

Defense-in-Depth Strategies

LIKE TRADITIONAL WARFARE, cyberwarfare is fought both offensively and defensively. Defending against cyberattacks is a complex task because of the broad range of attackers, the number of potential targets, and the huge variety of ways in which attacks are conducted. When you factor in that people also create vulnerabilities that can be leveraged to attack computers and networks, cyberdefense can seem nearly impossible.

The adversaries that network and systems defenders face vary. Nation-states, corporations, insurrectionists, and hacktivists each have goals in cyberwarfare. They may choose different targets and different methods. They also bring different levels of capability. They might attack using targeted malware or massive brute-force network attacks. Or they might attack via subtler methods that leverage human factors in addition to technological means over weeks or months. If they succeed, they may quietly gather data; continue their attacks to gain greater access; or immediately use their access to damage systems, networks, or infrastructure.

In traditional information security operations, security professionals warn their employers that there is no way to be perfectly secure. If a system is usable and useful, it has the potential to be attacked—no matter how well defended the networks, systems, and other cyberassets are. Worse, organizations have a finite amount of resources to spend on cyberdefense, and cyberdefense can often only defend effectively against threats that are known and understood. With technology's complexity and rate of progress, staying abreast of an organization's defensive needs is a challenge. When you consider the potential to have far more attackers than defenders, and for those aggressors to have far greater resources than your own organization possesses, defense can feel like a losing battle.

Despite these challenges, you can use methods to effectively defend assets, to reduce the chances of compromise, to detect those attacks that do occur, and to provide a competent response. Computer network defense (CND) strategies

attempt to first identify likely opponents, then to enumerate the threats and risks that an organization will face from those attackers. Once an organization has a good understanding of what it may face, it can design strategies to counter them using policies, procedures, technology, training, and a variety of other defensive options.

Since the creation of cyberwarfare as a concept, one of the key concepts for many CND strategies has been defense in depth. *Defense in depth* is the idea that defenses should have more than a single layer of protection between an attacker and the protected systems, data, or networks. Defense in depth in cyberwar is much like defense in depth in conventional military operations. It employs layers that use different methods to stop attackers so that a single attack or technology cannot succeed simply by penetrating a single system or layer of protection. In addition, it offers real advantages to those who are defending, as they can use simpler, easier-to-understand, and sometimes less-expensive defenses in each layer. Network defense in depth often starts with a strong design that involves network security devices. These include firewalls, intrusion detection and prevention systems, antivirus, authentication, logging, response, and restoration capabilities. Network defense in depth can also include the policies, procedures, training, and knowledge of the staff who use and support computer networks and systems.

This chapter looks at how modern computer networks and systems implement defense in depth by using a variety of strategies and technologies. It explores U.S. Department of Defense and National Security Agency strategies and concepts, as well as civilian know-how regarding the way in which people, technology, and operations influence defense strategies. You'll also learn where defense in depth can fail and why some experts have begun to claim that defense in depth is no longer the strong cyberwarfare defense strategy it once was.

Chapter 9 Topics

This chapter covers the following topics and concepts:

- What defense in depth is
- What the defense-in-depth strategies and concepts are
- Where and why defense in depth fails
- What the design elements of a modern defense-in-depth strategy are

Chapter 9 Goals

When you complete this chapter, you will be able to:


- Describe defense in depth and why it is important
- Explain common elements of defense-in-depth strategies
- Describe how and why defense in depth can fail
- Explain the concept of dynamic defense
- Describe common elements in a modern defense-in-depth design

Defense in Depth

From ancient Roman fortifications to medieval castles, the concept of providing defense in depth by layering protective capabilities has been in use for thousands of years. The earliest motte and bailey fortifications used by the Norman invaders in England in the eleventh century are recognizable as the predecessors of the mighty medieval castles you are probably familiar with. This design layered ditches, mounds of earth, and wooden palisades around a central multistory, defensible house (keep). The Normans ruled the recently conquered countryside from these very early castles. They relied on the multiple layers to keep them safe even if attackers successfully crossed the ditch and burned the palisade down.

Over the next 200 years, those early fortifications evolved as technology and strategies for attacking castles changed. Castles became increasingly more complex as attackers became more organized and the technologies used to attack them became more effective. Stone replaced wood to avoid fire, and layers of defenses became deeper and stronger to combat larger, more organized armies. By the thirteenth century, concentric castles like those shown in Figure 9-1 had layers of stone walls, strong towers, and heavily fortified gatehouses with drawbridges, strong doors, strong internal gates that could divide invading groups, and a myriad of ways to attack enemies trapped inside. These concentric castles are a common sight when describing defense in depth because they so clearly show the layered defenses available to a medieval lord, and thus are a useful metaphor for how to layer modern defenses.

The weapons and strategies used in warfare have never stood still for long, and changes were already beginning to occur even as these mighty stone castles were being built. By the middle of the fifteenth century in Europe, cannons and gunpowder had begun to change the balance of power in warfare. Traditional castles, keeps, and city walls with their tall stone construction were particularly vulnerable to this new form of warfare.

**FIGURE 9-1**

Concentric castles provided defense in depth using stone walls, moats, gates, and terrain features like hilltops and raised earthen mounds. Note the layered walls, strongpoints near entrances, and narrow pathway to the castle.

For example, a cannon-equipped army could reduce the mighty fortifications to rubble from a distance.

Designers realized this, and they developed an updated castle design that specifically addressed the new world of cannon and siege warfare. They recognized that traditional defenses were no longer relevant, and that a new type of layered defense was necessary. Their fortification style, known as star forts (see Figure 9-2), changed how fortifications were designed and remained in use until the nineteenth century.

The constant change in both the weapons and technologies attackers use, and the ways in which defenders attempt to counter them, is the same challenge faced by information system defenders today. In fact, the experts assigned to defend modern computer networks and individual computers have often adopted similar strategies for the same reasons that fortress builders and defenders have throughout history: Enemies can often breach one layer of the defense. Layered defenses make it less likely that a single attack can completely compromise a network or system. They also allow for weaknesses and mistakes on the parts of both defenders and those who provide defenders' software, hardware, and devices.

Modern defense-in-depth strategies still use layers, but the layers are no longer stone walls and ditches. Figure 9-3 shows the U.S. Computer Emergency Readiness Team's recommended practice for defense in depth against vulnerabilities like a buffer overflow attack. Here, the strong outer layer relies on humans who are trained and know policies that will help prevent behaviors that allow attacks to succeed. Successive layers implement a combination of technologies, as well as human knowledge, to prevent and detect attacks.




FIGURE 9-2

Star forts like the Italian fort in Nicosia, Cyprus, marked a major change in defensive strategies due to technological change. Note the multiple layers of low angled walls to defeat cannon fire, the dry moat and ditches to prevent foot-soldier assaults, and the angled projections that allowed defenders to fire sideways at attackers.

