Chapter 4

Authentication and Access Control

Three-quarters of all businesses in North America are connected to the Internet. In Australia/Oceana, this figure falls to 60 percent, and in Europe (East and West combined) it is 50 percent (source: *http://www.internetworldstats.com/stats.htm*). There is little doubt that businesses are increasingly seeing online transactions as a way to improve business efficiency.

As organizations embrace the use of electronic transactions, the "information velocity" (a term made famous by Bill Gates in his book *Business @ the Speed of Thought*) increases, and the speed at which they make decisions increases. This phenomenon, along with the overall greater access to information that the Internet affords, improves business decision-making within the firm, which in turn increases revenue and decreases costs.

The result is that more organizations transacting business on the Internet need to assure themselves that the entities with whom they are doing business are who they purport to be and can legitimately do business with them. Failure to do so might result in unauthorized ordering of goods, illegal transfer of funds, or malicious alteration of data.

The act of verifying the credentials (which could be identity, qualifications, or authorization level) of an entity (it could be a person or a business entity) is called *authentication*. The core activity of any identity management environment is to provide authentication services. Authentication, as the word implies, is the act of verifying a person's identity as the person tries

to access restricted resources. This process most commonly refers to the log-on procedure that users must complete before being granted access to a company's computing resources. (The terms "account log-on" and "network log-on" are used synonymously.)

Authentication differs from *authorization*, which is the act of granting access to a specific computer application or maybe to just one or two of the application's features. This process is often referred to as *access control*, which is a somewhat broader term that encompasses physical access to buildings as well as logical access to computer systems. Either way, a user's credentials are compared with an access control list that determines the level of access the user is entitled to receive.

Authentication, then, is the act of confirming that users are who they purport to be before granting them access to corporate resources. Once a user is authenticated, authorization provides access to computer programs (applications) commensurate with the user's authenticated identity. This activity is a critical one for any organization, but it becomes particularly acute for a company with high-security requirements. All companies have security issues; for instance, they don't want external entities to gain access to their price lists, inventory levels, or strategic direction statements. Some businesses, such as pharmaceutical companies, defense-related organizations, or companies working in sensitive areas, must protect their resources to a higher degree. The higher the security requirement, the higher the cost to implement a mechanism that protects corporate resources.

Equally, it makes little sense for a company to spend a lot of money implementing elaborate firewalls and monitoring facilities if there is little reason for anyone to try to gain access to the company's facilities in the first place. Before an authentication mechanism is put in place, it is a good idea to conduct (and document) a risk assessment that identifies the degree to which resources need to be protected. Such an assessment should include the reasons for selecting the preferred authentication mechanism.

Another word requiring definition in this discussion is *validation*. Typically, the validation stage refers to the check of identity source documents as part

of an enrollment process. Before gaining access to protected resources, a person must produce identity documents to validate his or her identity claims. This *evidence of identity (EoI)* check is an integral part of the validation process. Validation is undertaken once, whereas authentication occurs whenever the user logs on to the network. Figure 4.1 summarizes the points at which validation, authentication, and authorization come into play.

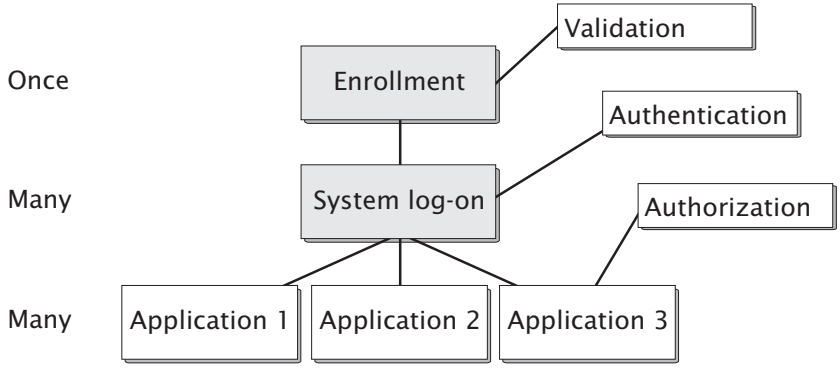

Figure 4.1: Validation, authentication, and authorization

Methods of Authentication

By far, the most common authentication method is user name and password. Approximately 95 percent of identity management systems use passwords to authenticate users, leaving but 5 percent for all the other mechanisms. This fact isn't surprising, because passwords are usually quite satisfactory for the purpose to which they are put. Remember that there are two reasons for authentication: identification and protection. For these purposes, passwords are generally sufficient.

