DOMAIN 1

SSCP®

Access Controls

ACCESS CONTROL IS CONCERNED with determining the allowed activities of legitimate users, mediating every attempt by a user to access a resource in the system. Access controls permit the security practitioner to specify what users can do, which resources they can access, and what operations they can perform on a system. Access controls provide the security practitioner with the ability to limit and monitor who has access to a system and to restrain or influence behavior on that system. In some systems, complete access is granted after successful authentication of the user, but most systems require more sophisticated and complex control. In addition to the authentication mechanism such as a password, access control is concerned with how authorizations are structured. Access control systems define what level of access an individual has to the information contained within a system based on predefined conditions such as authority level or group membership. Access control systems are based on varying technologies, including passwords, hardware tokens, biometrics, and certificates, to name a few. Each access control system offers different levels of confidentiality, integrity, and availability to the user, the system, and stored information.

TOPICS

The following topics are addressed in this chapter:

- **Implement authentication mechanisms**
  - **Single/multifactor authentication**
  - **Single sign-on**
  - **Offline authentication**
  - **Device authentication**
- **Operate internetwork trust architectures (e.g., extranet, third-party connections, federated access)**
  - **One-way trust**
  - **Two-way trust**
  - **Transitive trust**
- **Administer identity management lifecycle**
  - **Authorization**
  - **Proofing**
  - **Provisioning**
  - **Maintenance**
  - **Entitlement**
- **Implement access controls (e.g., subject-based, object-based)**
  - **Mandatory**
  - **Non-discretionary**
  - **Discretionary**
  - **Role-based**
  - **Attribute-based**

OBJECTIVES

A Systems Security Certified Practitioner (SSCP) is expected to demonstrate knowledge in how different access control systems operate and are implemented to protect the system and its stored data. In addition, the security practitioner must demonstrate knowledge in the following:

- Account management
- Access control concepts
- Attack methods that are used to defeat access control systems

ACCESS CONTROL CONCEPTS

Security practitioners planning to implement an access control system should consider three constructs: access control policies, models, and mechanisms. Access control policies are high-level requirements that specify how access is managed and who may access information under what circumstances. For instance, policies may pertain to resource usage within or across organizational units or may be based on need-to-know, competence, authority, obligation, or conflict-of-interest factors. At a high level, access control policies are enforced through a mechanism that translates a user's access request, often in terms of a structure that a system provides. An access control list is an example of an access control mechanism. Access control models bridge the gap between policy and mechanism. Rather than attempting to evaluate and analyze access control systems exclusively at the mechanism level, the security practitioner should use security models, which are usually written to describe the security properties of an access control system. Security models are formal presentations of the security policy enforced by the system and are useful for proving the theoretical limitations of a system. Discretionary access control (DAC), which allows the creator of a file to delegate access to others, is one of the simplest examples of a model.

Access controls provide for the ability to control "who" can do "what" with respect to data, applications, systems, networks, and physical spaces. In the simplest of terms, an access control system grants system users only those rights necessary for them to perform their respective jobs. The following definitions of key terms will be helpful for the security practitioner:

- A *subject* is an active entity that requests access to an object or the data within an object. The subject is the actor.
- An *object* is a passive entity being accessed, or the item being acted upon.

- *Access* is the ability of a subject to do something, such as read, create, delete, or modify. Access is also considered the flow of information between a subject and object.
- *Access control* is focused on the security features that control how subjects and objects communicate and interact with each other and the flow of information.

Applying Logical Access Control in Terms of Subjects

An access control subject is an active entity and can be any user, program, or process that requests permission to cause data to flow from an access control object to the access control subject or between access control objects.

Access control subjects include

- Authorized users
- Unauthorized users
- Applications
- Processes
- Systems
- Networks

The authorization provided to the access control subject by an access control system can include but is not limited to the considerations shown in Table 1.1.

TABLE 1.1 Access Control Subject/Object Comparison

ACCESS CONTROL SUBJECT	ACCESS CONTROL OBJECT
Temporal—time of day, day of request.	Data content of the object.
Locale from where the access control subject was authenticated.	The access control subject may be restricted from accessing all or part of the data within the access control object because of the type of data that may be contained within the object.
Inside or outside of the network.	Transaction restrictions may also apply.
Password or token utilized.	
An individual access control subject may have different rights assigned to specific passwords that are used during the authentication process.	

The attributes of a subject are referred to as privilege attributes or sensitivities. When these attributes are matched against the control attributes of an object, privilege is either granted or denied.

In a typical access control system, there are additional subject-specific requirements:

- A secure default policy should be applied to any newly created subject.
- The attributes of the subject should not be expressed in terms that can easily be forged, such as an IP address.
- The system should provide for a default deny on all permissions for the subject, thereby requiring that access to any object be explicitly created by an administrator.
- In the absence of policy for a given subject, the default policy should be interpreted as default deny.
- A user ID should remain permanently assigned to a subject.

The configuration of privileges in access control for an individual subject affords maximum granularity to the security practitioner. In systems with perhaps hundreds or thousands of users, this granularity can quickly become a management burden. By incorporating multiple subjects with similar permissions within a group, the granularity is thereby coarsened and the administration of the access control system is simplified. For example, look at Figure 1.1. Notice that the access control entry for Student\NHM_E4 has five permissions associated with it. Managing these permissions for a single user is not very difficult, nor does it present the security practitioner with a situation that would be too challenging to document and manage over the lifecycle of the SSCP Access Control Example document. However, even with just a single user and the permissions associated with their access to the document, there are a minimum of 10 different possible outcomes that the security practitioner will have to keep in mind as potential access levels for the user with regards to the document if the *standard* permissions are considered only. When the *special* permissions are added as well, the number jumps to a minimum of 26 potential outcomes if all permissions were employed.

The total number of permissions available for use in a Windows operating system such as Windows 7 or Windows 8 that uses the NTFS file system would be 14 if all possible standard and special permission options were included for potential use. This would include the five standard permissions, the additional eight special permissions available, as well as the 14th permission, which would be **no access** (full control = DENY). The security practitioner always needs to keep in mind what permissions have been assigned

to a resource, either explicitly or implicitly, and, by extension, which permission(s) have not been assigned. A complete listing of the NTFS special permissions is as follows:

- Full control
- Traverse folder/execute file
- List folder/read data
- Read attributes
- Read extended attributes
- Create files/write data
- Create folders/append data
- Write attributes
- Write extended attributes
- Delete
- Read permissions
- Change permissions
- Take ownership

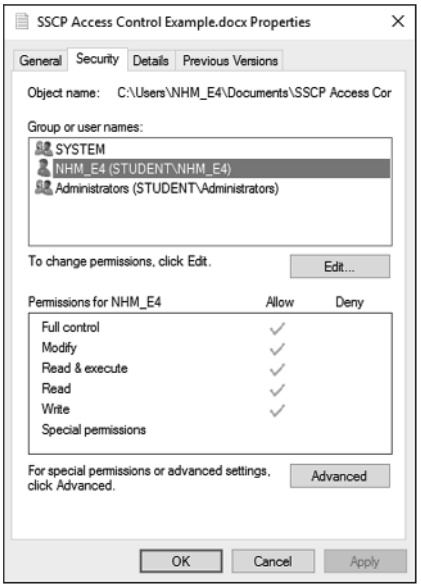

FIGURE 1.1 **Subject Group Access Control—User**

The security practitioner needs to keep in mind that permissions can be assigned to the user, or set, as either ALLOW or DENY, as shown in Figure 1.2.

When Figure 1.3 is examined, one will notice that there are access control entries for multiple users. Each user has the potential to have different permissions assigned to them by the owner of the SSCP Access Control Example document. As a result, the security practitioner now has a situation that will require them to manage and document permissions assigned to multiple users. Managing these permissions for multiple users is more challenging, as there are a minimum of 10 different possible outcomes multiplied by the four users that the security practitioner will have to keep in mind as potential access levels for the user concerning the document if the standard permissions are considered only. This means that the security practitioner will now have to keep track of a potential minimum of 40 different user/permission combinations. When the special permissions are added as well, the number jumps to a minimum of 26 potential outcomes multiplied by the four users, which is a minimum of 104 outcomes, if all permissions were employed.

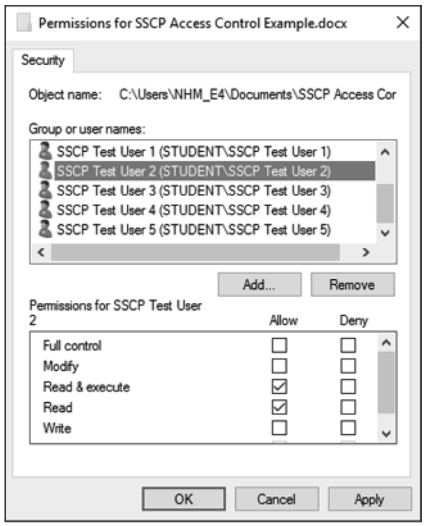

FIGURE 1.2 Subject Group Access Control—User permissions Allow and Deny

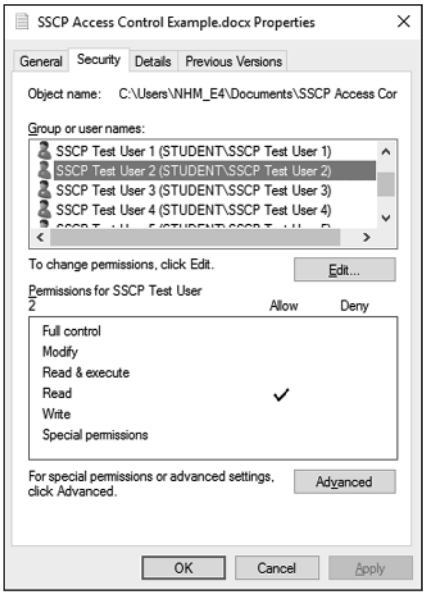

FIGURE 1.3 Subject Group Access Control—Multiple Users

In Figure 1.4, the access control entry for the Student\Administrators group has five permissions associated with it. On the surface, this group presents the same scenario to the security practitioner that the Student\NHM_E4 user from Figure 1.1 does, and the same minimum number of outcomes for both the standard and special permissions. The key difference for the security practitioner is the ability to leverage the power of membership in the group in order to simplify the management overhead involved with assigning, documenting, and tracking permission combinations. By placing users with similar access needs into a single group, the security practitioner will be able to use the power of the group to ***assign once and manage many***, resulting in two key advantages. The first advantage is that the security practitioner will be able to streamline the permission provisioning process for the users requiring access to the SSCP Access Control Example document, resulting in less management overhead as more users require access over the lifetime of the document. The second advantage is that the likelihood of an incorrect permission assignment being made for one or more users, leading to either too little or too much access to the SSCP Access Control Example document, is greatly reduced if the security practitioner is focused on ensuring that the group permissions are assigned based on job role or access need, and as a result, that membership in the groups are managed the same way. The security practitioner should always strive to use group membership as the basis for assigning access to resources when planning access control solutions, as it offers more flexibility and forces the data owner to carefully consider the requirements for data access *prior* to assignment.

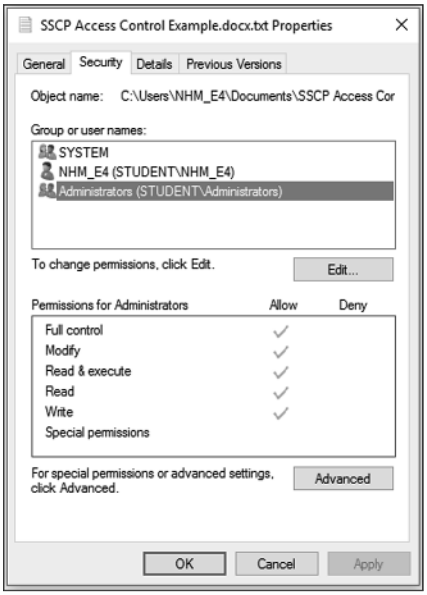

FIGURE 1.4 **Subject Group Access Control—Group**

Applying Logical Access Control in Terms of Objects or Object Groups

An access control object is a passive entity that typically receives or contains some form of data. The data can be in the form of a file, can be in the form of a program, or may be resident within system memory.

Access control objects include:

- Data
- Applications
- Systems
- Networks
- Physical space, for example, the data center

Typical access control object considerations can include but are not limited to the following:

- Restrict access to operating system configuration files and their respective directories to authorized administrators.
- Disable write/modify permissions for all executable files.
- Ensure that newly created files inherit the permissions of the directory in which they were created.
- Ensure that subdirectories cannot override the permissions of parent directories unless specifically required by policy.
- Log files should be configured to only permit appending data to mitigate the risk of a log file's contents being purposely deleted or overwritten by a malicious user or process.
- Encryption of data at rest can afford additional security and should be a consideration in the determination of the policies for access control objects.

The configuration of privileges to access an individual object affords maximum granularity. It is common today for the number of objects within an access control system to number in the tens or even hundreds of thousands. While configuring individual objects affords maximum control, this granularity can quickly become an administrative burden. It is a common practice to assign the appropriate permissions to a directory, and each object within the directory inherits the respective parent directory permissions. By incorporating multiple objects with similar permissions or restrictions within a group or directory, the granularity is thereby coarsened and the administration of the access control system is simplified. Figure 1.5 shows the permission entries for the SSCP_1 folder, a child object of the parent SSCP folder object. As a child object, the SSCP_1 folder

automatically upon creation is set to accept inheritable permissions from the object's parent as indicated by the button with the text "Disable inheritance." This setting ensures that *all* objects created within the SSCP_1 folder will inherit the existing access control settings already in place at the parent object, the SSCP folder, in addition to whatever new settings are assigned once the object is created by the object owner.

The "Replace all child object permission entries with inheritable permissions from this object" setting is never set by default and must be manually selected to be used. This setting indicates that the object owner has decided to break the original hierarchical inheritance chain between the parent and child objects and, as a result, all additional hierarchical generations that are created below the child as well. Further, the breaking of the hierarchical inheritance chain at this point will result in all new objects that are created being blocked from inheritance of the parental object's existing access control settings, thus ensuring that these newly created child objects are not bound by *any* of the access control settings in place at the parent object.

Figure 1.6 illustrates this exact outcome, as the language "This will replace explicitly defined permissions on all descendants of this object with inheritable permissions from" indicates. This action will also effectively promote the current child to the status of a parent for any/all newly created objects at this level, as well as all sublevels, ensuring that these objects inherit their access control settings from their newly created parent object, not the original parent object that they are now disassociated from due to the breaking of the inheritance chain.