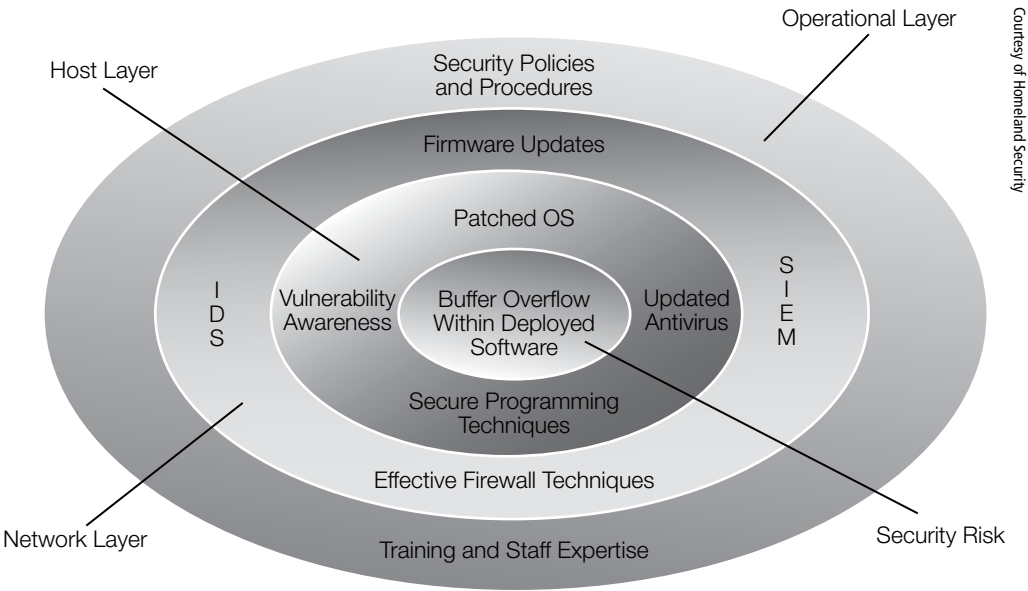



FIGURE 9-3

The US-CERT's defense-in-depth strategy for protecting individual systems (hosts) against a sample attack layers expert staff, firewalls, detection and monitoring, antivirus, and patching to prevent attacks.

FIGURE 9-4

The C-I-A triad shows the interaction between confidentiality, integrity, and availability when handling data and services.



The C-I-A Triad

The NSA uses a common information security conceptual model known as the C-I-A triad as part of its design. The **C-I-A triad** consists of confidentiality, integrity, and availability, as shown in Figure 9-4. It is often a key part of defense-in-depth designs, as well as throughout information security and cyberwarfare theory and practice.

The components of the C-I-A triad are:

- **Confidentiality** ensures that information is not accessible or disclosed to unauthorized systems or individuals.
- **Integrity** ensures that information has not been modified by unauthorized users or systems, and remains accurate and consistent.
- **Availability** ensures the system, data, network, or service is available and can be used or accessed.

The NSA and other information security practitioners also commonly add *authentication* or **authenticity**, which is the ability to validate that the system or user is who he or she claims to be, and **nonrepudiation**, which means that the sender cannot claim not to have sent the data or messages received.

Some sources refer to this as the A-I-C triad to distinguish it from the U.S. Central Intelligence Agency (CIA).

technical TIP

A *buffer overflow attack* attempts to overfill a memory location in a program. This causes the program or server to fail, or, in some cases, allows attackers to cause the system to run their program instead of the program it should be running.

Defense-in-Depth Strategies

Many groups and organizations have published defense-in-depth strategies. All have emphasized elements that are specific to their organizational goals, the technologies that they rely on, and the attackers and attack techniques they expect to face. The following sections examine publicly available defense-in-depth strategies from the U.S. National Security Agency (NSA), the NSA's Information Assurance Division, and the Department of Defense. You'll also examine how elements of the U.S. government's defensive strategies and priorities mirror those of the SANS Critical Controls, a popular list of network security design considerations that the business world often uses in network security implementations.

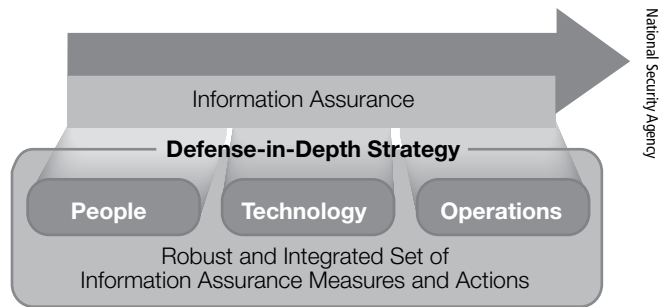
The NSA People, Technology, and Operations Defense Strategy

The National Security Agency's Information Assurance–based defense-in-depth strategy is based on the idea that people, technology, and operational security must be provided to ensure end-to-end defense. The NSA points to availability, confidentiality, integrity, authentication, and nonrepudiation services as key parts of the ability to protect against, detect, react to, and recover from attacks.

As with most high-level conceptual security designs, the NSA's Information Assurance strategy model is typically explained with a simple diagram. (See Figure 9-5.) Note the emphasis on robust and integrated measures and actions. The NSA realizes the importance of a strong defensive design, with multiple supporting layers that cover both the individuals who use technology; the technology itself; and how the daily operations that support, monitor, and maintain them work. Although the diagram looks simple, the underlying implementation can be quite complex, as you'll see in looking at how these elements interact.

FIGURE 9-5

The NSA's Information Assurance and defense-in-depth conceptual model combines people, technology, and operations into a defense-in-depth strategy in support of information assurance.



People

The NSA's people-based strategy relies on hiring talented staff, training and rewarding them, and penalizing unauthorized and unacceptable behavior. To do this, the NSA built a framework that includes policies and procedures, training and awareness, system administration, physical security, personnel security, and facilities countermeasures. Combining these elements provides depth by ensuring that (1) the staff knows the right thing to do based on policies and procedures, (2) they know how to do it because of their training, and (3) they are in an environment that helps to enforce those requirements with effective system administration and physical security. These elements also help ensure security through personnel security. To do this, they use background checks and other review of their staff, and then ensure that staff work facilities can provide appropriate levels of security, oversight, and separation.

Technology

The technology portion of the NSA's recommendations focuses on how technology is designed, acquired, configured, managed, and maintained. Technology focus areas emphasize the need to defend in multiple places at once, including:

- Defending the network and infrastructure
- Defending the enclave boundary
- Defending the computing environment
- Supporting infrastructure like key management, public key infrastructure (PKI), detection, and response

The NSA's defense-in-depth strategy uses layered defenses that work together to ensure that the failure of one layer of protection will not expose the data, service, or system to attack. The NSA emphasizes the need for each layer of defense to create unique barriers to access, so that a single attack cannot bypass multiple layers simultaneously. It also focuses on the need to detect and respond when a breach of a defense layer occurs.

technical TIP

Key management and public key infrastructure are parts of an encryption strategy. Keys are the part of a cipher that is used with the algorithm to specify how the cipher's encoding will transform the original unencrypted data. Public key infrastructure is the set of systems and software that make public key encryption work. PKI includes a certificate authority that issues and verifies certificates, registration authorities that verify whether the entity requesting a certificate is valid, directories of certificates, and the certificate-management system itself.

Operations

The operations leg of the three-part Information Assurance defense-in-depth strategy focuses on:

- Security policy
- Certification and accreditation
- Security management
- Key management
- Readiness assessments
- Attack sensing
- Warning
- Response
- Recovery and reconstitution

In essence, this is where the daily activities of the defense-in-depth strategy occur. Elements of this strategy include testing and validating configurations, systems, and software; ensuring that patching and updates occur; performing regular assessments; monitoring; and restoring normal functionality after a successful attack.

