© Rodolfo Clix/Dreamstime.com

CHAPTER

12

Defending Networks

COMPUTER NETWORKS ARE A CRITICAL PART of national infrastructure. They provide communications between computers and other devices around the world, allowing command and control, infrastructure monitoring and management, and all of the other parts of modern business, government, and society to transfer data effectively. Of course, this makes computer networks one of the most important targets in cyberwarfare, as attackers can use networks to access data, to disable systems and infrastructure, and to attack vulnerable endpoints and other networks. As defenders add wireless and mobile network technologies to network designs, the complexity of network defense increases. Complexity also increases as networked devices layer additional ways to connect to multiple networks, adding additional pathways to networks that defenders may believe are well defended through their expected entry and exit points.

Many network attacks focus on gaining access to the systems and devices that reside on that network. To get to those endpoints, attackers must make their way past the protective measures that defenders put into place. Defenders have a broad range of options to help endpoints defend themselves, ranging from network firewalls to systems that can detect and stop attacks as they occur. Network administrators and security specialists must carefully design defenses to balance the ability to defend a network against the operational needs and requirements of the network's users.

Attacks against systems are not the only challenge that networks can face. Denial of service attacks, kinetic attacks against critical network infrastructure, sabotage, or even specialized attacks against network devices are all possibilities in cyberwar. The U.S. Department of Defense (DoD) pays particular attention to the concept of *mission assurance*—the ability to provide continuous operations despite attacks, system failures, or other disruptions. Mission assurance requires additional effort in design, maintenance, and management of computer and other networks, including specialized network operation procedures designed to ensure the security and availability of critical defense networks.

The U.S. National Security Agency (NSA) also recommends a defensive posture using the Protect, Detect, React paradigm for networks based on its information assurance defense-in-depth strategy. The NSA's defense-in-depth strategy includes an emphasis on preserving the confidentiality and availability of data that crosses networks and the need to separate portions of networks based on the roles of users and the requirements of the data that is transmitted across them. The emphasis that both the U.S. DoD and the NSA put on network defense demonstrates the importance of defending networks to national security and the defense of national assets in cyberwar.

Network defenders must consider both existing frameworks, like those used by DoD and the NSA, and the actual attacks that occur both in cyberwar activities and in traditional computer and network attacks and exploits. The same technologies and designs used for network security are frequently found in nation-state network implementations. The defenses deployed to protect military platforms and weapons systems can combine network technologies found in the civilian world with customized or unique platforms or systems. Or they may be entirely distinct from the software and hardware available to the open market. This creates unique problems for network defenders in cyberwar, who cannot always rely on the existing body of knowledge and technology created to protect commercially available networked devices.

The network border is often the first line of defense in cyberwar. An effective network defense strategy, as well as an understanding of the technologies and concepts typically used in a well-designed network, is critical to a successful defense. Defense in depth combining network devices, procedures, policies, and standards is a key tool in computer network defense strategies.

Chapter 12 Topics

This chapter covers the following topics and concepts:

- What designing defense in depth into computer network defense entails
- What mission assurance and information assurance in network defense entails
- What critical concepts in network security design are
- What technologies commonly used for network defense are
- What active defense in network security is

Chapter 12 Goals

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

- Describe network defense in depth
- Describe the Protect, Detect, and React paradigm
- Explain mission assurance and information assurance as used by the U.S. DoD in network defense design and implementation
- Describe commonly used network defense technologies like firewalls, IPS, and SIEM devices

Network Defense in Depth

Effective defenses for networks typically require multiple layers of technologies, processes, procedures, awareness, and training, which operate in layers to provide defense in depth. Network defense in depth must focus on attacks from outside a network; attacks from trusted systems allowed to connect to the network or which may be compromised and then used to attack the network from the inside; and attacks from coworkers or trusted network or system administrators.

Layered defenses in a network design allow defenders to deal with a breach in one or more layers effectively by using other defenses to prevent the attack from making further progress. It allows them to detect the attack and react appropriately. In computer network defense activities, layered defenses are an expected part of almost any network design. They ensure that the network can survive attacks from nation-state-level adversaries who can bring significant resources to bear when they want to bypass network defenses.

FYI

The Stuxnet attack offers an excellent example of how a determined and skilled group of attackers can bypass strong network defenses. The Iranian nuclear facilities used an air gap, or physical separation between their computer networks, to help protect their production facilities. Physical separation of networks is usually considered one of the strongest protections defenders can put in place. In the case of the Stuxnet attack, the attack succeeded because of Iranian reliance on contracted engineers who used USB flash drives to transfer data between unprotected systems and the separated nuclear-facilities production network. The attack's success helps demonstrate the need for both technical and human factors in defense in depth. It also illustrates the need for a complete vision and ongoing testing when defending networks.

Cyberwarfare defenders may need to consider more than just the typical computer networks that most people envision when they hear the word *network*. In fact, cyberwarfare defenders must consider many types of networks when they design defenses. Among those networks, four major types of networks are most commonly found in many organizations. Their distinct needs can shape the choices defenders make. Those types are:

- **Local area networks (LANs)**. Defenders typically can control LANs relatively effectively, allowing them to choose technologies, implementation, practices, and policies to best suit the needs of their organizations.
- **Wide area networks (WANs)** between local networks. Wide area networks often use third-party service providers to interconnect local networks, resulting in risks based on the abilities of those service providers to ensure reliable, secure service.
- The Internet, which is a decentralized, open network that is not designed to be inherently secure. The Internet is a network of networks, and provides both an immense variety of options for organizations that want to communicate and just as many possibilities for attack.
- Proprietary networks, which are based on specialized technologies or which have some other distinguishing factor that makes them separate from other, more traditional networks.

This chapter explores both technologies and strategies for defending each type of network, and uses examples of cyberattacks that can provide insight into defensive strategies that have worked and those that have failed.

Network defense of potential cyberwar targets creates significant challenges for military and government defenders due to the open nature of the Internet. Utility companies, Internet service providers, and the supply chain for military and civilian operations all rely on the Internet for their day-to-day business. This means that defending the entire Internet is a concern for defenders.