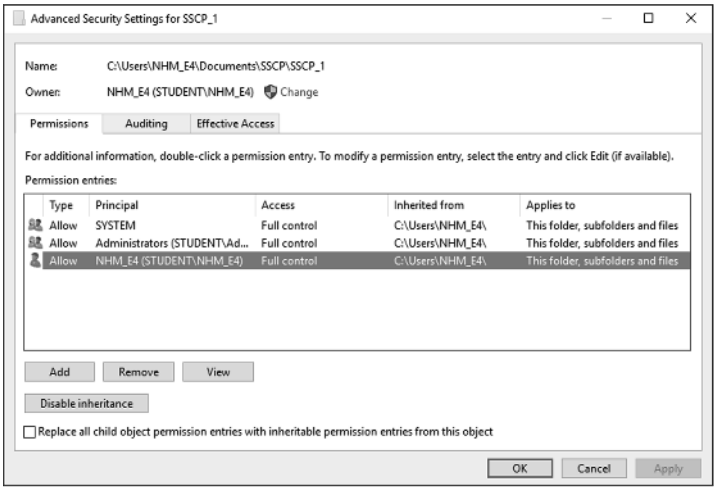

FIGURE 1.5 **Hierarchical permission inheritance**

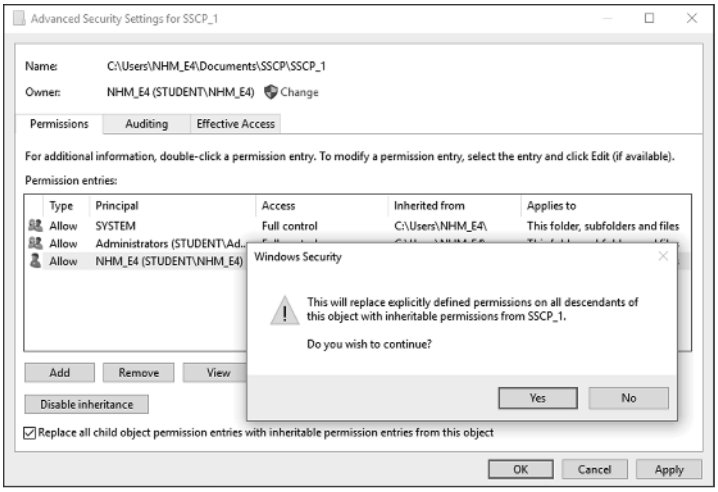

FIGURE 1.6 **Replacement of all child object permissions**

IMPLEMENTING ACCESS CONTROLS

Access controls are used in a system to ensure that authorization and authentication are properly implemented. Authorization is the process where requests to access a particular resource should be granted or denied. Authentication is providing and validating identity. The SSCP should be familiar with the different types of access control methods available, as well as how they work.

Discretionary Access Control

A *Discretionary Access Control* (DAC) policy is a means of assigning access rights based on rules specified by users. This class of policies includes the file permissions model implemented by nearly all operating systems. In Unix, for example, a directory listing might yield "... rwxr-xr-x ... SSCP File 1.txt," meaning that the owner of SSCP File 1.txt may read, write, or execute it, and that other users may read or execute the file but not write it. The set of access rights in this example is {read, write, execute}, and the operating system mediates all requests to perform any of these actions. Users may change the permissions on files they own, making this a discretionary policy. A mechanism

implementing a DAC policy must be able to answer the question: "Does subject Sayge have right Read for object SSCP File 1?" More practically, the same information could also be represented as an access control matrix. Each row of the matrix corresponds to a subject and each column to an object. Each cell of the matrix contains a set of rights. Table 1.2 shows an example of an access control matrix.

TABLE 1.2 An Access Control Matrix

	SSCP FILE 1	SSCP FILE 2
Aidan	Read | Write | eXecute	Read | eXecute
Sayge	Read	Read | Write

Systems typically store the information from this matrix either by columns or by rows. An implementation that stores by columns is commonly known as an access control list (ACL). File systems in Windows and Unix typically use such an implementation: Each file is accompanied by a list containing subjects and their rights to that file. An implementation that stores by rows is commonly known as a capability list. For example, it is easy in an ACL implementation to find the set of all subjects who may read a file, but it is difficult to find the set of all files that a subject may read.

The underlying philosophy in DAC is that subjects can determine who has access to their objects. In Discretionary Access Control (DAC), the owner of the access control object would determine the privileges (i.e., read, write, execute) of the access control subjects. In the DoD 5200.28-STD, Department of Defense Standard Department of Defense Trusted Computer System Evaluation Criteria, Discretionary Access Control is defined as "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 certain access permission is capable of passing that permission (perhaps indirectly) on to any other subject (unless restrained by mandatory access control)."[1]

This methodology relies on the discretion of the owner of the access control object to determine the access control subject's specific rights. Hence, security of the object is literally up to the discretion of the object owner. DACs are not very scalable; they rely on the decisions made by each individual access control object owner, and it can be difficult to find the source of access control issues when problems occur.

Rule Set–Based Access Controls

Rule Set–Based Access Controls (RSBAC) are discretionary controls giving data owners the discretion to determine the rules necessary to facilitate access. RSBAC is an open source access control framework for current Linux kernels, which has been in use since January 2000 (version 1.0.9a). RSBAC allows full fine-grained control over objects (files, processes,

users, devices, etc.), memory execution prevention (PaX, NX), real-time integrated virus detection, and much more. The RSBAC framework logic is based on the work done for the Generalized Framework for Access Control (GFAC) by Abrams and LaPadula.[2]

All security relevant system calls are extended by security enforcement code. This code calls the central decision component, which in turn calls all active decision modules (the different modules implementing different security models) and generates a combined final decision. This decision is then enforced by the system call extensions. Decisions are based on the type of access (request type), the access target, and the values of attributes attached to the subject calling and to the target to be accessed. Additional independent attributes can be used by individual modules. All attributes are stored in fully protected directories, one on each mounted device. Thus, changes to attributes require special system calls to be provided.

RSBAC works at the kernel level and affords flexible access control based on several modules:

- Mandatory Access Control (MAC) module
- Privacy module (PM)
- Function Control module (FC)
- File Flag module (FF)
- Malware Scan module (MS)
- Role Compatibility module (RC)
- Function Control module (FC)
- Security Information Modification module (SIM)
- Authentication module (Auth)
- Access Control List module (ACL)

Figure 1.7 illustrates the RSBAC access request process.

✔ Try It for Yourself—With a Live CD

Test RSBAC with a Debian-based live CD, or use it on a USB key/drive. This will allow full testing of RSBAC functionality without having to install it. Just insert the CD or USB key, reboot, and try it!

Download here:

```
https://www.rsbac.org/download
```

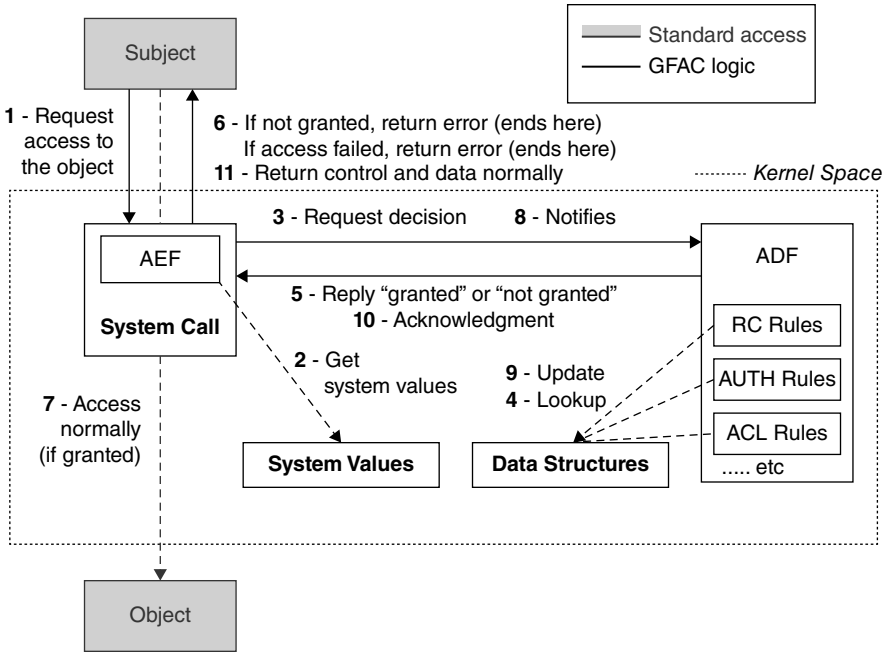

FIGURE 1.7 **The Rule Set Based Access Control (RSBAC) Generalized Framework for Access Control (GFAC) logic for data access request**

Role-Based Access Controls

With role-based access control, access decisions are based on the roles that individual users have as part of an organization. Users take on assigned roles (such as Backup Operator, Performance Log Users, and Administrators). The process of defining roles should be based on a thorough analysis of how an organization operates and should include input from a wide spectrum of users in an organization.

Access rights are grouped by role name, and the use of resources is restricted to individuals authorized to assume the associated role. For example, within a network the role of Performance Log User can include operations to open, read, save, and delete log files; and the role of Backup Operators can be limited to activities related strictly to the backing up of specified data, but not be designed to include the activities associated with restoring the data if required.

Under the RBAC framework, users are granted membership into roles based on their competencies and responsibilities in the organization. The operations that a user is permitted to perform are based on the user's role. User membership in roles can be revoked easily and new memberships established as job assignments dictate. Role associations can be established when new operations are instituted, and old operations can be deleted as organizational functions change and evolve. This simplifies the administration and management of privileges; roles can be updated without updating the privileges for every user on an individual basis.

Under RBAC, when a user is associated with a role, the user should be given no more privileges than are necessary to perform their role. This concept of least privilege requires identifying the user's job functions, determining the minimum set of privileges required to perform that function, and restricting the user to a role with those privileges and nothing more. In less precisely controlled systems, this is often difficult to achieve. Someone assigned to a job category may be allowed more privileges than needed because it is difficult to tailor access based on various attributes or constraints. Since many of the responsibilities overlap between job categories, maximum privilege for each job category could cause undesired or unlawful access.

Under RBAC, roles can have overlapping responsibilities and privileges; that is, users belonging to different roles may need to perform common operations. Role hierarchies can be established to provide for the natural structure of an enterprise. A role hierarchy defines roles that have unique attributes and that may contain other roles; that is, one role may implicitly include the operations that are associated with another role.

✔ Try It for Yourself—RBAC in a Box

Now you will interact with RBAC first hand.

What's Needed?

A Windows-based computer and a user account with administrative rights.

How to Do It

Use the following step-by-step guidance:

1. Open the Control Panel from the Windows desktop, or simply type **control panel** in the Run line and hit Enter. (Alternately, you can type **`compmgmt.msc`** directly in the Run line to bypass the Control Panel and go directly to the Computer Management Console.)
2. From the Control Panel open Computer Management.
3. From within the Computer Management Console go to the Local Users and Groups item and then select the Groups folder in the left window. You will see the various groups that are already present on the system displayed in the right portion of the window. (See Figure 1.8.)

CONTINUES

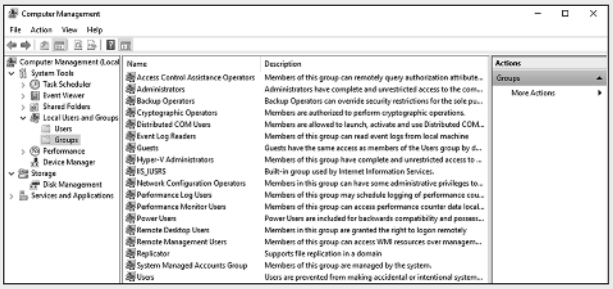

FIGURE 1.8 **Local Users and Groups in a Windows 7 computer**

4. Select a group to examine the permissions for, such as the Backup Operators group or the Power Users group.
5. Open the Windows Explorer window to examine the files and folders on the system.
6. Pick a file or folder from within the Windows Explorer window in order to examine the RBAC permissions for the group you chose in step 4. Any file or folder in the computer may be used, but it would be best if one were created specifically to test with, so existing file permissions are not mistakenly changed.
7. Once the file or folder has been selected, right-click on it from within the Windows Explorer window and choose Properties from the shortcut menu that pops up. When the Properties window for the file or folder has opened, click on the Security tab (second in line, moving left to right). Something similar to Figure 1.9 should be displayed.

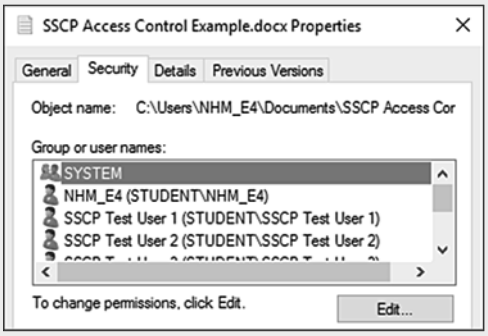

FIGURE 1.9 **File permissions *before* adding an RBAC example in Windows 7/Windows 8 computer**

8. Click the Edit button and then click the Add button on the Security tab that will appear once the Edit button has been clicked.
9. Use the group that was selected in step 4. Type the name of the group into the dialog within the Select Users or Groups screen that has appeared, as shown in Figure 1.10.

   Please note the following: Type the group name into the window in the format shown in Figure 1.10, which is **Machine Name\Group Name**. You can find the machine name listed under the "From this location" area, right above where the machine name\group name information will be typed.

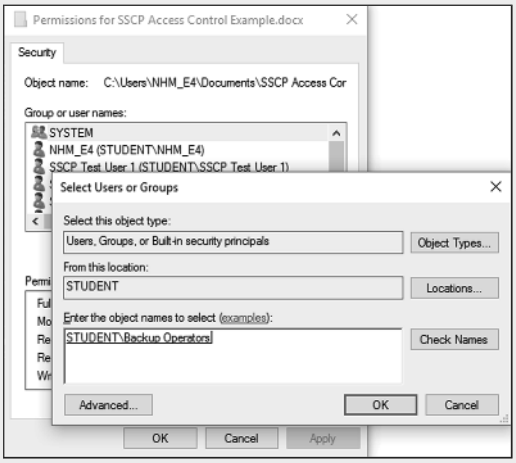

FIGURE 1.10 The Select Users or Groups screen that appears after the Edit button has been clicked

10. Once done entering the group information, click the OK button. Something similar to Figure 1.11 should be displayed.

CONTINUES

CONTINUED

11. Figure 1.11 shows the Backup Operators group and the Role Based Access Control permissions associated with the group. RBAC has been successfully demonstrated!

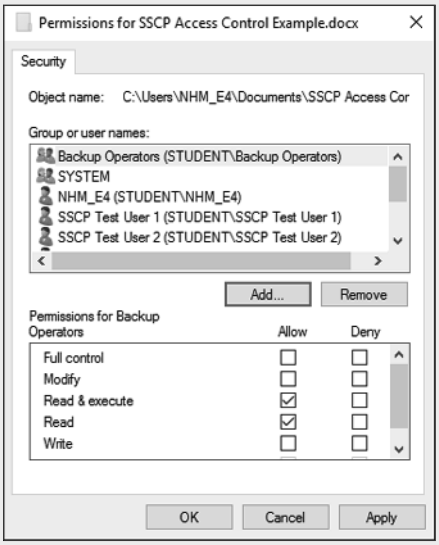

FIGURE 1.11 Folder permissions *after* adding an RBAC example in Windows 7/ Windows 8; resultant set of permissions for Backup Operators group

Constrained User Interface

Constrained User Interface (CUI) is a methodology that restricts the user's actions to specific functions by not allowing them to request functions that are outside of their respective level of privilege or role. One of the most common examples of a Constrained User Interface can be found in online banking applications and ATMs where the limited menus are not readily apparent until after the user has properly authenticated, thereby establishing their respective role/level of privilege.

