




24

THE SECURITY PERSPECTIVE

Desired Quality	The ability of the system to reliably control, monitor, and audit who can perform what actions on these resources and the ability to detect and recover from failures in security mechanisms
Applicability	Any systems with publicly accessible interfaces, with multiple users where the identity of the user is significant, or where access to operations or information needs to be controlled
Concerns	Policies, threats, mechanisms, accountability, availability, and detection and recovery
Activities	Identify sensitive resources, define the security policy, identify threats to the system, design the security implementation, and assess the security risks
Architectural Tactics	Apply recognized security principles, authenticate the principals, authorize access, ensure information secrecy, ensure information integrity, ensure accountability, protect availability, integrate security technologies, provide security administration, and use third-party security infrastructure
Problems and Pitfalls	Complex security policies, unproven security technologies, system not designed for failure, lack of administration facilities, technology-driven approach, failure to consider time sources, overreliance on technology, no clear requirements or models, security as an afterthought, security embedded in the application code, piecemeal security, and ad hoc security technology

Many factors drive today's need for information systems security, including the increasing trend to distribute systems, the use of public networks (particularly the Internet) as part of system infrastructure, the rising interest in interorganizational computing (such as that envisaged by Web services), and





other less technical reasons such as the increasing interest the media and the public have shown in computer security. All of these factors point to the fact that today your system's stakeholders are likely to be more interested in the security of the system than they would have been only a couple of years ago.

We define *security* as the set of processes and technologies that allow the owners of resources in the system to reliably control who can perform what actions on particular resources. The *who* refers to the people, pieces of software, and so on that form the set of actors in the system who have a security identity; security specialists normally call such actors *principals*. The *resources* are the parts of the system considered sensitive (i.e., those to which access must be controlled) such as data elements and operations. The *actions* are the operations that the principals in the system will want to perform on the resources (e.g., read them, change them, execute them, and so on), as shown in Figure 24-1.

The resources, principals, and actions that need to be considered are often very specific to the system. An Internet service provider is likely to have a totally different set of security concerns compared with those of a military intelligence organization, which will be different again from an enterprise implementing an internal information system that allows dialup access to its employees. However, in all these cases, security is still the business of allowing the right levels of access to the right resources to the right people.

It is also important to recognize that security is not a simple process of "being secure" or not. Rather than a binary state, security is really a process of risk management that balances likely security risks against the costs of guarding against them. Bear this in mind to help you set realistic assumptions in the minds of your stakeholders and to make intelligent tradeoffs that address the real security risks your system faces.

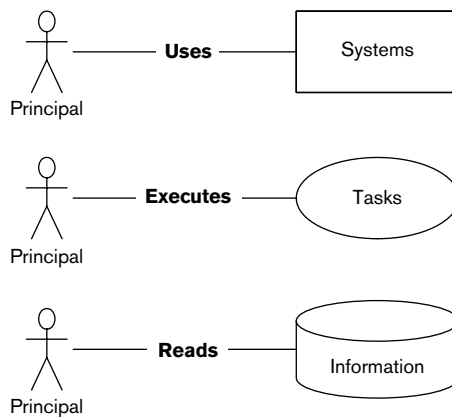


FIGURE 24-1 PRINCIPALS, ACTIONS, AND RESOURCES



**TABLE 24-1** APPLICABILITY OF THE SECURITY PERSPECTIVE TO THE SIX VIEWS

View	Applicability
Functional	The Functional view allows you to clearly see which of the system's functional elements need to be protected. Conversely, the functional structure of the system may be impacted by the need to implement your security policies.
Information	The Information view also helps you see what needs to be protected—in this case, the sensitive data in the system. Information models are often modified as a result of security design (e.g., partitioning information by sensitivity).
Concurrency	The Concurrency view defines how functional elements are packaged into runtime elements like processes. Security design may indicate the need to isolate different pieces of the system into different runtime elements, and if so, this will affect the system's concurrency structure.
Development	You may identify guidelines or constraints that the software developers will need to be aware of in order to ensure that the security policy is enforced. You need to include these guidelines or constraints in (or reference them from) the Development view.
Deployment	The security design may have a major impact on the system's deployment environment. For example, you may need security-oriented hardware or software, or you may need to change previously assumed deployment arrangements in order to address security risks.
Operational	Enforcing security policy is not just a matter of adding advanced technological features to a system. How the system is operated once it is in production will have a major affect on its security. The Operational view needs to make the security assumptions and responsibilities extremely clear, so that these aspects of the security implementation can be reflected in operational processes.


APPLICABILITY TO VIEWS

Table 24-1 shows how the Security perspective affects each of the views we discussed in Part III.

CONCERNS

Policies

The security policy for a system defines the system's security needs. It defines the controls and guarantees that the system requires for its resources and identifies which principals (or groups of principals) should be granted which types of access to each resource (or type of resource) within the system. A security





policy can be considered the security specification for a system, as it defines the set of security-related constraints that the system should be able to enforce.

A typical security policy defines the information access policy in terms of the different types of principals the system contains (e.g., clerks, managers, administrators) and, for each type of information (e.g., payroll records, customer details, invoice information), what sort of access each principal group requires (e.g., whether they can view the information, change it, delete it, share it). The policy also should define information integrity constraints that must be enforced, such as the integrity rules and checking required in data stores and protection of messages from unauthorized changes.

Threats

While the policy defines the security constraints the system requires, the threats the system faces are the possible ways the security constraints might be breached by an attacker who wishes to avoid them. Explicitly considering the security threats the system faces allows you to identify security enforcement mechanisms that can then counter these threats. Adding security mechanisms to a system is never free and often adds significantly to the complexity and cost of the system, while often reducing its usability and making it more difficult to operate in production. Therefore, it is important to select mechanisms according to the threats faced, to ensure that all of the chosen mechanisms are applied to counter realistic, credible security threats to the system.

Common threats to information systems include password cracking, network attacks that exploit software or configuration vulnerabilities, and denial-of-service attacks as well as nontechnical, social-engineering attacks that try to trick authorized users of the system to perform operations on behalf of the attacker.

Mechanisms

The security mechanisms in a system are the set of technologies, configuration settings, and procedures required to enforce the rules established by the security policy. Information systems security is a relatively mature field, and many proven security technologies already exist to act as mechanisms in an information system. Examples of commonly used security technologies include user name and password authentication, single-sign-on systems, virtual private networks to secure network links, database access control systems, and SSL/TLS encryption for client/server connections.