FYI

One question that is frequently asked when trying to build a defensible network is "Who owns the Internet?" In the context of cyberwarfare activities like attack and defense, the question often becomes "Who is responsible for the traffic on the Internet?" Because the Internet is intentionally designed to be decentralized and difficult for any country or corporation to shut down, the Internet is not the sole responsibility of any single organization. For many years, the United States held authority over the root zone file, which contained all of the names and addresses of the top-level domains for the Internet, providing it with greater power than other countries. Even that once centralized point of control is changing, as the United States announced in March 2014 that control will be handed over to the Internet Corporation for Assigned Names and Numbers (ICANN).

Unlike smaller countries that can easily identify a limited number of primary Internet connections to their national electronic infrastructure, large, first-world countries often have hundreds, if not thousands, of major connections to the Internet. Tactics like those used by Syria in 2013, Iran in 2012, and Egypt in 2011, when they shut down or strictly limited their Internet connections, are more difficult for heavily connected countries that have significant national business, security, and infrastructure requirements that make it necessary for them to stay connected to the Internet.

NOTE

The U.S. Department of Defense claimed to have 7 million networked devices and thousands of network enclaves in the February 2014 testimony to the Senate Committee on Armed Services.

The United States and other major players in cyberwar are generally heavily connected with significant infrastructure that requires Internet access to operate successfully and to leverage their own cyberwar capabilities. The major Internet powers are therefore investing in both military and government programs to foster stronger, more secure networks through cooperation with corporate and other organizations through information sharing, standards, and training. At the same time, their intelligence agencies and computer network attack teams are actively looking for vulnerabilities in computer networks, creating situations where national interests may create a desire for both secure networks, and to keep known issues secret so they can be used by the nation that discovered them for intelligence or cyberwarfare activities.

The March 2014 discovery of the Heartbleed bug in the popular OpenSSL encryption package, which provides secure connections for many Web sites on the Internet, provided an excellent example of a bug that may have been known and exploited by a nation-state prior to its discovery by the rest of the world. News reports indicate that the NSA may have been aware of and using the Heartbleed bug to view unencrypted Web traffic or to capture server security certificates for up to two years before the bug was patched.

Protect, Detect, React

The U.S. government provides many of the most accessible strategy and implementation documents available describing computer network defense. The NSA publishes strategy and technology implementation documents that are broadly adopted by organizations worldwide.

The technology portion of the NSA's People/Technology/Operations defense-in-depth strategy relies on a Protect, Defend, React paradigm. This requires the network to include protective mechanisms and abilities, as well as capabilities, processes, and technologies that can detect attacks. This assumption that attacks will occur and that they may succeed then drives the need for the ability to react. Thus, the NSA's requirement for defense in depth includes the ability to handle attacks and breaches, and not only to respond appropriately, but also to recover after the breach.

NOTE

The NSA's defense-in-depth strategy, found at *http://www.nsa.gov/ia/_files/support/defenseindepth.pdf* explains the Protect, Defend, React paradigm as part of an information assurance strategy.

FIGURE 12-1
The NSA's defense-in-depth strategy.

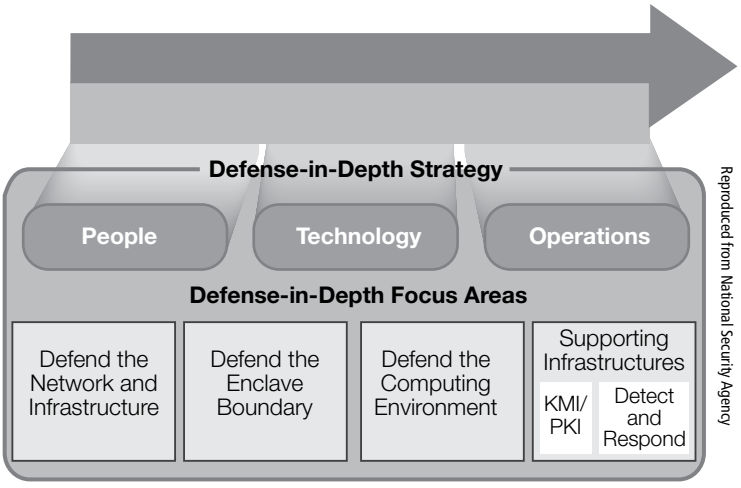

The technology for the NSA's defense-in-depth strategy, shown in Figure 12-1, focuses on network and infrastructure protection in three major areas:

- Protection of LANs and WANs against attacks like denial of service (DoS) attacks that attempt to prevent use of networks and systems on networks
- The requirement to protect both confidentiality and integrity, which ensures that data is not exposed and has not been modified when transmitted over networks
- The use of enclave boundaries or network segment protection by means of firewalls and intrusion detection systems to both limit traffic and detect attacks

The way that people are trained as well as the policies, standards, and processes they use all have an important role in the NSA's strategy—as do the operational elements such as log monitoring, certification, and day-to-day management. In fact, every part of the NSA's strategy is critical to an information assurance program.

Mission Assurance

In addition to the NSA's Protect, Defend, React paradigm, DoD uses a concept known as **mission assurance** as a key part of its defense-in-depth strategy. Mission assurance is defined as "a process to ensure that assigned tasks or duties can be performed in accordance with the intended purpose or plan ... to sustain ... operations throughout the continuum of operations." In short, mission assurance in cyberwarfare operations is the process by which systems and networks are kept online and functioning. DoD helps to ensure this by using a risk management process that assesses the functional requirements of networks and systems, then matches appropriate methods of control to them to ensure that the risks they face will not disrupt them.

DoD has steadily developed its mission assurance strategy over time, and in March 2014, it adopted a new standard as part of its overall **information assurance** strategy. The Risk Management Framework for DoD Information Technology is designed to provide mission assurance, and is aligned with the National Institute of Standards and Technology's (NIST) Risk Management Framework (RMF). The policy under which the RMF was adopted states that DoD will use enterprise-wide decision-making processes for cybersecurity risk management based on NIST standards and guidance.

