POSIX Security functionality

Introduction

This paper gives initial thoughts on the taxonomy and categorisation of security services provided by distributed systems, and attempts to integrate the security services specified in ISO 7498-2, ECMA TR-46, POSIX 1003.6 and the X/Open DSF. In particular a consistent view of security services across the communication models adopted by ISO and those appropriate to the platform view of security in POSIX is required. The model proposed here gives an initial taxonomy based on access control functionality based on attributes of the subject and target object, augmented by the concept of barrier defences in space and time.

Basic Concepts

The model is based on the traditional reference monitor model of active subjects (which may themselves be objects) accessing objects within the system environment. The access by subject to object is constrained by a reference monitor which performs access control based on the attributes of the subject and object. The interpretation of the attributes of the subject may involve appropriate mappings between the security/administrative domain in which the subject attributes are defined, and the domain in which the object resides.

The access control to objects is both spatial and temporal. In particular the traditional discretionary and mandatory access control mechanisms specified by the Orange book are spatial access controls which constrain access to an object in space irrespective of the time of access. Integrity controls are primarily temporal access controls which constrain access to an object in time. That is, any attempts to access an object are serialised to ensure that simultaneous access does not occur. Hybrid controls may also exist which constrain access in both time and space.

To enforce access control it is insufficient merely to provide for access control decisions, but is also necessary to segregate objects from subjects and other objects, and to ensure that all access to the object is arbitrated by the reference monitor. There is therefore a requirement for object segregation both in space and time. Spatial object segregation represents the traditional reliance placed on the platform to correctly separate the address/data space occupied by an object from adjacent objects. This includes traditional functions such as process context management, memory management, object allocation and deallocation. Temporal object segregation refers to the segregation of object contents in time and is traditionally addressed by object reuse controls.

Cryptographic security aims at implementing spatial object segregation using data encryption techniques, with access control being implemented by key management and distribution.

Object abstraction

The concept of an object may differ dependent on the level of abstraction of the defining process or application. In particular the POSIX 1003.22 framework permits service objects to be constructed which are hybrids of platform objects. The service objects may further be subject to alternate access control models which are mapped to the underlying platform access control policy.

The mapping of security policy also infers that a mapping of access control policy and object segregation policy be carried out. In this regard the access control policy applied to service objects must be mapped to an access control policy on platform objects, this mapping may include the use of privilege to bypass certain access control checks within the platform, and the enforcement of additional access controls in the services accessing service objects. To achieve temporal access control it will be necessary to extend object locking at the platform level through the definition of transaction management at the service level. A service level transaction represents an atomic means of manipulating a complex set of platform objects in a consistent manner.

Object segregation at the service level is also required both between service objects and between separate service domains. The former can be supported through the use of trusted services, although such trusted services should be segregated from objects which exist in domain which the service is not trusted to access. The platform should support the segregation of services through traditional memory management and context management.

The layering of services can be extended to support multiple levels of object and subject abstraction, with appropriate requirements being placed on the underlying objects and subjects for segregation. Furthermore the data required to support access control decisions in higher level subjects must itself be protected in an appropriate manner by the supporting services.

Authorisation and accountability

The concept of access control and segregation is augmented by two further security concepts, namely the acquisition and management of authority by subjects on which access control decisions can be based, and secondly ensuring that the subjects are considered accountable for their actions and the exercise of such authority. The concept of authority is traditionally embodied in the form of credentials.

Identification and authentication of users provides an initial route for acquisition of credentials by the subject acting on behalf of the user. Such credentials can be propagated, revoked or surrendered by a subject. Furthermore such credentials can be proffered in support of access control decisions on objects. The credentials may further be mapped from their issuing domain into credentials appropriate to a target domain.

Subjects are made accountable for their use of credentials through a variety of security mechanisms. Such mechanisms include auditing of access control decisions and the accounting of the use of system resources (which can be considered to be auditing of temporal access to objects). Further abstract concepts of accountability also exist in the form of non-repudiation and casuality relations. Non-repudiation establishes accountability for a transaction or message unit transferred via a communication bearer. Casual relations define a requirement for traceability of information flow within a communications network.