Your challenge as an architect is to select the right set of technologies from the wide array available and to apply them appropriately for the particular system you are building. In particular, it can be difficult to safely combine a number of these different technologies, along with system-wide security





configuration settings and a set of secure operational processes, to create a truly secure end-to-end system.

Accountability

Accountability is the means of ensuring that every action can be unambiguously traced back to the principal who performed it.

In centralized server-based systems, the notion of auditing is a common mechanism used to ensure accountability. In distributed systems, cryptographic message signing is typically used to prove that a message originated from a particular principal (a form of accountability also known as nonrepudiation).

Availability

Availability is often thought of as a purely operational consideration, but in fact it also has a relationship to security. Ensuring that potential attackers of your system cannot block its availability with denial-of-service attacks is an important part of your security design.



Designing a system for availability involves thinking about planned and unplanned outages that can occur due to the nature of the system itself; designing a system to avoid outages as a result of an external attack involves thinking about risks that can occur as a result of the environment in which the system resides (such as who can access the system and how). Availability security is an area that has rapidly gained importance as systems have started to be routinely connected to public networks like the Internet.

Detection and Recovery

System security is never perfect, and if a system is attractive enough to an attacker, it is almost inevitable that some sort of security breach will occur eventually. This means that an important security concern is the ability of the system to detect security breaches and recover from them. Addressing this concern is unlikely to be a purely technical matter but is more likely to involve people and processes as well as the technology required to spot a security breach and react to it appropriately.

ACTIVITIES: APPLYING THE SECURITY PERSPECTIVE

The activity diagram in Figure 24–2 illustrates a simple process for applying the Security perspective. In this section, we describe the activities in this process.



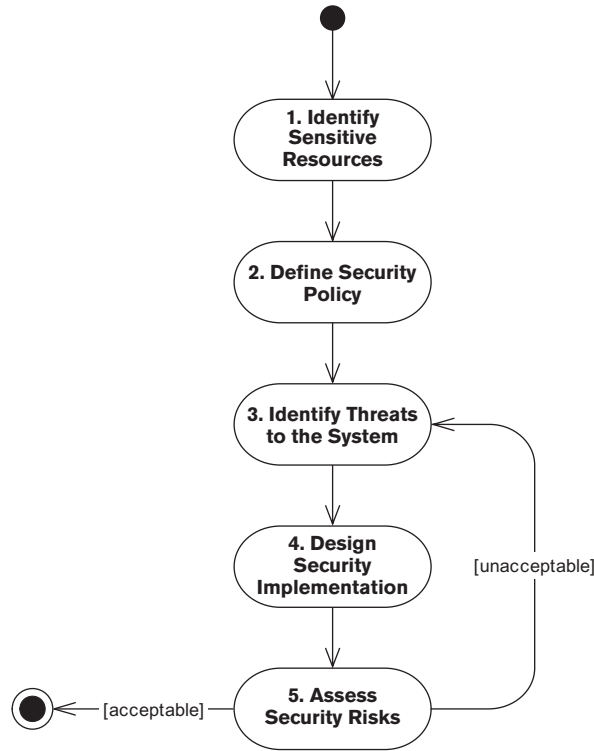


FIGURE 24-2 APPLYING THE SECURITY PERSPECTIVE

Identify Sensitive Resources

Before considering *how* to secure your system, you need to establish *what* needs to be secured. All of the security for your system needs to be driven from this key concern, which you should consider as early in the system’s lifecycle as possible.

The information you gather about the security-sensitive resources feeds into all of the other security-related decisions you make.

NOTATION A simple text-and-tables approach is normally sufficient, although you may choose to annotate one of the diagrammatic models of the architecture (typically in the Functional or Information views).

ACTIVITIES

Classify Sensitive Resources. Using the Functional and Information views as primary inputs, along with any security requirements information avail-

able, decide what the sensitive resources in the system are—typically, functional operations and data items. For each type of sensitive resource, define the reasons the resource is sensitive, who should be considered its owner, and the type of access controls required for the resource.



EXAMPLE Table 24–2 shows a simple example of documenting some of the sensitive resources that might appear in an e-commerce ordering system.

If you’re fortunate enough to have comprehensive security requirements defined for you, the sensitive resources may have already been identified. In many cases, though, you will need to do this yourself, in conjunction with

TABLE 24–2 EXAMPLE OF SENSITIVE RESOURCE IDENTIFICATION

Resource	Sensitivity	Owner	Access Control
Customer account records	Personal information of value for identity theft or invasion of privacy	Customer Care Group	No direct data access
Descriptive product catalog entries	Defines what is for sale and its description; if changed maliciously, could harm the business	Stock Management Group	No direct data access
Pricing product catalog entries	Defines pricing for catalog items; if maliciously or accidentally modified, could harm business or allow fraud	Pricing Team in Stock Management Group	No direct data access
Business operations on customer account records	Needs to be controlled to protect data access and integrity	Customer Care Group	Access to individual record or all records by authenticated principal
Descriptive catalog operations	Needs to be controlled to protect data access and integrity	Stock Management Group	Access to catalog modification operations by authenticated principal
Pricing catalog modification operations	Needs to be controlled to protect data access and integrity	Pricing Team	Access to price modification operations by authenticated principal, with accountability of changes
...



interested stakeholders (such as auditors and managers of the groups owning the information). Useful sources of input information for this exercise include organizational security policies and relevant external regulations, as well as the opinions of interested stakeholders.

When performing this exercise, be sure that everyone (including yourself) understands exactly what is included in each group of resources as well as the identity of the group's owner and the type of security required for the resources in the group.

Define the Security Policy

Having identified the system's sensitive resources and the threats against them, you should be able to define a security policy (also sometimes called a *trust model*) for your system. The security policy is the basis for the security implementation in the system. This model identifies *who* will be trusted with *what* access to *which* system resources (and any constraints on this access such as limiting it to certain times or days of the week), the *integrity* guarantees required within the system, and the *accountability* required when sensitive resources are accessed.

Ideally, you will have a complete, unambiguous, straightforward security policy defined for you by the system's main stakeholders that you can just check and review before proceeding. Alternatively, if it's like every system we've worked on, you'll have nothing of the sort, in which case you need to get a written, approved definition.

The policy definition should normally be defined in terms of groups of resources and principals (often based on organizational unit and role) rather than enumerating lots of specific cases. Also remember that this is a policy, not a design, so it needs to define what access will be provided to whom rather than defining how this will be achieved.