DoD uses three mission assurance levels, known as **Mission Assurance Categories (MAC)**. They are:

- MAC I, composed of systems that are critical to the completion of DoD's mission, and which therefore must have the highest possible level of integrity and availability, as well as extremely strong security controls.
- MAC II, which includes systems that support deployed and backup forces. Although not as critical as MAC I, MAC II systems could have a significant impact on DoD's ability to complete its mission. Loss of availability is only allowable for very short time periods, whereas loss of integrity through the modification of data is not acceptable for MAC II systems. DoD notes that MAC II–level systems must exceed common industry practices for security to provide high levels of integrity and availability.
- MAC III, which includes systems that are used for normal operations but are not used for operational support of deployed units or for mission-critical operations. Thus, although MAC III systems are important, the exposure or modification of data on MAC III systems, or issues with their availability, could be dealt with in a reasonable way. MAC III systems are expected to use common industry standards for security to provide integrity and availability assurance.

NOTE

NIST's guidance on how to apply its Risk Management Framework can be found at *http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf*.

The NIST RMF provides a process for assessing risk to systems and data, then selecting, implementing, using, and monitoring appropriate controls. Figure 12-2 shows the NIST RMF process flow for risk assessment. In computer network defense operations, the categorization process typically includes assessments of the classification of data and systems, as well as the mission assurance requirements of the systems and networks involved. Categorization also typically looks at relevant laws, standards, and strategic requirements, as well as what resources are available. NIST provides guidance on classification of systems for U.S. government agencies as part of FIPS 199, a standard written in 2004 for security categorization of federal information and information systems.

DoD's relatively quick changes in risk management and information assurance strategies demonstrate the pace at which defensive strategies and technologies have changed since DoD's Information Assurance Certification and Accreditation Process (DIACAP) was first adopted under the NSA's guidance. The level of reliance on computer networks,

FIGURE 12-2
The NIST risk management process.

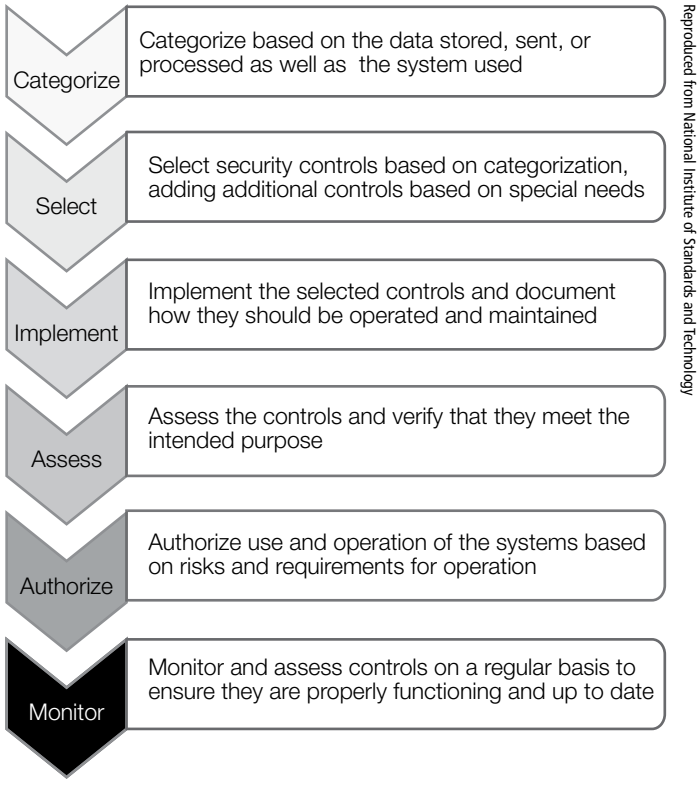

as well as their role in cyberwarfare and traditional kinetic warfare, has greatly increased in the seven years since DIACAP was fully approved, creating an ongoing need to update and replace existing standards. In addition to the changes made necessary by new technologies and threats, the move to DoD RMF for information technology (IT) also better aligns the department with other U.S. government agencies that have adopted the NIST Risk Management Framework for their own use. This allows shared training, easier auditing, and other advantages based on a shared knowledge base for military personnel and civilian contractors.

FYI

Prior to DoD RMF's adoption, department systems were certified and accredited under DoD Information Assurance Certification and Accreditation Process (DIACAP), adopted in 2007 as DoD Instruction 8510.01. DIACAP controls were based on a system or network's Mission Assurance Category (MAC) and its confidentiality level (CL). The predecessor to DIACAP was the Department of Defense Information Technology Security Certification and Accreditation Process (DITSCAP).

Surviving Attacks

The need for mission assurance and the ability to react mean that surviving network attacks is an important part of defending networks in cyberwar. Networks must be both resilient and adaptable to changing network conditions—like those found in large-scale denial of service attacks conducted from thousands of systems around the world—and yet still must be able to handle an attack that targets their component parts.

Providing a resilient network often combines a series of design elements that can include capabilities such as the following:

- Diverse geographic locations for major network components to make sure that natural disasters or direct enemy action cannot disable the entire network.
- Multiple service providers if commercial services are used to provide network connectivity.
- Careful analysis of the paths network traffic takes, to ensure that a single event can't disable multiple paths at one time.
- Selection of hardware and software to handle attacks that create excessive traffic or are aimed at a specific hardware or software version. This can require multiple brands or types of devices, as well as different testing than a typical commodity device might undergo.
- Certification and configuration testing to ensure that the device has been tested for security and design flaws. This doesn't prevent the flaws, but can help ensure that the devices are more secure by default.
- Use of commercial services that specialize in defense against network attacks like distributed denial of service (DDoS) attacks. These services use very high-bandwidth systems distributed closer to end users to ensure that an attack against a single location or data center cannot prevent the service from continuing to be available to most users.

Surviving network attacks requires both careful design and the ability to respond quickly and intelligently to attacks that do succeed. Restoring services must be carefully balanced with understanding what attackers did, and what the best route to recovering may be.

Network Operational Procedures

The procedures and processes used to operate a secure network are part of the human layers of protection that go into a well-designed defense-in-depth network defense strategy. The following procedures must be in place:

- Network device management, updates, and configuration must ensure that the devices making up the network are properly patched, do not use default passwords or vulnerable configurations, and that the configurations are backed up and consistent. In addition, the way that devices are managed can help to ensure that a trusted attacker cannot modify the device to breach the organization's security, and that attackers can't use a trusted system to make changes without those changes being noticed. This is done via automated configuration and validation tools or systems.