Three major types of restricted interfaces exist: menus and shells, database views, and physically constrained interfaces.

- **Menu and Shells**—When menu and shell restrictions are used, the options users are given are the commands they can execute. For example, if an administrator wants users to be able to execute only one program, that program would be the only choice available on the menu. This limits the users' functionality. A shell is a type of virtual environment within a system. It is the user's interface to the operating system and works as a command interpreter. If restricted shells were used, the shell would contain only the commands the administrator wants the users to be able to execute.
- **Database Views**—Database views are mechanisms used to restrict user access to data contained in databases.

- **Physically Constraining a User Interface**—Physically constraining a user interface can be implemented by providing only certain keys on a keypad or certain touch buttons on a screen. You see this when you get money from an ATM. This device has a type of operating system that can accept all kinds of commands and configuration changes, but it is physically constrained from being able to carry out these functions.

Another type of CUI is often referred to as View-Based Access Control (VBAC); it is most commonly found in database applications to control access to specific parts of a database. The CUI in VBAC restricts or limits an access control subject's ability to view or perhaps act on "components" of an access control object based on the access control subject's assigned level of authority. Views are dynamically created by the system for each user-authorized access.

Simply put, VBAC separates a given access control object into subcomponents and then permits or denies access for the access control subject to view or interact with specific subcomponents of the underlying access control object.[3]

Content-Dependent Access Control

Content-Dependent Access Control (CDAC) is used to protect databases containing sensitive information. CDAC works by permitting or denying the access control subjects access to access control objects based on the explicit content within the access control object. An example would be the use of CDAC in a medical records database application where a health-care worker may have been granted access to blood test records. If that record contains information about an HIV test, the health-care worker may be denied access to the existence of the HIV test and the results of the HIV test. Only specific hospital staff would have the necessary CDAC access control rights to view blood test records that contain any information about HIV tests.

While high levels of privacy protection are attainable using CDAC, they come at the cost of a great deal of labor in defining the respective permissions. It should be further noted that CDAC comes with a great deal of overhead in processing power as it must scan the complete record to determine if access can be granted to a given access control subject. This scan is done by an arbiter program to determine if access will be allowed.

Context-Based Access Control

Context-Based Access Control (CBAC) is used in firewall applications to extend the firewall's decision-making process beyond basic ACL decisions to decisions based on state as well as application-layer protocol session information. A static packet-filtering firewall is a good example of a firewall that does not use CBAC. It looks at each packet and compares the packet to an ACL rule base to determine if the packet is to be allowed or

denied. A stateful inspection firewall is a good example of a firewall that uses CBAC. The firewall also considers the "state of the connection"; i.e., if a packet arrives that is part of a continuing session that had previously been permitted to pass through the firewall, then subsequent packets that are part of that session are allowed to pass without the overhead associated with comparing the packet to the ACL rules. CBAC affords a significant performance enhancement to a firewall.[4]

CBAC is often confused with CDAC, but they are two completely different methodologies. While CDAC makes decisions based on the content within an access control object, CBAC is not concerned with the content; it is concerned only with the context or the sequence of events leading to the access control object being allowed through the firewall.

In the example of blood test records for CDAC in the previous section, the access control subject would be denied access to the access control object because it contained information about an HIV test. CBAC could be used to limit the total number of requests for access to any blood test records over a given period of time. Hence, a health-care worker may be limited to accessing the blood test database more than 100 times in a 24-hour period.

While CBAC does not require that permissions be configured for individual access control objects, it requires that rules be created in relation to the sequence of events that precede an access attempt.

Temporal Isolation (Time-Based) Access Control

Temporal Isolation (Time-Based) Access Control is used to enhance or extend the capabilities of RBAC implementations. This combined methodology is often referred to as Temporal Role-Based Access Control (TRBAC).[5] TRBAC supports periodic role enabling and disabling and temporal dependencies among such actions. Such dependencies expressed by means of role triggers (active rules that are automatically executed when the specified actions occur) can also be used to constrain the set of roles that a particular user can activate at a given time instant. The firing of a trigger may cause a role to be enabled/disabled either immediately or after an explicitly specified amount of time. Enabling/disabling actions may be given a priority that may help in solving conflicts, such as the simultaneous enabling and disabling of a role. As expected, the action with the highest priority is executed. TRBAC effectively applies a time limitation to when a given role can be activated for a given access control subject.

- A high-level *top secret* role would be assigned to a given access control subject during the normal 8 a.m. to 5 p.m. working hours.
- A lower-level *confidential* role would be assigned to the same access control subject during the 5 p.m. to 8 a.m. nonworking hours.

To decrease the effort associated with assigning TRBAC rules to many individual access control subjects, most implementations of TRBAC assign the temporal-based classification levels to the access control objects rather than to the access control subject. Hence, a given access control object would have a temporal-based classification level that is effective against all access control subjects.

Temporal extensions are also used to enhance other access control methodologies. It is common today to find access control devices that support time-based access control rules. The temporal enhancement of the access control rule only allows the rule to be effective during the specified time period.

Nondiscretionary Access Control

According to the United States National Institute of Standards and Technology (NIST), in general, all access control policies other than DAC are grouped in the category of non-discretionary access control (NDAC). As the name implies, policies in this category have rules that are not established at the discretion of the user. Non-discretionary policies establish controls that cannot be changed by users, but only through administrative action.[6]

Mandatory Access Control

Mandatory Access Control (MAC) is typically used in environments requiring high levels of security such as government or military systems. In MAC, the inherent problems of trying to rely on each system owner to properly control access to each access control object is eliminated by having the system participate in applying a mandatory access policy; the system owner applies the "need to know" element. This policy affords typically three object classification levels: *top-secret*, *secret*, and *confidential*. Each access control system subject (users and programs) is assigned clearance labels, and access control system objects are assigned sensitivity labels. The system then automatically provides the correct access rights based on comparing the object and subject labels. MAC allows multiple security levels of both objects and subjects to be combined in one system securely.

Mandatory access control (MAC) policy means that access control policy decisions are made by a central authority, not by the individual owner of an object, and the owner cannot change access rights. An example of MAC occurs in military security, where an individual data owner does not decide who has a top secret clearance, nor can the owner change the classification of an object from top secret to secret. The need for a MAC mechanism arises when the security policy of a system dictates that

1. Protection decisions must not be decided by the object owner.
2. The system must enforce the protection decisions (i.e., the system enforces the security policy over the wishes or intentions of the object owner). Usually a labeling mechanism and a set of interfaces are used to determine access based on the

MAC policy; for example, a user who is running a process at the secret classification should not be allowed to read a file with a label of top secret. This is known as the *simple security rule*, or *no read up*. Conversely, a user who is running a process with a label of Secret should not be allowed to write to a file with a label of Confidential. This rule is called the **-property* (pronounced "star property") or *no write down*. The *-property is required to maintain system security in an automated environment. A variation on this rule called the *strict *-property* requires that information can be written at, but not above, the subject's clearance level. Multilevel security models such as the Bell–LaPadula Confidentiality and Biba Integrity models are used to formally specify this kind of MAC policy.

Attribute-Based Access Control

The following is a high-level definition of ABAC, according to NIST Special Publication 800-162, *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*:[7]

> *Attribute Based Access Control (ABAC) is an access control method where subject requests to perform operations on objects are granted or denied based on assigned attributes of the subject, assigned attributes of the object, environment conditions, and a set of policies that are specified in terms of those attributes and conditions.*

Here are some vocabulary terms that will help the security practitioner understand and apply the definition:

- *Attributes* are characteristics of the subject, object, or environment conditions. Attributes contain information given by a name-value pair.
- A *subject* is a human user or NPE, such as a device that issues access requests to perform operations on objects. Subjects are assigned one or more attributes. For the purpose of this document, assume that subject and user are synonymous.
- An *object* is a system resource for which access is managed by the ABAC system, such as devices, files, records, tables, processes, programs, networks, or domains containing or receiving information. It can be the resource or requested entity, as well as anything upon which an operation may be performed by a subject including data, applications, services, devices, and networks.
- An *operation* is the execution of a function at the request of a subject upon an object. Operations include read, write, edit, delete, copy, execute, and modify.
- *Policy* is the representation of rules or relationships that makes it possible to determine if a requested access should be allowed, given the values of the attributes of the subject, object, and possibly environment conditions.

- *Environment conditions* represent the operational or situational context in which access requests occur. Environment conditions are detectable environmental characteristics. Environment characteristics are independent of subject or object and may include the current time, day of the week, location of a user, or current threat level.

Separation of Duties

This aspect of access control establishes guidelines that require that no single person should perform a task from beginning to end and that the task should be accomplished by two or more people to mitigate the potential for fraud in one person performing the task alone. Separation of duties is a key element in the Clark–Wilson formal model.

SECURITY ARCHITECTURE AND MODELS

Security architects often use established security models as points of reference in design work. Established, tested models identify the major components in a security solution and how they interact. Chief among these models are the Bell–LaPadula confidentiality model, and the Biba and Clark–Wilson integrity models.

Bell–LaPadula Confidentiality Model[8]

The Bell–LaPadula model was designed as an architectural reference for controlling access to sensitive data in government and military applications. The components of the model are subjects, objects, and an access control matrix. *Objects* (access targets) are classified into a hierarchy of security levels based on sensitivity, from low to high. If information has been previously classified (top secret, secret, etc.), then classification levels corresponding to the organization's policy are used. *Subjects* (actors)—which may be human actors, application programs, or system processes—are assigned security levels called *clearance levels*. The relation between the sensitivity level of objects and the clearance level of subjects is defined in the *access control matrix*. The access control matrix defines permissions (read-only, read/write, append, execute) for each clearance level and object classification. Each access operation is defined within the matrix by a subject, object, and access permission triple. The matrix provides assurance that the confidentiality of the system will remain stable despite transitions in state; that is, a system that is in a secure state before an operation will be in the same secure state at the conclusion of the operation.

The basic tenet of Bell–LaPadula is that a given subject can read objects at the same or lower sensitivity level, but not those at a higher sensitivity level; this is called the *simple*

security property and can be remembered as "no read up." The simple property is usually sufficient for implementing systems that control access to classified documents and files when the files have corresponding read-only attributes. However, it does not take into consideration the possibility that a subject may add, append, or transmit sensitive information to an area of lower sensitivity and thus create a channel that defeats the access control mechanism. Bell–LaPadula adds another property to counteract this called the star (*) property. The * property blocks the channel between areas of different sensitivities such that when a subject has accessed an object for a read operation, then objects at a lower sensitivity level cannot be accessed for create and modify operations ("no write down"). Covert channels, such as backup and monitoring channels and image capture utilities, still present a risk for systems designed using Bell–LaPadula confidentiality models as these processes may be used for legitimate as well as illegitimate purposes.

Bell–LaPadula is not without its limitations. It is concerned only with confidentiality and makes no mention of other properties (such as integrity and availability) or more sophisticated modes of access. These have to be addressed through other models. More importantly, it does not address important confidentiality goals such as need-to-know, or the ability to restrict access to individual objects based on a subject's need to access them. Since Bell–LaPadula does not provide a mechanism for a one-to-one mapping of individual subjects and objects, this also needs to be addressed by other models.

Biba[9] and Clark–Wilson Integrity Models[10]

Like Bell–LaPadula, Biba is also a lattice-based model with multiple levels. It uses the same modes of access (read, write, and read/write) and describes interactions between subjects and objects. Where Biba differs most obviously is that it is an integrity model: It focuses on ensuring that the integrity of information is being maintained by preventing corruption. At the core of the model is a multilevel approach to integrity designed to prevent unauthorized subjects from modifying objects. Access is controlled to ensure that objects maintain their current state of integrity as subjects interact with them. Instead of the confidentiality levels used by Bell–LaPadula, Biba assigns integrity levels to subjects and objects depending on how trustworthy they are considered to be. Like Bell–LaPadula, Biba considers the same modes of access but with different results. Table 1.3 compares the BLP and Biba models.

For example, consider a subject that wishes to add two numbers together. The subject needs information that is reasonably accurate to two decimal places and has different values to choose from. Some of these values are accurate to more than two decimal places. Some are less accurate. To prevent corruption, the subject must only use information that is at least as accurate as two decimal places; information that is accurate only to one decimal place must not be used or corruption may occur.

TABLE 1.3 BLP and Biba Model Properties

PROPERTY	BLP MODEL	BIBA MODEL
ss-property	A subject cannot read/ access an object of a higher classification (no read up).	A subject cannot observe an object of a lower integrity level (no read down).
*-property	A subject can save an object only at the same or higher classification (no write down).	A subject cannot modify an object of a higher integrity level (no write up).
Invocation property	Not used.	A subject cannot send logical service requests to an object of a higher integrity.

Source: Hare, C., "Policy Development," *Information Security Management Handbook*, 6th ed., Tipton, H.F. and Krause, M., Eds., Auerbach Publications. New York, 2007.

In the * integrity property, a given subject has the ability to write information to different types of objects with differing levels of integrity or accuracy. In this case, the subject must be prevented from corrupting objects that are more accurate than it is. The subject should then be allowed to write to objects that are less accurate, but not to objects that are more accurate. To allow otherwise may result in corruption. Biba also addresses the problem of one subject getting a more privileged subject to work on their behalf. In the invocation property, Biba considers a situation where corruption may occur because a less trustworthy subject was allowed to take advantage of the capabilities of a more trustworthy subject by invoking their powers. According to Biba, this must be prevented or corruption could occur.

David D. Clark and David R. Wilson developed their Clark–Wilson integrity model to address what they viewed as shortcomings in the Bell–LaPadula and Biba models.[11] While these models were useful for protecting classified information from unauthorized access or leakage to unclassified systems, they did not provide any framework to prevent corruption of data (either maliciously or unintentionally) during processing of the data. Clark–Wilson's model addresses this risk using the idea of a well-formed transaction operating on the data. The components of this model also form a triple: authenticated principals (users), programs acting on data (transaction processes), and the data items themselves. Each triple or relation between user, transaction, and data item must be maintained in the system.

Systems designed to enforce the Clark–Wilson integrity policy consist of well-formed transactions, that is, transactions that maintain a consistent level of integrity between the initial and end state. Integrity verification processes ensure the integrity of data items before, during, and after a transaction. Clark–Wilson also protects against malicious users

by requiring separation of duties between people who can create relations used in a process and those who can execute the process.

Additional Models

Bell–LaPadula, Biba, and Clark–Wilson are all useful frameworks for designing so-called multilevel security (MLS) systems, in which information with various sensitivities or integrity requirements can be processed concurrently in a single system by users or actors with multiple levels of clearance or need to know. Some additional models that the security practitioner will want to familiarize themselves with are mentioned in the following sections.

