

Guide to Understanding I&A

NCSC-TG-017
Library No. 5-235,479
Version 1

FOREWORD

A Guide to Understanding Identification and Authentication in Trusted Systems provides a set of good practices related to identification and authentication (I & A).

We have written this guideline to help the vendor and evaluator community understand the requirements for I & A, as well as the level of detail required of I & A at all classes, as described in the Department of Defense Trusted Computer Systems

Evaluation Criteria. In an effort to provide guidance, we make recommendations in this technical guideline that are not requirements in the Criteria.

The I & A Guide is the latest in a series of technical guidelines published by the National Computer Security Center. These publications provide insight to the Trusted Computer Systems Evaluation Criteria requirements for the computer security vendor and technical evaluator. The goal of the Technical Guideline Program is to discuss each feature of the Criteria in detail and to provide the proper interpretations with specific guidance.

The National Computer Security Center has established an aggressive program to study and implement computer security technology. Our goal is to encourage the widespread availability of trusted computer products for use by any organization desiring better protection of its important data. One of the ways we do this is by the Trusted Product Evaluation Program. This program focuses on the security features of commercially produced and supported computer systems. We evaluate the protection capabilities against the established criteria presented in the Trusted Computer System Evaluation Criteria. This program, and an open and cooperative business relationship with the computer and telecommunications industries, will result in the fulfillment of our country's information systems security requirements. We resolve to meet the challenge of identifying trusted computer products suitable for use in processing delicate information.

I invite your suggestions for revising this technical guideline. We will review this document as the need arises.

Patrick R. GALLAGHER, JR. 1 September 1991
Director
National Computer Security Center

ACKNOWLEDGMENTS

The National Computer Security Center extends special recognition and acknowledgment to James Anderson and to Lt. Col. Rayford Vaughn (USA) as coauthors of this document. Lt. Patricia R. Toth (USN) is recognized for the development of this guideline, and Capt. James A. Muysenberg (USAF) is recognized for its editing and publication.

We wish to thank the many members of the computer security community who enthusiastically gave their time and technical expertise in reviewing this guideline and providing valuable comments and suggestions.

INTRODUCTION

BACKGROUND

The principal goal of the National Computer Security Center (NCSC) is to encourage the widespread availability of trusted computer systems. In support of this goal the NCSC created a metric, the DoD Trusted Computer System Evaluation Criteria (TCSEC) [3], against which computer systems could be evaluated.

The TCSEC was originally published on 15 August 1983 as CSC-STD-001 -83. In December 1985 the Department of Defense adopted it, with a few changes, as a Department of Defense Standard, DoD 5200.28-STD. DoD Directive 5200.28, Security Requirements for Automatic Information Systems (A 155) [11], requires the TCSEC be used throughout the Department of Defense. The TCSEC is the standard used for evaluating the effectiveness of security controls built into DoD AISs.

The TCSEC is divided into four divisions: D, C, B, and A. These divisions are ordered in a hierarchical manner, with the highest division (A) being reserved for systems providing the best available level of assurance and security. Within divisions C and B are subdivisions known as classes, which are also ordered in a hierarchical manner to represent different levels of security in these divisions.

PURPOSE

This document provides guidance to vendors on how to design and incorporate effective identification and authentication (I & A) mechanisms into their systems. It's also written to help vendors and evaluators understand I & A requirements. Examples in this document are not the only way of accomplishing identification or authentication. Nor are the recommendations supplementary requirements to the TCSEC. The only measure of TCSEC compliance is the TCSEC itself.

SCOPE

Computer security is founded upon the notion of controlling access between AIS users and the data within the AIS. The concept of controlled access relies upon establishing identifying information for the AIS users, such that this information can be used to determine whether the user has the proper clearance or identity to access a given data object. In this manner, the I & A requirements are central to the system's identification and authentication of users, and thus the enforcement of the Mandatory Access Control (MAC) and Discretionary Access Control (DAC) policies. I & A also provides the audit mechanism the information it needs to associate actions with specific users.

CONTROL OBJECTIVE

Identification is part of the accountability control objective. The accountability control objective states: "Systems that are used to process or handle classified or other sensitive information must assure individual accountability whenever either a mandatory or discretionary security policy is invoked. Furthermore, to assure accountability the capability must exist for an authorized and competent agent to access and evaluate accountability information by a secure means, within a reasonable amount of time, and without undue difficulty." [3, p. 76]

The fundamental identification requirement states: "Individual subjects must be identified. Each access to information must be mediated based on who is accessing the information and what classes of information they are authorized to deal with. This identification and authorization information must be securely maintained by the computer system and be associated with every active element that performs some security-relevant action in the system." [3, p. 4]

OVERVIEW OF PRINCIPLES

Identification and authentication requirements are found together throughout all evaluation classes. They are directly related in that "identification" is a statement of who the user is (globally known) whereas "authentication" is proof of identification. Authentication is the process by which a claimed identity is verified. The I & A procedures of a system are critical to the correct operation of all other trusted computing base (TCIS) security features.