- Response processes for both breaches and vulnerabilities are needed to ensure that issues are dealt with promptly and fully. Internal disclosure and communication can help to minimize an issue's effect, whereas incorrect or incomplete information or lack of a process can lead to the problem becoming even more of an issue.
- Network entry and access management can help ensure that devices are trusted and secure. The ways in which devices are added to a network and the methods used to verify that systems or devices are allowed to connect are particularly important in a high-security environment. In DoD's compartmentalization systems, a system that is not rated for Top Secret data should not be allowed into a Top Secret network. Ensuring that a system cannot just plug into a network jack or connect to a wireless network that isn't compatible with its sensitivity rating helps to protect both the data and the network.
- Device acquisition policies and procedures are increasingly important in computer network defense. The manufacturer of devices, the code and hardware that are part of the devices, and even the shipping process they go through can be important to the security of a device.

The fear that nation-state-level organizations have influenced or modified the hardware and software of network devices has been a repeated news topic in recent years. In 2007, the NSA launched an investigation of Huawei, a major Chinese electronic manufacturer. The NSA's goal was to determine if it had ties to the Chinese military and had provided, or intended to provide, a back door into devices it built. Reports in 2012 noted that 9,539 banned businesses had sold technology to the U.S. government, and that about 10 percent of those incidents involved counterfeit parts or gear.

More recently, information released in 2014 by Edward Snowden included pictures purporting to show routers that were diverted to facilities used by the NSA, where additional software was allegedly loaded onto the routers. The code reportedly included "phone home" capabilities, which would activate once the routers were added to a live network. Such activities can lead to devices being compromised even before they are delivered, resulting in difficult-to-detect security issues in apparently fresh-from-the-factory hardware.

Many operational processes and policies are needed for a strong, secure network. A complete defense-in-depth design must first identify what may be needed, and then must prioritize which processes are most important without leaving gaps attackers can exploit.

Network Security Design

The initial design of a network forms the basis of its security model. A strong design needs to follow organizational policies and strategies, and must incorporate elements that will help handle the risks the organization faces. In cyberwar, those risks involve nation-state-level attackers, cybercriminals, organized crime, activists, insurrectionists, and even internal threats. Defenders must carefully design networks and plan additions and updates. Fortunately, many resources and design guides exist to help security designers.

FYI

The concepts and commonly used design techniques behind network defense in depth could fill an entire book. Organizations like SANS and vendors like Cisco provide frameworks for secure design, which can be useful references:

- Cisco's SAFE network design framework can be found at *http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg/chap1.html*.
- The SANS reading room has numerous white papers on secure network design at *http://www.sans.org/reading-room/whitepapers/bestprac/*.

Network security designs in cyberwar typically combine the need to meet classification standards and the operational requirements of supporting government or military infrastructure. They then use traditional information security models to provide a strong, secure foundation for the network.

Classification

Nation-state-level defense networks typically add another layer of complexity to design and implementation that traditional organizational networks don't. The need for networks that can handle data classified under governmental or defense classification schemes leads to additional network segmentation. That segmentation often results in physically distinct networks. These are managed to a level of security that matches the data classifications requirements for the information traveling across them.

An excellent example of a publicly announced network scheme for handling classified data can be found in the networks of the U.S. Defense Information Systems Agency (DISA). DISA provides multiple levels of IP network services based on the data classification the network transmits. They include the following:

- SBU IP Data, formerly known as NIPRNet, is an unclassified network that provides Internet connectivity and is used for e-mail, file transfers, and Web traffic. SBU IP Data is designed to funnel traffic through security gateways, allowing monitoring, Web traffic filtering, and traffic protection.
- Secret IP Data, or SIPRNet, is designed to provide a transport network for data falling under the U.S. government's Secret classification. The network is specifically designed to allow compartmentalization based on classification and supports traditional applications like e-mail, Web traffic, and other similar uses for that data. The network extends to DoD, intelligence, the civilian U.S. government, and U.S. allies. It provides centralized, monitored access to other networks.
- TC/SCI IP Data, formerly known as JWICS, provides secure communications between Top Secret and sensitive compartmented information (SCI) users who are part of the U.S. intelligence community.

The Defense Information Systems Network

The DISA operates a global telecommunications network for the U.S. military. The Defense Information Systems Network (DISN) is designed to meet specialized requirements, including security, transmission capabilities, and specialized needs based on geographic locations.

The DISN's physical network design is divided into three major areas. The first is the sustaining base, which includes the parts of the network that exist in long-term bases, facilities, and locations that host military units and other organizations. The Command, Control, Communications, Computers, and Intelligence (C4I) infrastructure on individual bases is typically maintained by each branch of the military, rather than being centrally maintained by the DISA.

The second area is the long-haul infrastructure, which provides communications and network services between deployed units and task forces and the sustaining bases they come from. The DISA operates the long-haul infrastructure, including capabilities like IP telephony, teleconferencing, and various levels of secure network capabilities designed around classification levels.

The third element of the DISN is deployed units and individuals, as well as their commanders. Deployed units are typically supported by their individual units. This design requires coordination between bases, the DISA, and deployed units and their support elements to ensure network security and a strong defense. Monitoring and sharing information about attacks and issues is critical, but is increasingly complex as units use more and more networked services and platforms.

For obvious reasons, the DISA provides far less information about the more sensitive networks it operates than it does for the unclassified SBU IP Data network.

Network Defense Technologies

The technologies used to defend networks have evolved from basic network firewalls and logging systems into advanced systems that can detect attacks and take action to protect against them as they occur. Modern network defense technologies can be integrated into the underlying fabric of a network, and can provide protection at any point along the network from the network edge to the port or wireless access point that individual devices connect to.

Network defense technology deployments require an understanding of the types of devices and technologies that exist and how they interact. It's easy to deploy expensive and complex network defenses without significantly strengthening a network. Attackers will exploit any gap in a network's defenses that they can find. Defenders must deploy a well-designed, monitored, and responsive network to defend against attacks that they are aware of, and to help to prevent the attacks they haven't seen before.

FYI

Network protocols are used to determine how traffic flows between and inside networks, how traffic errors are handled, and a huge variety of network functions. This section covers two examples out of many possible choices.

Protocols

Networks rely on protocols to control how they send and receive information, how that information is formatted, and a whole host of other elements that enable them to work. Protocols play a significant role in the security of networks by allowing data encryption, ensuring that data arrives intact and unchanged, and ensuring that attackers who have access to network traffic cannot view it, even if they possess data decryption capabilities.