Brewer–Nash (the Chinese Wall) Model

This model focuses on preventing conflict of interest when a given subject has access to objects with sensitive information associated with two competing parties. The principle is that users should not access the confidential information of both a client organization and one or more of its competitors. At the beginning, subjects may access either set of objects. Once, however, a subject accesses an object associated with one competitor, they are instantly prevented from accessing any objects on the opposite side. This is intended to prevent the subject from sharing information inappropriately between the two competitors even unintentionally. It is called the Chinese Wall Model because, like the Great Wall of China, once on one side of the wall, a person cannot get to the other side. It is an unusual model in comparison with many of the others because the access control rules change based on subject behavior.

Graham–Denning Model

Graham–Denning is primarily concerned with how subjects and objects are created, how subjects are assigned rights or privileges, and how ownership of objects is managed. In other words, it is primarily concerned with how a model system controls subjects and objects at a very basic level where other models simply assumed such control.

The Graham–Denning access control model has three parts: a set of objects, a set of subjects, and a set of rights. The subjects are composed of two things: a process and a domain. The domain is the set of constraints controlling how subjects may access objects. Subjects may also be objects at specific times. The set of rights govern how subjects may manipulate the passive objects. This model describes eight primitive protection rights called commands that subjects can execute to have an effect on other subjects or objects. The model defines eight primitive protection rights:

1. **Create Object**—The ability to create a new object
2. **Create Subject**—The ability to create a new subject

3. **Delete Object**—The ability to delete an existing object
4. **Delete Subject**—The ability to delete an existing subject
5. **Read Access Right**—The ability to view current access privileges
6. **Grant Access Right**—The ability to grant access privileges
7. **Delete Access Right**—The ability to remove access privileges
8. **Transfer Access Right**—The ability to transfer access privileges from one subject or object to another subject or object

Harrison–Ruzzo–Ullman Model

This model is very similar to the Graham–Denning model, and it is composed of a set of generic rights and a finite set of commands. It is also concerned with situations in which a subject should be restricted from gaining particular privileges. To do so, subjects are prevented from accessing programs or subroutines that can execute a particular command (to grant read access for example) where necessary.

IMPLEMENTING AUTHENTICATION MECHANISMS—IDENTIFICATION, AUTHENTICATION, AUTHORIZATION, AND ACCOUNTABILITY

The process flow involved in the implementation of authentication mechanisms is to identify, authenticate, and authorize. Identification is the process used to allow the access control subject to provide information as to their identity, which can be used to validate them. Authentication is the act of providing and validating identity within the access control system. Authorization is the process where requests to access a particular resource should be granted or denied, based on the outcome of the authentication process. One example of a technology used to provide authentication services within an access control system is Biometrics. The SSCP should be familiar with the identification, authentication, and authorization processes and how they work together to create accountability within access control systems.

Identification (Who Is the Subject?)

Identification asserts a unique user or process identity and provides for accountability. Identification of an access control subject is typically in the form of an assigned user name. This user name could be public information whether intentional or not. A good example is that in most networks, the user name that identifies the user for network access is also the identification used as the e-mail account identifier. Hence, all one would have to do to determine

the account holder's user name would be to know the account holder's e-mail address. An access control that relied on the user name alone to provide access would be an ineffective access control. To prove that the individual who presented the user name to the access control is the individual who the user name was assigned to, a secret is shared between the access control system and the respective user. This secret is the user's password and is used to authenticate that the user who is trying to gain access is in fact the user who owns the rights associated with the respective identification.

Methods (User ID, PIN, Account Number)

The three most common methods used to provide user identity in an access control system are

- User ID—User name and password combination assigned to the user
- PIN—Typically a four-digit numerical combination created by the user during a sign-up/on-boarding process
- Account number—Typically an eight- to sixteen-digit unique numerical sequence assigned to an individual by the owner of the system

Regardless of the method used (user ID, PIN, or account number), each one must be unique to be valid for any user. Further care must be taken so that users are not readily identifiable from that of another user's user ID. An example of this problem would be to simply use the user's first initial and last name as his user ID. Anyone knowing the user's first and last names would then easily know the user's user ID.

Registration of New Users

Manual user registration provides for the greatest granularity but is also regarded as having too high of an administrative burden to be effective. Today it is often replaced with an automated provisioning solution. Automated provisioning solutions (identity management) provide a framework for managing access control policies by role, interconnection with IT systems, workflows to guide sign-off, delegated administration, password management, and auditing.

Periodic Review of Access Levels

The periodic review of user access levels is no longer simply a best practice and has been incorporated into current regulations including Sarbanes–Oxley. The mandatory periodic review of user access levels is necessary to ensure that each user's privilege continues to be appropriate and reflects any changes in their access requirements as their role and or responsibilities within the enterprise change.

Clearance

The proper application of clearance is critical in systems where access controls are based on security labels such as implementations of access control using the Bell–LaPadula model. Access control systems using clearances typically do so using a trusted user directory. Access to the directory is available only after successful authentication, and the directory must be trusted. Clearance levels, like other general access levels, must routinely be verified against each user's actual requirements, designated access, and status.

Certificates play an important role today in improving trust within a user directory. Instead of simply looking up a user in a directory to determine the level of clearance, a certificate with additional attributes, such as clearance lifecycle, can be used to verify by its digital signature that the clearance is valid.

Authentication (Proof of Identity)

Authentication is the process of verification that the identity presented to the access control system belongs to the party that has presented it. The three common factors in authentication are something you know, something you have, and something you are. In network authentication, the identification of the user is authenticated using a secret password that only the user should know. This would be referred to as simple authentication. There are more complex authentication methodologies such as *dual factor authentication* that not only require the secret that the user knows but also require another layer of authentication in the form of something the user "has" in their possession (such as a security token) or something the user "is" (as in the case of biometric authentication, a fingerprint or retina scan). We will discuss complex authentication methodologies such as dual factor later in this chapter. Again, the objective of authentication is to prove the identity of the user who is asking for some type of access from the access control system.

Knowledge (Static Passwords)

Knowledge is something someone knows, such as a password. Static passwords can be a password, a PIN, a passphrase, a graphic, etc. Regardless of length and character construction, static passwords that are not frequently changed are inherently insecure.

Secure storage is a necessity as legacy encryption of passwords in storage is typically easy to crack and makes unauthorized use of accounts a trivial matter for a determined malicious hacker. Tools such as Cain & Able along with Rainbow Tables can defeat the most commonly used password encryption methodologies in seconds. There are also Linux distributions such as KALI Linux that have a much broader toolset and function than just password cracking and are specifically engineered to provide an arsenal of tools to the security professional, password crackers among them, for detailed penetration testing. (Find it here: `https://www.kali.org/`.)[12] Password resets when the user forgets

their password consume a large volume of time in most IT support departments and also provide an effective entry vector for social engineering attacks. All too often password lockout mechanisms are disabled to reduce the number of required password resets, further increasing the risk of potential compromise. Automated password reset mechanisms range from the user being required to answer a series of personal questions that they previously provided responses for to newer technology-based reset mechanisms that use voice recognition to further automate the process.

Mass lockouts of user accounts are an effective denial-of-service attack. If a malicious hacker learns that you are using a standard "not unique" user name format, making the user names for authentication easy to guess, and that your access control system will lock out a user account after a given number of failed login attempts, it is a simple matter to quickly script an attack that walks through a failed login attempt, creating a locked-out account for every user. An example of this behavior can be found in the eBay Account Lockout Attack. At one time, eBay displayed the user ID of the highest bidder for a given auction. In the final minutes of the auction, an attacker who wanted to outbid the current highest bidder could attempt to authenticate three times using the targeted account. After three deliberately incorrect authentication attempts, eBay password throttling would lock out the highest bidder's account for a certain amount of time. An attacker could then make their own bid and the legitimate user would not have a chance to place a counter-bid because they would be locked out of their account.

Ownership

Ownership is something the user has in his possession such as a smart card or a token.

Smart Cards

Typically, smart cards are credit card size, contain a tamper-resistant security system, are managed by a central administration system, and require a card reader device, such as the typical card reader on an ATM or fuel pump at a gasoline station. There are contact and contactless smart cards and readers.

A contact card reader requires physical contact with the card reader. There are two primary methodologies for contact card readers. A landing contact requires physical contact with the contacts (landing zone) on the card when it is placed within the reader. Typical standards for landing contact readers include ISO 7816.[13] Landing contact readers are popular in physical access applications. A friction contact requires that the card landing contacts are wiped against the contact reader. Typical friction card readers are those used in credit card transactions at merchants.

Contactless card readers are quickly gaining in popularity and typically rely on radio-frequency identification (RFID) technology to facilitate reading. The additional security

mechanisms found in contactless card applications can include challenge/response-based encryption safeguards to reduce the risk of *card skimming*, whereby the account information is stolen in an otherwise legitimate transaction. Smart cards are discussed in more depth later.

Dynamic Passwords

A dynamic password methodology, also known as a *one-time password*, is typically implemented by utilizing hardware or software token technology. The password is changed after each authentication session. This effectively mitigates the risk of shoulder surfing or password sniffing, as the password is valid for only one session and cannot be reused.

Tokens

While tokens are available in many different form factors, there are two basic types of tokens in use today: *synchronous* and *asynchronous*.

With a synchronous token, time is synchronized between the token device and the authentication server. The current time value is enciphered along with a secret key on the token device and is presented to the access control subject for authentication. A popular synchronous token from RSA called "SecureID" provides for a new six- to eight-digit code every 60 seconds; it can operate for up to 4 years and can be programmed to cease operation on a predetermined date. The synchronous token requires fewer steps by the access control subject to successfully authenticate:

- The access control subject reads the value from his or her token device.
- The value from the token device is entered into the login window along with the access control subject's PIN.
- The authentication server calculates its own comparative value based on the synchronized time value and the respective access control subject's PIN. If the compared values match, access is granted.

An asynchronous token, such as the event-driven asynchronous token from Secure Computing called the SafeWord eToken PASS, provides a new one-time password with each use of the token. While it can be configured to expire on a specific date, its lifetime depends on its frequency of use. The token can last from 5 to 10 years and effectively extends the time period typically used in calculating the total cost of ownership in a multifactor authentication deployment. In the use of an asynchronous one-time password token, the access control subject typically executes a five-step process to authenticate identity and have access granted:

1. The authentication server presents a challenge request to the access control subject.

2. The access control subject enters the challenge into his/her token device.
3. The token device mathematically calculates a correct response to the authentication server challenge.
4. The access control subject enters the response to the challenge along with a password or PIN.
5. The response and password or PIN is verified by the authentication server and, if correct, access is granted.

The use of a PIN together with the value provided from the token helps to mitigate the risk of a stolen or lost token being used by an unauthorized person to gain access through the access control system. Tokens are discussed in more depth later in the "Tokens" section.

Radio Frequency Identification (RFID)

RFID is the wireless non-contact use of radio-frequency electromagnetic fields to transfer data, for the purposes of automatically identifying and tracking tags attached to objects. The tags contain electronically stored information. Some tags are powered and read at short ranges, typically a few meters, via magnetic fields. Others use a local power source such as a battery, or else have no battery but collect energy from the interrogating EM field, and then act as a passive transponder to emit microwaves or UHF radio waves. Battery-powered tags may operate at hundreds of meters. Unlike a bar code, the tag does not necessarily need to be within line of sight of the reader and may be embedded in the tracked object.

According to Technovelgy.com, some common problems with RFID are reader collision and tag collision:

> "Reader collision occurs when the signals from two or more readers overlap. The tag is unable to respond to simultaneous queries. Systems must be carefully set up to avoid this problem; many systems use an *anti-collision protocol* (also called a *singulation protocol*). Anti-collision protocols enable the tags to take turns in transmitting to a reader. Tag collision occurs when many tags are present in a small area; but since the read time is very fast, it is easier for vendors to develop systems that ensure that tags respond one at a time."[14]

Characteristics

A characteristic is defined as a physical trait of the user, also referred to as "what a person does" or "what a person is," that allows for the confirmation of an individual's identity based on either a physiological condition such as a fingerprint or retina scan or a behavioral characteristic such as keystrokes, speech recognition, or signature dynamics.

Characteristics are generally identified by using biometrics. Biometrics is discussed at length in the "Biometrics" section.

Biometrics

Biometrics is the science and technology of measuring and analyzing biological data. In information technology, biometrics refers to technologies that measure and analyze human body characteristics, such as DNA, fingerprints, voice patterns, facial patterns, and hand measurements, for authentication purposes. Biometric data cannot be considered to be secret in the way that private keys or passwords can. In contrast with private keys, biometric data is given to possibly hostile hosts to which a user wishes to authenticate. As opposed to passwords, biometric data cannot be changed, and a user cannot conveniently choose different biometric data to present to different hosts in the way that one might use a different password for a webmail account or a bank account. Moreover, in contrast with keys and passwords, biometric data such as user's facial characteristics and fingerprints are in the public domain and can be captured without the user's consent or knowledge. For this reason, protocols for biometric authentication should rely on proof of freshness of biometric data and cannot rely on its secrecy.

The processes involved within a biometric authentication solution could be classified as two steps: enrollment and verification. During the enrollment process, the user's registered biometric code is stored either in a system or on a smart card that is kept by the user. During the verification process, the user presents their biometric data to the system so that the biometric data can be compared with the stored biometric code. User verification can be carried out either within the smart card, a process called on-card matching, or in the system outside the card, known as off-card matching. The on-card matching algorithm protects the user's stored biometric code. The biometric code is not necessarily transferred to the outside environment if using this type of matching. Even though the biometric data is not considered to be secret, the protocol should not reveal it without the user's agreement. When the biometric data is used for biometric authentication, it should not only be protected from disclosure to an attacker, but also its origin should be guaranteed; this prevents an attacker from presenting the previously captured biometric data to the system in order to authenticate himself as the authorized user.

Biometrics can be broken down into two main classifications: behavioral and physiological.

Behavioral Biometrics

Behavioral biometrics includes signature analysis, voice pattern recognition, and keystroke dynamics.

Signature Analysis

The handwritten signature is unique to each individual. Most access control signature analysis access devices use a 3D analysis of the signature, which includes both the pressure and form of the signature. Signature analysis dynamically measures the series of movements, which contain biometric characteristics such as acceleration, rhythm, pressure, and flow. Signature analysis access control devices have become popular with credit card merchants for authorization of credit card transactions (see Figure 1.12).

FIGURE 1.12 Electronic signature pad. PHOTO OF ELECTRONIC SIGNATURE PAD (MODEL TM-LBK755) COURTESY OF TOPAZ SYSTEMS

> ✔ **Experience It!**
>
> To discover more about signature analysis, and the systems used to implement a solution, here is a list of vendors that provide systems:
>
> - `http://www.topazsystems.com/Software/index.htm`
> - `http://www.kofax.com/products/kofax-signature-solutions`

Voice Pattern Recognition

Voice pattern recognition works by creating a database of unique characteristics of the access control subject's voice. The access control subject then simply speaks at or near a microphone, and the access control device compares the current voice pattern characteristics to the stored characteristics to determine if access is to be granted. Biology, not

technology, is the issue with voice recognition. As the subject ages, the characteristics of the voice naturally change. Voice characteristics can change under stress, and during an emergency situation the access control subject could be denied access simply because of the stress he/she was under at that moment. Further, it is possible to create an error through the altering of the inflection of a given phrase. Voice recognition is an inexpensive methodology to implement, but because of the high probability of error it is best used to complement another more accurate technology, such as iris scanning, and not to be relied on as a primary access control device.