PURPOSE OF I&A

The strength of I & A procedures directly impacts the ability of the other TCB mechanisms to fulfill their function. For example, the strength of an audit mechanism and the assurance of correctness in the audit is dependent upon the I & A mechanism. If I & A is successfully circumvented, then all audited actions become unreliable, because an incorrect ID could be associated with auditable actions. In this sense, the I & A requirement is the foundation for other TCIB functional requirements (figure 1).

A TCB, using security-relevant data structures, controls access to a system, determines authorized access to objects within a system, and associates audited actions* with specific individuals based on their identity.

In controlling access to a system, the TCIS acts as the first line of defense. Once the TCIS identifies the user as a unique entity or a member of a group, it then accurately determines what access privileges the user may be assigned with respect to the data protected by the system. If the system has a C1 rating, group membership provides sufficient granularity for enforcing the I & A requirements. In C2 or higher rated systems, I & A enforcement must be at the individual user level. Systems at the B and A levels enforce a mandatory access control policy and use the I & A mechanism to establish an authorized security level or levels at which the user will operate during a given session.

The user's identity determines which functions or privileges the user can exercise on a system. In some systems (e.g., transaction systems), functional access may be the predominant expression of security policy. A common case of functional access found in many systems is the control of access to the system security officer's functions based on his identity.

The purpose of a device also may determine functions or privileges the user can exercise on a system. For example, the same physical mechanisms protecting system hardware normally protect the device commonly called the "operator's console." The user of this device is ordinarily subject to stronger physical controls and administrative procedures than are other users of the system. In some cases, access to the device implies physical access to the data (all storage media) that the TCB is charged to protect. In B1 and lower rated systems, the I & A requirements may be relaxed for the user of the operator's console if the device is protected by the same physical security mechanisms protecting the system itself.

Auditing functions in trusted systems are a TCB responsibility. Certainly, accurate identification of an individual user is required for proper attribution of actions taken on behalf of the individual by the system. Again, the strength of the I & A mechanism directly affects the reliability and assurance of correct audit functions.

THE I&A PROCESS

The identification and authentication process (typically called "Login") starts with a user establishing communications with the TCB. (In B2 and above TCBS, this is done by invoking a Trusted Path, which is guaranteed by design to be an inviolable communication path between the user and the TCIS.) Once the user is communicating with the TCIS, the user identifies himself (i.e., claims an identity). Either as part of the transmission claiming an identity or in response to a prompt from the TCIS, the user supplies one or more authentication elements as proof of identity

*Discretionary access controls, enforced at the C level and higher, are dependent upon effective and reliable I & A.

**See Interpretations et -cl-02-83 and ci -cl-04-86 on the NCSC DockMaster Eva __Announcements forum. B2 and above require the operator to log on, but operator logon is optional at B1 and below

The TCB, using the claimed identity and authentication elements as parameters, validates the supplied information against an authentication database. If the information from the user satisfies the TCB validation process, the TCB completes the login by establishing a user session, and associating the user's identity and access control information with that session. In C1-C2 systems, this access control information may merely be a recording of the user identity to compare to access control information associated with files. In B1-A1 systems, the TCB also associates a specific security level (within the valid range for the user) with the user session for use in making mandatory access control decisions.

At C2 and above, the TCB is able to record (audit) the success or failure of a new login. If the login succeeded, the TCB then completes any necessary actions to establish the user session.

ASPECTS OF EFFECTIVE AUTHENTICATION

Users' identities are verified using one of three generic methods something they know (type 1), something they have (type 2), or something they are (type 3). While the requirements of the TCSEC can be met by any ONE of these different methods, systems using two or more methods can result in greater system security. Systems using only one method of authentication may be more vulnerable to compromise of the authenticator, thus multiple methods are preferred.

Examples of type 1 mechanisms include passwords, passphrases, PINs (Personal Identification Numbers), and data about one's self or family members. Type 2 mechanisms include real and electronic keys, challenge response generators, and magnetic-strip cards or badges. The final category, type 3, includes fingerprints, retinal patterns, DNA patterns, and hand geometry.