Border Gateway Protocol

The Border Gateway Protocol (BGP) is a protocol used between network routers to communicate information about how and where to send traffic. It provides the ability to verify that other routers on a network are the neighbors they are supposed to be. It provides protections that ensure that the network lifetime of data packets has not been modified, which would allow attackers to conduct specialized attacks. It also provides protection against attackers inserting information about other routers into the network, thus allowing them to have traffic sent through their own malicious routers.

Choosing a routing protocol like BGP means that administrators must fully understand the implications of their choice and how to properly secure a BGP router. Each network protocol requires careful implementation and a full understanding of its specialized needs.

Transport Layer Security

One of the most common protocols used to secure network traffic is **Transport Layer Security (TLS)**. It uses asymmetric cryptography to allow endpoints that have no prior relationship to securely share symmetric session encryption keys. This allows Web traffic to be secured between applications and servers, providing both confidentiality and integrity of messages on a network.

FYI

Secure Sockets Layer (SSL) was the predecessor to TLS, and TLS is often still referred to as SSL. SSL itself is now outmoded, and TLS is the standard of choice of secure Web traffic.

Network Access Control

An important part of network defense is controlling which devices and users can connect to the network. In addition to physical controls, networks can also use technologies that allow them to identify users and systems before allowing them to fully connect to the network. Equally useful is the ability to use the information that is provided to place a system into a network segment that is appropriate to the system's rights or status.

Network access control (NAC) allows networks to require users to authenticate to the network. It can also use software on systems to assess those systems' security and compliance status. If the user is validated properly, and his or her system meets the requirements that the NAC system is configured to require, the user can then be placed into the proper network segment based on his or her user rights. If the user fails, the network can reject the connection, or can put him or her into a virtual network penalty box that limits access to only the tools needed to make the system compliant.

NAC systems can be very thorough in their testing, with some NAC software systems integrating antivirus scanning to check for malware and configuration scanning and validation tools to ensure that systems have every setting right. NAC systems can even allow users to change the network they connect to based on their needed access.

Network Firewalls

Firewalls are one of the most commonly deployed network security technologies, and are widely used to provide network separation and defense. Firewalls use a set of rules about which traffic may travel through them based on the network addresses, ports, protocols, or details of the traffic that is sent. These lists of rules are known as **rulesets**, and firewall administrators manage their firewalls' rulesets carefully to ensure that only allowed, trusted traffic is passed through.

Three common types of firewalls are deployed ranging from very basic packet filters to stateful packet inspection firewalls to advanced, next-generation, application-aware firewall devices. Each firewall type has different capabilities. The basic features of each are as follows:

- **Packet filter firewalls** are the most basic type of firewall. They inspect the IP address that is sending traffic, the receiver's IP address, the port that the traffic is sent or received from, and the protocol by which the traffic is being sent. Packet filters look at each packet that sent, repeating their check to ensure that the traffic matches what the firewall allows. They have no understanding of the content of the traffic—that means that traffic sent from a valid sender, on a valid port, and via a valid protocol to a system that will accept the sender's traffic will go through the firewall, even if it is actually an attack.

FYI

A well-designed firewall's final rule will always be a rule that denies all traffic from passing through it. This ensures that any traffic that doesn't match a rule allowing it will be stopped. The opposite rule, which allows all traffic, sometimes known as an *ANY/ANY rule*, will make the firewall useless because it will match all traffic and allow it through.

- **Stateful packet inspection firewalls**, or stateful firewalls, apply additional intelligence beyond the simple individual packet inspection applied by packet filters. Stateful firewalls maintain a mapping in memory of systems that have participated in allowed communications, and use that information to allow conversations to continue between systems. This means that instead of each packet being inspected, the firewall only needs to verify that the following packets are part of the same network conversation before allowing them to pass. It also allows the firewall to use simpler rules, as responses to messages sent can be allowed back in as part of the conversation, rather than requiring specific inbound rules to permit the traffic in. This simplifies ruleset construction, and makes it less likely that firewall administrators will accidentally forget to allow traffic back through a firewall.
- **Application-aware firewalls** are even more advanced, and provide the ability to inspect the actual protocols and application traffic that flow through the firewall. This means that the firewall must spend significantly more resources looking at traffic that passes through it, but it also means that it can detect attacks against the application servers that it protects. Application-aware firewalls can also verify that attacks against Internet protocols are not occurring—preventing attackers from exploiting the network software on servers and devices.

Figure 12-3 shows a sample attack from a remote system through each of the common types of firewall devices. The packet filter firewall checks to see if the traffic is allowed, and because that system is allowed to send responses to a trusted system, it passes the traffic. The stateful packet inspection firewall makes a similar check, and notices that the system started the conversation, rather than responding to an ongoing conversation, and stops the traffic. The application-aware firewall detects that the communication does not match the requirements of the application protocol, and stops the traffic because it is a known attack. It would also stop the traffic if the traffic were not part of an existing conversation started by host A.

Firewalls are used in most secure network designs to ensure that only permitted traffic is sent through network control points. In most designs, they aren't intended to react rapidly to new threats. Instead, firewalls are typically changed through a well-managed change process. This change process is designed to ensure that new rules don't prevent legitimate traffic. It also ensures that changes that are made are well understood and only allow the intended traffic to pass.

FIGURE 12-3

Network firewalls handling an attack from a trusted system.

