1

Fundamental Security Concepts

CERTIFICATION OBJECTIVES

1.01	Describe Principles of Information Security
1.02	Explain Information Security Fundamentals and Define Good Security Architectures
✓	Two-Minute Drill
Q&A	Self Test

We'll begin Part I of the book with the discussion of fundamental concepts and principles of information security. These general concepts and principles are relevant in all computing environments and serve as the foundation upon which all security mechanisms and controls are designed and implemented, regardless of the particular hardware platform, operating system, or application.

CERTIFICATION OBJECTIVE 1.01

Describe Principles of Information Security

First, let's define information security. If ten different people were asked to define information security, we might well receive ten different answers, but what is surprising is that they might all be correct. Nevertheless, the universal, classic definition of information security is brief and simple:

> Information security is the confidentiality, integrity, and availability of information.

Indeed, all the principles, standards, and mechanisms you will encounter in this book are dedicated to these three abstract but fundamental goals of confidentiality, integrity, and availability of information and information processing resources—also referred to as the *C-I-A triad* or *information security triad.*

Confidentiality

In the context of information security, *confidentiality* means that information that should stay secret stays secret and only those persons authorized to access it may receive access. From ancient times, mankind has known that information is power, and in our information age, access to information is more important than ever. Unauthorized access to confidential information may have devastating consequences, not only in national security applications, but also in commerce and industry. Main mechanisms of protection of confidentiality in information systems are cryptography and access controls. Examples of threats to confidentiality are malware, intruders, social engineering, insecure networks, and poorly administered systems.

Integrity

Integrity is concerned with the trustworthiness, origin, completeness, and correctness of information as well as the prevention of improper or unauthorized modification of information. Integrity in the information security context refers not only to integrity of information itself but also to the origin integrity—that is, integrity of the source of information. Integrity protection mechanisms may be grouped into two broad types: preventive mechanisms, such as access controls that prevent unauthorized modification of information, and detective mechanisms, which are intended to detect unauthorized modifications when preventive mechanisms have failed. Controls that protect integrity include principles of least privilege, separation, and rotation of duties—these principles are introduced later in this chapter.

Availability

Availability of information, although usually mentioned last, is not the least important pillar of information security. Who needs confidentiality and integrity if the authorized users of information cannot access and use it? Who needs sophisticated encryption and access controls if the information being protected is not accessible to authorized users when they need it? Therefore, despite being mentioned last in the C-I-A triad, availability is just as important and as necessary a component of information security as confidentiality and integrity. Attacks against availability are known as *denial of service* (*DoS*) attacks and are discussed in Chapter 7. Natural and manmade disasters obviously may also affect availability as well as confidentiality and integrity of information, though their frequency and severity greatly differ—natural disasters are infrequent but severe, whereas human errors are frequent but usually not as severe as natural disasters. In both cases, business continuity and disaster recovery planning (which at the very least includes regular and reliable backups) is intended to minimize losses.

exam watch

Understanding the fundamental concepts of confidentiality, integrity, and availability of information and their interaction is crucial for this exam. Make sure you know their definitions, summarized here, and can give examples of controls protecting them:

- *Confidentiality* ***is the prevention of unauthorized disclosure of information.***
- *Integrity* ***aims at ensuring that information is protected from unauthorized or unintentional alteration, modification, or deletion.***
- *Availability* ***aims to ensure that information is readily accessible to authorized users.***

Now that the cornerstone concepts of confidentiality, integrity, and availability have been discussed, let's take a look at identification, authentication, and authorization processes and methods, which are some of the main controls aimed at protecting the C-I-A triad.

Identification

Identification is the first step in the identify-authenticate-authorize sequence that is performed every day countless times by humans and computers alike when access to information or information processing resources are required. While particulars of identification systems differ depending on who or what is being identified, some intrinsic properties of identification apply regardless of these particulars—just three of these properties are the *scope*, *locality*, and *uniqueness* of IDs.

Identification name spaces can be local or global in scope. To illustrate this concept, let's refer to the familiar notation of Internet e-mail addresses: while many e-mail accounts named *jack* may exist around the world, an e-mail address *jack@company.com* unambiguously refers exactly to one such user in the company .com locality. Provided that the company in question is a small one, and that only one employee is named Jack, inside the company everyone may refer to that particular person by simply using his first name. That would work because they are in the same *locality* and only one Jack works there. However, if Jack were someone on the other side of the world or even across town, to refer to *jack@company.com* as simply *jack* would make no sense, because user name *jack* is not *globally unique* and refers to different persons in different *localities*. This is one of the reasons why two user accounts should never use the same name on the same system—not only because you would not be able to enforce access controls based on non-unique and ambiguous user names, but also because you would not be able to establish accountability for user actions.

To summarize, for information security purposes, unique names are required and, depending on their scope, they must be locally unique and possibly globally unique so that access control may be enforced and accountability established.

Authentication

Authentication, which happens just after identification and before authorization, verifies the authenticity of the identity declared at the identification stage. In other words, it is at the authentication stage that you prove that you are indeed the person or the system you claim to be. The three methods of authentication are *what you* *know, what you have,* or *what you are*. Regardless of the particular authentication method used, the aim is to obtain reasonable assurance that the identity declared at the identification stage belongs to the party in communication. It is important to note that *reasonable assurance* may mean different degrees of assurance, depending on the particular environment and application, and therefore may require different approaches to authentication: authentication requirements of a national security–critical system naturally differ from authentication requirements of a small company. Because different authentication methods have different costs and properties as well as different returns on investment, the choice of authentication method for a particular system or organization should be made after these factors have been carefully considered.

What You Know

Among *what you know* authentication methods are passwords, passphrases, secret codes, and personal identification numbers (PINs). When using *what you know* authentication methods, it is implied that if you know something that is supposed to be known only by *X*, then you must be *X* (although in real life that is not always the case). *What you know* authentication is the most commonly used authentication method thanks to its low cost and easy implementation in information systems. However, *what you know* authentication alone may not be considered strong authentication and is not adequate for systems requiring high security.

exam watch

Strong authentication is the use of two or more different authentication methods, such as a smart card and a PIN, or a password and a form of biometrics such as a fingerprint or retina scan.