Trust hierarchies and capabilities can be represented as means to acquire appropriate authority, or means of proffering authority respectively.

Object methods

The form of access control exercised upon objects is itself undefined in the above model. Such access may be the traditional read, write or execute permissions, or they may be permissible transitions of object state in Clarke-Wilson models. In general access control mechanisms will become more abstract as object and subject abstraction occurs. Thus a database may specify table update permissions subject to complex integrity control conditions. The action of access or update may itself initiate many access requests to less abstract services.

Communication vs static objects

This model currently requires additional effort to model the behaviour of data in transit over communication links and its mapping to static data retained on local storage. The treatment of data in transit requires consideration. Possible models of communications links are:

> Data in transit across a communication link is itself an object. The communication protocol services within the host and intermediate network nodes act as subjects which act upon the object in transit. In such a case the object is assumed to carry attributes which constrain access to the object. This may be the case when the object is cryptographically sealed, or when the object carries a tightly bound data label. This model is appropriate to datagrams.

Data can be considered to be transferred between two nodes forming an association. In such a case mediation of access to data can be carried out by constraining access to the nodes forming each end of the association, and placing a requirement on the supporting operating system to provide cryptographic sealing of data in transit between end nodes. The association/channel model is commonly used to alleviate the requirement for access checking on all data objects.

Further consideration needs to be given to whether data objects in transit can be be represented as a single data object whose location varies in space, and whose accessibility in time is not continuous. This model can be applied also to non-migrant local objects. Ideally the eventual security model should support migrant objects and the requisite support services for the maintenance of object segregation and consistent object mitigation security policies during object migration. The fragmentation of objects in transit across communication networks, and their possible destruction and duplication must also be considered.

Security services

Object segregation

Object segregation embraces all security services responsible for ensuring the integrity of object contents in space and time. Such services refer to support for segregation and non-interference between objects, and for the reliable retention of object state. The object segregation services can therefore be divided into:

Spatial segregation

The requirement that objects are allocated, managed and deallocated in a manner which prevents interference with the content of an object by other subjects or objects in the system environment. Spatial segregation refers explicitly to concepts such as:

- Initialisation - ensuring that all objects of a particular type are initialised to a consistent initial state, or that they may not be accessed until they have been initialised to such a state; further the allocation of resource capacity for the object must ensure that the object is distinct and does not interfere with the state of existing objects. Initialisation does not require association of a valid object access path or paths.

- Non-interference - ensuring that access to objects will not cause modification of other system objects, and that the resources allocated to objects remain distinct.

- Non-penetrability - ensuring that all access to objects are mediated through the reference monitor implementing object mediation.

The concepts of non-interference and non-penetrability also infer that the state of objects will not alter unilaterally without action on behalf of the object mediation services, i.e. that the underlying hardware and operating system will correctly and reliably maintain the state of the object.

Spatial object segregation is implemented in the operating system environment through process segregation, object and memory management. Cryptographic sealing may be used to detect violation of non-interference or non-penetrability.

Temporal segregation

Temporal segregation refers to security services required to support the concept of object segregation in time. The requirements placed on such services are ensuring that storage and resources allocated to objects do not overlap in time. This service is traditionally implemented via object reuse mechanisms.

- Object reuse - ensuring that the content of objects do not overlap in time.

The modelling of covert timing channels may impose additional resource allocation constraints in time which constrain object reuse. Cryptographic sealing can provide temporal segregation of objects when associated with key management services.

Cryptographic object segregation services divide into encryption of object content to prevent compromise of data, signing of object content to detect violation of data integrity, and the associated verification and decryption services.

- Encryption - sealing of object content to ensure segregation for confidentiality
- Signature - sealing of object content to validate segregation for integrity
- Decryption - unsealing of object content
- Validate - validation of object signature to detect violation of integrity segregation

Object mediation (access control)

The services above can provide support for construction of barriers between system objects and subjects which constrain access to be via the reference monitor or object mediation mechanism. Access to objects can therefore be mediated both in time and space by access control mechanisms supplied by the mediation services. Mediation can be divided into control of access to the object in space, and in time. Spatial controls represent traditional access control mechanisms, and simple integrity controls such Biba. Temporal controls represent transaction and atomicity mechanisms such as object locking. Hybrid controls which constrain access to an object or objects in time and space can also be envisaged.