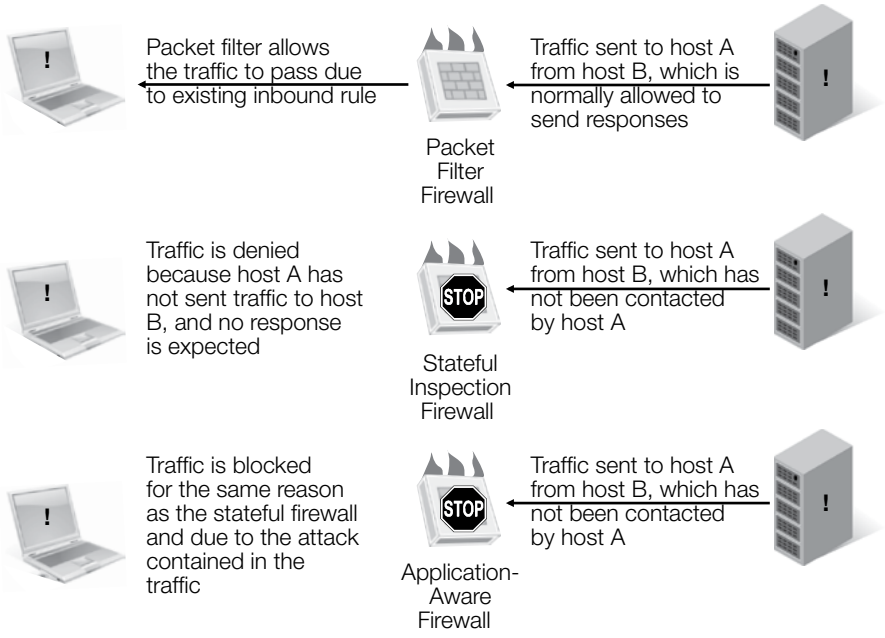

FYI

Host-based software firewalls also exist, and are built in to many modern operating systems. They provide protection only for the system they are installed on, rather than an entire network or groups of systems. This allows fine-grained protection for each device, but also requires ruleset maintenance for every machine that uses one. In most cases, host-based firewalls are either controlled centrally, or users are allowed control over the firewall so that they can allow applications to send or receive traffic.

NOTE

The term *enterprise* here refers to a professional-grade device of a type typically found in a large organizational network. Enterprise devices typically have additional functionality, may have redundant power supplies or other features, and provide more management and configuration options.

Routers and Switches

Networks rely on devices known as **routers** to send data between computer networks. Routers are typically found at network boundaries, or inside large networks to control traffic between network segments. Routers offer a wide range of features that can allow network defenders to provide access control and security for a network, including access control lists. Many modern enterprise routers provide the ability to add other network security devices to the router itself, providing advanced firewall features, as well as capabilities like intrusion detection or network monitoring.

Networks also use switches to distribute data to individual systems. Switches, much like routers, are available with a variety of capabilities, including some that have routing functionality built in to them, blurring the boundaries between a traditional switch and a router. Enterprise switches often can protect their ports from various attacks; they also have options and software that can help them support additional network-based authentication and control policies.

Access Control Lists. Most routers, more advanced switches, and some other network devices can use a tool known as an **access control list (ACL)**. An ACL lists what systems or network IP addresses are allowed to pass through a router, much like a firewall ruleset can limit traffic between systems. Typical ACLs provide less-advanced features than the more-advanced types of firewalls, and are therefore used for simple restrictions.

Most of the Internet uses a version of the IP protocol known as IPv4. There are a limited number of IPv4 network addresses in the world (4,294,967,296, or 2^{32} addresses), which led to worries about an impending disaster when the last IP address was allocated. A technique known as **network address translation (NAT)** was created to allow use of a range of IPs behind a router, with one or more addresses on the outside that those addresses are translated to. The router then tracks which traffic belongs to which internal system. NAT is sometimes used to help protect systems, as attackers normally can't directly target a device that is using an internal IP address behind a NAT device, because the actual address of the system isn't accessible.

Network Security Boundaries

Network defense inside networks often relies on the network security boundaries that network and security administrators can build between parts of the network or systems that have different security or access requirements. Boundaries in networks are typically created at network control points, such as the following:

- Firewalls or routers, which are used to divide networks
- Endpoint switches, which can limit access to wall jacks and individual systems, as well as tag the devices connected to a port into a given network
- Wireless access points, which can authenticate connections and can allow users and devices to connect to multiple distinct networks
- Remote access systems, which can authenticate users and place those users into separate network groups or address ranges

Administrators and defenders can both use these control points to separate systems based on the rights their users have, the configuration or state of the system, or other information they have about the systems based on software running on the system or an access control system providing information about the user's rights.

NIST recognizes the complexity of network boundaries and those found in more complicated network-based information systems. In NIST Special Publication 800-37,

FIGURE 12-4

NIST's use of network security boundaries.

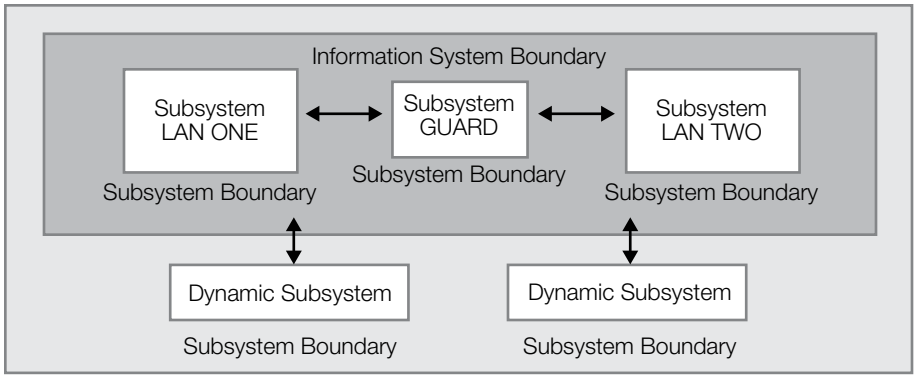

Reproduced from National Institute of Standards and Technology

a complex information system is broken down into two separate networks, as shown in Figure 12-4. In the NIST example, the subsystem guard provides separation between two networks while allowing controlled traffic to pass between them. The dynamic subsystems shown are systems that enter or leave the network over time, typically workstations or other devices that may need access to resources contained in the two separated networks.

Virtual LANs

Network security relies on two major types of separation: logical and physical. One of the most common ways to logically separate networks is using a virtual LAN (VLAN). VLANs are network segments that are logically separated using information known as a VLAN tag, which is applied to network traffic. VLANs are configured on network switches and routers, allowing networks to place systems into appropriate groups based on the tag provided.

Physical separation, on the other hand, is done by using completely separate hardware and wiring. Obviously, physical separation is typically both more complex and more expensive, but when networks have a very strong set of security requirements, physical separation can provide greater protection—unless your opponent can access the network in person or through social engineering.

VLANs are used when the security requirements for separation of networked systems and devices do not make physical separation necessary. They can provide a moderate level of separation, although they cannot guarantee separation because misconfiguration, access to the switches or routers, or VLAN hopping attacks could provide attackers with the ability to see traffic on other VLANs.

