3 Oct 1991]

threat
A potential violation of security.[ISO 7498-2:1989]

threat
an action or event that might prejudice security.[ITSEC Ver 1.2 1991]

traffic analysis
The inference of information from observation of traffic flows (presence, absence, amount, direction and frequency).[ISO 7498-2:1989]

traffic flow confidentiality
A confidentiality service to protect against traffic analysis.[ISO 7498- 2:1989]

traffic padding
The generation of spurious instances of communication, spurious data units and/or spurious data within data units.[ISO 7498-2:1989]

trap door
A hidden software or hardware mechanism that permits system protection mechanisms to be circumvented. It is activated in some non-apparent manner (e.g., special "random" key sequence at a terminal).[TCSEC Dec 1985]

trojan horse
Computer program containing an apparent or actual useful function that contains additional (hidden) functions that allow unauthorized collection, falsification or destruction of data. [NSTISSI 4009][FC Ver 1.0 Dec 1992]

trust
A relationship between two elements, a set of activities and a security policy in which element X trusts element Y if and only if X has confidence that Y will behave in a well defined way (with respect to the activities) that does not violate the given security policy.[ISO/IEC CD 10181-1:Dec 1992]

trusted functionality

That which is perceived to be correct with respect to some criteria, e.g., as established by a security policy.[ISO 7498-2:1989]

trusted path

Mechanism by which a person using a terminal can communicate directly with the TCB. [NSTISSI 4009]

Note: Trusted path can only be activated by the person or the TCB and cannot be imitated by untrusted software.[FC Ver 1.0 Dec 1992]

trusted third party

A security authority or its agent, trusted by other entities with respect to security-related activities.[ISO/IEC CD 10181-1:Dec 1992]

verification authentication information (verification AI)

Information used by a verifier to verify an identity claimed through exchange AI.[ISO/IEC DIS 10181-2 Jul 1991]

verifier

An entity which is or represents the entity requiring an authenticated identity. A verifier includes the functions necessary for engaging in authentication exchanges.[ISO/IEC DIS 10181-2 Jul 1991]

virus

Self replicating, malicious program segment that attaches itself to an application or other executable system component and leaves no external signs of its presence. [NSTISSI 4009][FC Ver 1.0 Dec 1992]

vulnerability

Weakness in an information system or components (e.g., system security procedures, hardware design, internal controls) that could be exploited to produce an information-related misfortune.[FC Ver 1.0 Dec 1992]