To be effective, authentication mechanisms must uniquely and unforgeably identify an individual. "Authentication by knowledge" (type 1) and "authentication by ownership" (type 2) mechanisms are limited in effectiveness through only being associated with a person by possession. A person possesses knowledge or some

Type 1 = Authentication by Knowledge (Something They Know)

Type 2 = Authentication by Ownership (Something They Have)

Type 3 = Authentication by Characteristic (Something They Are)

figure 2

identifying object, but since the user is not physically attached to the authentication information, the authentication information could be lost, stolen, or otherwise compromised.

What distinguishes the first two types is how effectively each can be protected; each has advantages and disadvantages. The principal weakness of type 1 is duplication. Not only is it usually very easy to learn something someone else knows, but it may be possible to guess the information without even having access to it. No special tools, skills, or equipment are required to duplicate "authentication by knowledge." One can often break it by a simple brute force guessing attack using automated techniques.

The most important advantage of this type of authentication item is this: the user can take it anywhere and change it whenever a compromise should occur. Another advantage is its simplicity, since such information tends to be easily represented to the TCB without special equipment. Any authentication data must ultimately be encoded in some form in the TCB's authentication database, and in this sense a copy of the information has to be kept by the TCB to be usable in authentication. Since a character string can usually represent type 1, it's easy to store it for later use by the TCB.

Although type 1 is easy to copy if it's genuinely unique, such as a nonsense word or number, it may be easier to guard it than a physical object. This is because an item of knowledge, although easily copied, is always fully in the possession of the person it identifies. Unlike a key, card, or other physical device, "authentication by knowledge" can't be stolen while temporarily left sitting on a desk, can't accidentally fall out of a pocket, and often can't be forcefully stolen unless the person stealing it can verify it is correct. Yet the ease of duplication makes type 1 always an imperfect form of authentication and thus requires conscientious protection to remain effective.

By comparison, "authentication by ownership" has its major strength in the difficulty of duplication. Type 1 is, in fact, an example of type 2. So when we speak of type 2 we will, by convention, mean a physical object rather than an item of knowledge. While such objects require more effort to guard from theft, they can be made using special equipment or procedures that are generally unavailable. Hence, duplicating them is more costly than the value of whatever is to be gained by falsifying them. This discourages their duplication, although it doesn't necessarily prevent duplication by a determined intruder.

The third type of authentication item, "authentication by characteristic," is much stronger than the first two. After all, the goal of authentication is to verify who you are, and type 3 is very closely tied to this.

A major obstacle with this type of authentication is the difficulty of building cost-effective peripherals that can obtain a complete enough sample of a characteristic to entirely distinguish one individual from another. Cost is also a factor in identifying type 2. But in type 3, the authentication item itself can be designed to simplify identification by the TCB. Conventional computer peripherals usually cannot easily encode "authentication by characteristic." While it may be easy to build devices that confirm someone's distinguishing features such as weight or finger length, the addition of more detail to completely distinguish one person from another can substantially increase the cost.

Fortunately, an adequately unique identification doesn't require every detail about a person. Thus, specific methods such as fingerprinting or eye retinal scans may be used alone, reducing costs in comparison with a total examination of all a person's physical attributes. Even these methods incur greater costs than simple use of a password, which requires no additional hardware at all, and they are not guaranteed to be infallible. Identical twins for instance would not be distinguishable by DNA readers, and might not be distinguishable by other specific tests of physical characteristics. In the imaginary case of entirely identical twins, the two individuals might be solely distinguished by things in the "authentication by knowledge" category.

Not only do the various types of authentication methods have cost and feasibility tradeoffs, but an adequate certainty of authentication may require several methods. (Each is subject to an amount of error at the most theoretical level.) One would expect greater assurance from a combination of type 1 and type 2 mechanisms than either used alone. Likewise, type 3 may provide more assurance than the combination of types 1 and 2 together. Potential approach choices are shown in figure 3, where type 12 represents the use of type 1 and type 2 mechanisms.

Direct comparisons of strength relationships are not possible unless one knows the exact implementation mechanism; however, one can theorize that some such relationships are likely. One might argue that type 2 is stronger than type 1 in terms of assurance, and type 123 is probably stronger than type 12. Singular mechanisms may offer the needed assurance at lower levels, whereas higher levels may require combinations to achieve adequate assurance.