Work to make the security policy as simple and general as you can, with as few special cases as possible. A straightforward policy model allows the simplest possible implementation of enforcement mechanisms and makes it more likely that the model will get enforced correctly.

The other vital quality of a security policy model is precision. This model is an important deliverable for your system, and it should be approved by all the important stakeholders. If the model isn't precise, each group of stakeholders will interpret it differently, according to their own interests, assumptions, and desires. This is bound to lead to problems when implementing the policy.

NOTATION A security policy is normally defined by using a simple text-and-tables approach to create a structured document. Tables are particularly valuable when defining groups of resources and principals and showing the types of access allowed.



ACTIVITIES

Identify the Principal Classes. In order to make defining the security policy more manageable, start by grouping your principals into classes you can treat as groups for security policy purposes. Partition the principals into sets based on the types of access they require to different types of sensitive resources.

Identify the Resource Classes. Partition the system's sensitive resource types into groups that you can treat together for access control purposes. (The resource types come from identifying the sensitive resources, as described earlier.)

Identify the Access Control Sets. For each resource class, define the operations that can be performed on members of that class and the principal classes that should be allowed to access each operation on the class.

Identify the Sensitive Operations. Consider any system-level operations that are independent of the system's managed resources (such as administrative operations), and define which principal classes should be allowed to access these operations.

Identify the Integrity Requirements. Consider any situations in the system where information is or could be changed, and identify the set of integrity guarantees required for that information.



EXAMPLE Table 24-3, on the following page, shows part of the result of defining the security policy for an e-commerce system.

Seven different classes of principals have been identified (including unauthenticated Web-site users) and six classes of sensitive resources. Only data administrators are allowed access to data directly; all other users can only execute business operations. Data administrators can also execute business operations (on the basis that they can simulate their occurrence anyway by changing data), but all of their activity is audited. The other classes of principals have been allocated just enough privileges in order to perform their roles in the system.

Identify Threats to the System

In a perfect world, we could now stop and simply publish the list of sensitive resources and the security policy, and everyone would voluntarily follow the policy. However, the reason we need security is that we can't necessarily trust everyone, so we need to identify the possible threats to the security policy. Identifying threats provides a clear definition of what in the system needs to be protected and what it needs to be protected from. This makes it clear which threats you are aware of and can try to guard against and also (implicitly) the threats you have not considered.

TABLE 24-3 EXAMPLE OF AN ACCESS CONTROL POLICY

	User Account Records	Desc. Catalog Records	Pricing Records	User Account Operations	Desc. Catalog Operations	Price Change Operations
Data administrator	Full with audit	Full with audit	Full with audit	All with audit	All with audit	All with audit
Catalog clerk	None	None	None	All	Read-only operations	None
Catalog manager	None	None	None	Read-only operations with audit	All	All with audit
Product price administrator	None	None	None	None	Read-only operations	All
Customer Care clerk	None	None	None	All	Read-only operations	None
Registered customer	None	None	None	All on own record	Read-only operations	None
Unknown Web-site user	None	None	None	None	Read-only operations	None

The result of this process is termed a *threat model*, which builds on the initial list of sensitive resources and should provide a thorough analysis of the threats you think the system is subject to, the impact of the threat being realized, and the likelihood of the threat occurring.

To create a threat model, you ask a number of key questions about your proposed system.

- Who is likely to try to infringe the security policy?
- How will they try to do so?
- What are the attacker's main characteristics (such as their sophistication, motivation, resources, and so on)?
- What are the consequences of the policy being breached in this way?

Explicitly identifying threats allows specialists from outside your project to provide assistance by reviewing the model and advising you about potential threats that have not been considered at all or threats that may have been incorrectly characterized in the model. Be sure to have it widely reviewed so that you can be confident you have considered all likely threats. The threat



model also allows you to systematically consider the set of security facilities your system needs.

NOTATION The threat model is normally presented by using a text-and-tables approach to create a structured document. A graphical alternative is to use an attack tree, which provides a structured, graphical notation to categorize and illustrate the threats a system faces and the likely probability of each occurring.



EXAMPLE An attack tree is a useful approach to representing the threat model for your system. The technique is based on the Fault Trees technique used in safety-critical systems design to analyze possible failure modes. In the security domain, you can take a similar approach to represent possible attacks on your system.

An attack tree represents the possible attacks your system may face in order for an attacker to achieve a particular goal. The root of the tree is the goal the attacker is trying to obtain, and the branches of the tree classify the different types of attacks the intruder could attempt in order to obtain the goal.

Attack trees can be represented graphically (as a tree structure with nodes and links) or textually as a “dotted decimal” hierarchy (like nested subheadings in a technical document). The latter form is often more practical because attack trees can get very large as all possible attacks are considered, which can make the graphical form difficult to draw and comprehend.

Here is a possible attack tree for the goal of extracting customer credit card details from an e-commerce Web site.

Goal: Obtain customer credit card details.

1. Extract details from the system database.
 - 1.1. Access the database directly.
 - 1.1.1. Crack/guess database passwords.
 - 1.1.2. Crack/guess operating system passwords that allow database security to be bypassed.
 - 1.2. Access the details via a member of the database administration staff.
 - 1.2.1. Bribe a database administrator (DBA).
 - 1.2.2. Conduct social engineering by phone/e-mail to trick the DBA into revealing details.



2. Extract details from the Web interface.
 - 2.1. Set up a dummy Web site and e-mail users the URL to trick them into entering credit card details.
 - 2.2. Crack/guess passwords for user accounts and extract details from the user Web interface.
 - 2.3. Send users a Trojan horse program by e-mail to record keystrokes/intercept Web traffic.
 - 2.4. Attack the domain name server to hijack the domain name and use the dummy site attack from 2.1.
 - 2.5. Attack the site server software directly to try to find loopholes in its security.
3. Find details outside the system.
 - 3.1. Conduct social engineering by phone/e-mail to get customer services staff to reveal card details.
 - 3.2. Direct a social-engineering attack on users by using public details from the site to make contact.

An attack tree should be created for each of the possible goals that an attacker may have for breaching your system's security. Once you have an attack tree, you can analyze each threat it contains to establish whether the system's security neutralizes the threat.

ACTIVITIES

Identify the Threats. Consider the security threats your system faces from the perspective of the sensitive resources within it, the possible access to these resources that potential attackers might wish to gain, the main characteristics of the potential attackers, and the types of attacks they are likely to carry out.