What You Have

Perhaps the most widely used and familiar *what you have* authentication methods are keys—keys we use to lock and unlock doors, cars, and drawers; just as with doors, *what you have* authentication in information systems implies that if you possess some kind of token, such as a smart card or a USB token, you are the individual you are claiming to be. Of course, the same risks that apply to keys also apply to smart cards and USB tokens—they may be stolen, lost, or damaged. *What you have* authentication methods include an additional inherent per-user cost. Compare these methods with passwords: it costs nothing to issue a new password, whereas per-user *what you have* authentication costs may be considerable.

What You Are

What you are authentication refers to biometric authentication methods. A *biometric* is a physiological or behavioral characteristic of a human being that can distinguish one person from another and that theoretically can be used for identification or verification of identity. Biometric authentication methods include fingerprint, iris, and retina recognition, as well as voice and signature recognition, to name a few. Biometric authentication methods are less well understood than the other two methods but when used correctly, in addition to *what you have* or *what you know* authentication, may significantly contribute to strength of authentication. Nevertheless, biometrics is a complex subject and is much more cumbersome to deploy than *what you know* or *what you have* authentication. Unlike *what you know* or *what you have* authentication methods, whether or not you know the password or have the token, biometric authentication systems say how much you are like the subject you are claiming to be; naturally this method requires much more installation-dependent tuning and configuration.

Authorization

After declaring identity at the identification stage and proving it at the authentication stage, users are assigned a set of *authorizations* (also referred to as rights, privileges, or permissions) that define what they can do on the system. These authorizations are most commonly defined by the system's security policy and are set by the security or system administrator. These privileges may range from the extremes of "permit nothing" to "permit everything" and include anything in between.

exam watch

Authorization is the process of ensuring that a user has sufficient rights to perform the requested operation and preventing those without sufficient rights from doing the same. At the same time, authorization is also the process which gives rights depending on the identity of the user—be it a human or another system.

As you can see, the second and third stages of the identify-authenticate-authorize process depend on the first stage, and the final goal of the whole process is to enforce *access control* and *accountability,* which is described next. User account management and access control in Solaris 10 are described in more detail in Chapters 9 and 10.

Accountability

Accountability is another important principle of information security that refers to the possibility of tracing actions and events back in time to the users, systems, or processes that performed them, to establish responsibility for actions or omissions.

A system may not be considered secure if it does not provide accountability, because it would be impossible to ascertain who is responsible and what did or did not happen on the system without that safeguard. Accountability in the context of information systems is mainly provided by *logs* and the *audit trail*.

Logs

System and application logs are ordered lists of events and actions and are the primary means of establishing accountability on most systems. However, logs (as well as the audit trail, which is described next) may be considered trustworthy only if their *integrity* is reasonably assured. In other words, if anyone can write to and/or erase logs or the audit trail, they would not be considered dependable enough to serve as the basis for accountability. Additionally, in case of networked or communication systems, logs should be correctly timestamped and time should be synchronized across the network so events that affect more than one system may be correctly correlated and attributed.

Audit Trail

The difference between the audit trail and logs is not clearly defined. However, we may say that logs usually show high-level actions, such as an e-mail message delivered or a web page served, whereas audit trails usually refer to lower-level operations such as opening a file, writing to a file, or sending a packet across a network. While an audit trail provides more detailed information about the actions and events that took place on the system, it is not necessarily more *useful*, in a practical sense of the word, than logs, simply because abundance of detail in an audit trail makes it more resource and time consuming to generate, store, and analyze. Another aspect by which logs and audit trails differ is their source: logs are usually and mostly generated by particular system software or applications, and an audit trail is usually kept by the operating system or its auditing module. Auditing and audit analysis in Solaris 10 are covered in detail in Chapter 5.

Functionality vs. Assurance

Having introduced the concept of accountability and how it is implemented on most systems, it's time to look at perhaps one of the most challenging issues of information security: the issue of *functionality versus assurance*. The best way to illustrate this is to refer to your own first-hand experience with computers: how many times has a computer failed to do something that you expected of it, and how many times did it do something you didn't want it to do? It is this difference between our expectations (as well as vendors' advertising of product features) and what happens in fact that is referred to as functionality versus assurance.

A particular system may claim to implement a dozen smart security features, but this is very different from being able to say with a high degree of confidence that it indeed implements them, implements them correctly, and will not behave in an unexpected manner. Another way of looking at the functionality versus assurance issue is that functionality is about what a system *can do* and assurance is about what a system *will not do*.

Although no quick and easy solutions are available in this case, we will discuss functionality and assurance issues in more detail in Chapter 3 of this book with regard to standards, certification, and accreditation.

Privacy

Privacy in the information security context usually refers to the expectation and rights of individuals to privacy of their personal information and adequate, secure handling of this information by its users. *Personal information* here usually refers to information that directly identifies a human being, such as a name and address, although the details may differ in different countries.

In many countries, privacy of personal information is protected by laws that impose requirements on organizations processing personal data and set penalties for noncompliance. The European Union (EU) in particular has strict personal data protection legislation in place, which limits how organizations may process personal information and what they can do with it. The U.S. Constitution also guarantees certain privacy rights, although the approach to privacy issues differs between the United States and Europe.

Since privacy is not only a basic human need but also a legally protected right in most countries, organizations should take necessary precautions to protect the confidentiality and integrity of personal information they collect, store, and process. In particular, organizations' information security policies should define how personal information is to be collected and processed. Because of these requirements, although not in the C-I-A triad, privacy is also an inseparable part of information security and must be addressed in all information security policies as part of the information security requirements.

Non-repudiation

Non-repudiation in the information security context refers to one of the properties of cryptographic digital signatures that offers the possibility of proving whether a particular message has been digitally signed by the holder of a particular digital signature's private key. Non-repudiation is a somewhat controversial subject, partly because it is an important one in this day and age of electronic commerce, and because it does not provide an absolute guarantee: a digital signature owner, who may like to repudiate a transaction maliciously, may always claim that his or her digital signature key was stolen by someone and that someone actually signed the digital transaction in question, thus repudiating the transaction. The following types of non-repudiation services are defined in international standard ISO 14516:2002, *Guidelines for the use and management of trusted third party services*.

Approval Non-repudiation of approval provides proof of who is responsible for approval of the contents of a message.