✔ Experience It!

To discover more about voice pattern recognition, and the systems used to implement a solution, here is a list of vendors that provide systems:

- `http://www.authentify.com/`
- `http://www.biovalidation.com/index.aspx`

Keystroke Dynamics

Keystroke dynamics rely on characteristics that are unique to an individual. Specifically, these are the characteristics of the access control subject's keystrokes as the user name and password are typed on the keyboard. The normal characteristics of the individual are learned over time and typically can be enrolled with six or eight samples. The individual characteristics used by the typical keystroke analysis device include but are not limited to

- The length of time each key is held down
- The length of time between keystrokes
- The typing speed
- The tendencies to switch between a numeric keypad and keyboard numbers
- The keystroke tendencies involved in capitalization

Figure 1.13 shows some standard aspects of keystroke dynamics that are measured.

The accuracy of keystroke dynamics can be affected by hand injuries, fatigue, arthritis, and perhaps temperature. In addition, the security of the keystrokes committed by the subject is open to compromise.[15] Hence, while keystroke dynamics is regarded as the lowest-cost authentication mechanism, it cannot yet be used reliably in a single-factor or perhaps two-factor (using passphrase) authentication methodology and is better suited to complement another technology such as iris scanning in a two-factor authentication scheme. It is important to note, however, that it does provide continuous authentication, if that is desirable.

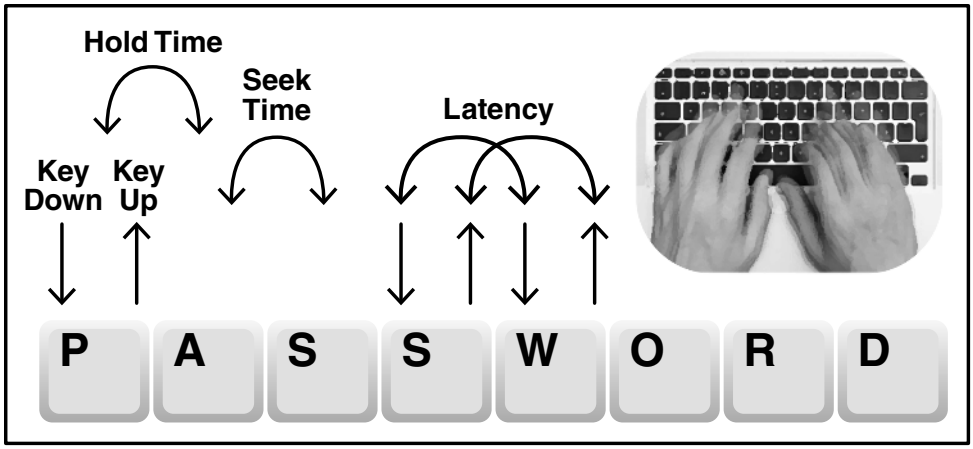

FIGURE 1.13 **Sample keystroke dynamics measurements**

✔ Experience It!

To discover more about keystroke dynamics, and the systems used to implement a solution, here is a list of vendors that provide systems:

- `http://www.biovalidation.com/index.aspx`
- `http://www.deepnetsecurity.com/authenticators/biometrics/typesense/`

Physiological Biometrics

There are several biometric devices that make use of the user's personal physiological data in access control applications. These apply fingerprint, hand, vascular, eye, or facial recognition technology.

Fingerprint Verification Technology

Fingerprint verification typically requires seven characteristics or matching points either to enroll a new access control subject or to verify an existing access control subject. The task is not as difficult as it may seem as the human finger contains 30–40 characteristics or matching points. The fingerprint reader does not store an image of the fingerprint. Rather, it creates a geometric relationship between the characteristics or matching points and stores and then compares that information. See Figure 1.14.

FIGURE 1.14 A fingerprint reader scans the loops, whorls, and other characteristics of a fingerprint and compares it with stored templates. When a match is found, access is granted.
PHOTO OF BIOMETRIC ID SIGNATURE PAD WITH OPTICAL FINGERPRINT SENSOR (MODEL TF-LBK464) COURTESY OF TOPAZ SYSTEMS.

One of the biggest challenges facing biometric technology in general, and fingerprint verification in particular today, is the ability to carry out performance evaluations unambiguously and reliably. One way this challenge is being addressed is through an innovative program called FVC-onGoing. FVC-onGoing is a web-based automated evaluation system for fingerprint recognition algorithms. Tests are carried out on a set of sequestered datasets, and results are reported online by using well-known performance indicators and metrics. While previous FVC initiatives were organized as "competitions," with specific calls and fixed time frames, FVC-onGoing is

- An "ongoing competition" always open to new participants
- An evolving online repository of evaluation metrics and results

Furthermore, FVC-onGoing performance evaluation is not only limited to fingerprint verification algorithms: ad hoc metrics and datasets for testing specific modules of fingerprint verification systems are available. This allows to better understand the limits

and the challenges not only of the whole recognition problem but also of its modules (e.g., feature extractor, matcher), with clear benefits for researchers and algorithms' developers. The aim is to track the advances in fingerprint recognition technologies, through continuously updated independent testing and reporting of performances on given benchmarks. The algorithms are evaluated using strongly supervised approaches to maximize trustworthiness.

FVC-onGoing is the evolution of FVC: the international Fingerprint Verification Competitions organized in 2000, 2002, 2004, and 2006. Find out more about FVC-onGoing at `https://biolab.csr.unibo.it/FVCOnGoing/UI/Form/Home.aspx`.

✔ Experience It!

To discover more about fingerprint verification, and the systems used to implement a solution, here is a list of vendors that provide systems:

- `www.supremainc.com`
- `http://usa.morpho.com`
- `http://www.zvetcobiometrics.com/`

Hand Geometry Technology

Hand geometry and geometry recognition technology is in broad use for access control as well as time and attendance applications (see Figure 1.15). An individual places their hand on a reader, and their identity is verified based upon the location of a number of key points on their hand (e.g., length of fingers, position of knuckles, etc.). Hand geometry technology measures the dimensions of hands and fingers, being mostly used in physical security applications. Applications include frequent traveler verification, identification of season pass holders for Walt Disney, and building security for hospitals. The advantage of hand geometry is that it provides a proven reliable verification even within difficult environments while being simple to operate. However, compared to other identification and verification methods, the method is less accurate and requires large and expensive equipment. Hand geometry verification is typically accomplished by building a five-element array of finger lengths determined from scanned matching points at the base and end of each finger. The stored five-element array is compared to a new hand scan, and a mathematical calculation is performed to determine the geometric distance between the respective arrays.

FIGURE 1.15 **Hand geometry reader**

✔ Experience It!

Discover more about hand geometry, and the systems used to implement a solution; here is a list of vendors that provide systems:

- `http://us.allegion.com/Products/biometrics/handkey2/Pages/default.aspx`
- `http://www.morphotrust.com/`

Vascular Patterns

This is the ultimate palm reader (see Figure 1.16). Vascular patterns are best described as a picture of the veins in a person's hand or finger. The thickness and location of these veins are believed to be unique enough to an individual to be used to verify a person's identity. The NTSC Subcommittee on Biometrics reports that researchers have determined that the vascular pattern of the human body is unique to the specific individual and does not change as people age. Claims for the technology include

- **Difficult to forge**—Vascular patterns are difficult to re-create because they are inside the hand, and for some approaches, blood needs to flow to register an image.
- **Contactless**—Users do not touch the sensing surface, which addresses hygiene concerns and improves user acceptance.

- **Many and varied uses**—It is deployed at ATMs, hospitals, and universities in Japan. Applications include ID verification, high security physical access control, high security network data access, and POS access control.
- **Capable of 1:1 and 1:many matches**—Users' vascular patterns are matched against personalized ID cards/smart cards or against a database of many scanned vascular patterns.

FIGURE 1.16 Vascular pattern reader. PHOTO OF PALMSECURE® ID MATCH COURTESY OF FUJITSU FRONTECH NORTH AMERICA

Eye Features/Retina Scan

The retina scan is one of the oldest and most accurate biometric authentication methodologies. Traditionally, the retina scan has been reserved only for the most secure application of physical access systems control. The retina scan simply maps the blood vessels in the back of the eye and only requires 10 or so seconds to complete a scan. There is no known technology that can forge a retina scan signature, and as the blood vessels quickly decay upon death, a retina scan on a dead individual will not create the same signature as that of the live individual.

How it works:

- The eye is read by a small green infrared light.
- Low-intensity infrared light is used because blood vessels on the retina absorb the infrared light faster than the surrounding eye tissues and the light is reflected back to a video camera.
- Initial scanning takes 10–15 seconds total, but verification scanning takes 2 seconds.
- Patterns of blood vessels are converted into mathematical patterns.

See Figure 1.17 for an overview of how retinal scanning works.

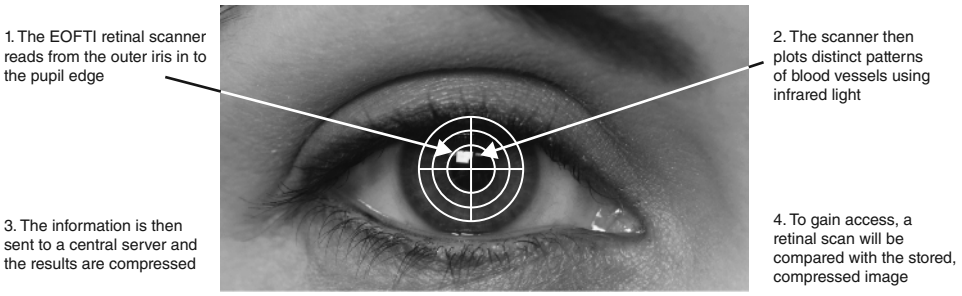

FIGURE 1.17 **How retinal scanners record identity source**

Eye Features/Iris Scan

Iris scanning is based on scanning the granularity of the richly detailed color bands around the pupil. The color bands are well defined at birth and change little over the subject's lifetime. The typical iris scanner maps nearly 247 variables in the iris and can do so at a distance of 19–20 inches. This makes the iris scanner potentially more accurate than a fingerprint, with only 40–80 characteristics, and is less obtrusive then a retina scanner as it does not require the same close proximity to the reading device or a light shining into the eye.

How it works:

- A person stands 1–3 feet away, and a wide-angle camera calculates the position of the eye.
- A second camera zooms in on the eye and takes a black-and-white image.
- The camera lays a circular grid on the image of the iris so the iris system can recognize patterns within the iris to generate points.
- The captured image or "eyeprint" is checked against previously stored reference template in the database.
- Software localizes the inner/outer boundaries of the iris and eyelid contours.
- Demodulation, or mathematical software, encodes the iris pattern.
- Then it captures the unique features of the iris, like a template (the IrisCode).

The template is immediately encrypted to eliminate the possibility of identity theft and maximize security.

Figure 1.18 shows a simplified overview of how iris scanning works.

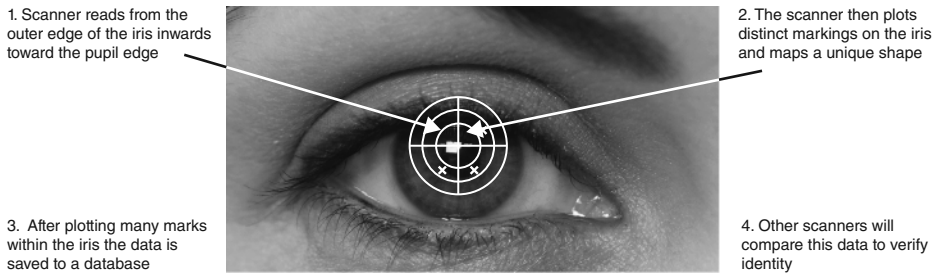

FIGURE 1.18 How iris scanners record identity

Here are some interesting facts about iris scan technology:

- No two irises are alike, not even with identical twins.
- The left eye and right eye are not the same on one person.
- The iris has six times more distinct identifiable features than fingerprints.
- The probability of having two irises that are alike is one in 10 to the 78th power (the population of the earth is approximately 10 to the 10th power).
- There is no known way to copy a retina, unlike an iris.
- A retina from a dead person would deteriorate too fast to be useful, so no extra precautions have been taken with retinal scanning to make sure the person is alive.

Security concerns have recently come up with regards to iris scans, which are considered to be one of the most secure biometric solutions currently in use. Announced during the annual Black Hat security conference in 2012, a team at the Universidad Autonoma de Madrid was able to re-create the image of an iris from digital codes of real irises stored in security databases. The researchers were able to print out synthetic images of irises, and as they tested their fake irises against one of the leading commercial recognition systems, they achieved an 80% false accept rate.[16] Another problem that has recently come to light is the effect that alcohol can have on iris scans. Research has shown that alcohol consumption causes recognition degradation as the pupil dilates/constricts, which causes deformation in the iris pattern. Experiments performed show that in matching pre- and post-alcohol consumption images, the overlap between genuine and impostor match scores increased by approximately 20%. This means that one in five subjects under alcohol influence potentially could be able to evade identification by iris recognition systems.[17]

The ISO/IEC standard 19794-6:2011, "Information technology—Biometric data interchange formats—Part 6: Iris image data," specifies two alternative image interchange formats for biometric authentication systems that utilize iris recognition. The first is based

on a rectilinear image storage format that may be a raw, uncompressed array of intensity values or a compressed format such as that specified by ISO/IEC 15444-1:2004/Amendments 1–5. The second format is based on a polar image specification that requires certain pre-processing and image segmentation steps but produces a much more compact data structure that contains only iris information. Data that complies with either one of the iris image formats specified in ISO/IEC 19794-6:2011 are intended to be embedded in a CBEFF-compliant structure in the CBEFF Biometric Data Block (BDB) as specified in ISO/IEC 19785-1:2006 and ISO/IEC 19785-1:2006 / AMD1:2010.

✔ Try It for Yourself— Open Source Iris Scan Project in a Box

Now you will work directly with an iris scan project.

What's Needed?

A Linux-based computer and a user account with administrative (root) privileges.

How to Do It

Use the following step-by-step guidance:

1. Download the source tarball from here (tarball name is **iris-0.1.tar.gz**):

   ```
   http://projectiris.co.uk/
   ```

2. To install for any Linux distribution:

   ```
   tar zxf iris-version.tar.gz
   cd iris-version/
   qmake
   make ./iris
   ```

✔ Experience It!

To discover more about iris scans, and the systems used to implement a solution, here is a list of vendors that provide systems:

- `http://www.sri.com/engage/products-solutions/iris-move-biometric-identification-systems`
- `http://www.crossmatch.com/seek-avenger/`
- `http://www.irisguard.com/index.php`

Facial Recognition

Like the fingerprint reader and hand geometry devices, facial recognition uses a mathematical geometric model of certain landmarks of the face such as the cheekbone, tip of the nose, and eye socket orientation, and measures the distance between them. There are approximately 80 separate measurable characteristics in the human face, but most facial recognition systems rely on only 14–22 characteristics to perform their recognition.