Characterize the Threats. Characterize each threat in terms of the resources that would be compromised if the attack were successful, the result of this compromise, and the likelihood of the attack occurring.

Design the Security Implementation

Once you understand the sensitive resources and threats, you can consider the technical security design for the system. The goal of this step is to design a system-wide security infrastructure that can enforce the system's security policy in the face of the risks identified in the threat model. In this step, you consider using specific security technologies such as single-sign-on systems, network firewalls, SSL communication link security, cryptographic technology, policy management systems, and so on.





This design process results in a number of design decisions you should incorporate in the architecture. These decisions affect a number of architectural structures, including those likely to be described in the Functional, Information, Deployment, and Operational views.


NOTATION As the outputs of this step are a set of design decisions to be reflected in the architectural views, the notation used for the technical security design depends on the notation used in each view. You may also produce some form of overall security design model, which is typically captured by using a software design notation such as UML (similar to its use in the Functional viewpoint).



EXAMPLE Based on the attack tree shown in the previous example, some of the security measures you might consider for the e-commerce system include the following:

- Isolating the database machines from the public network by using network firewall technology
- Isolating the security-sensitive parts of your system from the public network by using network firewall technology
- Analyzing the paths into your system to check them for possible vulnerabilities
- Arranging penetration testing to see if experts can find ways into your system
- Identifying an intrusion detection strategy that would allow security breaches to be recognized
- Training administration and customer service staff (in fact, probably all staff) to avoid social-engineering attacks and to abide by strict privacy protection procedures for customer information
- Designing your site so that a minimal amount of user information (ideally, none) is publicly viewable
- Designing your site so that sensitive information (e.g., credit card numbers) is never shown in full (e.g., display just the last four digits to allow legitimate users to identify their cards in lists)
- Constantly reminding users of security precautions they should take (e.g., not revealing passwords to anyone, including your staff; checking URLs before entering information; and so on)

It is interesting to note how much of a typical security implementation does not involve security technology directly but is about making sure that people act in a secure manner.



ACTIVITIES

Design a Way to Mitigate the Threats. For each of the previously identified threats to sensitive resources, design a security mechanism to address the threat. This may include modifying existing architectural decisions, applying one or more security technologies, and designing procedures and processes for the system's operation and use.

Design a Detection and Recovery Approach. Bearing in mind that the security mechanisms you identify are unlikely to be foolproof, design a system-wide approach to detecting violations of the system security policy and recovering from them. This will typically include technical intrusion detection solutions, internal system checks and balances to reveal unexpected inconsistencies, and a set of processes for regularly checking the system for intrusions and reacting to any discovered. Good solutions for intrusion detection are available as proven commercial and open source software packages; however, deciding how to use the packages and design the processes around them are still system-specific activities.

Assess the Technology. One way to address a threat is to use a piece of security technology to provide a security mechanism. This activity involves assessing which candidate security technologies are suitable for addressing a particular threat in a particular context. This includes generic assessment (e.g., checking reliability, ease of use, and so on) as well as context-specific assessment (e.g., checking that the proposed technology can be operated by the system's administrators and is efficient enough to meet the system's performance and scalability goals).

Integrate the Technology. Once you have decided which security technologies to apply to your system, the other important activity is deciding how you will integrate them with the primary system structure and with other security technologies. It is important to spend some time designing the integration approach carefully, to avoid possible security loopholes creeping in due to the unexpected side effects of an ad hoc integration approach.

Assess the Security Risks

No security system is perfect. The process of implementing system security is a balancing act: You balance the risks you believe you face against the cost of implementing solutions for those risks and the costs you may face if the risks occur.

Having designed a security infrastructure for your system, you now need to reevaluate the risks to consider whether your proposed security infrastructure has achieved an acceptable cost/risk balance. If so, your Security perspective is complete. If not, you need to return to a consideration of the threat model and redesign the security infrastructure in order to achieve an acceptable cost/risk balance.



TABLE 24-4 EXAMPLE OF RISK ASSESSMENT

Risk	Estimated Cost	Estimated Likelihood	Notional Cost
Attacker gains direct database access.	\$8,000,000	0.2%	\$16,000
Web-site flaw allows free orders to be placed and fulfilled.	\$800,000	4%	\$32,000
Social-engineering attack on a customer service representative results in hijacking of customer accounts.	\$4,000,000	1.5%	\$60,000
...

NOTATION The risk assessment delivers a record of the set of risks, the estimated likelihood of each risk occurring given the design of the system, and the notional cost (i.e., the estimated cost adjusted for the probability of occurrence) that each risk implies. This information is usually best presented with a simple tabular form.



EXAMPLE Table 24-4 shows an example, using a tabular presentation style, for a couple of the risks facing our e-commerce system.

A risk presentation in this form allows us to focus on the risks with the highest notional costs (i.e., those that have the highest risk of loss and greatest likelihood of occurrence).

ACTIVITIES

Assess the Risks. The single activity in this step is a process of risk assessment. For each risk in your threat model, reevaluate its likelihood of occurrence given the planned security infrastructure, and compare this against the security needs you established earlier in the process. Assess whether this is an acceptable level of security risk for your particular system's situation.

ARCHITECTURAL TACTICS

Apply Recognized Security Principles


The field of computer security is relatively mature for a rapidly evolving, high-technology discipline. There is an established body of knowledge, developed



and applied by an identifiable professional security community with its own principles, standards, norms, and culture.

This community of researchers and practitioners has established a number of widely accepted and commonly applied principles that are considered important to establishing security within a system. Some of the more important principles include the following.