Sending Non-repudiation of sending provides proof of who sent the message.

Origin Non-repudiation of origin is a combination of approval and sending.

Submission Non-repudiation of submission provides proof that a delivery agent has accepted the message for transmission.

Transport Non-repudiation of transport provides proof for the message originator that a delivery agent has delivered the message to the intended recipient.

Receipt Non-repudiation of receipt provides proof that the recipient received the message.

Knowledge Non-repudiation of knowledge provides proof that the recipient recognized the content of the received message.

Delivery Non-repudiation of delivery is a combination of receipt and knowledge, as it provides proof that the recipient received and recognized the content of the message.

There is also a difference between the legal concept of non-repudiation and non-repudiation as an information security/cryptographic concept. In the legal sense, an alleged signatory to a paper document is always able to repudiate a signature that has been attributed to him or her by claiming any one of the following:

- Signature is forged
- Signature is a result of fraud by a third party
- Signature was unconscionable conduct by a party to transaction
- Signature was obtained using undue influence by a third party

In the information security context, one should keep in mind that the cryptographic concept of non-repudiation may, and often does, differ from its legal counterpart. Moreover, in some countries there is a trend of moving the burden of proof from the party relying on the signature (which is applicable to regular on-paper signatures) to the alleged signatory party, who would have to prove that he or she *did not* sign something. Chapter 11 of this book looks at cryptography in more detail.

INSIDE THE EXAM

General Security Concepts

The Sun Certified Security Administrator for Solaris exam consists of 60 multiple-choice, drag-drop, and matching questions to be answered in 90 minutes. The passing score for the entire exam is 60 percent. Of these 60 questions, approximately ten questions cover "Section 1 – General Security Concepts" of the official exam objectives and include the following items:

1. Explain fundamental concepts concerning information security and explain what good security architectures include (people, process, technology, defense in depth).
2. Describe accountability, authentication, authorizations, privacy, confidentiality, integrity, and non-repudiation.
3. Identify the security life cycle (prevent, detect, react, and deter) and describe security awareness, security policies and procedures, physical security, platform security, network security, application security, and security operations and management.
4. Describe concepts of insecure systems, user trust, threat, and risk.
5. Explain attackers, motives, and methods.
6. Describe the benefit of evaluation standards and explain actions that can invalidate certification.
7. Describe how attackers gain information about the targets, and describe methods to reduce disclosure of revealing information.

The first two exam objectives in this list are covered in this chapter. The purpose of Section 1 of the exam objectives is to test your understanding of the general security concepts and principles. Unlike other sections of the exam, Section 1 tests your knowledge of abstract matters that are universally applicable and are not specific to the Solaris operating environment. To perform well on the exam, you must have a clear understanding of the material presented in this chapter; the self-test questions at the end of the chapter should help you to check and reinforce the most important concepts.

CERTIFICATION OBJECTIVE 1.02

Explain Information Security Fundamentals and Define Good Security Architectures

Now that we have armed ourselves with the fundamental concepts of information security, let's consider some of the universal security principles, such as principles of least privilege, minimization, and compartmentalization.

Least Privilege

The principle of *least privilege* stipulates, "Do not give any more privileges than absolutely necessary to do the required job." This principle applies not only to privileges of users and applications on a computer system, but also to other non-information systems privileges of an organization's staff. The principle of least privilege is a preventive control, because it reduces the number of privileges that may be potentially abused and therefore limits the potential damage. Like most good principles, the principle of least privilege is applicable in all information systems environments. Some examples of application of this principle include the following:

- Giving users only read access to shared files if that's what they need, and making sure write access is disabled
- Not allowing help desk staff to create or delete user accounts if all that they may have to do is to reset a password
- Not allowing software developers to move software from development servers to production servers

Defense in Depth

The principle of *defense in depth* is about having more than one layer or type of defense. The reasoning behind this principle is that any one layer or type of defense may be breached, no matter how strong and reliable you think it is, but two or more layers are much more difficult to breach. Defense in depth works best when you combine two or more different types of defense mechanisms—such as using a firewall between the Internet and your LAN, plus the IP Security Architecture (IPSEC) to encrypt all sensitive traffic on the LAN. In this scenario, even if your firewall is compromised, the attackers still have to break IP Security to get to your data flowing across the LAN.

Generally, different types of controls should be used together: first, preventive controls should be in place to try and prevent security incidents from happening at all; second, detective controls are necessary so that you can know whether preventive controls are working or have failed; and third, corrective controls are needed to help you respond effectively to security incidents and contain damage. However, the defense in depth principle does not mean that you should indiscriminately apply all the controls and security measures you can get your hands on: balance has to be found between security provided by the defense in depth approach and the financial, human, and organizational resources you are willing to expend following it. This balance is addressed by the cost-benefit analysis, introduced later on in this chapter.

Minimization

The *minimization* principle is the cousin of the least privilege principle and mostly applies to system configuration. The minimization principle says "do not run any software, applications, or services that are not strictly required to do the entrusted job." To illustrate, a computer whose only function is to serve as an e-mail server should have only e-mail server software installed and enabled. All other services and protocols should either be disabled or not installed at all to eliminate any possibility of compromise or misuse. Adherence to the minimization principle not only increases security but usually also improves performance, saves storage space, and is a good system administration practice in general.

Cost-Benefit Analysis

Although not strictly a principle, the *cost-benefit analysis* is a must when considering implementation of any security measure. It says that the overall benefits received from a particular security control or mechanism should clearly exceed its total costs; otherwise, implementing it would make no sense. Cost-benefit analysis directly affects return on investment (ROI). This may sound like simple common sense, and it probably is; nevertheless, this is an important and often overlooked concern. When doing cost-benefit analysis, one should consider all costs and all benefits over a period of time, for example from one to five years, to have a complete picture.

Risk-Control Adequacy

We will discuss risk analysis and management in more detail in Chapter 2. For now, suffice to say that controls should match the risks they are expected to control and should not be implemented just for the sake of having them. This is yet another common-sense principle that is often neglected.