Figure 1.19 shows how geometric properties are used for facial recognition.

Here are some interesting facts about facial recognition technology:

- Google's Picasa digital image organizer has a built-in face recognition system starting from version 3.5 onwards. It can associate faces with people so that queries can be run on pictures to return all pictures with a specific group of people together.
- Windows Live Photo Gallery includes face recognition.
- Sony's Picture Motion Browser (PMB) analyzes photos, associates photos with identical faces so that they can be tagged accordingly, and differentiates between photos with one person, many people, and nobody.
- OpenBR is an open source face recognition system and research platform for biometric algorithm development.

✔ Experience It!

To discover more about facial recognition, and the systems used to implement a solution, here is a list of vendors that provide systems:

- `http://www.morphotrust.com/`
- `http://www.luxand.com/index.php`
- `http://www.zoomyimages.com/`

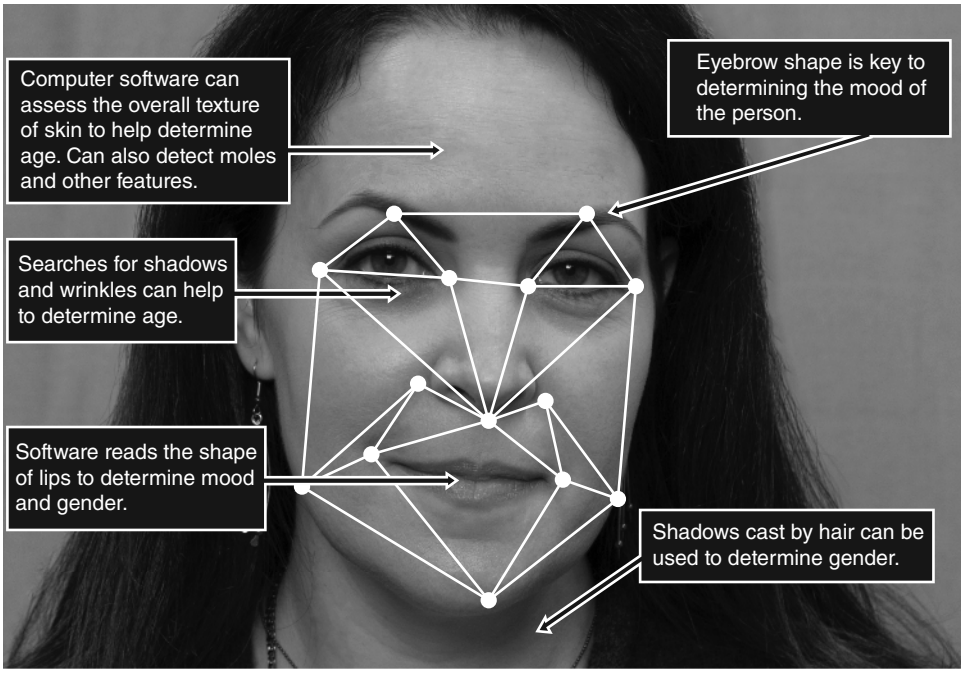

FIGURE 1.19 **Geometric properties of a subject's face used in facial imaging**

Biometric Implementation Issues

User acceptance is one of the most critical factors in the success of any biometric-based implementation. To minimize the risk of improper use, which can cause failed access, the device should not cause discomfort or concern and must be easy to use.

Biometric accuracy is measured by two distinct rates: the False Rejection Rate (FRR), referred to as a type 1 error, and the False Acceptance Rate (FAR), referred to as a type 2 error.

- **False Rejection**—Failure to recognize a legitimate user. While it could be argued that this has the effect of keeping the protected area extra secure, it is an intolerable frustration to legitimate users who are refused access because the scanner does not recognize them.
- **False Acceptance**—Erroneous recognition, either by confusing one user with another or by accepting an imposter as a legitimate user.

Failure rates can be adjusted by changing the threshold ("how close is close enough") for declaring a match, but decreasing one failure rate will increase the other.

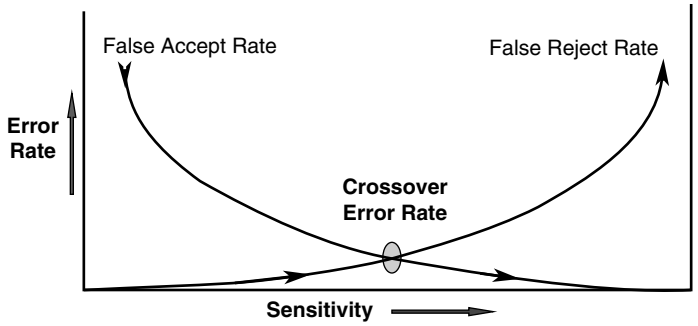

FIGURE 1.20 **Crossover error rate is one of three categories of biometric accuracy measurements.**

The actual methodologies of the measurement of accuracy may differ in each type of biometric device, but simply put, you can obtain a good comparative accuracy factor by looking at the intersection point at which the type 1 error rate equals the type 2 error rate as shown in Figure 1.20. This value is commonly referred to as the Crossover Error Rate, (CER). The biometric device accuracy increases as the crossover value becomes smaller, as shown in Table 1.4.

TABLE 1.4 **Biometric Crossover Accuracy**

BIOMETRIC CROSSOVER ACCURACY	
Retinal scan	1:100,00,000
Iris scan	1:131,000
Fingerprint	1:500
Hand geometry	1:500
Signature dynamics	1:50
Voice dynamics	1:50

A further comparison of biometric technologies is provided in Table 1.5.

In reusable password authentication, the access control subject had to remember a perhaps difficult password. In token-based authentication, the access control subject had to retain possession of the token device. In biometric, characteristic-based authentication, the actual access control subject *is* the authentication device.

TABLE 1.5 Comparison of Biometric Technologies

CHARACTERISTIC	FINGERPRINTS	HAND GEOMETRY	RETINA	IRIS	FACE	SIGNATURE	VOICE
Ease of Use	High	High	Low	Medium	Medium	High	High
Error Incidence	Dryness, dirt, age	Hand injury, age	Glasses	Poor lighting	Lighting, age, glasses, hair	Changing signatures	Noise, colds, weather
Accuracy	High	High	Very high	Very high	High	High	High
User Acceptance	Medium	Medium	Medium	Medium	Medium	Medium	High
Required Security Level	High	Medium	High	Very high	Medium	Medium	Medium
Long-Term Stability	High	Medium	High	High	Medium	Medium	Medium

Source: Liu, S., and Silverman, M., "A practical guide to biometric security technology," *IT Professional,* 3, 27–32, 2005. With permission.

Physical Use as Identification

Biometrics takes advantage of the unique physical traits of each user and arguably is the most effective methodology of identifying a user. It is important to note that in physical security, biometrics is often used as an identification mechanism, while in logical security biometrics is often used as an authentication mechanism. As biometric technologies evolve, accuracy rates are increasing, error rates are declining, and improved ease-of-use is increasing user acceptance.

Biometric Standards Development

Numerous activities regarding the interoperability of biometrics are ongoing at both the national and international levels. On the national level, ANSI INCITS 395-2005 specifies a data interchange format for representation of digitized sign or signature data, for the purposes of biometric enrollment, verification, or identification through the use of Raw Signature/Sign Sample Data or Common Feature Data. The data interchange format is generic, in that it may be applied and used in a wide range of application areas where electronic signs or signatures are involved. No application-specific requirements or features are addressed in this standard. At the international level, there are two corresponding

documents currently published: The first is ISO/IEC 19794-7:2014: Information technology—Biometric data interchange formats—Part 7: Signature/sign time series data, and the second is ISO/IEC 19794-11:2013/Amd.1:2014: Information technology—Biometric data interchange formats—Part 11: Signature/Sign Processed Dynamic Data.

The ISO JTC 1/SC 37 Biometrics working group homepage can be found here:

```
http://www.iso.org/iso/home/standards_development/list_of_iso_technical_
committees/jtc1_home/jtc1_sc37_home.htm
```

Tokens

Security tokens are used to prove one's identity electronically. There are four different ways in which this information can be used, according to Wikipedia:[18]

Static Password Token—The device contains a password that is physically hidden (not visible to the possessor) but that is transmitted for each authentication. This type is vulnerable to replay attacks.

Synchronous Dynamic Password Token—A timer is used to rotate through various combinations produced by a cryptographic algorithm. The token and the authentication server must have synchronized clocks.

Asynchronous Password Token—A one-time password is generated without the use of a clock, from either a one-time pad or a cryptographic algorithm.

Challenge Response Token—Using public key cryptography, it is possible to prove possession of a private key without revealing that key. The authentication server encrypts a challenge (typically a random number, or at least data with some random parts) with a public key; the device proves it possesses a copy of the matching private key by providing the decrypted challenge.

Smart Cards

A smart card, typically a type of chip card, is a plastic card that contains an embedded computer chip—either a memory or microprocessor type—that stores and transacts data. This data is usually associated with either value, information, or both and is stored and processed within the card's chip. The card connects to a reader with direct physical contact or with a remote contactless radio frequency interface. With an embedded microcontroller, smart cards have the unique ability to store large amounts of data, carry out their own on-card functions (e.g., encryption and mutual authentication), and interact intelligently with a smart card reader. Smart card technology conforms to international standards (ISO/IEC 7816 and ISO/IEC 14443) and is available in a variety of form factors, including plastic cards, fobs, subscriber identity modules (SIMs) used in GSM mobile phones, and USB-based tokens.

There are two general categories of smart cards: contact and contactless. According to the Smart Card Alliance,[19]

> A contact smart card must be inserted into a smart card reader with a direct connection to a conductive contact plate on the surface of the card (typically gold plated). Transmission of commands, data, and card status takes place over these physical contact points.
>
> A contactless card requires only close proximity to a reader. Both the reader and the card have antennae, and the two communicate using radio frequencies (RF) over this contactless link. Most contactless cards also derive power for the internal chip from this electromagnetic signal. The range is typically one-half to three inches for non-battery-powered cards, ideal for applications such as building entry and payment that require a very fast card interface.
>
> Two additional categories of cards are dual-interface cards and hybrid cards. A hybrid card has two chips, one with a contact interface and one with a contactless interface. The two chips are not interconnected. A dual-interface card has a single chip with both contact and contactless interfaces. With dual-interface cards, it is possible to access the same chip using either a contact or contactless interface with a very high level of security.

Improvements to Authentication Strategies

For many years knowledge-based authentication in terms of passwords was the most common methodology in use in access control systems. Weaknesses in the implementation of encryption (hashing) for passwords has effectively rendered these knowledge-based methodologies obsolete.

Multifactor Authentication

In October 2005, the Federal Financial Institutions Examination Council provided a recommendation to U.S. banks that included, in part, a requirement to replace passwords and single-factor authentication with multifactor authentication.[20] The recommendation clearly pointed out that passwords alone were simply no longer a secure methodology for authenticating users in the current Internet environment.

The best practice in access control is to implement at least two of the three common techniques for authentication in your access control system:

- Knowledge-based
- Token-based
- Characteristic-based

Two-Factor vs. Three-Factor Authentication

In two-factor authentication, typically the mechanism used provides for something the user has in the form of a physical token that generates a one-time password and something the user knows in the form of a PIN that is appended to the one-time password that is generated by the token. This methodology is regarded as more secure than historical single-factor methodologies such as traditional passwords; however, it does little to definitively identify the user. This can be significantly improved upon by incorporating a third factor in the form of a biometric that in fact identifies the user. An example of a three-factor authentication solution is the RSA AuthenTec Fingerprint device from Privaris. It incorporates a fingerprint reader to identify the user, as well as being something the user "has," and also incorporates the traditional one-time password and PIN combination found in common two-factor authentication tokens.

Dual Control

Dual control, also referred to as "split-knowledge," is built on the principle that no one person should have access to information that would allow the person to determine the encryption key used to encrypt protected information more quickly than a brute-force attack of the entire key-space. Effectively, the determination of any part of the encryption key would require collusion between at least two different trusted individuals. Encryption—splitkeys—is just one example of dual control. It has been said that because of its inherent complexity, dual control is not difficult to accomplish but is easy to get wrong.

Continuous Authentication

While traditional one-time authentication, otherwise known as transactional authentication, takes place only once before granting access, continuous authentication takes place both before granting access and then continuously through the entire duration of the user's connection to maintain the granted access.

Periodic Authentication

The most common use of periodic authentication first provides for traditional challenge/response authentication requiring user interaction and then begins periodically to issue challenge/response authentication queries with the user's token to determine if the user has physically left the area where he had authenticated. This methodology aids in reducing the risk that a user would walk away from a device or system he has authenticated access to before properly logging out.

Time Outs

If the user leaves the proximity of the device authenticated after a specific time period, the user is automatically logged off and the authentication process would start over, requiring user intervention to accomplish initial authentication before continuous authentication could again resume. Naturally, the shorter the timeout period, the higher the security that can be provided; however, as always, it comes at the cost of being intrusive to the user.

Reverse Authentication

With the advent of phishing it is no longer enough to simply authenticate the user in web-based transactions. Today, it is necessary to also authenticate the website/page to the user as part of the authentication process. Bank of America was a pioneer in reverse authentication with its roll-out of PassMark, a reverse authentication system that relies on a series of pictures that the user could identify and use to accomplish the authentication of the Bank of America website. Some had believed that the picture approach of PassMark was too simplistic and raised doubts about the technology. However, PassMark quickly grew in acceptance and was adopted by more than 50% of the online banking market.

Certificate-Based Authentication

Certificate-based authentication relies on the machine that the user authenticates from having a digital certificate installed that is used in part along with the encrypted user's password to authenticate both the user and the device the user is authenticating from. Effectively, the use of a certificate in the authentication process adds an additional element in security by validating that the user is authorized to authenticate from the device they are using because of the presence of the digital certification within the device. Great care must be taken in the management of the digital certificates by the certificate authority to ensure that the use of certificates is properly controlled and certificate renewal and revocations are accomplished in a timely and effective manner.

Authorization

What a user can do once authenticated is most often controlled by a reference monitor. A reference monitor is typically defined as the service or program where access control information is stored and where access control decisions are made. A reference monitor will typically decide if access is to be granted based on an ACL within the reference monitor. Once access is granted, what the subject can then do is controlled by the authorization matrix or table (see Table 1.6).

TABLE 1.6 Authorization Table—Matrix of Access Control Objects, Access Control Subjects, and Their Respective Rights

ACCESS CONTROL SUBJECTS	ACCESS CONTROL OBJECTS					
	Procedure A	Procedure B	File A	File B	File C	File D
Bob	Execute		Read	Read/Write		Read
Tom		Execute			Read	
Mary		Execute			Read	
Process A			Read/Write			Write
Process B			Write			Read/Write

Access to Systems vs. Data, Networks

Defining ACLs that only address access to systems can facilitate unintended user access to data that perhaps the user should not have had access to. Including access controls to specific data within a given system increases overall security. Consideration must also be given to ensuring that users only have access to intended networks, systems, and data.