- *Grant the least amount of privilege possible:* Always grant security principals the smallest set of privileges they require in order to perform their tasks. Consider varying the set of privileges a principal has over time if certain sensitive tasks are executed only intermittently.
- *Secure the weakest link:* The security of a system is only as strong as its weakest element, so understanding the weakest link in your security is an important step toward understanding how secure your system really is. The weakest link could be technological (an unsecured network link), procedural (allowing easy access to a data center), or human (people who write down their passwords). Identify and secure the weakest links in your system's security until you achieve an acceptable level of security risk.
- *Defend in depth:* If you examine physical security systems, you'll find that they rarely rely on just one security measure. Just as medieval castles had moats, drawbridges, and strong walls, banks have alarms, vaults, security guards, surveillance systems, and multiple locks on important doors. These are examples of the principle of defense in depth, where a series of defenses provides a greater level of security than a single technique could. Defense in depth is particularly relevant to computer systems, given that many of the security technologies we use may themselves have hidden flaws and that we are also susceptible to human and procedural failures. Rather than relying on one security measure to counter each threat to your system, consider possibilities for layering defenses to provide greater protection.
- *Separate and compartmentalize:* Attempt to clearly separate different related responsibilities so that authority for each can be assigned to different principals if required, and compartmentalize responsibilities for different parts of the system so that they can be controlled individually. This makes it easier to control access securely and means that a successful attack on one part of the system does not immediately compromise it entirely. Good examples are separating the "security override" privilege from the "alter audit trail" privilege (with the audit trail being used to record use of the "security override" privilege) and implementing separate security configuration and mechanisms for each major subdivision of a system.
- *Keep security designs simple:* Security professionals often say that complexity is the enemy of security. Complexity in a system is difficult to deal with and makes it very difficult to analyze the system to assess its secu-





rity, which makes problems more difficult to understand and rectify when you find them. Systems with stringent security requirements need to be simple enough to make it possible for them to be secured and verified.

- *Don't rely on obscurity:* Some years ago, it was common for the details of secure systems to be kept secret in an attempt to make them more difficult to attack. An example of this was hiding security secrets, such as cryptographic keys, in obfuscated (and undocumented) computer code for which the source was kept secret. The problems with this approach are that it assumes the attackers aren't smart enough to work past the obscurity and that it may prevent external experts from assessing the real level of security provided. The conventional wisdom in today's security community is total disclosure, where the system is designed assuming that potential attackers know it as well as its implementers do. This principle avoids the system's security becoming dependent on hiding secrets, which is often extremely difficult to do.
- *Use secure defaults:* Many systems we have encountered include inherently insecure default settings and behaviors, including empty default passwords, permissive default access control lists, network ports open by default, and so on. Such behavior is very likely to result in real security threats, particularly when exhibited by packaged software products, which are often installed by users who are not familiar with them. Ensuring secure default behavior makes a real contribution to practical systems security.
- *Fail securely:* Another common problem found in many real systems is insecure failure mode behavior, where the system is reasonably secure during normal operation but becomes insecure when things go wrong. Examples include audit trails that suspend auditing if the audit logs run out of space, systems that drop back to insecure mode processing if security negotiation fails, unprotected recovery consoles that appear after crashes, and so on. System failures are inevitable occasionally, so make sure that if your system fails, it does so securely.
- *Assume external entities are untrusted:* Within your system you should be able to exert a great deal of control over the security environment and the principals within it. This is not the case for external entities who access the system. Ensure that all external entities are totally untrusted until proven otherwise to avoid accidental security breaches in unexpected cases.
- *Audit sensitive events:* Most systems include a number of key security-related events that, if abused, could compromise the security system. Common examples include resetting passwords, assigning powerful roles, and manipulating audit trails. These sensitive events need to be securely audited so that their use can be monitored. This principle often places a number of constraints on your deployment structure to ensure that a reliable audit trail can be implemented.





Make sure that you understand these commonly accepted principles and apply them as appropriate to your system's security. In the Further Reading section at the end of this chapter, we recommend a number of background references that can provide much more detail on these subjects.

Authenticate the Principals

Authentication is the reliable identification of each of the principals who can use the system. These principals could be people, computers, pieces of software, or anything else in the system that has a security identity. Nearly all system security depends on being able to identify the various principals in the system reliably.

Implementing authentication involves reliably binding a unique name of some sort to each principal who can use the system. A wide variety of authentication technologies exist to allow you to identify the principals in your system, including user names with passwords, public/private key systems (such as X.509 digital certificates), and hardware token technologies (such as smart cards). Off-the-shelf products known as single-sign-on systems allow you to delegate the entire system-wide authentication process to them (which is particularly useful in mixed technology environments such as those with both UNIX and Windows authentication). Each of these technologies has its place, but each tends to be suited to different environments and security needs, so you need to select the right ones for your system.

The key decision to be made is how every principal in your system is uniquely identified, via a mechanism that is secure enough to address the particular risks your system faces. You should bear in mind that you may need to use different principal authentication technologies for different principal types, and in some cases, principals may need to be identified by more than one authentication technology (e.g., users in a corporation may have separate digital certificates, Windows networking logins, Oracle database logins, and Lotus Notes logins). Where multiple logins are necessary, we strongly recommend some form of single-sign-on technology as a unifying layer on top of the different underlying systems. People simply can't remember more than a couple of passwords, and forcing them to remember more will result in passwords written on notes stuck to terminals!

What is critical at the architectural level is that every principal can be reliably identified when required and that the system you use is simple and usable enough that the system's users will not try to work around it.

Authorize Access

Once you have identified the principals, authorization involves restricting and enforcing what those principals are permitted to do within the system. Concep-



tually, you define lists of principals and their permitted actions for each controlled resource in the system. This usually involves assigning principals to roles and groups so you can treat them as homogeneous populations. Then you can define permitted actions on resources for the roles and groups in the system.

Most enterprise software technologies already include some form of access control that allows you to limit access to resources based on user identity and credentials (such as group membership). In addition, software developers can code explicit security checks into their software when they implement the system.

In many systems, your key problem as the architect is how to create a coherent access control system from this patchwork of separate, often incompatible authentication and authorization technologies. For system security to be manageable and effective, the system architecture needs to include a clearly defined approach to enforcing principal authorizations, particularly if multiple candidate authorization systems are available in the system's environment.

A solution to this problem can be the use of purpose-built standard access control products. These software packages (such as IBM Policy Director or Entrust GetAccess) allow access to many sorts of resources to be controlled through a single product that interfaces to a variety of authentication systems. At the other extreme, including a simple homegrown access control manager in your system architecture may be sufficient.

Ensure Information Secrecy

Secrecy is the means of ensuring that only owners of information (or those with whom they choose to share it) can read that information.

Secrecy can be partially achieved by using authorization appropriately, to prevent access to stored information that a principal is not authorized to access. However, access control is rarely sufficient, and it usually needs to be combined with other techniques in order to achieve the secrecy required.

Traditional information systems were highly centralized, with huge information vaults held on central computers. The information was tightly controlled by access control systems (typically as part of database management systems) and was accessible only via custom applications running on terminals directly connected to the central computer.

In contrast, today's information systems are highly distributed, with many computers communicating in order to implement the functions of the system. This means that information can be much less secure than with a traditional centralized system—particularly if public networks are used for communication and the different parts of the system are within different organizations.