Compartmentalization

Compartmentalization, or the use of compartments (also known as zones, jails, sandboxes, and virtual areas), is a principle that limits the damage and protects other compartments when software in one compartment is malfunctioning or compromised. It can be best compared to compartments on ships and submarines, where a disaster in one compartment does not necessarily mean that the entire ship or submarine is lost. Compartmentalization in the information security context means that applications run in different compartments are isolated from each other. In such a setup, the compromise of web server software, for example, does not take down or affect e-mail server software running on the same system but in a separate compartment. Zones in Solaris 10 implement the compartmentalization principle and are powerful security mechanisms.

Keep Things Simple

Complexity is the worst enemy of security. Complex systems are inherently more insecure because they are difficult to design, implement, test, and secure. The more complex a system, the less assurance we may have that it will function as expected. Although complexity of information systems and processes is bound to increase with our increasing expectations of functionality, we should be very careful to draw a line between avoidable and unavoidable complexity and not sacrifice security for bells and whistles, only to regret it later. When you have to choose between a complex system that does much and a simple system that does a bit less but enough, choose the simple one.

Fail Securely

Although *fail securely* may sound like an oxymoron, it isn't. Failing securely means that if a security measure or control has failed for whatever reason, the system is not rendered to an insecure state. For example, when a firewall fails, it should default to a "deny all" rule, not a "permit all." However, *fail securely* does not mean "close everything" in all cases; if we are talking about a computer-controlled building access control system, for example, in case of a fire the system should default to "open doors" if humans are trapped in the building. In this case, human life takes priority over the risk of unauthorized access, which may be dealt with using some other form of control that does not endanger the lives of people during emergency situations.

Secure the Weakest Link

To people new to information security, many information security principles and approaches may sound like little more than common sense. Although that may well be the case, it doesn't help us much, because very often we still fail to act with common sense. The principle of securing the weakest link is one such case: look around and you will likely see a situation in which instead of securing the weakest link, whatever it may be, resources are spent on reinforcing already adequate defenses.

Use Choke Points

Security is very much about control, and control is so much more effective and efficient when you know all ways in and out of your systems or networks. *Choke points* are logical "narrow channels" that can be easily monitored and controlled. An example of a choke point is a firewall—unless traffic can travel only via the firewall, the firewall's utility is reduced to zero. Consider the example of controlled entrances to buildings or facilities of high importance, such as perimeter fencing and guard posts.

Leverage Unpredictability

Just as states don't publicize the specifics of their armaments, exact locations, or numbers of armed forces, you should not publicize the details of your security measures and defenses. This principle should not be seen as contradicting deterrent security controls—controls that basically notify everyone that security mechanisms are in place and that violations will be resisted, detected, and acted upon. The important difference here is that deterrent controls don't provide details of the defenses but merely announce their existence so as to deter potential attackers without giving them detailed information that later may be used against the defenders. In practical terms, this means you can, for example, announce that you are using a firewall that, in particular, logs all traffic to and from your network, and these logs are reviewed by the organization—there is no need to disclose the type, vendor, or version number of the firewall; where it is located; how often logs are reviewed; and whether any backup firewalls or network intrusion detection systems are in place.

Segregation of Duties

The purpose of the segregation (or separation) of duties is to avoid the possibility of a single person being responsible for different functions within an organization, which when combined may result in a security violation that may go undetected. Segregation of duties can prevent or discourage security violations and should be practiced when possible. Although the actual job titles and organizational hierarchies may differ greatly, the idea behind the principle of separation of duties stays the same: no single person should be able to violate security and get away with it. Rotation of duties is a similar control that is intended to detect abuse of privileges or fraud and is a practice to help your organization avoid becoming overly dependent on a single member of the staff. By rotating staff, the organization has more chances of discovering violations or fraud.

Types of Controls

Central to information security is the concept of controls, which may be categorized by their *functionality* (preventive, detective, corrective, deterrent, recovery, and compensating, in this order) and *plane of application* (physical, administrative, or technical). Physical controls include doors, secure facilities, fire extinguishers, flood protection, and air conditioning. Administrative controls are the organization's policies, procedures, and guidelines intended to facilitate information security. Technical controls are the various technical measures, such as firewalls, authentication systems, intrusion detection systems, and file encryption, among others.

Preventive Controls

Preventive controls are the first controls met by the adversary. Preventive controls try to prevent security violations and enforce access control. Like other controls, preventive controls may be physical, administrative, or technical: doors, security procedures, and authentication requirements are examples of physical, administrative, and technical preventive controls, respectively.

Detective Controls

Detective controls are in place to detect security violations and alert the defenders. They come into play when preventive controls have failed or have been circumvented and are no less crucial than detective controls. Detective controls include cryptographic checksums, file integrity checkers, audit trails and logs, and similar mechanisms.

Corrective Controls

Corrective controls try to correct the situation after a security violation has occurred. Although a violation occurred, not all is lost, so it makes sense to try and fix the situation. Corrective controls vary widely, depending on the area being targeted, and they may be technical or administrative in nature.

Deterrent Controls

Deterrent controls are intended to discourage potential attackers and send the message that it is better not to attack, but even if you decide to attack we are able to defend ourselves. Examples of deterrent controls include notices of monitoring and logging as well as the visible practice of sound information security management.

Recovery Controls

Recovery controls are somewhat like corrective controls, but they are applied in more serious situations to recover from security violations and restore information and information processing resources. Recovery controls may include disaster recovery and business continuity mechanisms, backup systems and data, emergency key management arrangements, and similar controls.

Compensating Controls

Compensating controls are intended to be alternative arrangements for other controls when the original controls have failed or cannot be used. When a second set of controls addresses the same threats that are addressed by another set of controls, the second set of controls are compensating controls.

Access Control Models

Logical *access control models* are the abstract foundations upon which actual access control mechanisms and systems are built. Access control is among the most important concepts in computer security. Access control models define how computers enforce access of subjects (such as users, other computers, applications, and so on) to objects (such as computers, files, directories, applications, servers, and devices). Three main access control models exist: the discretionary access control model, the mandatory access control model, and the role-based access control model.

Discretionary Access Control (DAC)

The discretionary access control model is the most widely used of the three models. In the DAC model, the owner (creator) of information (file or directory) has the discretion to decide about and set access control restrictions on the object in question—which may, for example, be a file or a directory. The advantage of DAC is its flexibility: users may decide who can access information and what they can do with it—read, write, delete, rename, execute, and so on. At the same time, this flexibility is also a disadvantage of DAC because users may make wrong decisions regarding access control restrictions or maliciously set insecure or inappropriate permissions. Nevertheless, the DAC model remains the model of choice for the absolute majority of operating systems today, including Solaris.

