INFORMATION PRIVACY ENGINEERING AND PRIVACY BY DESIGN

Understanding Privacy Threats, Technology, and Regulations Based on Standards and Best Practices

WILLIAM STALLINGS

Information Privacy Engineering and Privacy by Design

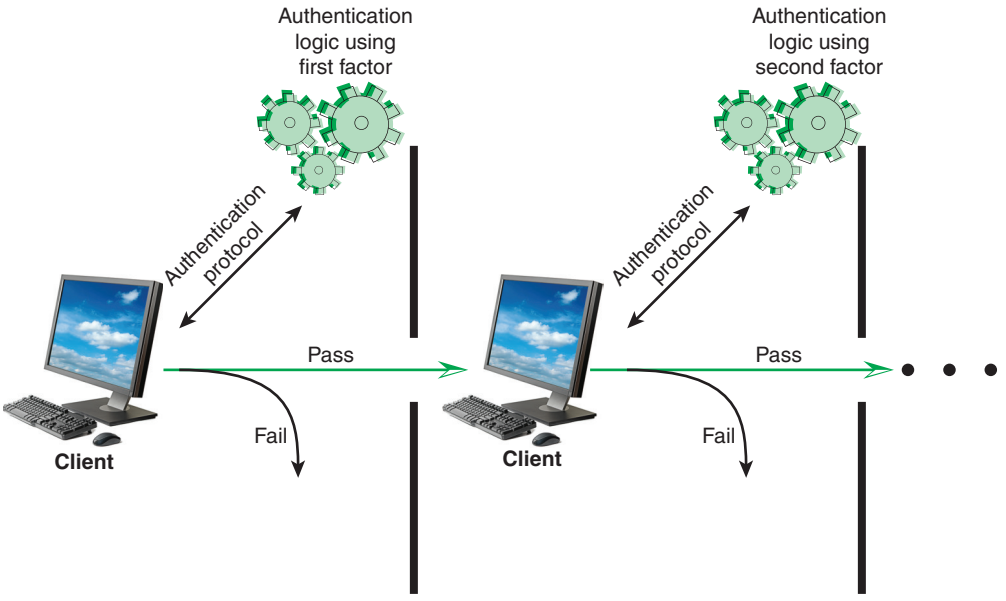

FIGURE 5.2 Multifactor Authentication

A Model for Electronic User Authentication

NIST SP 800-63 (*Digital Identity Guidelines*) defines a general model for user authentication that involves a number of entities and procedures, as shown in Figure 5.3.

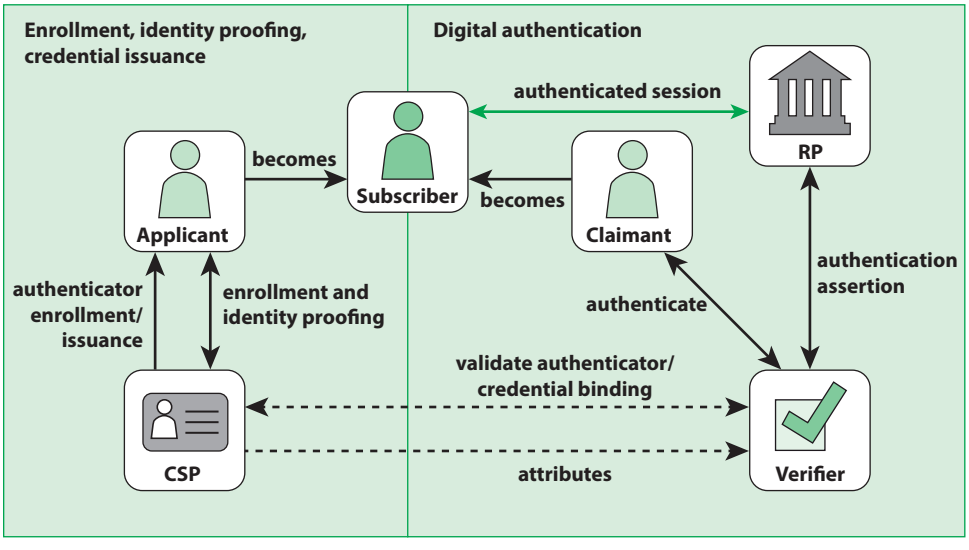

FIGURE 5.3 The NIST 800-63 Digital Identity Model

Three concepts are important in understanding this model:

- **Digital identity:** The unique representation of an individual, generally referred to as a subject, engaged in an online transaction. The representation consists of an attribute or a set of attributes that uniquely describe a subject within a given context of a digital service but do not necessarily uniquely identify the subject in all contexts.
- **Identity proofing:** The process of establishing that a subject is who they claim to be to a stated level of certitude. This process involves collecting, validating, and verifying information about a person.
- **Digital authentication:** The process of determining the validity of one or more authenticators used to claim a digital identity. Authentication establishes that a subject attempting to access a digital service is in control of the technologies used for authentication. Successful authentication provides reasonable risk-based assurances that the subject accessing the service today is the same as the subject who previously accessed the service.

Six entities are defined in Figure 5.3:

- **Credential service provider (CSP):** A trusted entity that issues or registers subscriber authenticators. For this purpose, the CSP establishes a digital credential for each subscriber and issues electronic credentials to subscribers. A CSP may be an independent third party or may issue credentials for its own use.
- **Verifier:** An entity that verifies the claimant's identity by verifying the claimant's possession and control of one or two authenticators using an authentication protocol. To do this, the verifier may also need to validate credentials that link the authenticator(s) to the subscriber's identifier and check their status.
- **Relying party (RP):** An entity that relies on the subscriber's authenticator(s) and credentials or a verifier's assertion of a claimant's identity, typically to process a transaction or grant access to information or a system.
- **Applicant:** A subject undergoing the processes of enrollment and identity proofing.
- **Claimant:** A subject whose identity is to be verified using one or more authentication protocols.
- **Subscriber:** A party who has received a credential or an authenticator from a CSP.

The left-hand portion of Figure 5.3 illustrates the process whereby an applicant is enrolled into the system for purposes of accessing certain services and resources. First, the applicant presents to the CSP evidence of possession of the attributes to be associated with this digital identity. Upon successful proofing by the CSP, the applicant becomes a subscriber. Then, depending on the details of the overall authentication system, the CSP issues some sort of electronic credential to the subscriber. The ***credential*** is a data structure that authoritatively binds an identity and additional attributes to one or more authenticators possessed by a subscriber, and it can be verified when presented to the verifier in an authentication

transaction. The authenticator could be an encryption key or an encrypted password that identifies the subscriber. The authenticator may be issued by the CSP, generated directly by the subscriber, or provided by a third party. The authenticator and credential may be used in subsequent authentication events.

Once a user is registered as a subscriber, the authentication process can take place between the subscriber and one or more systems that perform authentication (right-hand portion of Figure 5.3). The party to be authenticated is called a *claimant*, and the party verifying that identity is called a *verifier*. When a claimant successfully demonstrates possession and control of an authenticator to a verifier through an authentication protocol, the verifier can verify that the claimant is the subscriber named in the corresponding credential. The verifier passes on an assertion about the identity of the subscriber to the relying party (RP). That assertion includes identity information about a subscriber, such as the subscriber name, an identifier assigned at registration, or other subscriber attributes that were verified in the registration process. The RP can use the authenticated information provided by the verifier to make access control or authorization decisions.

In some cases, the verifier interacts with the CSP to access the credential that binds the subscriber's identity to his or her authenticator and to optionally obtain claimant attributes. In other cases, the verifier does not need to communicate in real time with the CSP to complete the authentication activity (e.g., with some uses of digital certificates). Therefore, the dashed line between the verifier and the CSP represents a logical link between the two entities.

An implemented system for authentication will differ from or be more complex than this simplified model, but this model illustrates the key roles and functions needed for a secure authentication system.