One of implications of this shift is that sensitive information needs to be protected once it moves outside the authorization controls of the database servers where it resides. Ensuring secrecy of information outside databases



normally involves the use of cryptography, to encrypt information in such a way that only principals who have access to a particular key can decrypt it and read it. In reality, the inherent complexity of cryptography means that it is typically implemented by securing network links through the use of communications protocols that include cryptographic protection, like SSL/TLS.

From the architect's point of view, the snag with this neat picture is that extensive use of cryptography in your system is not without cost. It will add significant complexity to many aspects of the system (cryptographic key management alone is a complex problem), and it also tends to have high computational costs, burning up precious computing resources that could be dedicated to useful functional processing. It can be difficult to apply cryptography correctly so that the system is actually secure. It is significantly more complex than just using server certificate-based SSL to identify your servers because this leaves your system open to any number of alternative attacks (such as intercepting information elsewhere in the system once it has been decrypted).

In general, use your threat model to identify where information needs to be protected, and use the minimum amount of cryptography that leaves your system acceptably secure.

Ensure Information Integrity

Integrity is the name given to the concern of ensuring that information is protected from unauthorized change (particularly during message transmission).

Implementing integrity in most information systems also involves the application of cryptography. We would like the recipients of messages sent over a network to be able to satisfy themselves that they have received each message unchanged and that no one has tampered with it. Ensuring integrity of information usually involves the application of cryptographic hash functions. Hash functions are cryptographic operations that use a secret key to compute a large numeric value for a block of information. They allow the recipient of that block to ensure that the information has not been changed since the hash value was computed. However, useful as this is, the same caveats mentioned for secrecy also apply to integrity, in particular that it adds both complexity and runtime overhead to the system and so needs to be used judiciously.

Ensure Accountability

Many systems require some or all of their users to be accountable for their actions. In some systems (such as certain financial and medical systems), legal requirements mandate user accountability for key operations. There tend to be two distinct forms of accountability required in information systems: *auditing* and *nonrepudiation* of messages.





Auditing is usually used in situations where secure central servers execute the system's primary operations. These servers can record logs of operations performed by the system's users, which can later be used to establish how a particular situation occurred.

In distributed systems without centralized servers, auditing is less relevant (due to the lack of a secure store for the audit trail). The related requirement in such systems is nonrepudiation, the ability to definitively identify a message's creator in a way that she cannot plausibly deny. A common solution to this problem is to use digital signatures on messages that can be generated only by someone with access to the corresponding digital certificate's private key.

Both of these approaches come with their own costs. Auditing has a runtime performance penalty associated with it, as well as management and storage overhead for the records created. Digital signatures are computationally expensive to create and verify, and they can be complex to administer because they require all of the principals in the system to have a unique digital certificate. We are also ignoring other important complexities such as certificate management and revocation and the difficult problem of ensuring that technically sophisticated users do not compromise other users' digital certificates.

Once again, your role as an architect is to balance the cost of these accountability mechanisms against the risks you think your system faces, applying these techniques where the threat model indicates they are really needed.

Protect Availability

When thinking about system availability, it is natural to focus on hardware reliability, software replication, failover, and so on. Such approaches are an important piece of the availability puzzle, as we discuss in Chapter 26 on the Availability and Resilience perspective. However, there is an easily overlooked security aspect related to availability: protecting your system from hostile attacks that aim to reduce its availability.

Such attacks, known as *denial-of-service (DoS) attacks*, have been the subject of increasing interest in the last few years as systems have started to expose their network interfaces to public networks like the Internet. DoS attacks can range from something as simple as forcing a user to be locked out by repeatedly using an incorrect password with her login,¹ to sophisticated

1. This doesn't sound like a terribly serious attack until you consider time-sensitive systems such as online auctions. This attack has allegedly been used on a number of occasions at Internet auction sites to prevent users bidding on (and so raising the prices of) items that the attacker wishes to buy. Other systems can also be vulnerable to similar DoS attacks—for example, a battlefield system isn't much use if an attacker can lock out all of its users simply by bombarding it with credible user names and random passwords.



network attacks using custom software that overloads systems, to physical attacks such as attempting to interrupt a system's power supply.

Work to understand the possible DoS threats your system faces and the impact of any of the threats occurring. This will probably require working with security specialists who understand this area in detail. Once you understand the threats, their costs, their likelihoods, and the potential approaches to protect against them, you need to decide the right level of protection for your system by balancing these factors. Be aware, however, that protecting your system from network DoS attacks can be tremendously difficult, and for most systems you may simply have to accept some level of risk of such an attack occurring.

Integrate Security Technologies

Almost inevitably, when you design the security implementation, you will need a number of different technologies in order to fulfill all of the requirements. Similarly, security normally needs to be implemented across a number of distinct parts of your system. Given this situation, you need to make an early design decision about how to achieve end-to-end security in the system. Implementing a large number of unrelated security technologies is highly undesirable because it introduces complexity and the likelihood of creating security vulnerabilities at the boundaries of each technology.

Part of your role in the security design of the system is to ensure that security is implemented consistently and that the different pieces of the technology are put together to form a complete, integrated security system. This will be particularly important if you're using security consultants from a number of sources (such as different product vendors) who will all focus on their own part of the security puzzle.

Provide Security Administration

Administration is a weakness of many existing security systems, particularly where complex policies need to be enforced or where there is a need to cope with the sort of large user populations found in today's global organizations. This problem can become even more awkward when you need to combine a number of security technologies in order to implement the security policy.

As part of architecture definition, you need to ensure that the planned security implementation can be administered effectively (this will form part of the Operational view). Without good administration, security policy is likely to be ignored simply to "get the job done." Complex security administration facilities can also lead to security loopholes caused by administrators who are unaware of the full ramifications of their actions. Remember to involve key stakeholders when considering security administration. Your system administrators will be able to provide a lot of feedback about the acceptability of the





administration facilities available—testers too often have a perspective on this because they need to reproduce large security configurations automatically to allow reliable, repeatable testing of the security facilities.

Use Third-Party Security Infrastructure

Many of today's infrastructure technologies (such as J2EE and .NET servers, enterprise directories, and e-mail servers) provide a number of standard security functions. Also, specific security infrastructure technologies (such as enterprise access control products) are available that can provide security services to applications. These technologies offer an alternative to the traditional approach to implementing application security—namely, embedding policy enforcement in the application code.