Mandatory Access Control (MAC)

Mandatory access control, as its name suggests, takes a stricter approach to access control. In systems utilizing MAC, users have little or no discretion as to what access permissions they can set on their information. Instead, mandatory access controls specified in a system-wide security policy are enforced by the operating system and applied to all operations on that system. MAC-based systems use data classification levels (such as public, confidential, secret, and top secret) and security clearance labels corresponding to data classification levels to decide, in accordance with the security policy set by the system administrator, what access control restrictions to enforce. Additionally, per-group and/or per-domain access control restrictions may be imposed—that is, in addition to having the required security clearance level, subjects (users or applications) must also belong to the appropriate group or domain. For example, a file with a confidential label belonging only to the research group may not be accessed by a user from the marketing group, even if that user has a security clearance level higher than confidential (for example, secret or top secret). This concept is known as *compartmentalization* or *need to know*.

Although MAC-based systems, when used appropriately, are thought to be more secure than DAC-based systems, they are also much more difficult to use and administer because of the additional restrictions and limitations imposed by the operating system. MAC-based systems are typically used in government, military, and financial environments, where higher than usual security is required and where the added complexity and costs are tolerated. MAC is implemented in Trusted Solaris, a version of the Solaris operating environment intended for high-security environments.

Role-Based Access Control (RBAC)

In the role-based access control model, rights and permissions are assigned to roles instead of individual users. This added layer of abstraction permits easier and more flexible administration and enforcement of access controls. For example, access to marketing files may be restricted to the marketing manager role only, and users Ann, David, and Joe may be assigned the role of marketing manager. Later, when David moves from the marketing department elsewhere, it is enough to revoke his role of marketing manager; no other changes would be necessary. When you apply this approach to an organization with thousands of employees and hundreds of roles, you can see the added security and convenience of using RBAC. Solaris has supported RBAC since release 8.

Centralized vs. Decentralized Access Control

Further distinction should be made between centralized and decentralized (distributed) access control models. In environments with centralized access control, a single, central entity makes access control decisions and manages the access control system; whereas in distributed access control environments, these decisions are made and enforced in a decentralized manner. Both approaches have their pros and cons, and it is generally inappropriate to say that one is better than the other. The selection of a particular access control approach should be made only after careful consideration of an organization's requirements and associated risks.

Information Security Architectures

In the rest of this chapter we will discuss information security architectures and best practices. You will see that information security is not only a technological challenge but a human challenge as well and needs human solutions first and foremost. We will also try to identify the most common shortcomings and pitfalls that result in inefficient or insufficient information security and see what can be done to minimize their impact and the rate of occurrence.

We begin with an overview of information systems governance and how good governance practices improve information systems security during planning, design, implementation, use, and maintenance stages of information systems.

Information Systems Governance

Information security is a part of *information systems governance*. With our exponentially increasing dependence on information systems in all areas of human life, information systems affect nearly everyone. How we build and use information systems affects our national security, our competitiveness, and our economy. Information systems governance is the foundation that determines whether these information systems are aligned with our objectives and serve our needs. Therefore, good information systems governance practices are vital for every organization, and these practices or their absence directly affect security of information systems.

Information systems governance is mainly concerned with two responsibilities: making the most of available information systems resources and at the same time managing risks associated with use of these information systems. Ultimately, information systems governance is the responsibility of an organization's highest governing bodies, which are usually the board of directors or trustees. The board of directors is responsible for understanding the role and impact of information systems on the organization, defining high-level policies, measuring performance of information systems, and delegating management of information systems risks. Regretfully, information security is often mistakenly viewed as a technology issue only, with little consideration given to business priorities and requirements, although research indicates that it is first and foremost a business issue. Responsibility for governing and managing the improvement of information security has been mostly limited to technical management and IT staff. However, for information security to be managed properly, understanding and involvement of the board of directors and executive management is necessary. In particular, the board of directors should do the following:

- Be informed about information security
- Set policy and strategy
- Provide resources for information security
- Assign responsibilities to management and set priorities

Executive management, in turn, should assume responsibility for the following aspects of information systems governance and be proactive in their management:

- Setting information security policy
- Assigning responsibilities to staff
- Assessing, analyzing, and managing risks associated with information systems
- Defining the information security management framework
- Implementing security awareness training of all staff

To assess information systems governance effectively in an organization, an information security governance maturity model defined by the Information

Technology Governance Institute (ITGI) as part of the Control Objectives for Information Technology (COBIT) framework may be used. COBIT defines the following as information systems governance criteria:

- Effectiveness
- Efficiency
- Confidentiality
- Integrity
- Availability
- Compliance
- Reliability

The information technology includes the people, technology, applications, facilities, and data of the organization. This model defines six levels of information systems and security governance, ranging from nonexistent to optimized. Although surveys show that very few organizations find themselves at these polar levels, the absolute majority of organizations have a way to go to reach higher levels of information systems governance.

The six levels of the information systems governance maturity model are as follows:

- **Level 0** Nonexistent
- **Level 1** Initial
- **Level 2** Repeatable
- **Level 3** Defined
- **Level 4** Managed
- **Level 5** Optimized

Nonexistent Information systems governance and security management processes are nonexistent. The board of directors and management are breaching their fiduciary duties to the organization's stockholders by not fulfilling their responsibilities with regard to the governance and protection of the organization's information systems and assets.

Initial Although existing, information systems governance and security management are mostly ad hoc and not well organized. Policies are not well founded. Standards and procedures are either unknown or are not being applied in an appropriate manner. Governance and management may be said to be in their embryonic form.

Repeatable At this level, although it is understood that information systems governance and security management are important to the organization, the actual implementation of information systems governance practices is not adequate. Security awareness is fragmented, and information security management is not effective. This is the level of information systems governance and security management at most organizations.

Defined Information systems governance processes are documented and they are communicated to all relevant parties. Risk-management policy exists and it is communicated to responsible staff. Information security policies exist and they are communicated to staff. Security awareness is facilitated, and information security responsibilities are defined and delegated.