When taken together, the People, Technology, and Operations design philosophy reflects common practices for most mature information security operations. It should come as no surprise that the threats the NSA faces mirror those that are found elsewhere, even though the NSA may face them in the form of cyberwarfare activities.

FYI

The Information Assurance Directorate (IAD) of the National Security Agency provides standards and a technical framework at <https://www.iad.gov/iad/index.cfm>. The Common Criteria Protection Profiles provide both configuration and testing certification information for systems, software, and other products tested to meet the Common Criteria. The **Common Criteria** make up an international standard for computer security certification and testing. You can find them at <http://www.commoncriteriaportal.org/ppsl/>.

The National Security Agency Information Assurance Directorate

In addition to the Information Assurance plan discussed previously, the National Security Agency's Information Assurance Directorate (IAD) provides a brief, highly focused "Confidence in Cyberspace" guide intended to provide guidance on how to fight attacks throughout their life cycle. To do so, the IAD identifies four major goal areas:

- **Device integrity** helps to ensure that attackers have not modified or changed systems and devices. This includes ensuring that even difficult-to-detect attacks like those used by advanced persistent threats are not allowed to take over a device.
- **Damage containment** helps when a compromise or intrusion does occur. It has a goal of limiting the damage done from the loss or modification of data, retaining functionality, and ensuring that successful attacks don't lead to further compromises or damage.
- Defense of accounts ensures that credentials are not exposed or misused.
- Secure and available transport allows data to be sent and ensures that it isn't modified or accessed during transit.

These are obviously broad goal areas, and you can approach each of them in many ways. Given the scope of government activities and systems, nearly infinite combinations of threats and corresponding defenses exist. Thus, the IAD has defined a set of top strategies to fit these goals. They are:

1. Use **application whitelisting**—a technological solution that uses known, allowed programs to run on trusted systems. It also prohibits unknown programs not on the whitelist from running. Administrators must approve each program, preventing individual users from downloading and running new applications, plug-ins, and other software.
2. Control **administrative privileges** by limiting when administrator privileges are used and who has access to them. One goal is to ensure that users have the lowest level of privilege needed to perform their jobs. A second goal is to prevent attackers, a variety of technical and policy means, from being able to exploit administrative accounts.
3. Limit workstation-to-workstation communications by ensuring that workstations generally cannot monitor or send traffic to one other. This prevents attacks in which attackers compromise one workstation, and then use that workstation to compromise other workstations until they get the credentials or access they need to move up through the network.
4. Use antivirus (AV) file reputation services that use centralized antivirus company data to determine whether files and Web sites are malicious.
5. Implement antiexploitation features, which are built in to many modern operating systems. These features help prevent buffer overflow and other attacks that target vulnerabilities in the operating system.

6. Implement **host intrusion prevention system (HIPS)** rules. A HIPS is an intrusion prevention system that monitors traffic to and from a system and blocks known attacks, or attacks which match behavioral rules that the HIPS enforces. By enabling and configuring a HIPS, the system can attempt to defend itself by blocking traffic before the operating system or services are allowed to see the attack.
7. Set a secure **baseline configuration**. A wide variety of system security standards is available. These can be used to ensure that deployed systems match expected configurations, preventing the use of default or dangerous settings.
8. Use Web Domain Name System (DNS) reputation. This technology leverages commercial or private services that maintain information about whether a given **domain** or address can be trusted. The domain or address is trusted—or not trusted—based on behavior it has displayed and reports from others using the reputation service.
9. Take advantage of software improvements (patches) to ensure that known vulnerabilities and issues are fixed in a timely manner. Although zero-day vulnerabilities are often exploited before they are known and before a patch is available, keeping systems up to date with known patches prevents widely available exploits from being used against them.
10. Segregate networks and functions to ensure that the compromise of one section of the network or system does not allow compromise of other sections. Segments should be separated by role and function, thus ensuring that higher-security or more-valuable assets are not in the same segment as lower-security or less-valuable systems and data.

These strategies are broad, and allow a range of implementations based on their general concepts. For example, to meet the strategy to segregate networks and functions, an organization could choose to segregate networks in a variety of ways. It could use firewalls, deploy virtual networks to separate systems based on their function or trust levels, or choose complete physical separation. Thus, the IAD's list makes a lot of sense when applied to the many environments in which U.S. federal government computing exists. Flexibility in implementation helps to prevent blind spots created by centralization of requirements. This can be a real benefit. Of course, flexibility also means that some organizations may make poor or uninformed choices when designing and selecting their solutions.

It is interesting to note the relative comfort that the IAD appears to have with community or commercially sourced services like AV file reputation and DNS reputation services. In some high-security environments, third-party services have historically been viewed as creating more risk than they are worth. This list appears to acknowledge that commodity software and services are a fact of life for the organizations the IAD advises.

Comparing the SANS Top 20 and the NSA IAD Recommendations

Federal agencies aren’t the only groups that produce lists like this. In fact, one of the most frequently cited lists is the SANS Top 20 Critical Security Controls. As you might expect, those top controls match many on the IAD’s list. Table 9-1 compares the items on the IAD’s list with their matches in version 5 of the SANS list. The NSA’s Information Assurance Directorate’s top strategies map well to the SANS Top 20 Critical Controls, showing the common choices and recommendations for cyberwar and traditional cybersecurity defense strategies.

TABLE 9-1 Mapping the NSA strategies to the SANS Critical Controls.

IAD RECOMMENDATION	SANS TOP 20 CRITICAL CONTROLS
Administrative privilege control	12. Controlled use of administrative privileges 16. Account monitoring and control
Use AV file reputation services	5. Malware defenses
Implement HIPS	Not specifically mentioned
Use DNS reputation services	13. Boundary defense 11. Limitation and control of network ports, protocols, and services
Segregate networks and functions	19. Secure network engineering

The Department of Defense and Defensive Design

The U.S. Department of Defense (DoD) also provides public information about cybersecurity and DoD defensive design. Interestingly, DoD Instruction 8500.01E, the most recent DoD instruction addressing cybersecurity at the time of this writing, does not mention defense in depth in so many words. It does, however, provide a significant amount of information about how the DoD expects cybersecurity to be implemented. In addition, many of the same elements that are found throughout the NSA's People, Technology, and Operations strategy, the IAD's recommendations, and the SANS Top 20 can be found in the DoD's instruction.

Much like the designs you've already looked at, the DoD Instruction requires operational resilience, or availability. In fact, it requires:

- Information and services to be available to authorized users “whenever and wherever required according to mission needs...”
- Security posture to be monitored and available to those responsible for it at all times
- Technology components to be able to self-defend, reconfigure, and optimize with little or no human intervention

NOTE

You can find all publicly released DoD instructions on the DoD Issuances Web site at <http://www.dtic.mil/whs/directives>.