In general, using external security infrastructure is simpler than coding it directly; after all, someone else has done a lot of the work already. This approach is often much more flexible, too. Rather than needing to change program code when the security policy changes, it is often possible to change the configuration of the infrastructure element—a process that can be achieved by the administrators without involving software developers. Most importantly, a security infrastructure product is likely to provide more reliable security than application-specific code because it is more widely used and has been written by software developers with specialized interest and training in system security.

We strongly recommend adopting a security approach of pushing as much of your policy enforcement into the underlying infrastructure of your system as possible. We suggest this for three reasons: simplicity, flexibility, and reliability. Security infrastructure products hide much of the complexity of integrating different security systems. They should also provide some degree of abstraction over individual mechanisms, allowing the mechanisms to be changed more easily. Finally, widely deployed infrastructure products have been widely tested, analyzed, and understood, increasing their reliability when compared with custom solutions.

PROBLEMS AND PITFALLS

Complex Security Policies

Security policies have a habit of starting out simple and well-defined but becoming complex, full of special cases, and poorly understood as they are reviewed by more and more stakeholders. Implementing system security is difficult during the best of times. If the security policy is not a simple and regular rule set, this job becomes much more difficult during both development and system operation, and the likelihood of flaws in the security implementation is much higher.



RISK REDUCTION

- Make sure that the security policy is as simple as it can be—make functional, performance, and other tradeoffs if appropriate in order to achieve this goal.
- If policies appear to be very complex, consider partitioning the security-sensitive resources or the principals in different ways to make the policy simpler.

Unproven Security Technologies

There are many security technologies available for implementing your security infrastructure. They range from simple, very well understood technologies such as user names and passwords to large-scale, sophisticated systems in their own right such as many public/private key cryptography implementations.

However, whatever technology you use, your system's security is only as strong as its weakest link. A single weak element in your security infrastructure has the potential to make the rest of the infrastructure irrelevant. The weakness in an element can come from an inherent design or implementation flaw in the underlying product, a mistake made in its application, or indeed a mistake made in operating the technology.

RISK REDUCTION

- Err on the side of caution when selecting your security technology—such technology needs to be well understood and proven in operation before you can consider it a sound choice for a generic piece of infrastructure.
- If you don't have practical experience with a particular technology, find someone who does have such experience to guide you. (And don't assume that everything the vendor says is true!)

System Not Designed for Failure

As any architect knows, things go wrong in the process of designing, building, and operating information systems. Therefore, system elements are often created to fail in certain ways in order to minimize the impact of a particular failure. This is often a sound approach. It allows the system to continue being used (possibly with reduced function) while one or more pieces of it are being fixed. However, when considering security, this is not always the best approach.

To illustrate this, consider the failure of a security element that controls access to a sensitive resource. If the element is unexpectedly unavailable, the high-availability approach to design might suggest that the default behavior in such cases is to simply allow access until the security element is fixed. However, this obviously isn't the right thing to do from a security perspective.





Similar but much more subtle examples of this problem can happen when unexpected errors occur and security has not been considered while designing the relevant error-handling code.

RISK REDUCTION

- Design your security infrastructure to cope with failure in a safe way from the beginning so that unexpected element failures do not open security loopholes in the system.
- Ensure that your security infrastructure is configured to fail securely if unexpected situations occur.
- When reviewing your architecture in failure scenarios, be sure to check what impact the failures have on the system's security.

Lack of Administration Facilities

An important but often neglected part of designing security infrastructure is to ensure that it can be effectively administered when the system is in production. Your testers may also have administration needs to ensure that they can reproduce large-scale test scenarios effectively.


Two common problems in this area include administration tools that can be used only for simple test cases (rather than for the thousands of users who can be found in real systems) and the use of a patchwork of administration tools that need to be carefully combined in order to manage the system. Such a situation will likely result in loopholes in security policy enforcement as administrators struggle to use inadequate tools to control access.

RISK REDUCTION

- Review your architecture to ensure that its administration facilities are adequate for the expected size of the user population and the complexity of the security policy.
- Make sure that the administration facilities you are planning to provide are acceptable to the administrators and operators who will have to use them.
- If possible, make sure that administrative security operations can be performed by using one task-driven interface, rather than requiring the use of a number of tools to perform a single task. If many steps are required, the likelihood of omission or error is much higher.

Technology-Driven Approach

A problem we have seen with the security design of some systems occurs when the available technology drives the security design process. You can tell that you have this sort of problem when you inquire whether the designers



have considered security yet and you get the reply, “Oh yes—we’re secure, we use SSL.” In such situations, although security technology is deployed, you often have no idea whether it addresses the system’s security needs.

RISK REDUCTION

- Drive your security design by the resources that need to be protected, the security model that needs to be implemented, and the security threats that the system faces.
- Avoid designing your security architecture around specific pieces of security technology. Don’t incorporate security technology that is not justified by the security needs of your system.

Failure to Consider Time Sources

A number of security mechanisms rely on checking the passage of time (e.g., product license timeouts and password expirations). These mechanisms assume that a reliable source of time is readily available. Although the specifics of the mechanisms and the time source they need differ, they all share the characteristic that if the time source is compromised, the mechanism is ineffective.




RISK REDUCTION

- Identify the security mechanisms that require accurate time and the characteristics of the time they need (such as accuracy, absolute or relative, universal or time zone, and so on).
- Incorporate enough secure time sources in your system to meet the needs of your security mechanisms.
- Make sure you understand what will happen if a secure time source is unavailable or is compromised.
- Protect secure time sources against likely threats (using mechanisms such as calling back to secure servers, using external time servers or dedicated hardware devices, using operating system access controls to the time source, and so on).

Overreliance on Technology

It is often said by security luminaries like Bruce Schneier² and others that security is a process, not a product. Yet you can often find system designers

2. *Cryptogram*, February 15, 2002, www.schneier.com/crypto-gram-0202.html—in reference to Microsoft security.





placing great reliance on particular security products to keep their systems secure. Although you certainly do need to use good third-party security products to secure your system, you need to use them intelligently as part of an overall security design that encompasses all of the different aspects of your system.

At a recent security conference, an executive of a security product company pointed out that “there is no way you can buy anything, subscribe to anything and say you are 100% secure.”³ This is sound advice, coming from someone actually in the IT security industry.

RISK REDUCTION

- Use your threat model to drive your security design. Addressing these tangible threats will help you design a system that is secure, rather than a system that just uses security technology.
- Design a sound set of operational procedures to avoid human error exposing the system to security threats.