Possible Authentication Approaches

Type-1 Type 12 Type 2

Type 13 Type 23

Type 3

Type 1 = Authentication by Knowledge

Type 2 = Authentication by Ownership

Type 3 = Authentication by Characteristic

FIGURE 3

SECURITY OF AUTHENTICATION DATA

Identification and authentication data (like most transmissions) is vulnerable to interception (e.g., eavesdropping, spoofing) by an intruder interposed between the user and the TCB. As a consequence, the connection between a user and the trusted computer requires protection commensurate with the sensitivity of the data the system processes. Due to government regulations, systems used to process classified data must meet stringent physical security standards that include protection of the connection of a terminal to the TCB. (This protection can involve putting the computer and its terminals in a physically secure area, protecting the wireways between the terminals and the computer, or using cryptography to protect the transmissions.) Unclassified systems may require similar protection. Additionally, 192 and higher rated systems must have trusted paths.

Networks provide many opportunities for intruders to intercept the I & A data.

One-time passwords can help protect against that possibility.

The user authentication data must be stored, protected, and maintained by the TCB. It should be accessible only to the System Security Officer (550). However, even the 550 should be barred from seeing the actual plain text version of the data (for example, the passwords used by the users.) To assure only the 550 can access and manipulate the I & A database, a unique and possibly extended special 550 identification and authentication procedure must be embedded in the TCB. The TCB should use this procedure to verify the identity of the 550 (and perhaps the device used) when that individual maintains the I & A database.

Besides interception, operator misfeasance or unauthorized physical access could compromise I & A data. This may be done by mishandling off-line versions of the data in such forms as system backup files, fault-induced system dumps, or listings. To protect I & A data from this kind of unauthorized disclosure, the data could be stored in encrypted form. Several so-called one-way transformations (non-invertible functions) of authentication data exist that could serve the function (see Purdy [10]). However, if one uses one-way transformations, it's important it be a true non-invertible transformation. For an example of how an ad hoc one-way transformation was broken, see Downey's paper [4].

The authentication database needs protection from general access whether the database is encrypted or not. Even in an encrypted form, a database may be subject to a "catalog attack." (Such an attack was highly successful during the November 1988 INTERNET attack by a network worm program. [5]) A catalog attack is conducted by encrypting a dictionary of probable authentication data (e.g., passwords). The attacker then matches the ciphertexts of the authentication database with the ciphertext of the dictionary to discover legitimate authenticators. If the database stores both the user's identification and authentication in encrypted form, the attacker must find a user's identification AND authenticator (e.g., the password) in COMBINATION. However, the need to protect the transformation mechanism remains.

MEETING THE TCSEC REQUIREMENTS

This chapter describes the I & A requirements from the TCSEC and their interaction with related TCSEC requirements.

C1 REQUIREMENTS

I & A Requirements

"2.1.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall use a protected mechanism (e.g., passwords) to authenticate the user's identity. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user."

Other Related Requirements

"2.1.1.1 discretionary Access Control

The enforcement mechanism (e.g., self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals or defined groups or both."

Comments

The related Discretionary Access Control (DAC) requirement, allowing users to control sharing by defined groups, means it's sufficient to identify users as a member of a group. Individual identity is NOT required at C1. Thus, it's acceptable to have a collective identity such as "Purchasing," authenticated with a password controlling access to a purchasing file. The TCSEC requires no additional information. And, without individual accounting, auditing isn't possible or required.

What is sufficient authentication? This is a difficult question, since it interacts critically with how the Trusted System is used, combined with the assurances and design features associated with the ratings. (See Guidance for Applying the Department of Defense Trusted Computer System Evaluation Criteria in Specific Environments [7].) The C1 requirement specifies only the use of a protected mechanism and gives, as an example, passwords. As discussed elsewhere in this guideline, passwords can be an effective authentication mechanism if conscientiously applied.

C2 REQUIREMENTS

I & A Requirements

“2.2.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall use a protected mechanism (e.g., passwords) to authenticate the user's identity. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADI° system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.”

Other Related Requirements

“2.2.1.1 Discretionary Access Control

The enforcement mechanism . . . shall allow users to specify and control sharing of . . . objects by . . . defined groups of individuals These access controls shall be capable of including or excluding access to the granularity of a single user. .

“2.2.2.2 Audit

... The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, . . . actions taken by computer operators and system administrators and/or system security officers For each recorded event, the audit record shall identify: date and time of the event, user, type of event, and success or failure of the event;. For identification AND Authentication events the origin of request (e.g., terminal ID) shall be included in the audit record. . . . The ADP system administrator shall be able to selectively audit... one or more users based on individual identity.”

Comments

Beginning at the C2 level, individual users are identified. The access control requirement mandates using individual identity for access decisions. If group-based access control is available, membership in the group is based on the identity of the individual rather than a user providing a group name with an authenticator. This is an important distinction. With a group identifier, a collective name and shared authenticator is valid. With individual identifiers, a unique individual ID, verified through unique authentication, is used to determine membership in the group, with group identification then used to access the data. In this latter implementation, the system can audit the actions of the individual, thus ensuring individual accountability.