NOTE

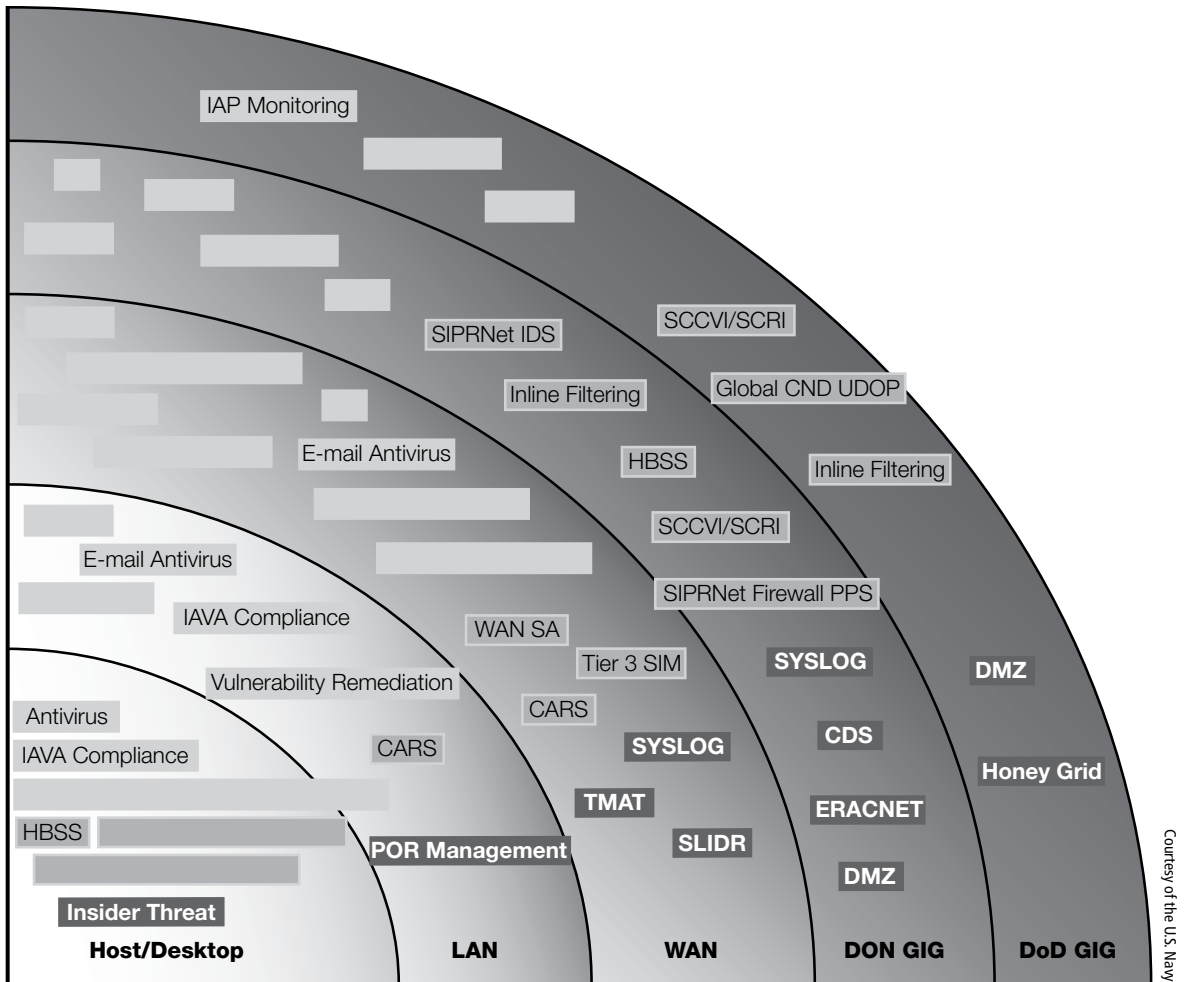
The term *self-defending network* implies that the network is able to respond to attacks by changing rules, modifying how it is configured, and otherwise responding to problems. You'll learn more about the concept of a self-defending network later in this chapter.

Computer Network Defense and Defense in Depth

At the beginning of this chapter, you learned that defense in depth is part of the defensive side of cyberwar operations known as computer network defense. To best understand where defense in depth fits in a complete CND strategy, you need to understand all of the elements that CND typically entails, and how current network defense strategies came about.

The U.S. Department of Defense and the U.S. National Security Agency's early Information Assurance designs commonly cited defense in depth. However, the term *computer network defense* was first used in the late 1990s. The Joint Task Force for Computer Network Defense was created in 1998 as part of U.S. Space Command (USSPACECOM). This creation, and its later growth to cover computer network attacks, is one of the first highly visible uses of the term CND.

Over time, the U.S. government grew CND into a complete discipline with dedicated cyberwarfare support organizations. Those organizations now implement a top-to-bottom CND strategy with elements similar to those shown in Figure 9-6, which depicts the U.S. Navy's CND strategy.

**FIGURE 9-6**

The Department of the Navy CND defense-in-depth strategy combines elements at the host, network, network edge, and policy layer.

NOTE

The Defense Information System Agency of the U.S. Department of Defense provides unclassified information on CND at <http://iase.disa.mil/policy-guidance/index.html#cnd>.

As you dig into current CND documentation like CJCSI 6510.01E, you will still find defense-in-depth language like “Implement a defense-in[-]depth strategy for ISs and supporting infrastructures through an incremental process of protecting critical assets or data first. The defense-in[-]depth strategy must establish protection and trust across various network layers (e.g., application, presentation, session, transport, network, data link, or physical).”

Where and Why Defense in Depth Fails

Defense in depth can fail, however, in a number of scenarios. Much like the castles that come to mind when you diagram defense-in-depth designs, changes in technology, flaws in design, and trusted insiders can all breach or betray even the strongest defenses. But attacks aren't the only reason that defense in depth can have problems. Some of the biggest problems with defense in depth result from tradeoffs it creates simply because of the way it must be implemented.

Recall the three main concepts from the C-I-A triad: confidentiality, integrity, and availability. If you think about these three concepts in the context of defense in depth, it quickly becomes obvious that the more you protect confidentiality, the harder it becomes to provide provable integrity. This is because proving that something has not been changed requires access to it so you can compare it with a known good version of itself. Providing access means that you have another component that must be trusted. That addition of trust makes it harder to have confidentiality. This simple issue can drive significant costs, or it can create unexpected holes in your layers of security as you increase either confidentiality or integrity at the expense of the other.

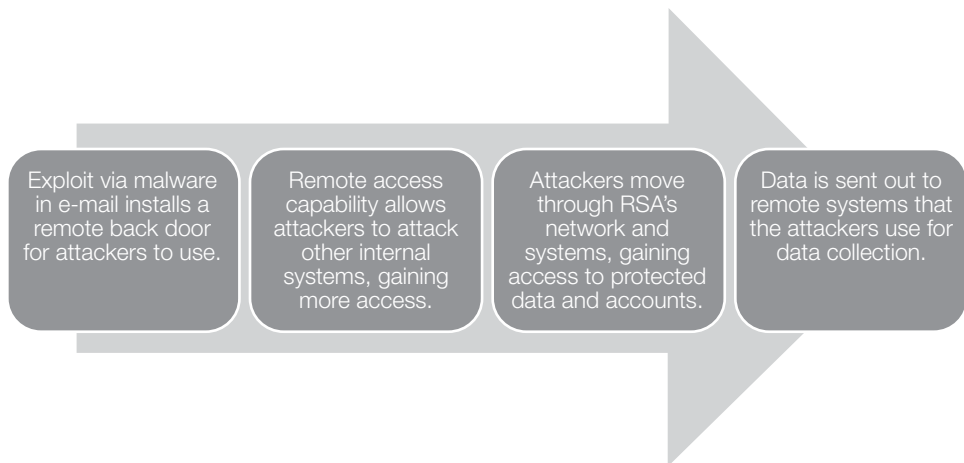
Availability can be similarly affected. As environments grow increasingly complex, the infrastructure and overhead required to maintain high levels of availability typically increases. Thus, as you increase the number of security layers, or increase their complexity, availability often suffers. Modern cloud solutions attempt to fix this by creating multiple regional data centers with hundreds or thousands of systems. However, they face the same issue: If you increase the number of places your data can be, you face the issue of securing it in all of those places at the same level of confidentiality and integrity.