Identification

We want to identify the user accessing our resources. When someone accesses articles over the Internet, the publisher likely wants to know who that person is so it can track who, and how many, people are downloading from its Web site. The publisher also might want to follow up with marketing material. So it needs a basic level of identification, but it has no need for a high degree of authentication. In this case, a password

authentication mechanism is sufficient to enable tracking the user's interest in the publisher's articles.

If, however, a liability is associated with the service being accessed, more than just a password may be necessary. This point brings up another word requiring definition: *repudiation*. For companies that provide a service for which they must be sure the user is who he or she purports to be, passwords may not be enough. If something goes wrong with the service provision, or with payment for it, it is important for the service provider to be able to go back to the user and ensure that the user can't "repudiate," or refute the validity of, the transaction or claim it was in fact someone else who undertook the transaction. If the possibility exists that the authentication system might have been compromised, the user could in fact repudiate the transaction. (This consideration is important for transactions involving credit card payments because over the Internet such transactions are not "card-present" transactions; they are conducted under "money order, telephone order," or MOTO, rules and in most countries can be repudiated.)

Protection

We want to protect our resources from inappropriate or harmful use. If we find that a user is misusing our service, we want to be able to go back to the user to rectify the situation. For this purpose, password protection is likely satisfactory.

If a greater level of authentication is required, we have multiple options:

- *One-time password* — In this scenario, users are issued a hardware token that is synchronized with the organization's back-end systems. A display on the token shows a number that changes approximately every minute. When users log in, the system prompts them to enter the current number to substantiate that they are who they purport to be. If a token is lost, the organization must be advised immediately so that the device can be taken out of service.
- *Challenge response* — This method is widely used in password self-service applications. The Achilles heel of a user-name/password

authentication mechanism is the problem of users forgetting their passwords. Most organizations can attest to the high number of help-desk calls to renew passwords. Most businesses now use a password reset facility that requires users to establish one or more challenge questions and their response. When a user wants to change his or her password, a challenge in the form of one or more questions is issued; upon receiving the correct response, the system updates the password.

- *Digital certificate* — Issuing a digital certificate to a user requires evidence of the completion of an identity step, in which the user is required to produce one or more forms of identification before the certificate is issued. Accompanying the certificate is a private key that must be safeguarded. Often, this key is provided on a token storage device, such as a smartcard or a USB memory stick.
- *Biometrics* — Another form of authentication that is generally considered more secure uses biometric identification. A high level of confidence can be provided with the storage of users' biometric detail. Popular biometrics include fingerprints, facial image templates, and iris scans. These authentication methods obviously require the installation of hardware that users can access, and they are not generally used by organizations with a controlled population. One area in which biometrics are being used with members of the public is electronic passports.

Combining Authentication Methods

By combining authentication methods, organizations can increase the security, and therefore the protection, that an authentication scheme provides.

One-factor Authentication

Single-factor authentication mechanisms typically rely on "something you know," and this something is usually a password. (Passwords fall into a category of authentication known as *shared secret* methodology. This mechanism is widely used for over-the-phone authentication and self-service password resets.) Under such a method, if you can enter your user name and password, you will be granted access to the system. This approach provides a relatively weak form of authentication because one user might give his

or her password to someone else, allowing the second individual to fraudulently access the system in question.

Varying strengths are associated with passwords. Many systems require a password to be a combination of letters and numbers and to include at least one case change. Some systems require the use of at least one special character in the password.

Unfortunately, one simple mechanism to ensure a strong password is often obviated by the system itself. Permitting a user to select a phrase as the password reduces the possibility that the user will forget the password and lessens the likelihood that a brute-force attack will be successful, but many systems restrict password length to 15 characters, and some legacy systems support only eight characters.

Some single-factor, or "shared information," authentication systems use a more sophisticated challenge-response methodology that includes multiple "questions" that the user must answer correctly to be authenticated. In this situation, users establish one or more questions to which they, and only they, would be expected to know the answer. The system stores the responses, enabling persons (or systems) to verify that they are who they say they are because they know the answers. A typical question is "What is your mother's maiden name?" or "Where did you first go to high school?" or "What is your favorite color?" Systems may ask multiple questions and accept a combination of correct answers.