Access Control Lists/Matrix

An authorization table is a matrix of access control objects, access control subjects, and their respective rights, as shown in Table 1.6. The authorization table is used in some DAC systems to provide for a simple and intuitive user interface for the definition of access control rules. While an authorization table provides for an increase in ease of use, it does not solve the inherent issue of DAC in that the system is still relying on the access control object owner to properly define the access control rules. Further, the use of an authorization table does not decrease the instance of errors or violations that may occur when changes are made within the authorization table.

An access control matrix is used in a DAC system to provide for a simple user interface to implement an ACL. The access control matrix determines the access rights for access control objects to access control subjects, as shown in Table 1.7. Like the authorization table mentioned earlier, the access control matrix does not decrease the instance of errors or violations that may occur when changes are made within the access control matrix.

TABLE 1.7 Access Control Matrix Determines the Access Rights for Access Control Objects to Access Control Subjects

ACCESS CONTROL SUBJECTS	ACCESS CONTROL OBJECTS															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	X		X		X	X		X					X	X		
2	X		X	X						X	X					
3	X	X				X	X	X						X		X
4			X	X	X											
5		X			X		X	X	X					X	X	
6												X				
7						X	X	X	X	X	X					
8		X	X										X	X		

Directories

Lightweight Directory Access Protocol (LDAP) is an application protocol used for querying and modifying directory services over TCP/IP. An LDAP directory is a logically and hierarchically organized group of objects and their respective attributes using an LDAP directory tree. An LDAP directory tree typically starts with domain names at the top of the hierarchy followed by organizational boundaries, then groups followed by users and data, such as groups of documents.

X.500 relies also on the use of a single Directory Information Tree (DIT) with a hierarchical organization of entries that are distributed across one or more servers. Every directory entry has what is referred to as a Distinguished Name (DN), which is formed by combining its Relative Distinguished Name (RDN), one or more attributes of the entry itself, and the RDN of the superior entries reaching all the way up to the root of the DIT.

The Microsoft Active Directory Domain Services (ADDS, originally called NT Directory Services) stores data and information within a central database, is highly scalable, and provides a wide variety of other network services including LDAP-like directory services, authentication, and Domain Name System–based naming. While ADDS is primarily used for assignment of policies because of its many attributes, it is commonly used by separate services to facilitate software distribution within a network.

Think of directory structures like LDAP, X.500, and ADDS as telephone directories where all entries are based on an alphabetical order and have attached addresses and telephone numbers.

Single Sign-On

Single sign-on (SSO) can best be defined as an authentication mechanism that allows a single identity to be shared across multiple applications. Effectively, it allows the user to authenticate once and gain access to multiple resources.

The primary purpose of SSO is the convenience of the user. With that in perspective, SSO can also help in mitigating some of the inherent risks of access control subjects using a different password or authentication mechanism for each of the many systems they access in a large network. Simply put, the chances of a security breach naturally increase as the number of passwords and or authentication mechanisms increase. This must, of course, be balanced against the additional risk of using SSO in that once it's implemented, a malicious hacker now only has to obtain a single set of authentication credentials and then has access to all of the systems that the respective access control subject was permitted to access. The advantages as well as disadvantages of SSO must also be considered (Table 1.8).

TABLE 1.8 Advantages and Disadvantages of SSO

ADVANTAGES OF SSO	DISADVANTAGES OF SSO
More efficient log-on process.	Difficult to implement across the enterprise.
Easier administration.	Many systems use proprietary authentication systems that will not work well with standard SSO systems.
When a new employee is hired, all of the accounts on all of the systems the new employee needs to access can be quickly added from a single administration point.	Time-consuming to implement properly.
When an existing employee is terminated, all access can be quickly and simultaneously restricted at a single administration point.	Many underestimate the amount of time necessary to properly implement SSO across all systems in the enterprise.
If an existing user loses their token or forgets their password, the administrator can quickly update the user's authentication credentials from a single administration point.	Expensive to implement.
Can mitigate some security risks.	Because of the difficulty and time involved to properly implement SSO, it is expensive. A redundant authentication server is required to avoid a single point of failure.
Reduces the inherent risk of a user having to remember passwords for multiple systems, within the enterprise.	
Because only a single password is used, the user is more apt to use a much stronger password.	

ADVANTAGES OF SSO	DISADVANTAGES OF SSO
Timeout and attempt thresholds are enforced consistently across the entire enterprise. SSO generally offers a good return on investment for the enterprise. The reduced administrative costs can often pay for the cost of implementing SSO in a short period of time. However, it should be noted that if scripting is used to facilitate the implementation of SSO, the typical reduced administration costs associated with SSO could in fact be negated because of the effort required to maintain numerous scripts.	Proprietary authentication systems may need expensive custom programming to be used in an SSO implementation, and more often than not this cost is not considered in the original estimates and results in SSO implementation cost overruns. In some cases the original authentication system for a difficult to implement system has to be weakened in an effort to get it to work reliably in an SSO system.

There are a couple of significant risks inherent with SSO:

- **Single point of failure**—With all of the users' credentials stored on a single authentication server, the failure of that server can prevent access for those users to all applications that it had provided authentication services for.
- **Single point of access**—Because SSO affords a single point of access, it is more prone to mass denial-of-service attacks whereby entire groups of users can be denied access to systems by attacking the single point of access.

Authentication Using Kerberos

Kerberos, described in RFC 1510, was originally developed by the Massachusetts Institute of Technology (MIT) and has become a popular network authentication protocol for indirect (third-party) authentication services.[21] It is designed to provide strong authentication using secret-key cryptography. It is an operational implementation of key distribution technology and affords a key distribution center, authentication service, and ticket granting service. Hosts, applications, and servers all have to be "Kerberized" to be able to communicate with the user and the ticket granting service.

Like the previously discussed indirect authentication technologies, Kerberos is based on a centralized architecture, thereby reducing administrative effort in managing all authentications from a single server. Furthermore, the use of Kerberos provides support for

1. **Authentication**—A user is who they claim to be.
2. **Authorization**—What can a user do once properly authenticated?
3. **Confidentiality**—Keep data secret.
4. **Integrity**—Data received is the same as the data that was sent.
5. **Nonrepudiation**—Determines exactly who sent or received a message.

The process in the use of Kerberos is substantially different from those indirect authentication technologies previously reviewed and is considerably more complex. The following is a simplified explanation of the Kerberos process that was adapted for use here from *Applied Cryptography: Protocols, Algorithms, and Source Code in C* by Bruce Schneier (New York, NY: Wiley, 1993).

1. Before an access control subject can request a service from an access control object, it must first obtain a ticket to the particular target object; hence, the access control subject first must request from the Kerberos Authentication Server (AS) a ticket to the Kerberos Ticket Granting Service (TGS). This request takes the form of a message containing the user's name and the name of the respective TGS.
2. The AS looks up the access control subject in its database and then generates a session key to be used between the access control subject and the TGS. Kerberos encrypts this session key using the access control subject's secret key. Then, it creates a Ticket Granting Ticket (TGT) for the access control subject to present to the TGS and encrypts the TGT using the TGS's secret key. The AS sends both of these encrypted messages back to the access control subject.
3. The access control subject decrypts the first message and recovers the session key. Next, the access control subject creates an authenticator consisting of the access control subject's name, address, and a time stamp, all encrypted with the session key that was generated by the AS.
4. The access control subject then sends a request to the TGS for a ticket to a particular target server. This request contains the name of the server, the TGT received from Kerberos (which is already encrypted with the TGS's secret key), and the encrypted authenticator.
5. The TGS decrypts the TGT with its secret key and then uses the session key included in the TGT to decrypt the authenticator. It compares the information in the authenticator with the information in the ticket, the access control subject's network address with the address the request was sent from, and the time stamp with the current time. If everything matches, it allows the request to proceed.
6. The TGS creates a new session key for the user and target server and incorporates this key into a valid ticket for the access control subject to present to the access control object server. This ticket also contains the access control subject's name, network address, a time stamp, and an expiration time for the ticket—all encrypted with the target server's secret key—and the name of the server. The TGS also encrypts the new access control subject target session key using the session key shared by the access control subject and the TGS. It sends both messages to the access control subject.

7. The access control subject decrypts the message and extracts the session key for use with the target access control object server. The access control subject is now ready to authenticate himself or herself to the access control object server. He or she creates a new authenticator encrypted with the access control subject target session key that the TGS generated. To request access to the target access control object server, the access control subject sends along the ticket received from Kerberos (which is already encrypted with the target access control object server's secret key) and the encrypted authenticator. Because this authenticator contains plaintext encrypted with the session key, it proves that the sender knows the key. Just as important, encrypting the time of day prevents an eavesdropper who records both the ticket and the authenticator from replaying them later.
8. The target access control object server decrypts and checks the ticket and the authenticator, also confirming the access control subject's address and the time stamp. If everything checks out, the access control object server now knows the access control subject is who he or she claims to be, and the two share an encryption key that they can use for secure communication. (Since only the access control subject and the access control object server share this key, they can assume that a recent message encrypted in that key originated with the other party.)
9. For those applications that require mutual authentication, the server sends the access control subject a message consisting of the time stamp plus 1, encrypted with the session key. This serves as proof to the user that the access control object server actually knew its secret key and was able to decrypt the ticket and the authenticator.

To provide for the successful implementation and operation of Kerberos, the following should be considered:

1. Overall security depends on a careful implementation.
2. Requires trusted and synchronized clocks across the enterprise network.
3. Enforcing limited lifetimes for authentication based on time stamps reduces the threat of a malicious hacker gaining unauthorized access using fraudulent credentials.
4. The key distribution server must be physically secured.
5. The key distribution server must be isolated on the network and should not participate in any non-Kerberos network activity.
6. The AS can be a critical single point of failure.

Kerberos is available in many commercial products, and a free implementation of Kerberos is available from MIT.[22] Table 1.9 shows the ports used by Kerberos during the authentication process.

TABLE 1.9 Network Ports Used During Kerberos Authentication

SERVICE NAME	UDP	TCP
DNS	53	53
Kerberos	88	88

User/Device Authentication Policies

Security and authentication policies are often unique to a given organization; effective security is never a one-size-fits-all proposition. A basic security policy—defining what information is sensitive, who can have access to this information and under what circumstances, and what to do in the event of a breach—is a must. Simple and obvious elements, like requiring PIN codes on mobile devices and regular password changes, are essential. Policy can go further to explain what a given user/device combination can do based on credentials and context. Only after policies are set and tested in an isolated or pilot setting should specific user/device authentication technologies be considered.

Authentication can take many forms. The security practitioner should familiarize themselves with the methods listed here:

Computer Recognition Software Using the computer as a second authentication factor is accomplished by installing a small authentication software plug-in that places a cryptographic device marker onto the user's computer, which can then be verified as a second factor during the authentication process. The authentication process would then include two factors: a password (something you know) and the device marker on the user's computer (something you have). Because the device marker is always present on the user's computer, the user only has to enter their username and password to log in.

Biometrics Using biometrics as a second factor is accomplished by verifying physical characteristics such as a fingerprint or eye using a dedicated hardware device. We discussed biometrics at length earlier in this chapter.

E-mail or SMS One-Time Password (OTP) Using e-mail or SMS OTP as a second factor is accomplished by sending a second one-time use password to a registered e-mail address or cell phone. The user must then input that second one-time password in

addition to their normal password to authenticate. This method is generally considered too cumbersome for everyday logins because there is a time lag before users get the OTP they need to login but is often used for the initial enrollment before providing another form of authentication.

One Time Password (OTP) Token Using an OTP token as a second factor is accomplished by providing users with a hardware device that generates a constantly changing second password that must be entered in addition to the normal password. OTP tokens require the user to carry the token with them to login.

Out of Band Using an out-of-band verification for authentication involves the target system calling a registered phone number and requesting that the user enter their password over the phone prior to allowing the user to log in. Similar to e-mail or SMS OTPs, this requirement introduces a time lag and requires that the user be at the location of the registered phone number during the login sequence.

Peripheral Device Recognition Using peripheral device recognition as a second factor is accomplished by placing a cryptographic device marker on a user's existing device such as a USB flash drive, an iPod, or smart phone memory card and then requiring that device to be plugged into the computer when the user logs in. This can be a good alternative to the OTP token because it provides a hardware-based second factor but does not require the user to carry an additional device. In addition, device markers from multiple systems can reside on a single hardware device.

COMPARING INTERNETWORK TRUST ARCHITECTURES

Computers are connected together using networks, and different types of networks provide different levels of trust. Primarily, there are four types of trust architectures: the Internet, an intranet, an extranet, and a demilitarized zone (DMZ, or perimeter network). The security practitioner is expected to understand all of them.

Internet

The Internet is a global system of interconnected computer networks that use the standard Internet protocol suite (TCP/IP) to link several billion devices worldwide. It is an international network of networks that consists of millions of private, public, academic, business, and government packet-switched networks, linked by a broad array of electronic, wireless, and optical networking technologies. The terms Internet and World

Wide Web are often used interchangeably in everyday speech; it is common to speak of "going on the Internet" when invoking a web browser to view web pages. However, the World Wide Web or the Web is just one of a very large number of services running on the Internet. The Web is a collection of interconnected documents (web pages) and other web resources, linked by hyperlinks and URLs. In addition to the Web, a multitude of other services are implemented over the Internet, including e-mail, file transfer, remote computer control, newsgroups, and online games. All of these services can be implemented on any intranet, accessible to network users.

Intranet

An intranet is a network based on TCP/IP protocols (an internet) belonging to an organization, usually a corporation, accessible only by the organization's members, employees, or others with authorization. Intranets utilize standard network hardware and software technologies like Ethernet, Wi-Fi, TCP/IP, web browsers, and web servers. An organization's intranet typically includes Internet access but is firewalled so that its computers cannot be reached directly from the outside.

Extranet

An extranet is a computer network that allows controlled access from the outside for specific business or educational purposes. Extranets are extensions to, or segments of, private intranet networks that have been built in many corporations for information sharing and ecommerce. In a business-to-business context, an extranet can be viewed as an extension of an organization's intranet that is extended to users outside the organization, usually partners, vendors, and suppliers, in isolation from all other Internet users. An extranet is similar to a DMZ in that it provides access to needed services for channel partners, without granting access to an organization's entire network.

Demilitarized Zone (DMZ)

A *DMZ* is a computer host or small network inserted as a "neutral zone" between a company's private network and the outside public network (see Figure 1.21). It prevents outside users from getting direct access to a server that has company data.

In addition to the four types of trust architectures, the security practitioner should be familiar with the trust types, which are discussed in the following section.