Spatial controls

Examples of spatial controls are discretionary access control mechanisms based on p-bit models, access control lists or capabilities; mandatory access control mechanisms based on object/subject label comparison; and Biba integrity models based on integrity labels. Spatial controls also constrain the allocation and deallocation of resources to objects.

Temporal controls

Examples of temporal controls are object lock mechanisms supporting serialisation of access and transaction based mechanisms supporting atomicity of access to clusters of objects.

Cryptographic services may support object mediation in both space and time through the interaction with key management functions implement authority management.

Object identification

To permit mediation of access to an object it is necessary to identify the object in some manner. Such identification may be by internal memory or disk address, or via the association of a numeric, alphabetic or other identifier. Objects may be referred to by a number of identifiers, furthermore mediation may be exercised over a single identifier or simultaneously over all identifiers.

Identifier allocation and management are subject to mediation in the same manner as object allocation and deallocation, such mediation may segment name spaces by security classification, randomise name allocation or constrain covert channel exploitation relating to name space allocation.

The object mediation and object naming services are themselves subject to object segregation and object mediation. Objects supporting object mediation such as access control lists will themselves be subject to object mediation and segregation.

Temporal segregation and mediation require the existence of a time service that will itself be protected by the object mediation and segregation mechanisms. Such a time service may require temporal segregation to implement fuzzy time covert channel concealment.

Subject identification

In an analogous manner to object identification, it is necessary to provide a mechanism for identifying all subjects active within the system. Subjects will be subject to segregation and mediation based on such identifiers.

Identification of the subject may be provided intrinsically by the underlying security services by examination of the environment in which a service request has been made, or through the proffering of an identifier by the subject, or by engaging in a dialogue sufficient to identify the subject. Services may be provided to support the acquisition of identity, propagation of identity, and revocation of identity that are analogous to the authorisation services detailed below.

Subject authorisation

Object mediation is carried out on the basis of the attributes of the object being accessed, and the attributes of the subject attempting the access. The means by which a subject acquires attributes are referred to here as authorisation. The process of authorisation may involve acquisition of authority by interaction with the security services of the system and the propagation of authority from other subjects. The initial propagation of authority from a user to the subject sponsor can be modelled as establishment of user identity, establishment of user authority, followed by propagation of such authority to the subject sponsor. In this regard I&A does not differ from the general case of authority propagation between subjects which exist in diverse domains. In such a case the requesting subject will identify itself to the remote domain, acquire and then propagate authority to a subject in the remote domain.

The management of authority can be divided into:

- Acquisition - the process by which a subject interacts with the security service to acquire authority to carry out a series of functions. This will normally include the proffering of an identity by the subject, and the mapping of the identity to an appropriate authority dependent on domain.

- Propagation - the process by which a subject which has authority can selectively propagate all, or part of its authority to another subject. Such propagation may be complete, limited to specific actions, limited to specific times or intervals, or constrained by limits on further propagation.

- Combination - the combination of authority received from multiple subjects to perform an action. This typically models separation of duty concepts such as a two man rule. A range of set operations can be conceived including disjunction and conjunction.

- Revocation - the deliberate revocation of authority by the granter of authority from a subject to which such authority has been propagated. This may include partial or

complete revocation, or modification of the constrains under which authority may be utilised.

- Resignation - the release of authority on a permanent or semi-permanent basis by a subject holding authority, such release may be partial or may modify the constrains under which the subject may utilise the remaining authority.

Cryptographic key management can be modelled as the management of certificates of authority to access data, where the certificates are cryptographic keys which when presented to the object mediation function (decryption module) will bypass object segregation (cryptographic security). The object mediation function may constrain the circumstances in which a cryptographic key permits access to the object.

Subject accountability

The sections above have addressed the segregation of objects and their access mediation based on subject authority. The user (as identified by the subject identification services) is made accountable for the use of his/her authority. The accountability traditionally takes the form of audit and accounting information recorded by the system, although other forms of accountability are now manifesting themselves in the area of digital signature of non-repudiable transactions.