Two-factor Authentication

Two-factor authentication mechanisms typically rely on "something you know" and "something you have." Users are required not only to know a password (or PIN) but also to have something, such as a security dongle that plugs into the USB port or a smartcard that must be inserted into a card-reader receptacle, to gain access to the system.

One-time passwords also fall into this category because they rely on the possession of a hardware device that displays the required password. When prompted by the system being accessed, the user reads the password

currently displayed on the hardware device and enters it into the system. The user can be identified because the system knows the password being displayed on each device at any point in time. Passwords typically change every minute or so to ensure that the user has the device at the exact time he or she is authenticating to the system.

Three-factor Authentication

Three-factor authentication mechanisms require users to display "something you know," "something you have," and "something you are." In this instance, a user might be required to carry a smartcard with a biometric feature on it. Typically, biometrics are fingerprints or facial templates that carry the unique characteristics of the user's fingerprint or facial features.

In a typical three-factor authentication system, a user plugs the smartcard into a reader (something you have), types in a PIN (something you know), and has a facial recognition system verify the facial template (something you are).

Choosing a Methodology That's Right for You

Although the preceding discussion indicates the normal selection of one-, two-, and three-factor authentication mechanisms, in reality you can combine these authentication methods in any way to meet the required protection:

- Something you know (shared secret)
- Something you have (dongle, token card, signed and verified certificate)
- Something you are (biometrics)

These methods can be used together in any number of ways. Each factor is something verifiable, and as they are combined, they provide stronger authentication. Even if a biometric method (considered quite strong) is selected on its own, it is still a single-factor authentication schemes.

Similarly, a two-factor authentication scheme isn't always something you know plus something you have; a two-factor scheme might combine a password with biometrics.

Also note that while more factors are generally associated with greater security, each additional authentication method represents more inconvenience for users. It is important for authentication schemes not to impose an authentication method just because they can; the scheme must match the level of authentication required to provide the desired level of protection and security.

Levels of Authentication

The levels of authentication cover a continuum from a simple password system to an elaborate public key infrastructure (PKI) installation. Companies and governments typically recognize four levels of authentication. (We exclude the additional "no authentication" level here. Organizations should not install an authentication mechanism unless it is really required.)

As Figure 4.2 depicts, as the level of risk (gauged by the severity of consequence in the event that the risk is triggered) increases, the authentication mechanism must change appropriately.

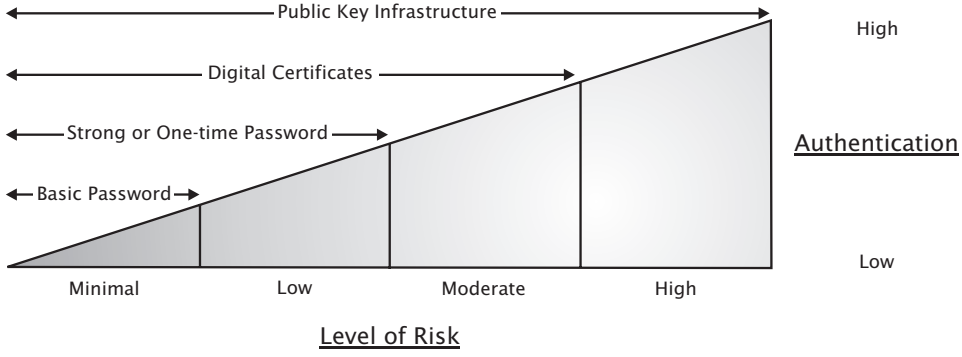

Figure 4.2: Levels of risk vs. authentication

At a basic, or minimal, risk level, a user-name/password authentication mechanism will suffice. At a somewhat higher level of risk, it might be necessary to implement a strong password format with a minimum length and the inclusion of mandatory character types. At a moderate level of risk, the organization might need to issue participants a digital certificate, kept

either on their PC or on a token storage device such as a smartcard. At a high risk level, it will be necessary to implement a public key infrastructure whereby each participant is issued an asynchronous private-public key-pair with a public key certificate.