If you have assessed your organization's threats and risks, and have spent your time and resources well, it's worth exploring the common places that defense in depth fails: neglecting layers, trusted attackers, human factors, and changes in technology.

Neglecting Layers: Getting Past the Shell

Much like a medieval castle attacked through its postern gate (a secondary entrance used for messengers and others who needed to come and go without being noticed), modern networks are vulnerable to attacks that are allowed to bypass their layered defenses. One of the most common mistakes made in defense-in-depth strategies occurs when defenders build strong outer layers of firewalls and network defenses, and then neglect the systems that inhabit the core of their defensive rings.

A great example of an attack that did exactly this occurred in 2011 when attackers targeted RSA, a major security company (and part of EMC, a large technology company). The RSA attackers crafted an e-mail and sent it to specific users in the company. The e-mail appeared to come from Beyond.com, a job-hunting site. The e-mail included an infected spreadsheet as an attachment, and that spreadsheet included an exploit that

**FIGURE 9-7**

The attack process at RSA started because users had the Flash plug-in installed, allowing attackers to use a zero-day exploit delivered via e-mail to compromise systems and to move through its network.

NOTE

Interestingly, none of the RSA employees who received the e-mail were high-profile employees, and the e-mail was actually flagged as spam. This crosses over into human factors: The employees who opened the e-mail had to retrieve it from their spam folder to open it!

targeted the popular Flash player. When the user opened it, the exploit compromised the user's system and turned it into a remote-controlled gateway into the RSA network.

Once attackers had an entryway into RSA's network, they moved laterally, targeting accounts and systems with greater privileges and access. Eventually, they found the data they were looking for. In the case of RSA, that data was key information about RSA's core business: one-time passwords, also known as two-factor authentication software and hardware. In cyberwar, it's a huge win to successfully target the technology used to secure an enemy's authentication systems that prevent simple passwords from being targeted. The attackers who went after RSA had the keys to the kingdom for major companies and military users around the world.

In Figure 9-7, a strong defensive layer of intrusion prevention systems, e-mail spam and antivirus filters, firewalls, monitoring systems, and patching existed at RSA. However, users were allowed to receive attachments from the outside world. Their systems used the Flash plug-in, and they may not have used application whitelisting, which would have helped prevent the malware from running. Once a user opened an Excel spreadsheet with its embedded Flash malware, his or her system was compromised and allowed access to other systems in that network segment and beyond. If there isn't sufficient internal defense, the attackers can **pivot**, attacking systems inside the heavily defended outer shell.

System Administrators: Trusted Attackers

System administrators provide another potential way in for attackers. Edward Snowden was a trusted contractor. Snowden and his collection of NSA data provide a case study of a privileged user leveraging his access to gather significant amounts of data without being detected. Snowden did not compromise NSA systems to allow outsiders in, but he still left with a massive amount of data about NSA programs. This cyberintelligence has been a huge coup for foreign intelligence agencies, but access to the systems he was trusted with would have been an even bigger win for them.

Publicly released information about Snowden's attack indicates that it was not an incredibly technical exploit, or one that compromised systems to succeed. Instead, he used a commonly available tool used to download the content of Web sites for later viewing, as shown in Figure 9-8. Using the tool from a trusted position in the network allowed him to gather huge amounts of data and explain the network traffic as part of his normal system administration job. Once he had the data, it was a relatively simple process to find a way to leave with that data, whether it was via a USB thumb drive or some other method. System administrators often carry around drives or send large files as part of their jobs, and this too would likely go unnoticed in many cases.

For those who want to have truly secure networks, it is an unfortunate fact that security systems require at least one trusted staff member to administrate them. This fact creates a potential problem with trusted attackers. It also creates a target for attackers who want the administrator's privileges. Thus, system administrators are a dual threat:

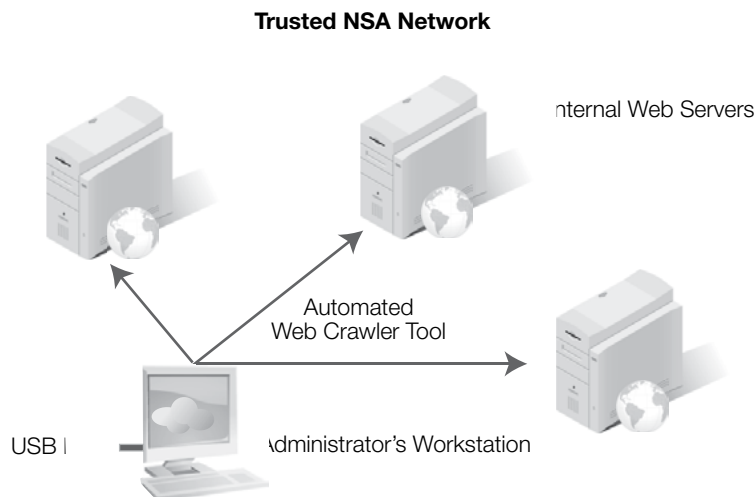


FIGURE 9-8

The trusted administrator's system has access to systems inside of a secure network. Because the administrator's system is expected to contact those machines, the Web crawler's data gathering is likely to go unnoticed, or at least be easily explained.

They could choose to become attackers themselves, and their accounts and privileges are targets for attackers, even if the administrators can be trusted.

In the defense-in-depth strategies discussed earlier in this chapter, system administrator privileges are one of the most heavily protected and monitored assets in any successful design. Knowing when those privileges are used, that the person using them is the authorized user, and what they are doing with those privileges is critical to preventing attacks and detecting any that succeed.

Attacking the User: Human Factors


Humans are in many ways the most vulnerable part of technological systems, and are the hardest to fix. Human instinct drives people to be helpful. Further, people are trained to assist others, to trust, and to sympathize. The problem is that attackers gladly use these natural instincts to their advantage. People tend to behave in ways that attackers can manipulate to access systems and data. If an attacker can predict how people will act, he or she can target the attack to use them to bypass security.

Using human factors to attack systems was an important part of the Stuxnet malware attack on Iranian nuclear facilities. The facilities used an air gap, or physical separation of systems that were network accessible from the control systems in the design of nuclear-enrichment facilities. This is a common feature in power plants and industrial facilities, as well as in high-security military operations where Internet-accessible systems could lead to fatalities or infrastructure failure. Air gaps are designed to prevent users who have access to the protected area from transferring materials from an untrusted area across the gap without taking appropriate security measures.

In Figure 9-9, you can see how a critical part of the Stuxnet attack leveraged expert technical staff. The staff members unwittingly carried malware on USB thumb drives that they used to transfer files from Internet-connected laptops to the protected network. Once the air gap had been crossed with human help, the Stuxnet malware could directly target the hardware it was looking for.