Strengthening the requirement for individual identity is the audit requirement.

This means the system administrator can audit the actions of any one or more users, based on individual identity. I & A must distinguish operators, system administrators, and system security officers from ordinary users in order to record security related events as actions initiated by the individuals performing those roles. Since individuals performing those roles may also be ordinary users of the system, it's necessary to distinguish the people when acting as ordinary users.

As an example, in one (network) system with many individuals performing administrator or security officer tasks, the system established an identifier associated with the role being performed (e.g., System Administrator (SA)). In an extended log-on, a two-step identification and authentication occurred; first as the SA, and then as the individual performing that role. If the individual wasn't recognized as one of those authorized the SA functions, the logon ended and the system audited the event.

Audit records taken of actions done by the SA incorporated the individual's identity. Later examination of the audit records permitted collection of all actions by the SA in time-sequenced order. Within the SA function, the system identified individuals performing in the SA role. Since this was a very large international time-sharing system, two or more people might be doing SA functions totally independently of each other. The system recorded all their activities under the SA identity, and within each record were the identities of the individuals actually performing the function.

Finally, the related audit requirement calls for identification of the origin of I & A events. The example given is from a terminal, but, in some systems, it may be a stored batch procedure (PROC on some systems) activated by an operator from the operator's console. In this case, the audit record should contain both the operator's console ID and an indication the operator ran the job for some individual identified in the batch procedure. The origin of a connection is often joined with the user's identity to insure the terminal's location is approved to handle classified material at the user's authorized level.

B1 REQUIREMENTS

I & A Requirements

“3.1.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identity of individual users (e.g., passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorizations of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.”

Other Related Requirements

“3.1.1.4 Mandatory Access Control

... Identification and authentication data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorization of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user.”

Comments

B1 is the first rating level in which access and data movement control based on sensitivity levels of subjects and objects takes place. For an unprivileged user, the I & A data is used to determine the user's current authorization level,

which the TCB compares with its user database containing authorization ranges for each user. If the logon information is correct and his level is valid, the TCB lets him on the system. Then the TCB moderates the access to files based on the user's current level and the label on the file or object the user is trying to get to. Since the sensitivity level is represented by the clearance and category of access, and object access permission is determined by the sensitivity associated with both the subject (outside of the TCB) and the object, the authorization of a subject becomes a component of the authentication requirement.

The meaning of the term "authorizations" in this section includes functional roles assigned to individuals. The authorizations associated with user roles (e.g., SA, 550, operator) define modes of access that may or may not be controlled by a label-processing (or MAC) policy, depending upon the particular system. One can have a system where a user, acting as an authorized 550, may add new users, delete users, or modify their authentication data to increase or decrease their authorized access, all without any sensitivity label associated with the records or the 550's actions. Such actions are, of course, subject to audit. The better approach uses MAC mechanisms to provide additional support for administrative least privilege. Here the user, as the 550, must still log onto the system at whatever level is necessary to do his work.

B2 REQUIREMENTS

I & A Requirements

3.2.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identity of individual users (e.g., passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorizations of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual."

3.2.2.1.1 Trusted Path

The TCB shall support a trusted communication path between itself and [the] user for initial login and authentication. Communications via this path shall be initiated exclusively by a user."

Other Related Requirements

3.2.3.1.4 Trusted Facility Management

The TCB shall support separate operator and administrator functions."

Comments

The trusted path requirement is the principal addition at this level. The B2 level is the first rating level providing sufficient architectural support for trusted paths in an operating system. This requirement ensures that at the B2 level and above, the individual user logging in is in unforgeable communication with the TCB, and not some program masquerading as a TCB. Otherwise, the user may be spoofed into divulging his authentication data to an application program.

The requirement to support separate operator and administrator functions could place an additional burden on the I & A function to distinguish individuals acting in those roles. To this end, the (functional) authorizations associated with the authentication data from the B1 requirement can be effectively used.

B3 (AND A1) REQUIREMENTS

I & A Requirements

“3.3.2.1 Identification and Authentication

The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identity of individual users (e.g., pass words) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user's identity and to ensure that the security level and authorizations of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.”

“3.3.2.1.1 Trusted Path

The TCB shall support a trusted communication path between itself and users for use when a positive TCB-to-user connection is required (e.g., login, change subject security level). Communications via this trusted path shall be activated exclusively by a user or the TCB and shall be logically isolated and unmistakably distinguishable from other paths.”