5.4 Access Control

This section provides an overview of important aspects of access control. It is useful to begin by defining the following terms:

- **Access:** The ability and means to communicate with or otherwise interact with a system, to use system resources to handle information, to gain knowledge of the information the system contains, or to control system components and functions.
- **Access control:** The process of granting or denying specific requests (1) for obtaining and using information and related information processing services and (2) to enter specific physical facilities.
- **Access control mechanism:** Security safeguards (i.e., hardware and software features, physical controls, operating procedures, management procedures, and various combinations of these) designed to detect and deny unauthorized access and permit authorized access to an information system.
- **Access control service:** A security service that protects against a system entity using a system resource in a way not authorized by the system's security policy.

Subjects, Objects, and Access Rights

The basic elements of access control are subject, object, and access rights. A ***subject*** is an entity capable of accessing objects. Generally, the concept of subject equates with that of process. Any user or application actually gains access to an object by means of a process that represents that user or application. The process takes on the attributes of the user, such as access rights.

A subject is typically held accountable for the actions he or she has initiated, and an audit trail may be used to record the association of a subject with security-relevant actions performed on an object by the subject.

Basic access control systems typically define three classes of subject, with different access rights for each class:

- **Owner:** This may be the creator of a resource, such as a file. For system resources, ownership may belong to a system administrator. For project resources, a project administrator or leader may be assigned ownership.
- **Group:** In addition to the privileges assigned to an owner, a named group of users may also be granted access rights, such that membership in the group is sufficient to exercise these access rights. In most schemes, a user may belong to multiple groups.
- **World:** The least amount of access is granted to users who are able to access the system but are not included in the categories owner and group for this resource.

An ***object*** is a resource to which access is controlled. In general, an object is an entity used to contain and/or receive information. Examples include records, blocks, pages, segments, files, portions of files, directories, directory trees, mailboxes, messages, and programs. Some access control systems also encompass bits, bytes, words, processors, communication ports, clocks, and network nodes.

The number and types of objects to be protected by an access control system depend on the environment in which access control operates and the desired trade-off between security on the one hand and complexity, processing burden, and ease of use on the other hand.

An ***access right*** describes the way in which a subject may access an object. Access rights could include the following:

- **Read:** User may view information in a system resource (e.g., a file, selected records in a file, selected fields within a record, some combination). Read access includes the ability to copy or print.
- **Write:** User may add, modify, or delete data in system resource (e.g., files, records, programs). Write access includes read access.
- **Execute:** User may execute specified programs.
- **Delete:** User may delete certain system resources, such as files or records.

- **Create:** User may create new files, records, or fields.
- **Search:** User may list the files in a directory or otherwise search the directory.

Access Control Policies

An access control policy dictates what types of access are permitted, under what circumstances, and by whom. Access control policies are generally grouped into the following categories:

- **Discretionary access control (DAC):** Access control based on the identity of the requestor and on access rules (authorizations) stating what requestors are (or are not) allowed to do. 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.
- **Mandatory access control (MAC):** Access control based on comparing security labels (which indicate how sensitive or critical system resources are) with security clearances (which indicate that system entities are eligible to access certain resources). This policy is termed *mandatory* because an entity that has clearance to access a resource may not, just by its own volition, enable another entity to access that resource.
- **Role-based access control (RBAC):** Access control based on user roles (i.e., a collection of access authorizations a user receives based on an explicit or implicit assumption of a given role). Role permissions may be inherited through a role hierarchy and typically reflect the permissions needed to perform defined functions within an organization. A given role may apply to a single individual or to several individuals.
- **Attribute-based access control (ABAC):** Access control based on attributes associated with and about subjects, objects, targets, initiators, resources, or the environment. An access control rule set defines the combination of attributes under which an access may take place.

DAC is the traditional method of implementing access control. MAC is a concept that evolved out of requirements for military information security. Both RBAC and ABAC have become increasingly popular.

These four policies are not mutually exclusive. An access control mechanism can employ two or even all four of these policies to cover different classes of system resources.