Authentication Assurance Levels

It should be obvious by now that an identity management environment cannot be designed without a good understanding of the risks associated with access to the resources being managed by the selected authentication mechanism. It is necessary, therefore, for the level of assurance provided by that mechanism to match the protection need. An assessment of the required level of assurance will determine the selection of the most appropriate authentication mechanism.

The four levels of risk identified in Figure 4.2 can be mapped onto the level of assurance as set out in Table 4.1.

Table 4.1: Risk levels and associated consequences

Risk level	Consequences of compromise
Minimal	• Insignificant inconvenience to either party • No release of private or sensitive information • No threat to commercial or government interests • No opportunity for associated criminal activity
Low	• Possible inconvenience to either party • No release of private or sensitive information • Minor threat of financial loss to either party • No threat to government interests • No opportunity for associated criminal activity
Moderate	• Significant inconvenience to either party • Possible release of private or sensitive information • Threat of significant financial loss to either party • Threat to non-national security government interests • Possible opportunity for associated criminal activity
High	• Major inconvenience to either party • Release of private or sensitive information • Significant financial loss to either party • Threat to government interests • Threat to national security • Opportunity for associated criminal activity

Registration Assurance Levels

While the initial registration process of an authentication mechanism normally will match the mechanism's assurance level, it is worth noting that this process is an important part of any identity management facility. Put bluntly, a full-blown public certificate infrastructure will technically provide a bulletproof solution to most authentication requirements, but if you can drive a truck through the registration process, your solution is a total waste of effort and money.

The registration process for any authentication mechanism should match the level of assurance that the mechanism purports to provide. Again, the registration rigor should be evaluated on a four-level scale, as Table 4.2 describes.

Table 4.2: Registration process requirements

Confidence level	Registration process description
Low	Self-registration provision of basic identity data (name, address, and contact details) is conducted, but no validation of documentation.
Medium	Some validation of identity details is performed with self-registration (e.g., ZIP code check, email address validation), but manual validation is typical.
High	A recognized evidence of identity check is performed with sighting of appropriate identity documents.
Very high	A substantial in-person evidence of identity check is performed, with a formal validation of identity documents and retention of proof.

It is important that the evidence of identity check performed as part of the identity validation matches the requirement of the authentication mechanism.

The 100-point check conducted by financial institutions in Australia is a popular one but should be reviewed before adoption to ensure that the correct attributes are being verified. This check classifies identity documents according to their veracity and credibility. Documents such as birth certificates and passports are typically category A documents, worth 70 points. Documents of a less robust nature, such as driver's licenses, mortgage documents, or student cards, are category B documents and are of less value typically 25, 35, or 40 points. To satisfactorily complete an

EoI check, an applicant must show one category A document and sufficient category B documents to compile the requisite 100 points.

Table 4.3 shows the risk matrix that combines the authentication assurance rating and the identity registration assurance rating.

Table 4.3: Authentication risk matrix

Registration assurance	Authentication assurance			
	Minimal	Low	Moderate	High
Low				
Medium				
High				
Very High				

The minimal authentication assurance level will command either an unvalidated registration process or a medium registration process in which there is a basic level of validation. The low assurance level will need at least a basic level of identity validation. The moderate level will require a check of identity credentials. The high level will require a substantial check, possibly via a third-party registration authority. Remember, the high level often will require non-repudiation and be associated with a significant financial liability. For this reason, the registration agent may be required to retain copies of identity documentation.

Access Control

Authentication is the basic mechanism for restricting access to a company's corporate resources. These resources are typically computer resources but also can include physical access to the company's buildings or equipment. If someone has been issued a password, digital certificate, or other authentication mechanism, that person has been given the "keys to the kingdom." He (or she) can access whatever he has been authorized for, can request extended access rights, and will retain that access until it is rescinded.

Identity management is crucial to managing this access and protecting the corporation's assets.

Identities and Access Control

Authorization or access control is the "raison d'être" for most identity management deployment. While there may be some benefit inherent in effectively and efficiently managing the identities within an organization, these activities are usually conducted for the purpose of granting access to restricted facilities, both virtual and physical.

Access control, by definition, must be real-time. As a user attempts to gain access to a computer application, the access control system must provide the user credentials to enable the user to gain the appropriate access. For instance, an account clerk might get access to the company's financial system to allow the entry of a customer transaction. The finance manager, however, will require far greater access to be able to create reports and monitor all activity in the system. It is the access control mechanism that will provide this differentiation.