FIGURE 9-9

Targeted malware seeks to infect systems like those that frequently hired technical consultants to the Iranian nuclear program use. Once infected, those systems deliver malware via any USB drive plugged into them.



This type of attack takes careful planning and required the attackers to cast a broad net. The attackers had to infect many systems belonging to potential technical advisors and workers who might work with the Iranian facilities. The Stuxnet malware had to be able to bypass their antivirus software. In addition, the attackers had to rely on the workers using removable drives to transfer files. The sequence of events that Stuxnet relied on was likely to happen eventually, but setting up the chain of infections to get it to the right place at the right time was also incredibly complex.

Changes in Technology

When designing defense in depth, technological change is one of the hardest challenges to defend against. Last week's secure design might be an outdated relic this week if a technological breakthrough is made or if a critical flaw is found in an old technology. Even before the modern idea of cyberwar existed, technological progress was responsible for success in breaking cryptographic systems. The German World War II Enigma device, once considered unbreakable, drove the development of computerized code-breaking systems, which made Enigma and similar systems obsolete and led to entirely new types of encryption systems.

Modern advances in code-breaking technology, known as rainbow tables, have made the use of MD5 and SHA1 hashes for password storage similarly obsolete. Rainbow tables allow attackers to quickly retrieve passwords if the attacker knows the password's hash. Unfortunately, many Web sites store their passwords as MD5 or SHA1 hashes because both are commonly available in toolkits for developing Web applications. Thus, organizations that are unaware of the ease with which rainbow tables can allow password recovery have lost the race to protect their users' data. If organizations' storage for hashed passwords is breached, attackers frequently discover the actual passwords.

technical TIP

Hashes are cryptographic functions that take a block of data and perform operations on it to produce a fixed-length string of characters. If the same data is hashed, it will always return the same string. However, modifying the data will always change the hash. Hashes aren't intended to be reversible, which makes them a great way to compare data like passwords without exposing the data itself.

Rainbow tables are a prebuilt list of all the hashes that strings can hash to. Thus, if an attacker has the MD5 hash of a password that was created by simply hashing the password with no additional steps taken to protect it, the attacker can simply perform a lookup against the rainbow table and identify the original password.

Using rainbow tables to test against sample password lists available from previously hacked sites has resulted in exposing more than 50 percent—and in some cases more than 90 percent—of all the hashed passwords.

Unfortunately, defenders are often unaware of the technological changes that have made their previously secure technologies obsolete. Some defenders unwittingly make this mistake because the vulnerable technology is embedded in their infrastructure, leaving them vulnerable without their knowledge. Thus, attackers frequently hold the advantage, and defenders need to stay vigilant and update their technology before it reaches that point. Defense in depth does offer some remedies for this by layering technologies with different weaknesses, but that defensive measure usually comes at a high cost.

One of the biggest problems for defenders is that changes in technology are typically far easier for attackers to exploit than they are for defenders to update. Defenders often have existing infrastructure and businesses they cannot disrupt without creating losses. Thus, they have to balance security against changes in technology and may end up on the losing end of the battle. The star forts mentioned at the beginning of this chapter are a perfect historical example of this. They were created because cannons were capable of shattering the huge existing stone castles of the day, leaving the defenders unable to protect themselves. By the time that defenders built star forts in large numbers, technology had again advanced, and new techniques were in place to overcome them.

Designing a Modern CND Strategy

Modern computer network defense strategies must be designed to defeat both nation-state level capabilities and those of asymmetric actors like hackers, organized crime, and even corporate employees seeking to engage in corporate espionage. Thus, a CND strategy makes use of many traditional information security design elements layered with government- and military-specific procedures, policies, and technologies. The additional considerations these designs take into account reflect the needs of military operations, such as chain of command, classification, and resistance to electronic warfare (as opposed to cyberwarfare). The hardware, operating systems, and software they are deployed on may be significantly different from their civilian counterparts, or may simply be civilian versions configured to meet CND requirements. This part of the chapter explores the concepts of dynamic defense and some of the common elements of a secure network.

Dynamic Defense

A primary criticism of defense in depth is that adding layers makes using the defended assets more difficult. Each successive layer adds overhead to manage, monitor, and access the environment. Those layers can also cause failures if they themselves have problems. When organizations consider layer upon layer of fixed defenses, they often wonder if there is a way to have a reactive defense mechanism that uses situational knowledge, a range of capabilities, and an adaptive mechanism to handle unexpected events. Thus, the concept of *dynamic defense*, or defense that can change in reaction to threats and new risks, has become more popular. In fact, dynamic defense is often what organizations and authors are describing when they propose self-defending systems and networks that can react to threats and modify their defense schemes to meet them.

The idea of dynamic defense isn't new. J. R. Winkler, C. J. O'Shea, and M. C. Stokrp outlined the basic concept in a 1996 paper, noting that operational security analysis should allow dynamic defenses to:

- Add additional countermeasures that can reconfigure to handle new threats
- Implement interoperability between defensive systems, allowing them to coordinate their defenses
- Provide both centralized and decentralized components
- Use a database of operational security information to support the components of the defense system

The Self-Defending Network and Self-Defending Hosts

One increasingly popular concept for defense in depth is that of a *self-defending* network or host. These self-defending systems are intended to adapt to prevent attacks by monitoring systems, users, and network traffic. When they detect something that should not be occurring, or that does not fit their security posture, they block the traffic, disable the user or service, or otherwise take what is hopefully appropriate action. As with any automated system, a system that can modify itself to handle problems can also likely cause itself harm by responding inappropriately to those threats.

A few common elements show up in most literature about self-defending networks and hosts:

- They use cloud-hosted or Big Data information to monitor for attackers, known malicious sites, and malware.
- Software-defined networks can redefine their boundaries as needed, moving systems into much more granular enclaves than a traditional network design typically provides.
- User- and system-contextual awareness can grant permissions to users based on behavior and their role in the network at the time they use it. Rather than giving users all of their rights all of the time, this contextual awareness focuses on giving users only the rights they need for the time they need it.
- Direct, secure interconnection between endpoints allows end-to-end security of data without systems and networks between them being able to see or modify the traffic between them.
- System posturing relies on more than just the user or system connected to the network. Often systems are granted rights based on the logged-on account, or where in the network the host resides. In a fully self-defending network, the host would have a more complex posture based on an up-to-date profile of the system, the user's rights, and where the device is.

As you might imagine, self-defending networks and hosts are a challenge to pay for, build, configure, and maintain. In many cases, organizations are using some of these concepts and capabilities, but true end-to-end, self-defending networks that use all of these elements remain largely theoretical.

Despite almost two decades of conceptual development, well-integrated dynamic defenses largely remain a dream for most organizations. More-advanced security systems, however, have begun to implement these ideas within their own bounds. Some defense systems are designed to allow custom-built modules to control other systems, but those are rarely built using a shared control standard.

CND and Defense-in-Depth Design


System and network security design is an incredibly complex topic with a wide range of options and strategies available to defenders. Often the best place to start is with a strong policy and procedural background, as well as with awareness and training to support the design. Figure 9-10 shows a conceptual layered design that uses many of the elements and concepts discussed in this chapter, built on a foundation of policy and procedure. Keep this design in mind as you review the common elements of modern security designs.