No Clear Requirements or Models

It is common for systems to have no clear, well-defined security requirements and no formal security models at all. The problem with these systems is that you simply don’t know if they are secure because you don’t know what “secure” means. Security is an area where it isn’t possible to be confident that you have met the needs of the stakeholders without firm requirements and formal models. Security can’t be seen, and most stakeholders won’t test it. You’ll find out that you have problems only when the security is breached—at which point you can be pretty certain that the stakeholders won’t think their needs are being met.

RISK REDUCTION

- Drive the security design process by using threat and security policy models. Developing these models will help define requirements because they focus the stakeholders’ minds on what is valuable, what is likely to be attacked, and what the impact of such attacks would be.
- Use plenty of concrete examples when discussing security with stakeholders. Security is an abstract area that requires some lateral and imaginative thinking in order to be effective. Examples will help your stakeholders think clearly about what they need.

3. Arthur Wong, CEO and founder of Security Focus, speaking at RSA Conference 2002, as reported on [cnn.com](http://archives.cnn.com/2002/TECH/internet/02/25/2002.security.idg/), February 25, 2002, <http://archives.cnn.com/2002/TECH/internet/02/25/2002.security.idg/>.



- When identifying your architectural scenarios, consider those related to security as well as to functionality, performance, evolution, and so on. Thinking through security-related scenarios can help identify and clarify important security requirements.

Security as an Afterthought

Because stakeholders often don't think about security requirements explicitly, these requirements don't always get mentioned in the initial requirements analysis. This can lead to the problem where security has to be added to the system at some point during (or even after) development. At best, this is likely to be an expensive and painful process involving a lot of rework. At worst, it won't be possible to introduce the required security without changing the system in some way that upsets a stakeholder group.

RISK REDUCTION

- Start considering the system's security as soon as you start developing its functional and information structure. This will allow you to understand the security needs early in the lifecycle and make sure that the system you design can be secured.

Security Embedded in the Application Code




A problem observed in a number of systems occurs when the code that enforces the security model is found sprinkled throughout the application itself. As explained earlier, the problems with this approach include reduced reliability, the difficulty of changing the security model being enforced, and the likelihood of introducing security errors in your system.

RISK REDUCTION

- Push as much of the security technology as possible into your underlying infrastructure elements.
- If security does need to be in application code, apply good software engineering judgment to encapsulate as much of it as possible in a single place.

Piecemeal Security

To be effective, security needs to be considered on a holistic basis, throughout the system. A problem seen in many systems is security being applied to parts of the system but not to others. For example, highly sensitive data is encrypted during transmission but not when stored. This may or may not be a





problem depending on the sensitivity of the data and the threats the system is likely to have to withstand. Your role is to make sure that security is implemented everywhere it is needed as revealed by the threat model—not just in the places that immediately spring to mind.

RISK REDUCTION

- The use of an architecture-driven development process will help address this risk. Make sure that you keep considering the architecture and particularly its security as a whole, rather than as a set of separate parts.

Ad Hoc Security Technology

Computer security is a specialized field with its own culture, standards, processes and background. Security engineers (the people who build security technology and secure systems) tend to have a lot of specialized training and experience. Similarly, cryptographers (the people who study and create ways to encrypt and decrypt data) tend to have advanced degrees in cryptography. Despite this, many software developers without this specialized background fancy themselves as amateur security engineers or cryptographers and decide to create some or all of the security technology in their systems.

In general, we suggest that this is a bad idea—creating truly secure technology is harder than it looks, and without specialized training, most of us can't do it reliably. Given the integrated nature of security, the consequences of one piece of weak security technology can be pretty catastrophic.


RISK REDUCTION

- Use proven, widely accepted security technology from established providers wherever possible, and get expert help with its use and deployment.
- Make sure you find out what previous users of the possible technologies thought of it and how the security community as a whole rates it.
- If you have to create your own technology, engage expert assistance to help you systematically develop it.

CHECKLISTS

Checklist for Requirements Capture

- Have you identified the sensitive resources contained in the system?
- Have you identified the sets of principals that need access to the resources?






- Have you identified the system's needs for information integrity guarantees?
- Have you identified the system's availability needs?
- Have you established a security policy to define the security needs for the system, including which principals are allowed to perform which operations on which resources and where information integrity needs to be enforced?
- Is the security policy as simple as possible?
- Have you worked through a formal threat model to identify the security risks your system faces?
- Have you worked through example scenarios with your stakeholders so that they understand the planned security policy and the security risks the system runs?
- Have you reviewed your security requirements with external experts?

Checklist for Architecture Definition

- Have you addressed each threat identified in the threat model to the extent required?
- Have you used as much third-party security technology as possible?
- Have you produced an integrated overall design for the security solution?
- Have you considered all standard security principles when designing your security infrastructure?
- Is your security infrastructure as simple as possible?
- Have you defined how security breaches will be identified and how to recover from breaches?
- Have you applied the results of the Security perspective to all of the affected views?
- Have external experts reviewed your security design?

FURTHER READING

You can find a short but thorough and very readable introduction to the main concepts of information systems security in a recent magazine article [LAMP04]. A much deeper but still quite readable introduction to the process of building secure systems and processes appears in Anderson [ANDE01]; this book provides a comprehensive introduction to many important security topics, as well as being an entertaining read, full of interesting stories from the security





field. A more specific background text aimed at information systems architects is Hardman et al. [HARD01], which explains how to build secure enterprise information systems using Enterprise Java Beans and CORBA technology. In addition to detailed information on security specifics for these technologies, the book includes some useful background and puts the facilities in the context of an overall security process.

A well-known, generally respected, and always colorful figure in the security field is Bruce Schneier. Two of his books that offer useful background are his introduction to cryptography [SCHN95], which provides a good nuts-and-bolts introduction to cryptographic technology for software engineers, and his later book [SCHN01], which moves beyond his initial technology-focused approach and explains how technology will never provide a complete solution to security problems. This second book is worth reading to understand just how complex the security field is and how broad effective security solutions need to be.

Two good, practical security books aimed at software developers are Viega and McGraw [VIEG02] and Howard and LeBlanc [HOWA02]. Both books explain how to construct software that is secure by design rather than by accident or buzzword compliance. Chapter 5 in the former book also presents and explains a practical and simple set of security principles. The Fault Trees approach, which forms the basis of the attack trees we introduced in this perspective, is explained in Leveson [LEVE95].