Controlling access to computer applications is becoming more important for organizations as the focus on properly managing access to documents and files increases. It is important that access be available only based on a proven identity validated by a trusted entity. This access must be integrated with the organization's identity management environment. In too many companies, the access control mechanism is independent and open to discrepancy.

Single Sign-on

One of the biggest issues with a disassociated access control mechanism is the potential for multiple sign-ons. Once a user has logged on to the company system, he or she must then individually log on to various applications, retyping user names that often differ between applications and entering passwords that are not synchronized (i.e., when one system forces a password change, it is not copied to other systems). Often, the password change frequency isn't synchronized either, with some applications requiring changes every month, some every 90 days, and some never. The result

is that users are forced to remember multiple user names and passwords and often resort to unsafe practices such as keeping written records of passwords or not changing their passwords at regular intervals.

For these reasons, there is a growing emphasis on *single sign-on (SSO)*—integrating the access control mechanisms for multiple applications.

Enterprise SSO

Enterprise SSO refers to the integration of the main corporate applications. This integration is typically quite difficult to achieve because the corporate applications are often spread over multiple types of systems. For example, the main enterprise resource planning system might be on an IBM mainframe, the corporate financial system might be an Oracle application on an IBM i system, and the corporate email might be Microsoft Exchange Server 2007. Each of these applications typically will have a separate access authorization mechanism populated individually by separate administrators. This is a common but very costly environment. Not only does it require multiple administrators to keep the access control lists up-to-date, but it also means that inevitably there will be differences between the applications (e.g., some users not removed when they leave the organizations, some with their names spelled differently in different applications) and the need for users to remember multiple passwords.

To integrate this environment is not trivial. While some applications operating in a Microsoft environment can be managed with Windows Integration, non-Windows applications remain hard to integrate into a single authorization environment. Sometimes the best that can be accomplished is the synchronization of the underlying identity stores.

Web SSO

One area in which integration of applications is typically easier to achieve is the Web environment. Users may be connecting to multiple applications, but they are all operating in a common environment. Once a user has logged on, the access control credentials are more easily passed between applications. There are multiple ways to achieve this integration, depending

on the ways in which each application grants access to users. As Figure 4.3 illustrates, some applications maintain their own identity repositories.

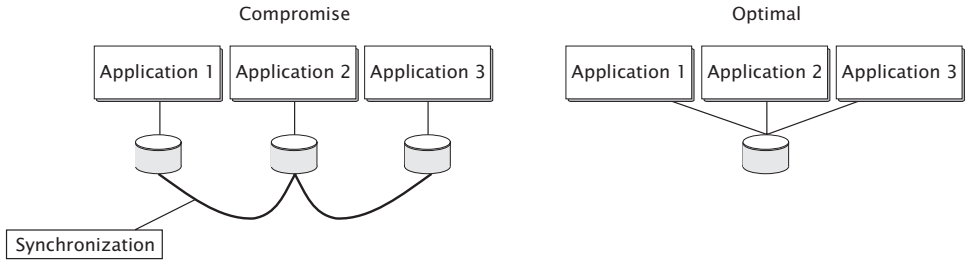

Figure 4.3: Multiple vs. single identity stores

There are two requirements for moving to the optimal solution:

- The data store must have the capability of handling the sophistication required in the dependent applications.
- The applications must be capable of communicating with the central directory.

In many cases, the applications are legacy systems that were built to operate only with their data stores, and it would be foolhardy to try to modify them. In such cases, synchronization is the preferred strategy. However, the enterprise architecture of an organization should mandate the approach for any further application development or acquisition. Increasingly, all applications are required to be LDAP-capable, and in some cases in which federated authentication is required, companies are mandating Security Assertion Markup Language (SAML) compliance.

While support for the same version of SAML between two applications does not guarantee interoperability, it does heighten the chances of achieving it. SAML will look after the sharing of messages and ensure they are intelligible to both parties; the sender and receiver must agree on the content and meaning of the message component. The Extensible Access Control Markup (XACML), discussed later in the chapter, helps.

Fine-grained Access Control (a.k.a. Entitlement Management)

With the growing focus on compliance, be it forced or self-imposed, there is an increasing call on IT infrastructure to provide *fine-grained access control*, also called entitlement management.