Managed Information systems governance processes are not only documented and communicated to all parties but are also monitored to ensure correct application and enforcement. Responsibilities for information security are clearly assigned, managed, and enforced. Information security risk analysis is consistently performed. Security policies and practices are complete with specific security baselines. Standards are used. Security awareness training is mandatory. Security certification of staff is encouraged and recognized.

Optimized Best practices are followed in all aspects of information systems governance and information security management. Risk assessment has developed to the stage at which a structured process is enforced and well-managed. Information security is a joint responsibility of business and information systems management and is aligned with business objectives. Security functions are integrated in applications at the design stage, and end users are increasingly accountable for managing security. The existing processes provide for continuing self-improvement. Third-party certification or accreditation attests to the deployment of best practices in information systems governance and security.

As you can see, we have a way to go before we reach the more advanced levels of this model, but every step toward more effective and efficient information systems governance is a worthwhile investment. We will cover fundamentals of information security management in Chapter 3.

The Weakest Link

It is a widely held opinion among information security professionals that we, the humans, are the weakest link in the chain of information security. Nevertheless, we continue to suffer from our own errors and omissions, despite increasing information security risks. Until and unless we realize and accept this reality, it would be very difficult to rectify this unfortunate situation. Everyone involved with information systems—including their owners, managers, designers, administrators, and of course users—needs to do their part in being a responsible citizen of the information society. Owners and managers should accept responsibility for the overall management and governance of information systems; designers and administrators should adhere to security principles and best practices; and users should not compromise security for minute convenience.

The Information Security Process

Bruce Schneier, one of the world's most well-known experts on security, once wrote that "security is a process, not a product." Indeed, with all the changing variables and players, security is a never-ending evolutionary process, wherein defenses change in response to new threats and new threats emerge with the introduction of new systems and defenses. This view only reinforces something that has been known for a while—that you cannot buy security off the shelf. In the case of Solaris, the operating environment provides a good set of proven security features and technologies, but they are useless if not correctly used and administered as part of an encompassing security process that tries to prevent security violations, detect them when they occur, apply corrective controls, recover from the incident, and improve itself. The strength of protection against information security risks lies in the way this process is practiced: whether sufficient resources are allocated, whether qualified professionals do their jobs well, and whether management actually cares about the whole issue.

Know Your Enemy

Earlier in this chapter, we mentioned cost-benefit analysis and risk-control adequacy. It is worth revisiting this topic from a slightly more specific angle, that of knowing your enemy. No one would argue that before building defenses and getting ready for the battle, one should have at least a general idea of who may attack and, therefore, what can and cannot be done to prevent or fight against the attack. Meaningful cost-benefit and risk-control analyses are impossible or would yield incorrect results if one doesn't know his enemies and is unaware of their means and intentions. An adversary's profile determines to a great extent the type and quantity of controls that should be applied and therefore affects the cost-benefit equation in a nonlinear way. We will be discussing the attackers and their motives in more detail in the next chapter; for the purposes of this chapter, let's note that before erecting defenses, you should understand from whom you are trying to protect your organization's information assets and resources.

CERTIFICATION SUMMARY

In this chapter, we explained the fundamental information security concepts and principles, looked at what constitutes good security architectures and practices, and learned that good practices include people, processes, and technology working in concert. We also discussed the concepts of accountability, authentication, authorization, privacy, confidentiality, integrity, and non-repudiation, as well as types and functionalities of information security controls and the importance of information systems governance.