Risk and Threats

CND designs must respond to the threats and risks they face. Standardized designs that do not address the realities of the attackers and defenders and the capabilities of both will fail quickly. Modern designs typically start by identifying the risks to the organization through some form of risk assessment methodology, which also identifies the threats and actors that the organization will face.

FIGURE 9-10

A sample CND design starts with policies and procedures. Physical security then follows the layers of the technical environment from the data resident on the hosts to the software and applications, hosts, network, and the border of the network.



FYI

Risks are defined as the potential that a threat or actor will exploit vulnerabilities, causing harm to an organization. There are many different risk assessment methods, but the key element of each is that they identify what could happen, what the result would be, how likely it is, and how much harm it would cause.

Threats are possible dangers that could cause a security breach or exploit resulting in harm. Understanding the threats an organization faces helps to inform risk assessments and security design.

Once the risks the organization faces have been assessed and prioritized, best practices and standardized design concepts like those discussed earlier in this chapter are applied to counter them. Layered defenses are typically designed where possible to ensure that a single failure cannot bring down the organization's entire security architecture.

Secure Networks

The technology elements of computer network defense rely on having a secure network, including the systems on the network, the network devices, and the technologies used to connect them. Although this is a broad topic, this section briefly explains a few of the key technologies and concepts associated with current network security designs, including compartmentalization, defensive technologies, monitoring, cryptography, defense against malware, endpoint security, and physical security.

Network Enclaves and Properties

The concept of protected network enclaves has increased in popularity as networks have become increasingly complex, and the uses for them have expanded. A **network enclave** is a logically or physically separated portion of a network that isolates systems and devices from others based on rules, identity, location, or purpose. Enclaves are a particularly useful concept when organizations need to separate sensitive data or systems from other parts of their infrastructure.

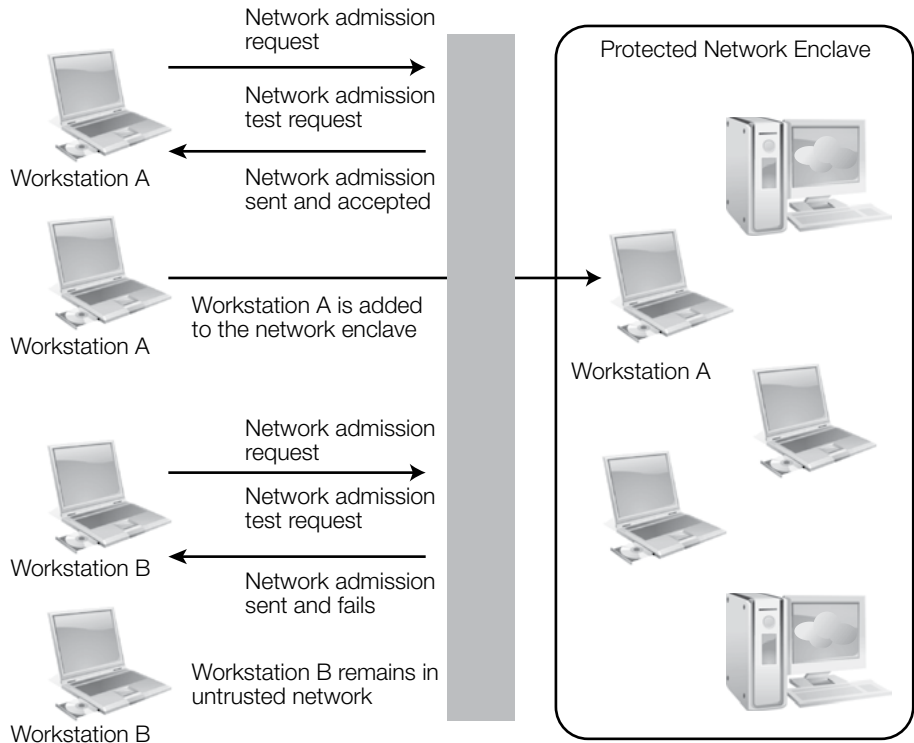
FYI

The U.S. Department of Defense describes *enclaves* as follows:

Enclaves provide standard cybersecurity, such as boundary defense, incident detection and response, and key management, and deliver common applications, such as office automation and electronic mail. Enclaves may be specific to an organization or a mission, and the computing environments may be organized by physical proximity or by function independent of location. Examples of enclaves include local area networks and the applications they host, backbone networks, and data processing centers.

FIGURE 9-11

A sample network enclave admission process.



Network enclaves are often built using internal firewalls using virtual network segmentation—a **virtual LAN (VLAN)**—or using a **virtual private network (VPN)** that allows creation of virtual networks of systems. Network enclaves often use a technology known as **network admissions control (NAC)** that requires systems to authenticate and verify that they match required settings and policies before they can join a protected network.

In Figure 9-11, Workstation A attempts to enter the protected network enclave by authenticating and performing tests required by the NAC system. It succeeds and is added to the protected network. Workstation B performs the same process, but is not properly configured and fails the test, leaving it in an untrusted network.

Network Defense Technologies and Devices. Technical network defenses are typically selected based on the security requirements of the network. The defense-in-depth designs explored in this chapter combine many different network security technologies. Some of the most common defensive devices and systems are as follows:

- **Firewalls**—These are network security devices that allow or deny traffic based on rules. The rules can specify ports and protocols. More advanced systems apply additional filters intended to detect attacks and abuse of systems.

- **Intrusion detection systems (IDSs) and intrusion prevention systems (IPs)**—Along with their local system sibling the host intrusion prevention system (HIPS), IDSs, and IPs are designed to detect and/or prevent attacks. They rely on either behavior-based detection, which looks for attacks based on behavior, or signature-based detection, which uses a fingerprint of known attacks.
- **Proxies and gateways**—These allow access either to or from secure networks, providing the ability to verify users and to filter what they send or receive.
- **Authentication and authorization systems**—These verify user identity and the permissions they are granted.
- **Virtual private networks (VPNs)**—VPNs provide encrypted secure networks virtually through existing networks.
- **Security Information and Event Management (SIEM) systems**—With their siblings, Security Information Management (SIM) and Security Event Management (SEM) systems, SIEM systems combine event logs, network traffic data, and other security and log information to detect and take action on security events, attacks, and even service issues.

Monitoring

Networks, devices, and software all typically offer some method of providing status information in the form of logs. In addition to logs, you can add additional instrumentation to validate and monitor what is occurring. This can help provide a layer of defense against both attackers and system errors.

Monitoring inside of a CND environment typically takes on a number of forms, including:

- Network traffic monitoring, such as network flows, bandwidth utilization, and other information about network traffic levels and types.
- System monitoring, such as error logs, administrative and privileged user access logs, authentication logs, and a host of other details about what systems are doing and what problems they encounter.
- Software and application logs, which provide information about how software and applications are performing and whether they have encountered errors.
- A host of other items. Almost anything that happens on a computer or a network can be monitored.

Due to the broad variety of potentially monitored items, many organizations use automated log analysis and SIEM systems. It can be difficult to balance the desire to gather logs in case they are needed with the sheer amount of data that in-depth logging can create.

technical TIP

Network flows are a useful type of network monitoring, and can be thought of like a traditional phone bill's call log: They show which systems communicated, via which port and protocol, and how much traffic they sent. This can be a treasure trove of useful information when you are looking for a system compromise or attempting to identify attackers.