Accountability divides into:

- Accountability for spatial access to system objects - traditionally reflected in the system audit trail which records all decisions made by the object mediation services.

- Accountability for temporal access to system objects - reflected in accounting functions which record a subjects use of system resources and associated financial costs.

While these mechanisms ensure that a subject is made accountable for his actions on the local system, it is increasingly important to provide mechanisms to support accountability within a wider forum. To support this concept mechanisms are being provided for providing certificates linking the exercise of authority by identifiable subjects to the transactions or access to objects carried out by the subject. Such certificates support non-repudiation, examples include non-repudiation of mail sending and non-repudiation of mail receipt.

Inter-domain support services

To support the concept of subject identification, subject authorisation and the hybrid concept of non-repudiation it is necessary to allow for the mapping of identity and authority between management domains in which such identity and authority are defined. The services required to support such inter-domain services include:

- Mapping of identity - mapping of the identity of a subject between a domain originating a request for manipulation of an object, and the domain in which such manipulation is carried out.

- Mapping of authority - mapping of the authority of a subject between a domain originating a request for manipulation of n object, and the domain in which such manipulation is carried out.

- Mapping of non-repudiable transaction - provision of a common definition of the nature of a non-repudiable transaction initiated by a remote domain.

Support services also include provision of unified object name spaces, identification of host systems, synchronisation of time services, and provision of distributed locking and object mediation services.

Inter-domain services are also required on a single platform to map between levels of subject and object abstraction. In such instances inter-domain services will support identity, authorisation and object mapping.

Summary

Object segregation
Object mediation
Object identification

Subject identification
Subject authorisation
Subject accountability

Inter-domain services

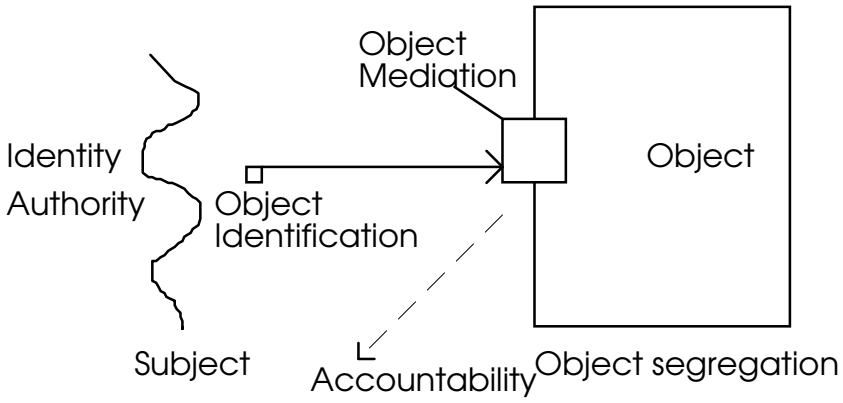

Object access by subject

A subject applies authority which is interpreted by an object mediation services to determine the rights of the subject to access an object as defined by the object segregation services, the exercise of such rights is recorded by subject accountability services using the identity supplied by the subject identification service.

Typical UNIX DAC access

Objects are files. The segregation is provided by the kernel file system and memory management mechanisms. Mediation is performed by the DAC controls in the "open" system call. The subject identity and authority is retrieved from the process environment. The subject is made accountable for its actions by an entry in the system audit trail.

Considering the time access... the segregation is provided by the kernel object reuse controls, mediation is provided by the kernel lock manager (if mandatory locking is in use), identity and authority are as before, accountability is provided by the system audit trail and file space accounting mechanisms.

Finally, to do...

The taxonomy above should be powerful enough to map all known forms of security mechanism or service in the ISO OSI, IETF, ECMA and POSIX frameworks. The taxonomy also can be extended to identify some of the holes... particularly in the provision of temporal security services such as object atomicity and transaction mechanisms. The concept of a network link or association requires modelling, as does the implication of unreliable object storage on local machines and how data corruption relates to object segregation.

Suggestions...

This taxonomy should be aligned with appropriate terminology. Each network and host service should be mapped against the taxonomy. We should also seek to identify where holes exist, and consider whether any such holes should be plugged prior to adopting existing practice which implements a subset of the security services.

D Ferbrache
DRA