Other Related Requirements

“3.3.1.1 Discretionary Access Control

The TCB shall define and control access between named users and named objects (e.g., files and programs) These access controls shall be capable of specifying, for each named object, a list of named individuals and a list of groups of named individuals with their respective modes of access to that object. Furthermore, for each Such named object, it shall be possible to specify a list of named individuals and a list of groups of named individuals for which no access to the object is to be given....”

Comments

The trusted path's use is generalized to “when[ever] a positive TCB-to-user connection is required,” not just for login. In 193 systems, other TCB-to-user communications may be going on, hence the requirement to logically isolate and to distinguish the TRUSTED path from all other paths. Note that the TCB can start the trusted path if necessary. The distinction between trusted path at B3 and trusted path at B2 hinges on whether the TCB needs to be aware of a previous context. In the B2 case, the only requirement for trusted path is at login. In the B3 case, a trusted path may be re-quired for a user to change security levels or to initiate another process at a different security level from the one he is now in. An example of the TCB starting the trusted path could be telling a user his security level has been changed as requested.

The principal impact of this requirement is the establishment of a trusted path between a user (i.e., an individual) and the TCB, not a process-subject, nor any other user-surrogates. It's as important for the person connecting to the TCB to be assured of the identity of the TCB as it is for the TCB to be assured of the identity of the individual. Earlier work in computer security suggested re-authentication as an assurance mechanism to cover cases of this kind. If the system has such a capability, the time between (re)authentications should be a configuration parameter.

The related Discretionary Access Control requirement has two components; first, each named object (already controlled by the Mandatory Access controls) shall also be under Discretionary Access controls. Second, lists of named individuals or groups authorized or specifically denied access must be maintained for each named object.

POSSIBLE METHODS OF IMPLEMENTATION

There are a wide variety of implementation methods possible for identifying and authenticating people. The challenge to the TCB designer is how to integrate the method chosen into the rest of the TCB design in such a way that the chosen technique cannot be bypassed or penetrated. The most frequent method of identifying individuals is still the claimed identity authenticated with a password. A reasonable discussion of the issues is carried in Hoffman [8], although it's not complete with respect to all the problems. The following discussion serves as a set of good examples or a general guideline only. There are many other acceptable methods of I & A.

SOMETHING A USER KNOWS (TYPE 1 METHOD)

This method is almost entirely focused on passwords. In the past, eight character passwords have been more or less the standard on most systems providing user identification and authentication, although there is no specified standard for password length. The Department of Defense Password Management Guideline [2] recommends a minimum of six characters. Two appendices in that guideline discuss the parameters affecting the choice of password length.

The guideline also strongly recommends automating password generation. A description of a pronounceable password generator can be found in Gasser's paper [6].

As indicated above, simple passwords are sufficient for the lower rating classes.

As one moves up in the ratings, the password or passphrase system should become more sophisticated. In the higher classes the password scheme should be some variation of the one-time password schemes or a combination of techniques, as depicted in figure 3.

In a one-time password scheme, the system provides the user with an initial password to authenticate his claimed identity. During the initial logon, the user receives a new password for the next logon. In the earliest conception of this approach, no one was concerned for wiretapped lines intercepting the users' passwords, since all lines were within protected wireways. The only concerns were for someone masquerading as another user without being discovered or for users writing down their passwords so they wouldn't forget them.

In modern applications, particularly with personal computers used as terminals, the TCB could encrypt the next password for the user; The user would receive his next password, decrypt it with his (personal, unique) decrypt key, and save it for his next session.

Other proposals include storing personal data about an individual, such as your grandmother's maiden name, the age of your father when he was married, the middle name of your 3rd sister, etc. This method falls short of the TCSEC requirements for a variety of reasons. It's difficult to administer for any reasonable number of users, and, even if one randomizes the challenge, the total number of available answers is too small. Personal information relevant to an individual is normally available from public sources and not protected.

SOMETHING A USER HAS (TYPE 2 METHOD)

This type contains several artifact approaches to providing positive identity of users. The schemes span a wide range, from magnetic strip readers to various forms of ignition keys (some with cryptographic subsystems, others merely alternate forms of the magnetic strip reader). An interesting form of "something one has" that combines the artifact with a password scheme is typified by a one-time password (challenge and response) device. The device is a calculator-sized unit that, if keyed correctly by a user with a personal identification number (PIN), generates a correct one-time response to a password challenge issued by a server host.

Possession of the device alone does not let one obtain the correct response to a random challenge. One must have the artifact and know the PIN to use it. There's no known way to read the PINs set in this particular product, so loss of the device may be an inconvenience rather than a security breach of major magnitude if reported immediately.