FYI

Attacks against VLANs known as *VLAN hopping attacks* can allow attackers to send traffic to other VLANs. Fortunately, modern network devices provide protection against VLAN hopping, although it is not always enabled.

Intrusion Detection and Prevention Systems

The ability to detect attacks as they occur is very useful when defending a network. **Intrusion detection systems (IDSs)** or devices were created to detect attacks based on their distinct network signatures or the behaviors they exhibit. After creating IDSs, the next step was to put the IDS between the attackers and their potential victims and allow the system to both detect and then prevent the attacks. This resulted in the creation of **intrusion prevention systems (IPSs)**.

Both IDS and IPS software and devices rely on two major types of detection:

- Signature-based detection looks for specific attack signatures that malware or attacks use. Signature-based detection requires that the defender know what the attack looks like before detecting or stopping it. This means that signature-based systems tend to be reactive instead of proactive. Properly built signatures will catch all examples of a known attack that they can see, but can't stop unknown attacks.
- Heuristic or behavior-based detection systems look for patterns or types of action that match attacks. They can stop unknown attacks that fit those patterns, but can also stop legitimate traffic that happens to match the same pattern.

Most IDS and IPS systems can combine signatures and heuristics, and administrators have to choose which rules and methods they want to use to balance the security requirements of their network against the detection schemes' strengths and weaknesses.

The Dangers of Detection

DoD's mission assurance strategy creates a challenge for IPS operators. One of the frequent challenges for IPS operators is balancing the need to block attacks against the need to allow legitimate traffic. Unfortunately, the balance between identifying attacks and not stopping legitimate traffic can be difficult. Often, new attacks will appear to be very similar, if not identical, to normal traffic. As administrators and network defenders write new rules to identify attacks, it is easy to make mistakes or to create an overly broad rule.

Once a rule is in place and active, it can result in significant network outages. A single IPS placed in a critical location that blocks important traffic, or that identifies a common network protocol as a source of attack, can bring network operations to a halt. Much like a hardware device failure, a misconfigured IPS or one with a bad rule in place can knock an organization offline.

This doesn't mean that an IPS isn't worth the risk. But it does mean that defenders whose primary goal is to provide mission assurance must carefully consider the balance between the risk of attack and the risk of network outages. Some defenders may opt for an IDS that will make them rely on manual responses to attacks. Others may design a strong testing process that will slow down deployment of potentially important new protection rules.

Security Event and Information Management Systems

Network devices can provide information about their status using automated logging capabilities that send details of what they are doing, whether errors have occurred, and what their configuration is. This data, as well as traffic data from the network, firewall, IDS, IPS, and other device status, and user actions are often gathered centrally. Once that huge amount of data is securely copied to a central location, defenders must then use it. Traditionally, that required huge amounts of time and manual searches through data. In many cases the data was simply ignored or used only after a compromise or attack occurred.

Fortunately, software and devices that can automatically manage the data and identify events and issues have become more common and capable. There are a few distinct types of devices, including the following:

- Security Event Managers (SEMs) provide event management and threat analysis, and often have an incident-response focus for security operations. Traditional SEM systems are not designed to handle massive volumes of logs, and instead focus on real-time analysis.
- Security Information Monitors (SIMs) monitor security logs from devices, looking for security events that indicate breaches or violations of policy. They often focus on reporting and analysis of massive amounts of log and event data, as well as long-term storage of the data they gather.
- **Security Information and Event Management (SIEM)** combines the capabilities of both SIM and SEM devices into a single, integrated solution. Many SIM and SEM systems are becoming SIEM products. SIEM devices often include the ability to monitor network traffic and traffic flows, logs, security information provided by machines, vulnerability scanning data, and other security information. They match the information to known attacks, as well as patterns that may indicate attack, and allow administrators to create rules to set off alarms or to take preprogrammed actions when those rules are matched.

Managing the overall security posture and status of network devices can allow defenders to determine if a compromise has occurred or whether an attack is in progress. It can also allow them to see abnormal traffic or behaviors that might help show that a trusted user is acting outside his or her role. SIEM systems help to sort through the massive amounts of data that a well-configured and logged network provides.

Physical Network Protection

Physical access to networks provides an ideal situation for attackers. Once they have physical access to the network or its component devices, they can do anything from monitor the traffic on the network to insert their own devices into it. Monitoring miles of cables and interconnected devices spread across huge distances can be a real challenge, which means that most physical network defense is done at the building or installation level, with traditional physical security boundaries like walls, fences, and security guards employed to ensure that attackers cannot gain access.

Access to the physical cables that data travels over is harder to prevent. The United States deployed listening devices into undersea cables using submarines in the early 1970s as part of Operation Ivy Bells. It tapped undersea fiber-optic cables in the 1990s, demonstrating the lengths that nation-states will go to in acquiring intelligence. Detecting devices that can simply read traffic passing through a cable can be nearly impossible. This makes it a necessity to strongly encrypt traffic that passes through unsecured network infrastructure, wireless connections, or where traffic has to be carried through commercial providers.

Network defenders are not always responsible for the physical-security design of their network. But they must be aware of the threats physical access to the network and its component devices can create. Nation-state-level resources can make otherwise unlikely scenarios a reality when the traffic that passes through a network is sufficiently valuable.

Wireless Network Security

Wireless networks create additional security issues for defenders due to the inability to restrict their signals. Wireless signals can be accessed from far-greater ranges than users may expect by attackers using directional antennas and other tools. This makes data sent over wireless connections much more easily accessible than physical network traffic in a well-secured building.

In addition to the threats that defenders must take into account with the spread of wireless signals, a number of other wireless threats exist, such as:

- Fake wireless networks, which lure unsuspecting users into sending their data through untrusted access points
- Tools that exploit unencrypted wireless connections to capture data or to modify data sent by users
- Fake cellular towers, which capture data seemingly sent through to legitimate carriers

With a broad range of attacks possible against wireless networks, physical wires remain a popular solution for the highest-security environments in use today. Where wireless is used, defenders often layer encryption and additional tools that help to prove the network identity of both the wireless infrastructure and the systems that connect to it using tools like digital certificates.