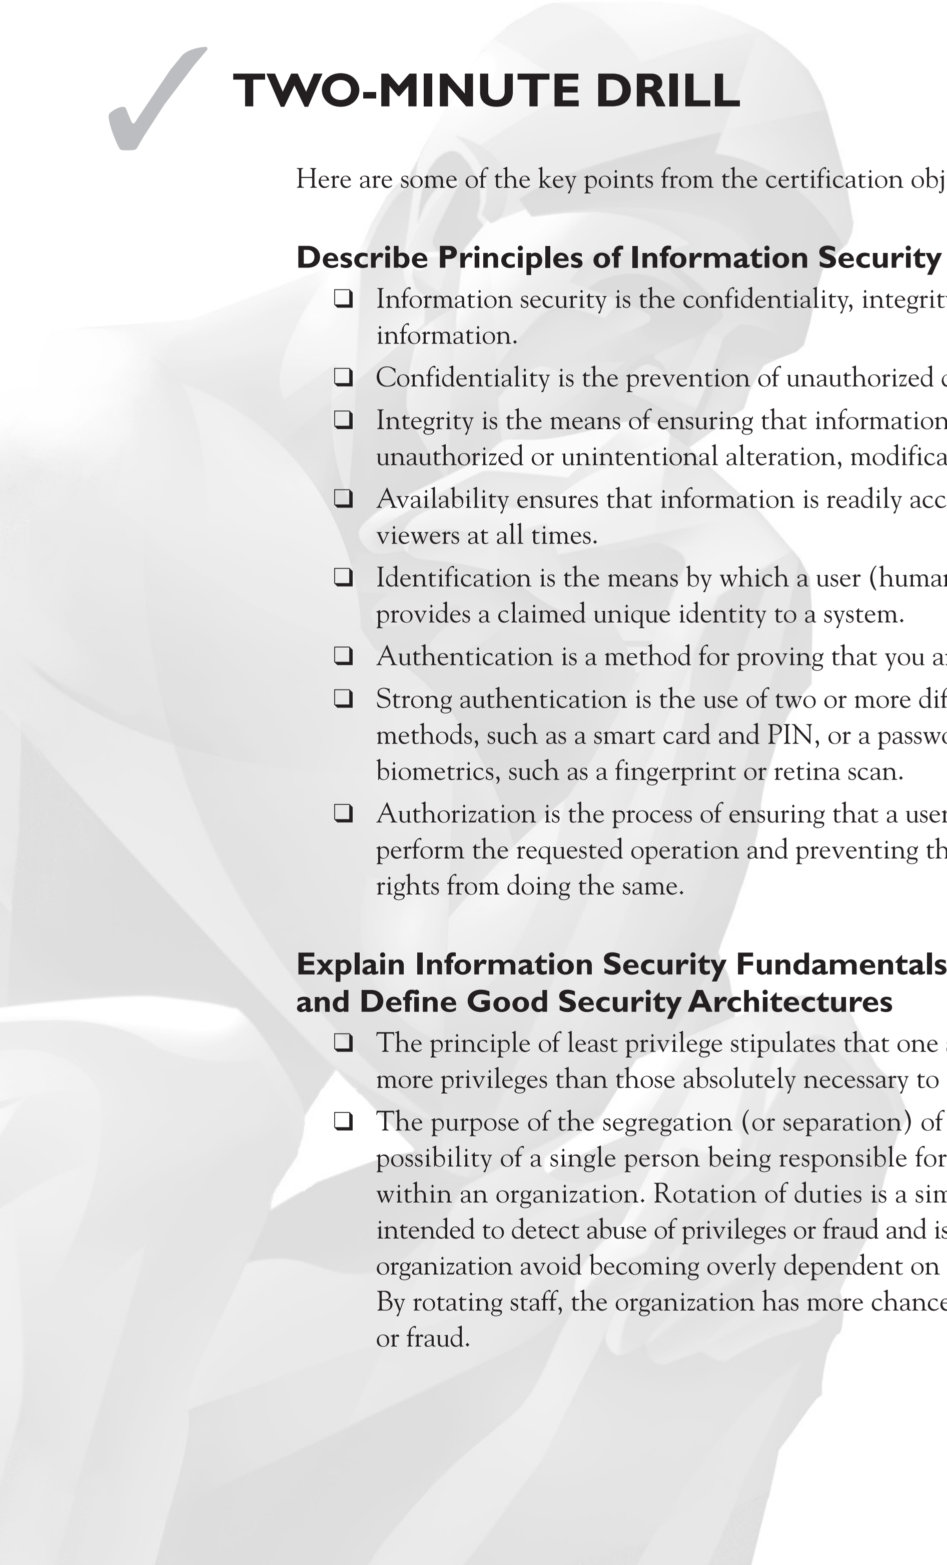

✓ TWO-MINUTE DRILL

Here are some of the key points from the certification objectives in Chapter 1.

Describe Principles of Information Security

- □ Information security is the confidentiality, integrity, and availability of information.
- □ Confidentiality is the prevention of unauthorized disclosure of information.
- □ Integrity is the means of ensuring that information is protected from unauthorized or unintentional alteration, modification, or deletion.
- □ Availability ensures that information is readily accessible to authorized viewers at all times.
- □ Identification is the means by which a user (human, system, or process) provides a claimed unique identity to a system.
- □ Authentication is a method for proving that you are who you say you are.
- □ Strong authentication is the use of two or more different authentication methods, such as a smart card and PIN, or a password and a form of biometrics, such as a fingerprint or retina scan.
- □ Authorization is the process of ensuring that a user has sufficient rights to perform the requested operation and preventing those without sufficient rights from doing the same.

Explain Information Security Fundamentals and Define Good Security Architectures

- □ The principle of least privilege stipulates that one should not be assigned any more privileges than those absolutely necessary to do the required job.
- □ The purpose of the segregation (or separation) of duties is to avoid the possibility of a single person being responsible for a variety of functions within an organization. Rotation of duties is a similar control that is intended to detect abuse of privileges or fraud and is a practice that helps the organization avoid becoming overly dependent on a single member of staff. By rotating staff, the organization has more chances of discovering violations or fraud.

SELF TEST

The following questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully because there might be more than one correct answer. Choose all correct answers for each question.