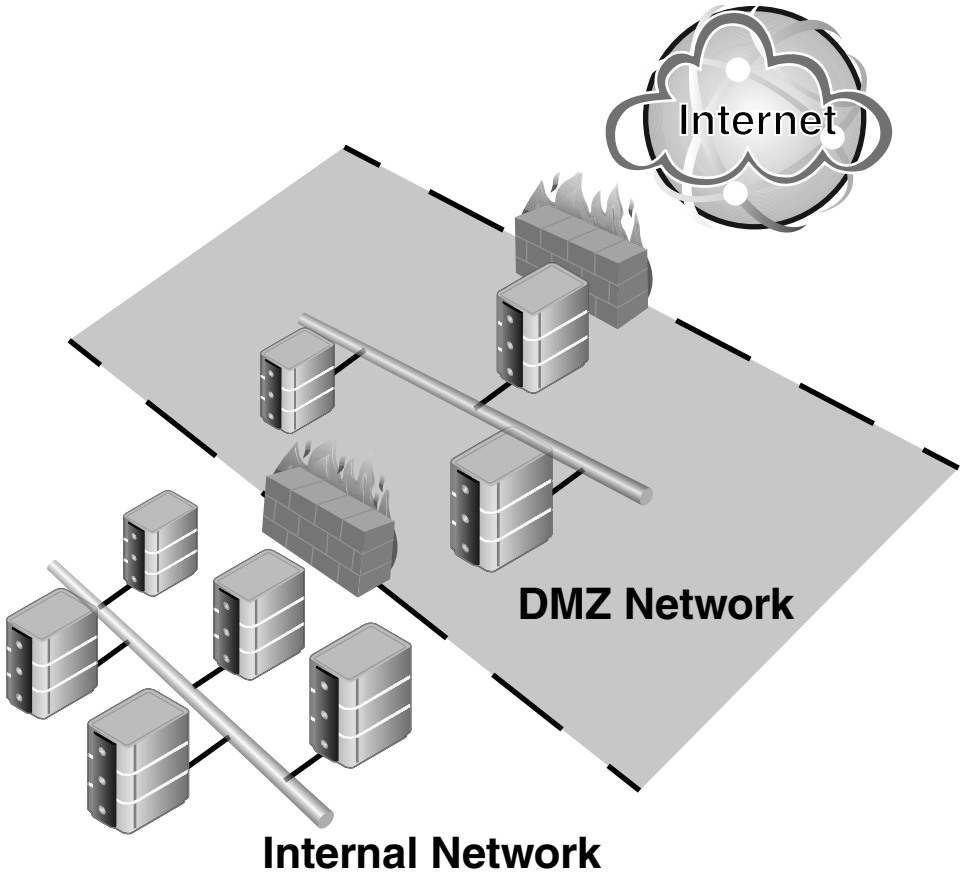

FIGURE 1.21 **A typical DMZ design**

TRUST DIRECTION

The trust type and its assigned direction affect the trust path that is used for authentication. A trust path is a series of trust relationships that authentication requests must follow between domains. Before a user can access a resource in another domain, the security system on domain controllers must determine whether the trusting domain (the domain that contains the resource that the user is trying to access) has a trust relationship with the trusted domain (the user's logon domain). To determine this, the security system computes the trust path between a domain controller in the trusting domain and a domain controller in the trusted domain.

One-Way Trust

A one-way trust is a unidirectional authentication path that is created between two domains. This means that in a one-way trust between Domain A and Domain B, users in Domain A can access resources in Domain B. However, users in Domain B cannot access resources in Domain A. Some one-way trusts can be either a non-transitive trust or a transitive trust, depending on the type of trust that is created.

Two-Way Trust

In a two-way trust, Domain A trusts Domain B, and Domain B trusts Domain A. This means that authentication requests can be passed between the two domains in both directions. Some two-way relationships can be either non-transitive or transitive, depending on the type of trust that is created.

Trust Transitivity

Transitivity determines whether a trust can be extended outside the two domains between which the trust was formed. You can use a transitive trust to extend trust relationships with other domains. You can use a non-transitive trust to deny trust relationships with other domains.

ADMINISTERING THE IDENTITY MANAGEMENT LIFECYCLE

There are five areas that make up the identity management lifecycle:

- Authorization
- Proofing
- Provisioning
- Maintenance
- Entitlement

Authorization

Authorization determines whether a user is permitted to access a particular resource. Authorization is performed by checking the resource access request against authorization policies that are stored in an Identity Access Management (IAM) policy store. Moreover, authorization could provide complex access controls based on data or information or policies

including user attributes, user roles/groups, actions taken, access channels, time, resources requested, external data, and business rules.

Proofing

According to Gartner's IT Glossary, identity-proofing services, which verify people's identities before the enterprise issues them accounts and credentials, are based on "life history" or transaction information aggregated from public and proprietary data sources. These services are also used as an additional interactive user authentication method, especially for risky transactions, such as accessing sensitive confidential information or transferring funds to external accounts. Identity-proofing services are typically used when accounts are provisioned over the Web or in a call center. However, they can also be used in face-to-face interactions.[23]

Provisioning

According to the Encyclopedia of Cryptography and Security, provisioning is the automation of all procedures and tools to manage the lifecycle of an identity: creation of the identifier for the identity, linkage to the authentication providers, setting and changing attributes and privileges, and decommissioning of the identity.

Maintenance

This area is comprised of user management, password management, and role/group management. User management defines the set of administrative functions such as identity creation, propagation, and maintenance of user identity and privileges.

Entitlement

According to the Open Group, entitlement is a set of rules, defined by the resource owner, for managing access to a resource (asset, service, or entity), and for what purpose. The level of access not only is conditioned by your identity but is also likely to be constrained by a number of further security considerations such as your company policy, your location (i.e., are you inside your secure corporate environment, connected via a hotspot, or working from an Internet café, etc.), or time of day.[24]

SUMMARY

Controlling physical access to IT assets is an important element in protecting the availability and integrity of services provided by the assets. Many considerations factor into the selection and implementation of physical access controls. Using multiple layered controls

such as deterrence and detection safeguards and response and recovery plans provide the greatest span of asset protection. Different types of assets may require specific access control systems to be able to effectively limit access to authorized personnel. Since information systems are comprised of various components, the availability of the information system as a whole is one of the main goals of an effective access control. The availability of the entire information system is only as strong as the weakest control afforded each component. The security practitioner should always be focused on the systems that are deployed to achieve access control within the environment, as well as how they are being monitored and maintained.

SAMPLE QUESTIONS

1. What type of controls are used in a Rule Set–Based Access Control system?
   - **a.** Discretionary
   - **b.** Mandatory
   - **c.** Role Based
   - **d.** Compensating
2. What framework is the Rule Set–Based Access Controls logic based upon?
   - **a.** Logical Framework for Access Control
   - **b.** Specialized Framework for Access Control
   - **c.** Technical Framework for Access Control
   - **d.** Generalized Framework for Access Control
3. View-Based Access Controls are an example of a(n):
   - **a.** Audit control
   - **b.** Constrained User Interface
   - **c.** Temporal constraint
   - **d.** Side Channel
4. Which of the following are supported authentication methods for iSCSI? (Choose two.)
   - **a.** Kerberos
   - **b.** Transport Layer Security (TLS)
   - **c.** Secure Remote Password (SRP)
   - **d.** Layer 2 Tunneling Protocol (L2TP)

5. According to the following scenario, what would be the most appropriate access control model to deploy?

   Scenario: A medical records database application is used by a health-care worker to access blood test records. If a record contains information about an HIV test, the health-care worker may be denied access to the existence of the HIV test and the results of the HIV test. Only specific hospital staff would have the necessary access control rights to view blood test records that contain any information about HIV tests.

   a. Discretionary Access Control

   b. Context-Based Access Control

   c. Content-Dependent Access Control

   d. Role-Based Access Control

6. Which of the following is *not* one of the three primary rules in a Biba formal model?

   a. An access control subject cannot request services from an access control object that has a higher integrity level.

   b. An access control subject cannot modify an access control object that has a higher integrity level.

   c. An access control subject cannot access an access control object that has a lower integrity level.

   d. An access control subject cannot access an access control object that has a higher integrity level.

7. Which of the following is an example of a firewall that does not use Context-Based Access Control?

   a. Static packet filter

   b. Circuit gateway

   c. Stateful inspection

   d. Application proxy

8. Where would you find a singulation protocol being used?

   a. Where there is a Radio Frequency ID system deployed and tag collisions are a problem

   b. Where there is router that has gone offline in a multi-path storage network

   c. Where there is a Radio Frequency ID system deployed and reader collisions are a problem

   d. Where there is switch that has gone offline in a multi-path storage network

9. Which of the following are not principal components of access control systems? (Choose two.)
   a. Objects
   b. Biometrics
   c. Subjects
   d. Auditing
10. Which of the following are behavioral traits in a biometric device?
   a. Voice pattern and keystroke dynamics
   b. Signature dynamics and iris scan
   c. Retina scan and hand geometry
   d. Fingerprint and facial recognition
11. In the measurement of biometric accuracy, which of the following is commonly referred to as a "type 2 error"?
   a. Cross-over error rate (CER)
   b. Rate of false rejection—False Rejection Rate (FRR)
   c. Input/output per second (IOPS)
   d. Rate of false acceptance—False Acceptance Rate (FAR)
12. What is the difference between a synchronous and asynchronous password token?
   a. Asynchronous tokens contain a password that is physically hidden and then transmitted for each authentication, while synchronous tokens do not.
   b. Synchronous tokens are generated with the use of a timer, while asynchronous tokens do not use a clock for generation.
   c. Synchronous tokens contain a password that is physically hidden and then transmitted for each authentication, while asynchronous tokens do not.
   d. Asynchronous tokens are generated with the use of a timer, while synchronous tokens do not use a clock for generation.
13. What is an authorization table?
   a. A matrix of access control objects, access control subjects, and their respective rights
   b. A service or program where access control information is stored and where access control decisions are made
   c. A listing of access control objects and their respective rights
   d. A listing of access control subjects and their respective rights

14. What ports are used during Kerberos Authentication?
    a. 53 and 25
    b. 169 and 88
    c. 53 and 88
    d. 443 and 21
15. What are the five areas that make up the identity management lifecycle?
    a. Authorization, proofing, provisioning, maintenance, and establishment
    b. Accounting, proofing, provisioning, maintenance, and entitlement
    c. Authorization, proofing, provisioning, monitoring, and entitlement
    d. Authorization, proofing, provisioning, maintenance, and entitlement

NOTES

[1] See the following for the full DoD 5200.28 TCSEC Standard: `http://csrc.nist.gov/publications/history/dod85.pdf`

[2] See the following for the original paper by Leonard J. LaPadula, "Rule-Set Modeling of a Trusted Computer System": `http://www.acsac.org/secshelf/book001/09.pdf`

[3] See the following for RFC 3415—View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP): `http://tools.ietf.org/html/rfc3415`. This document describes the View-based Access Control Model (VACM) for use in the Simple Network Management Protocol (SNMP) architecture. It defines the Elements of Procedure for controlling access to management information. This document also includes a Management Information Base (MIB) for remotely managing the configuration parameters for the View-based Access Control Model.

[4] Firewalls make context-based access decisions when they collect state information on a packet before allowing it to transit the network. A stateful firewall understands the necessary steps of communication for specific protocols. For example, in a TCP connection, the sender sends a SYN packet, the receiver sends a SYN/ACK, and then the sender acknowledges that packet with an ACK packet. A stateful firewall understands these different steps and will not allow packets to go through that do not follow this sequence. Therefore, if a stateful firewall receives a SYN/ACK and there was not a previous SYN packet that correlates with this connection, the firewall "understands" that the context for this packet is not right and disregards the packet.

[5] See the following for the original research paper that defined TRBAC as a model. "TRBAC: A Temporal Role-Based Access Control Model" by Elisa Bertino, Piero Andrea Bonatti, and Elena Ferrari. *ACM Transactions on Information and Systems Security*, Vol. 4, No.3, August 2001, 191-223.

[6] See the following for the NIST Interagency Report 7316: "Assessment of Access Control Systems." http://csrc.nist.gov/publications/nistir/7316/NISTIR-7316.pdf

[7] http://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.sp.800-162.pdf

[8] Read more about Bell–La Padula here: http://www.acsac.org/2005/papers/Bell.pdf

[9] Read more about the Biba Integrity Model here: http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA166920 (Page 27).

[10] Read more about the Clark–Wilson model here: http://www.cs.clemson.edu/course/cpsc420/material/Policies/Integrity%20Policies.pdf

[11] As it turns out, Biba addresses only one of three key integrity goals. The Clark–Wilson model improves on Biba by focusing on integrity at the transaction level and addressing three major goals of integrity in a commercial environment. In addition to preventing changes by unauthorized subjects, Clark and Wilson realized that high-integrity systems would also have to prevent undesirable changes by authorized subjects and to ensure that the system continued to behave consistently. It also recognized that it would need to ensure that there is constant mediation between every subject and every object if such integrity was going to be maintained.

[12] See the following for an overview of the major Linux Security Distros: http://www.itproportal.com/2016/02/02/the-top-10-linux-security-distros/

[13] For an overview of the entire ISO 7816 Standards set, see the following: http://www.smartcardsupply.com/Content/Cards/7816standard.htm

[14] See the following: http://www.technovelgy.com/ct/Technology-Article.asp?ArtNum=20

[15] Although it requires a higher level of skill, keystrokes can be hacked. At DEFCON 17, Andrea Barisani and Daniele Bianco demonstrated how to sniff keystrokes using unconventional side channel attacks. Wires in PS/2 keyboards leak information from the data wire into the ground wire, which acts like an antenna. The leaked information about the keyboard strokes can be detected on the power outlet, as well as other wires on the same electrical system. By slicing open one of these lines, cutting the ground wire, and attaching a probe, the line can be monitored and the signal isolated by filtering out the noise using software such as Scilab. The waves from the oscilloscope and the data can be streamed to the hacker's computer where additional software is used to extract the victim's keystroke information. In addition, a research team from the Ecole Polytechnique Federale de Lausanne was able to pick up electromagnetic radiation that is generated every time a computer keyboard is tapped by using an oscilloscope and an inexpensive wireless antenna; the team was able to pick up keystrokes from virtually any keyboard, including laptops, with 95 percent accuracy. See the following for more information on each of these instances:

1. DEFCON 17: Sniff Keystrokes With Lasers/Voltmeters—YouTube video: http://www.youtube.com/watch?v=xKSq9efXmh8

2. Robert McMillan. "A Way to Sniff Keystrokes From Thin Air," *PCWorld*. http://www.pcworld.com/article/161166/article.html

[16] See the following for the BBC News article describing the events at Black Hat 2012: http://www.bbc.co.uk/news/technology-18997580#

[17] See the following for the complete report on the research findings: Arora, S.S.; Vatsa, M. ; Singh, R. ; Jain, A. "Iris recognition under alcohol influence" (Conference Publications). 978-1-4673-0397-2. Biometrics (ICB), 2012 5th IAPR International Conference, pp. 336–341.

[18] See the following: http://en.wikipedia.org/wiki/Security_token

[19] See the following: http://www.smartcardalliance.org/smart-cards-intro-primer/

[20] See the following for the Authentication in an Internet Banking Environment recommendation: http://www.ffiec.gov/pdf/authentication_guidance.pdf

[21] See the following for RFC 1510, The Kerberos Network Authentication Service (V5): http://www.ietf.org/rfc/rfc1510.txt

[22] See the following for the MIT Kerberos home page: http://web.mit.edu/kerberos/

[23] See the following: http://www.gartner.com/it-glossary/identity-proofing-services

[24] See the following: http://blog.opengroup.org/2012/08/07/entities-and-entitlement-the-bigger-picture-of-identity-management/