Discretionary Access Control

A general approach to DAC, as exercised by an operating system or a database management system, is an ***access matrix***. One dimension of the matrix consists of identified subjects that may attempt data access to the resources. Typically, this list consists of individual users or user groups, although access could be controlled for terminals, network equipment, hosts, or applications instead of or in addition to users. The other dimension lists the objects that may be accessed. At the greatest level of detail,

objects may be individual data fields. More aggregate groupings, such as records, files, or even the entire database, may also be objects in the matrix. Each entry in the matrix indicates the access rights of a particular subject for a particular object.

Part a of Figure 5.4 shows a simple example of an access matrix. Thus, user A owns files 1 and 3 and has read and write access rights to those files. User B has read access rights to file 1, and so on.

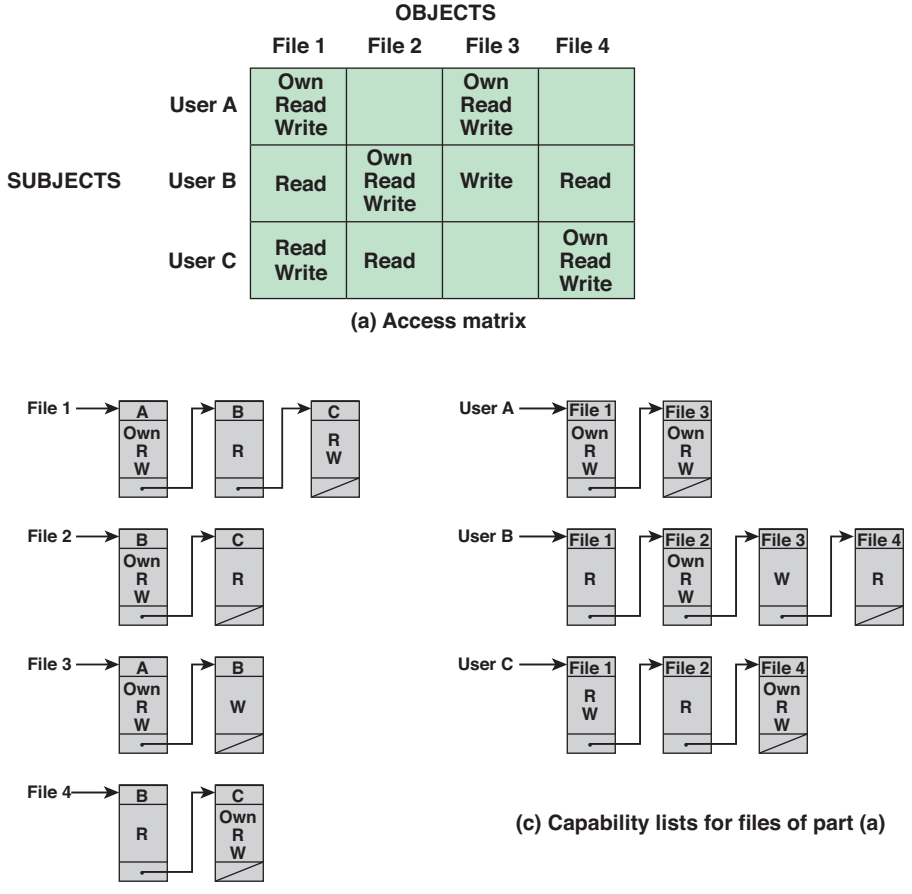

FIGURE 5.4 Example of Access Control Structures

In practice, an access matrix is usually sparse and is implemented by decomposition in one of two ways. The matrix may be decomposed by columns, yielding ***access control lists*** (ACLs); see part b of Figure 5.4. For each object, an ACL lists users and their permitted access rights. The ACL may contain a default, or public, entry. This allows users who are not explicitly listed as having special rights to a default set of rights. The default set of rights should always follow the rule of least privilege or read-only access, whichever is applicable. Elements of the list may include individual users as well as groups of users.