Describe Principles of Information Security

1. What is the purpose of audit trails and logs?

   A. They record events as they happen.

   B. An audit trail can be used in court proceedings but logs cannot.

   C. They serve to establish accountability.

   D. They may be used in place of deterrent controls.

   E. All of the above

2. Fingerprints can be used for

   A. *What you have* authentication

   B. *What you are* authentication

   C. Biological identification

   D. Keeping things simple

   E. All of the above

3. What type of control is intended to offset deficiencies of other controls?

   A. Preventive

   B. Defensive

   C. Compensating

   D. Recovery

   E. All of the above

4. What is strong authentication?

   A. Strong authentication uses long passwords.

   B. Strong authentication requires smart cards.

   C. Strong authentication requires the use of at least two different authentication methods.

   D. Strong authentication is provided via biometrics.

   E. All of the above

5. The principle of least privilege applies only to user accounts.

   A. True

   B. False

   C. True, but only on non-Solaris systems.

   D. True, provided users use good passwords.

6. The principle of isolating process spaces from each other is known as

   A. Virtualization

   B. Separation

   C. Defense in depth

   D. Compartmentalization

   E. All of the above

7. Surveys show that most organizations are at which level of the information security maturity model?

   A. Nonexistent

   B. Defined

   C. Detective

   D. Repeatable

   E. All of the above

8. Privacy is a concern in which of the following industries?

   A. Financial services

   B. Financial services and government

   C. Telecommunications

   D. All of the above

9. What is assurance?

   A. It is a type of insurance against security violations.

   B. It is the written security policy.

   C. It is about the trustworthiness of a system.

   D. It is provided by the mandatory access control (MAC).

   E. All of the above

10. Information security policies and procedures are a(n)

A. Technical control

B. Administrative control

C. Form of access control

D. Operational control

E. All of the above

11. In information security context, names must be

A. Unique locally

B. Unique globally

C. Standardized

D. Secret

E. All of the above

12. What risks apply to *what you have* authentication methods? (Choose all that apply.)

A. Same risks as with *what you are* authentication

B. Same risks that apply to regular keys

C. Risks that apply to all authentication methods

D. Certain non-assurance–related risks

E. All of the above

Explain Information Security Fundamentals and Define Good Security Architectures

13. Who must be ultimately responsible for information security within organizations?

A. Information security professionals

B. Information systems auditors

C. Top management

D. Stockholders

E. All of the above

14. Fundamental security principles

A. Do not apply in all situations

B. Apply to most information systems

C. May be used only in enterprise systems

D. Are system-dependent

E. All of the above

15. Information systems governance is about what?

A. Information security

B. Effective and risk-aware use of information systems

C. Risk management

D. Corporate responsibility

E. All of the above

16. What is the advantage of Role-Based Access Control (RBAC) over Discretionary Access Control (DAC)?

A. RBAC has no advantages over DAC.

B. RBAC is an improved version of DAC.

C. RBAC improves management of access control and authorizations.

D. RBAC is one level below Mandatory Access Control (MAC).

E. All of the above

17. Which authentication method is the most complex to administer?

A. *What you know*

B. *What you have*

C. *What you are*

D. *Who you are*

E. All of the above

18. What is the purpose of choke points?

A. Choke points are used to isolate firewalls.

B. Choke points protect confidentiality of information.

C. Choke points may be used only on TCP/IP networks.

D. Choke points are for control and monitoring of data flow.

E. All of the above

19. What is the purpose of authentication?

A. To obtain proof of claimed identity

B. To implement access control

C. To establish accountability

D. To allow use of different authorizations

E. All of the above

20. What is the benefit of cost-benefit analysis? (Choose all that apply.)

A. It is necessary because organizations cannot reduce all risks to zero.

B. It increases an organization's return on investment.

C. It prevents denial of service attacks.

D. It is a good governance practice.

E. All of the above

SELF TEST ANSWERS

Describe Principles of Information Security

1. ☑ **C.** The purpose of the audit trail and logs is to provide accountability in information systems.
☒ **A** is correct but is not the best answer; choices **B** and **D** are wrong. The issue of whether audit trails and logs can be used in court proceedings would depend on particular jurisdiction and is outside the scope of this book; audit trails and logs are detective controls but may function as deterrent controls as well when their existence is known to potential attackers.

2. ☑ **B.** Fingerprints can be used for *what you are*, or biometric, authentication.
☒ **A** is wrong because *what you have* authentication refers to token-based authentication mechanisms. **C** is wrong because there is no such term as biological identification in information security. **D** is wrong because use of fingerprints does not simplify authentication or identification since this requires additional configuration and tuning.

3. ☑ **C.** Compensating controls offset deficiencies of other controls.
☒ There is no such term as defensive controls in information security, so that rules out **B**. Choices **A** and **D** are incorrect because preventive controls aim to prevent security violations and recovery controls are not intended to offset deficiencies of other controls.

4. ☑ **C.** At least two different authentication methods are necessary for strong authentication.
☒ Long passwords do not provide strong authentication on their own, so answer **A** is not correct. Strong authentication does not necessarily require use of smart cards, as stated in **B**. And **C** is wrong because biometrics does not necessarily provide strong authentication on its own.

5. ☑ **B.** The principle of least privilege does not only apply to user accounts but is a universally applicable principle.
☒ The answers are incorrect because the principle of least privilege has no relation to use of good passwords and is not dependent on a particular operating system or environment.

6. ☑ **D.** Compartmentalization is the isolation of process spaces from each other in order to minimize the effect of security violation in one compartment on another.
☒ Answer **A,** virtualization, is a related concept but is not the correct answer. **B** is wrong because compartmentalization is the correct term. **C** is wrong because defense in depth is about using several types and/or layers of defense.

7. ☑ **D.** Most organizations are at the repeatable level of the information security maturity model.
☒ **C** is inappropriate because it refers to a type of control. Other choices are wrong because surveys show that most organizations are at the repeatable level.

8. ☑ **D.** All of the above. Privacy is a concern in all industries, because organizations in all industries collect, process, and store personal information of employees, clients, and partners.

9. ☑ **C.** Assurance is about the trustworthiness of a system.
☒ **A** is wrong because there is no such type of insurance. **B** is wrong because, although written security policy is always required, it is not a guarantee of assurance. **D** is wrong because the use of MAC does not guarantee assurance.

10. ☑ **B.** Information security policies and procedures are an administrative control.
☒ **A** is wrong because policies and procedures are not a technical control. **C** is wrong because policies and procedures are not a form of access control. **D** is wrong because, although policies and procedures address operational controls, choice **B** is a better answer.

11. ☑ **A.** Names *must* be unique locally.
☒ **B** is wrong because names *may* be unique globally, but it's not necessary. **C** is wrong because names may be standardized, but that is not mandatory. **D** is wrong because names are not necessarily secret.

12. ☑ **B** and **C** are correct because *what you have* authentication methods are subject to the same risks (such as theft and damage) as regular keys, and they are subject to the same general risks that apply to all authentication methods (such as unauthorized access).
☒ **A** is wrong because risks of *what you are* and *what you have* authentication methods are different, and **D** is wrong because it doesn't make sense.

Explain Information Security Fundamentals and Define Good Security Architectures

13. ☑ **C.** Top management must be ultimately responsible for information security within an organization.
☒ **A** is incorrect because information security professionals advise management and implement management's decisions. **B** is wrong because information systems auditors report on the organization's security to the board of directors and/or the stockholders. **D** is incorrect because stockholders appoint management and are not involved in day-to-day management.

14. ☑ **B.** Fundamental security principles apply to most information systems.
☒ **A** is wrong because it is not the best available answer. **C** is wrong because fundamental security principles do not apply only in enterprise systems, and **D** is wrong because fundamental security principles are not system dependent.

15. ☑ **E.** All of the answers are correct.

16. ☑ **C.** RBAC improves management of access control and authorizations by introducing the concept of roles distinct from individual users.
☒ **A** is wrong because RBAC has advantages over DAC; **B** is wrong because RBAC is not an improved version of DAC; **D** is wrong because it doesn't make sense.

17. ☑ **C.** *What you are* (biometrics) is inherently more complex to administer than *what you have* or *what you know* authentication methods.
☒ **A, B,** and **D** are incorrect because none of these methods is as difficult to administer as *what you are*.

18. ☑ **D.** Choke points are logical "narrow channels" that can be easily monitored and controlled.
☒ **A** is wrong because choke points are not used to isolate firewalls. Choke points do not affect confidentiality of information, so **B** is wrong. And **C** is not the answer because choke points are not protocol-dependent.

19. ☑ **E.** All of the above. Authentication is needed to obtain proof of claimed identity, to implement access control, to establish accountability, and to allow for different users with different authorizations.

20. ☑ **A, B,** and **D.** Cost-benefit analysis is necessary because organizations cannot reduce all risks to zero, it increases an organization's return on investment, and it is a good governance practice.
☒ **C** is wrong because cost-benefit analysis is not related to, and does not prevent, denial of service attacks.