Cryptography

The ability to securely store and transmit data is important for computer network defense. Encrypting data in transit can ensure that attackers cannot capture what is being sent. That same data must be protected when it is at rest, waiting to be accessed on a hard drive or other storage device. Thus, encryption capabilities and their implementation are part of the CND plan:

- In-place encryption is often based on full-drive encryption, or encryption of specific files or data stores that are sensitive. In-place encryption is often stronger than in-transit encryption because the data is not transient, and thus must resist potential long-term attacks.
- In-transit encryption includes the encryption used to protect and encapsulate virtual private networks, allowing them to operate inside of and through other networks without having their data exposed. It also includes the encryption used for secure Web traffic (SSL/TLS) and other technologies that help to ensure that network and other transmissions remain secure. Because in-transit encryption is typically short-lived, the encryption used is often weaker than that found in in-place encryption systems.

Smart attackers look for moments when data is exposed and target locations where they can capture that data. Often, this is when the data is transmitted or is momentarily unencrypted for access, such as when a credit card or password has to be processed by the system receiving it.

Defense Against Malware

Malware is a major concern for organizations—particularly state-sponsored threats, such as the very complex and well-supported advanced persistent threats that have appeared in the past few years. Malware can provide access into your network, can allow control of your systems, and targeted malware like Flame and Stuxnet can attack specific systems and capabilities for cyberwarfare purposes. Thus, protection against malware is a major part of defense in depth.

Anti-malware packages provide a number of noteworthy abilities that should be considered when designing a defensive strategy:

- Whitelisting capabilities allow only trusted software to run, and whitelisting software checks to ensure that the software running matches a known good version of the software. Unfortunately, whitelisting also requires a lot of work to verify each package every time it is updated, and limits flexibility by requiring every piece of software to be added to the list.
- Heuristic or behavior-based anti-malware capabilities look for behaviors associated with attackers, such as unexpected data transfers, scans of other systems, or access to memory or data that is atypical.
- Signature-based systems attempt to gather fingerprints of known malware and then compare files and applications with those fingerprints. Signature-based detection is becoming increasingly difficult as malware packages frequently change themselves using a variety of techniques every time they are copied or infect a new system. Thus, signatures don't match, and anti-malware vendors must track millions of similar packages.

Anti-malware technologies can be deployed on networks and on systems and devices. Unfortunately, modern malware concealment techniques can make advanced malware nearly impossible to detect. It has been shown that advanced persistent threat actors have been able to exploit the detection of their malware infections to upgrade them, using what they learn to defeat the tools that detected them. This results in an ongoing game of cat and mouse that pits defenders against attackers.

Endpoint Security Design

Security design and implementation for the endpoint is sometimes one of the hardest elements of a defense-in-depth plan. Endpoints are typically laptops, desktops, tablets, mobile devices, and other individual systems. They can also be network devices, cameras, and other parts of the technology infrastructure.

Endpoint security design is typically based on a few important elements:

- Securing of the operating system or software that controls the device's basic functionality
- User accounts and privileges
- Control of the software or other key functions of the device or system
- Configuration baselines that ensure that the system matches expected settings
- Monitoring and exception reporting
- Data protection, often in the form of encryption, backups, and integrity monitoring to ensure that data is not modified without permission

Endpoint security can be incredibly complex. Modern operating systems have hundreds and sometimes thousands of possible security settings when loaded with applications and services. This is where configuration baselines come in to help organizations ensure that they have common settings properly in place.

Selecting a Configuration Baseline

One of the key parts of a defense-in-depth strategy is baseline security configuration of systems and devices. This can be a complex task, as modern operating systems often have hundreds of possible security and configuration options. Organizations that are attempting to build a security standard are faced with a massive amount of work to design and test a security baseline for each of the operating systems or devices they intend to deploy.

Fortunately, a variety of configuration baselines are available to start from, and the best of the baselines provide useful commentary and information about the settings themselves. This allows security administrators and system administrators to make informed decisions about which settings they may need to modify due to the requirements of their business or organization.

Commonly used baseline standards are available from a variety of sources, including:

- Microsoft at <http://technet.microsoft.com/en-us/library/cc677002.aspx>
- Apple at <http://www.apple.com/support/security/guides/>
- The National Security Agency at http://www.nsa.gov/ia/mitigation_guidance/security_configuration_guides/operating_systems.shtml
- The Center for Internet Security at <https://benchmarks.cisecurity.org/>
- The Defense Information Systems Agency (DISA) Security Technical Implementation Guides at <http://iase.disa.mil/stigs/>

Important decisions that your organization might face when adopting a baseline include whether you should create your own baseline or modify one or more of these, which changes you should make, and how you will manage systems to ensure that they meet the baseline standard.

Unfortunately, vendors often lag behind their own releases of new operating systems, devices, and software. It is not uncommon to have a new operating system available for months or even years before a baseline is released. This leaves administrators scrambling to build a new baseline using the knowledge they have from older versions.

Physical Security

Although it might seem odd in the context of CND, physical security is actually one of the most important parts of defending a computer network and its systems. If an attacker has physical control of a system, he or she can almost always eventually compromise it in some way. Thus, a well-designed defense-in-depth plan must account for how computers, network devices, and the physical network itself are protected.

Physical security designs typically start with an understanding of the environment to be protected and how that environment will be used. You can use the same methods you used to assess risks and threats for the computer-based environment to assess physical

Travel Danger

The authors of this book have spoken with corporate security officers who have discovered physical bugs placed in laptops returning from overseas. The security staff said that they now issue disposable laptops that are physically disassembled and inspected when they return from overseas, rather than allowing travelers to take their daily work systems with them. They also provide a clean operating system and limit remote access to corporate networks for the systems they send overseas so that corporate data cannot be copied off the laptops while they are out of the country.

threats and risks. Unlike network defenses, physical security must also take into account nature itself, as the availability portion of physical security can be greatly harmed by events like earthquakes, hurricanes, floods, and tornadoes.

Physical security efforts also have to include policy and procedure. A heavily secured location that protects systems, but that allows travelers with laptops to have remote access can find itself leaking data through theft or because the laptop was physically bugged while it was out of the traveler's hands.



CHAPTER SUMMARY

This chapter examined the concept of defense in depth in cyberwarfare. Defense in depth is a key component of published cyberwar computer network defense strategies provided by the U.S. National Security Agency, Department of Defense, and civilian organizations that specialize in information and network security. These defense-in-depth strategies focus on people, operations, and technological strategies that blend how networks and systems are designed, built, managed, and monitored, as well as the handling of responses to and recovery from cyberattacks.

Although a key concept in CND strategy, defense in depth can create problems. In fact, defense in depth involves tradeoffs between confidentiality, integrity, and availability, the three common precepts in information security designs. Defense in depth is also vulnerable to design issues, trusted administrators and users acting in bad faith, and technological change. These issues don't mean that defense in depth doesn't make sense. But they do require you to think clearly about which defenses you use, and both where and why you use added layers—because these have costs and potential flaws themselves.



KEY CONCEPTS AND TERMS

Administrative privileges	Damage containment	Network enclave
Application whitelisting	Device integrity	Nonrepudiation
Authenticity	Domain	Pivot
Availability	Host intrusion prevention system (HIPS)	Virtual LAN (VLAN)
Baseline configuration	Integrity	Virtual private network (VPN)
C-I-A triad	Network admissions control (NAC)	
Common Criteria		
Confidentiality		