While it's possible such devices could be stolen from the rightful user, the security breach might be manageable, unless the user doesn't report the loss immediately or carelessly writes the PIN down. Furthermore, if one augments

the mechanism with the standard password approach to form a Type 12 method, one gets much greater assurance. There's a tendency to require total security for simple devices (locking them in safes or restricting where they may be carried). Many times all that's needed is the ability to detect the loss or compromise of the device.

SOMETHING A USER IS (TYPE 3 METHOD)

This category of authentication includes all the biometric schemes, such as fingerprint readers, lip print readers, retinal scanners, DNA analyzers, and dynamic signature readers. Some claim these devices have substantial resolution powers, virtually eliminating false acceptance, while keeping the false rejections at a reasonable level*. However, the statistical nature of the acceptance or rejection is something to consider. We noted earlier one could double up the authentication mechanisms for higher rated systems so authentication is based on two independent elements, for example a fingerprint reader and a password (type 13 method). Such a scheme would virtually eliminate false acceptance of the I & A procedure. In a doubled up authentication scheme, the system shouldn't accept either one of the elements unless the other element is also correct.

The vendors of biometric devices have a harder time than those who are content with simple passwords, since it's virtually impossible to change the biometric parameter being measured. However, it may be possible to copy the biometric parameter in such a way to gain access to a system as though one is the actual user. If an intruder interposes himself between the measuring device and the TCB, absent a protected path, he can copy the reading for replay at a later time. As this may be possible regardless of the authentication technique used, this suggests either making the authentication element be one-time, or protecting the path between authentication entry (measurement) and the TCB against interception.

COMMENTS

The requirements for identification and authentication are stressed heavily in the direction of authentication. One claims an identity, then must prove it to the operating system. It's "theoretically" possible to have self-identifying authentication (or it might be called self-authenticating identification). Examples of such might be a fingerprint reader or a DNA analysis of cells scraped from the skin. Of course, it's always (and probably incorrectly) assumed one can maintain the integrity of one's skin or fingerprints. However, it might be possible to copy fingerprints onto a latex finger (or fingers), and obtaining a patch of skin might not be as difficult as one might imagine. Even for self-identifying authentication, the protection of the authentication mechanism is key to its successful application.

*Martin, in his book, reported false acceptance of a voice recognizer at 1%, and a hand geometry reader at 0.5 to 0.05%, depending on the measurement tolerances. The voice recognizer had a false rejection of 1%, while the hand geometry reader was "very low." [9]

As indicated above, if applied in situations calling for application of lower rating classes, such methods that combine identification and authentication could very well be sufficient. In situations where higher classes are called for, the methods could be combined with another generic authenticator. In effect this erects a second access barrier for systems used in higher risk situations.

How much authentication is enough has not been adequately addressed to date and is a matter for discussion between the site accreditor or site certifying officer and the vendor. One could require multiple authentication mechanisms for any claimed identity for higher level systems. However effective such a requirement might be, it would be quite expensive for those myriads of systems NOT at special hazard but which use B2 or higher ranked systems in quasi-system high environments because of their label processing provisions. (Quasi-system high environments are those where all users are given all clearances and categories within the organization, but the data must be properly labeled to exercise control on where it's exported.) Perhaps this is a point of application for an I & A subsystem, to augment one built into a TCB. The principal recommendation is it be a different basic mechanism. If the built-in authenticator is based on authentication by characteristic, the I & A subsystem could be either authentication by ownership (type 2) or authentication by knowledge (type 1). Conventional wisdom says password systems can be good enough, but better means of authentication may be required.

In general, one would expect the lower rated systems to use simpler mechanisms than the higher rated systems. At C1, simple type 1 mechanisms or any of the simplest artifact schemes (e.g., lock combinations or keys to locked

doors to terminal rooms for a user population known to the system only as defined groups) might be sufficient, based on the site where the system is used.

At C2, any of the random pronounceable password systems or any of the simpler artifact (e.g., a challenge response table) or biometric systems (e.g., measurement of hand geometry, inter-ocular distance, or other Bertillon measures) seem appropriate.

At B1, passphrases, random pronounceable passwords of at least 8 characters, challenge-response schemes, one-time passwords, advanced artifacts (e.g., terminal logon "keys" or magnetic striped cards), or biometric systems (e.g., fingerprints, retinal images, or voice images) might be appropriate.

For higher levels (B2, O3, and A1), authentication could be based on at least two of the three generic ways of verifying identity. A magnetic striped card and a password, a PIN used to start a challenge-response device, or a biometric device and a password could be used in combination to increase the "work factor" of attempting to subvert or diagnose the authentication parameter(s).