While access control provides protection in that users are granted access only to the resources specifically enabled by the permissions they are granted, fine-grained control goes further, restricting access to specific times of day or to specific features available in the facility. For instance, students may be granted access to a laboratory only during class time; outside of class time, the facility will not be available. To accomplish this fine-grained level of control, we need a more sophisticated authorization mechanism.

A fine-grained access control environment will typically divide the specific functions required into discrete points, as depicted in Figure 4.4. This level of abstraction provides a better design that allows for the selection of the best product for each component of the environment.

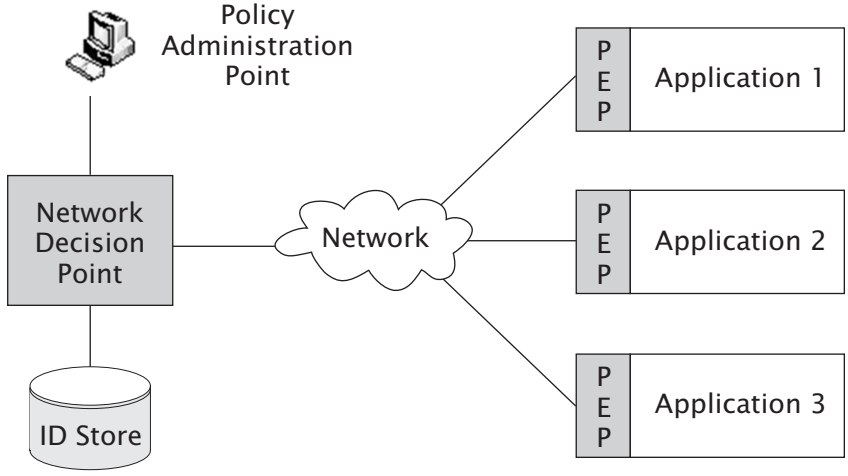

Figure 4.4: Fine-grained access control

The *policy decision point (PDP)* is the component of the environment that matches access requests with the criteria for access and notifies the *policy*

enforcement point (PEP). The PEP is the point at which access requests are granted or refused. The *policy administration point (PAP)* is the facility used to manage policy definitions.

Much effort has been expended over the past few years to define a standard for fine-grained access control and for the transmission of entitlement messages between PDP and PEPs. The result is the XACML standard, whereby entitlement decisions can be communicated between PDPs and PEPs.

XACML

The Extensible Access Control Markup language is a policy language that lets application administrators set the access control parameters for users of their applications. XACML combines both a data schema and associated language to combine complex rules and associated logic to make real-time decisions about a user's access rights to the application.

The policy enforcement point passes the policy set to the policy decision point, which passes the user information through its logic and returns the authorization decision. The PEP then responds to the user's request with the appropriate response.

A policy set consists of a target, a rule, and an obligation. The target contains the conditions that the user (subject) must meet to access the resource and the action required to meet the policy set. Successfully meeting the policy set or rule will return a "permit" decision to the PEP.

Discussion Questions

1. In what way do the attributes required for authentication differ from those required for authorization? Which are more volatile?
2. Why does authentication typically use a directory and authorization typically use a database? Discuss.
3. Both MasterCard and Visa have introduced higher-level authentication mechanisms with their MasterCard Secure and Verified by Visa programs. If the purpose of authentication is primarily to

identify and protect, who are these programs identifying and who are they protecting? Discuss.

4. Why is an integrated identity data store more difficult to attain than a distributed system with synchronization between data stores? Discuss which approach is preferable.
5. Why is Web SSO easier to achieve than enterprise SSO?
6. Think of a situation in which fine-grained access control would be beneficial. What are the attributes that a policy decision point would need to know before it could grant access to a user?

Case Study

Refer to the case study in Appendix A in answering the following questions.

1. A large part of the enrollment process is out of the university's control. It is conducted by a government university admissions office, with files of prospective students periodically passed to the university. What are the advantages of such a system to the university? Are there any disadvantages?
2. When a new staff member joins the university, that person is entered in the HR system, and, when approved by the HR manager, a nightly batch process populates the authentication directory (Active Directory). Discuss why this process is efficient or inefficient. How might you redesign the authentication process?
3. Would you consider the university systems to be high security or low security? What systems within the university might warrant digital signatures? What level of assurance do digital signatures relate to?