Although not specifically addressed in the TCSEC, the evaluation process must consider the strength of the I & A mechanism in relation to the evaluation class. The assurance associated with a chosen mechanism must be appropriate for the evaluated class.

As an aside, the strength of the I & A mechanism should also be based on the environment the system will be used in and the risk of losing the data on the system. Remember, it's possible that a C2 system running at system high with very sensitive data would need a high assurance I & A mechanism just as an A1 system would.

It's interesting to observe that password systems have rarely failed to perform their function on the systems protected. The bulk of password failures is due to misfeasance, sharing of passwords with an otherwise unauthorized individual, or careless handling of passwords (at least as serious as equivalently careless handling of safe combinations). Some agencies treat careless handling of passwords with the same degree of seriousness as leaving safes unlocked and unattended. For multilevel systems handling classified data, the password is classified at the highest level of information authorized the user to whom it belongs.

One can manage a password scheme properly with frequent changes of passwords and a pronounceable random password generator used to eliminate some of the simpler guessing attacks. It's true people can misuse the system by not treating their passwords properly (e.g., by writing them on their terminals, by deliberately giving them away, or allowing them to be observed by others when used). Nevertheless, the low cost and high degree of effectiveness make passwords the authentication method of choice for most systems.

If users are allowed to pick their own specific authenticators, their behavior is stereotypical enough to permit diagnosis and recovery of the selected authentication. This is especially true of systems permitting users to pick their own passwords. As a consequence, the technique of a system administrator making the (initial) selection of the authenticator is better security practice than it appears at first glance.

GLOSSARY

ASSURANCE

A measure of confidence that the security features and architecture of an automated information system accurately mediate and enforce the security policy.

AUDIT TRAIL

A chronological record of system activities that is sufficient to enable the reconstruction, reviewing, and examination of the sequence of environments and activities surrounding or leading to an operation, a procedure, or an event in a transaction from its inception to final results.

AUTHENTICATE

Verify a claimed identity as legitimate and belonging to the claimant.

AUTHORIZATION

An individual's right to access or use an object.

CRYPTOGRAPHY

The principles, means, and methods for rendering information unintelligible, and for restoring encrypted information to intelligible form.

DISCRETIONARY ACCESS CONTROL (DAC)

A means of restricting access to objects based on the identity of subjects and/or groups to which they belong. The controls are discretionary in the sense that a subject with a certain access permission is capable of passing that permission (perhaps indirectly) on to any other subject.

DISCRETIONARY ACCESS PRIVILEGE

Access granted to objects based on the identity of a subject and/or the groups to which they belong.

DOMINATE

Security level S₁ dominates security level S₂ if the hierarchical classification of S₁ is greater than or equal to that of S₂ and the non-hierarchical categories of S₁ include all those of S₂ as a subset.

IDENTITY

A unique name or number assigned to an individual using a system.

LEAST PRIVILEGE

The principle that requires that each subject be granted the most restrictive set of privileges needed for the performance of authorized tasks. The application of this principle limits the damage that can result from accident, error, or unauthorized use.

MANDATORY ACCESS CONTROL (MAC)

A means of restricting access to objects based on the sensitivity (as represented by a label) of the information contained in the objects and the formal authorization (i.e., clearance) of subjects to access information of such sensitivity.

OBJECT

A passive entity that contains or receives information. Access to an object potentially implies access to the information it contains. Examples of objects are: records, blocks, pages, segments, files, directories, directory trees, programs, bits, bytes, words, fields, processors, video displays, keyboards, clocks, printers, and network nodes.

SENSITIVITY LABEL

A piece of information that represents the security level of an object and that describes the sensitivity (e.g., classification) of the data in the object. The TCB uses sensitivity labels as the basis for mandatory access control decisions.

SPOOFING

An attempt to gain access to a system by posing as an authorized user. Synonymous with impersonating, masquerading, or mimicking.

SUBJECT

An active entity, generally in the form of a person, process, or device that causes information to flow among objects or changes the system state. Technically, a process/domain pair.

TRUSTED COMPUTING BASE (TCB)

The totality of protection mechanisms within a computer system -including hardware, firmware, and software - the combination of which is responsible for enforcing a security policy. It creates a basic protection environment and provides additional user services required for a trusted computer system. The ability of a TCB to correctly enforce a security policy depends solely on the mechanisms within the TCB and on the correct input by system administrative personnel of parameters (e.g., a user's clearance) related to the security policy.

TRUSTED PATH

A mechanism by which a person at a terminal can communicate directly with the Trusted Computing Base. This mechanism can only be activated by the person or the Trusted Computing base and cannot be initiated by untrusted software.

VERIFY

To prove the truth of by presenting evidence or testimony; substantiate.