FYI

The U.S. National Security Agency and other high-security organizations sometimes design or modify their buildings to limit, or entirely prevent, the spread of wireless signals through the walls and floors of the buildings themselves. This is an expensive solution, but one that can be necessary to prevent wireless signals from being monitored by those outside the building or even between parts of the same structure. The NSA uses the code name Tempest for technologies related to both capturing and protecting information from capture through the use of leaked emanations like vibrations, radio frequency emissions, and sound waves.

Remote Access and Administration

Attackers can cause major damage if they can gain remote access to a network or take over the network's administrative systems. Network management and remote access capabilities are both important to the network's defensive strategy.

Network Management. The tools used to manage network devices play a major role in the maintenance, configuration management, and monitoring of a network. Network management tools help ensure that routers, switches, wireless access points and controllers, and other network devices have the correct configurations, that the devices are functioning properly, and that administrators can properly control them.

TOR: The Onion Router

Many recent cyberwar news articles have mentioned TOR, the Onion Router. TOR combines elements of a virtual private network with routing and encryption technologies to create an anonymizing network that permits the origin, destination, and content of traffic to be concealed from observers.

TOR passes encrypted traffic through TOR entry nodes, which then pass that traffic through other TOR nodes, each stripping off a layer of encryption to learn where to send the traffic next. The final node doesn't know where the traffic came from, only where it is going, and where along the chain to send responses. This conceals the full path from any participating node, and provides anonymity to TOR users. For this reason, TOR is used for both legitimate and illegitimate uses, including cyberattacks and cyberwarfare activities.

TOR also uses its own internal Domain Name System, allowing the creation of systems that are only accessible via the TOR network. TOR addresses have been used for a variety of purposes, including online drug and weapons sales, secret message boards, and other criminal enterprises. Due to the combination of criminal activity and anonymity, the TOR network has become a prominent target for intelligence and law enforcement agencies.

Various nation-state organizations, including the NSA and United Kingdom Government Communications Headquarters (UK GCHQ), have created tools designed to defeat TOR. The toolkits they have employed include compromising the originating systems via browser-based tools, using fingerprinting to uniquely identify browsers, and even adding government owned and monitored TOR nodes. TOR remains resistant to purely network-based attacks, as attackers would have to own a majority of the TOR nodes in the world to have a reasonable chance of identifying both the sender and receiver in any given communication. Attacks that focus on the originating system or the server that it communicates with have generally been more successful, as they allow unique identification of the traffic. Attacks against the server don't reveal the source of the traffic without additional data, but attacks against the sender have the potential to reveal everything sent and where it is sent.

Network management tools can also provide a useful means for managing standards. Many manufacturers provide standard configuration guidance that can ensure that devices meet security baselines, and their management tools can then use those prebuilt profiles. Network defenders who use administration tools can often use the time they save by using those tools to improve the security of their network in other ways.

Virtual Private Networks. One of the most common ways to remotely access a network is through the use of a virtual private network (VPN). A VPN uses software and network protocols to build a private channel through other networks. VPNs are often, but not always, encrypted, allowing a separate, secure network path through otherwise public or untrusted networks. VPNs are also used to connect remote networks together through third-party infrastructure such as the Internet. A VPN set up between two sites can allow those sites to appear to be directly connected, despite the fact they are not actually both connected to the same physical network.

VPNs are an attractive target for attackers, as VPN-connected users appear to be on the network they are connecting just like a local physically connected user would be. Defenders often choose to place VPN users in a distinct zone with different trust or monitoring rules. They may also make VPN users employ additional methods to verify their identity, such as two-factor tokens or additional logons.

Active Defense

The U.S. Department of Defense has increasingly required its networks to use an **active defense** paradigm, as shown in its 2011 strategy document. The basic design requirement is that a network must adapt to attacks. If an attacker is using a specific attack, the network should respond by blocking that attack. If a system is compromised, the network must detect it, remove the system from its trusted network zone, notify administrators that the system needs investigating, and then allow normal operations for the systems that were not part of the incident.

Active defenses require a carefully constructed network and provide many challenges for defenders. The more active defense capabilities a network has, the more likely it is that the network could inadvertently take actions that interfere with normal network operations. The risks are matched by advantages that drive DoD's push for networks that actively defend themselves. A well-built network with active defense capabilities can in theory respond to most attacks faster than human monitoring and management teams can. It can adapt to attacks that might overwhelm the network before it is taken down. In a cyberwar scenario, a **self-defending network** that responds to attacks as they occur and changes itself to prevent, defeat, or at least minimize the assault would be a major advantage against an attacker who did not have the same capabilities on his or her own network.

FYI

Self-defending networks have been proposed for years; however, a network with true end-to-end, self-defense capabilities doesn't yet exist. Still, networks often have some self-defense capabilities, and the advent of SIEM systems that integrate well into network management tools has brought the idea closer to reality in recent years.

DoD's active defense strategy envisions:

- Real-time detection of threats so that they can be handled as they occur
- Attribution of threat actors, providing the ability to deal with the actual adversaries as they attack, and expanding the potential for effective, real-time responses
- Flexibility of response actions with choices available to defenders about how to respond and on what scale
- Intelligence dissemination

Honeypots, Honeynets, and Darknets

In addition to self-defending networks, defenders can benefit from understanding the attacks their networks face. Nation-state-level network security organizations must remain constantly aware of new threats. One way to do this as part of an active defense design is to use a **honeypot**. A honeypot is a vulnerable system that is monitored and equipped with logging capabilities. When an attacker compromises the honeypot system, defenders are provided with copies of any files or software the attacker used, as well as complete information about how the system responded to the attack.

Honeypots can provide defenders with up-to-date information about attacks by showing the complete process that an attacker uses (see Figure 12-5). **Honeynets** take the concept of a honeypot to the Network Layer, and provide a network honeypot for attackers. Honeynets can combine monitored network devices with honeypot systems, often in a virtual environment. Defenders can easily restore compromised systems to their original state while retaining all of the information about how the system or network device was attacked.

Organizations that have unused network space sometimes also set up darknets. A **darknet** relies on the fact that unused space should have no legitimate traffic sent to it. Monitoring unused networks allows defenders to see what scans and attacks occur against systems, and can help identify attackers who are targeting broad ranges of network space.

NOTE

The term *darknet* is sometimes applied to TOR networks because they operate in an otherwise difficult-to-detect part of the Internet.

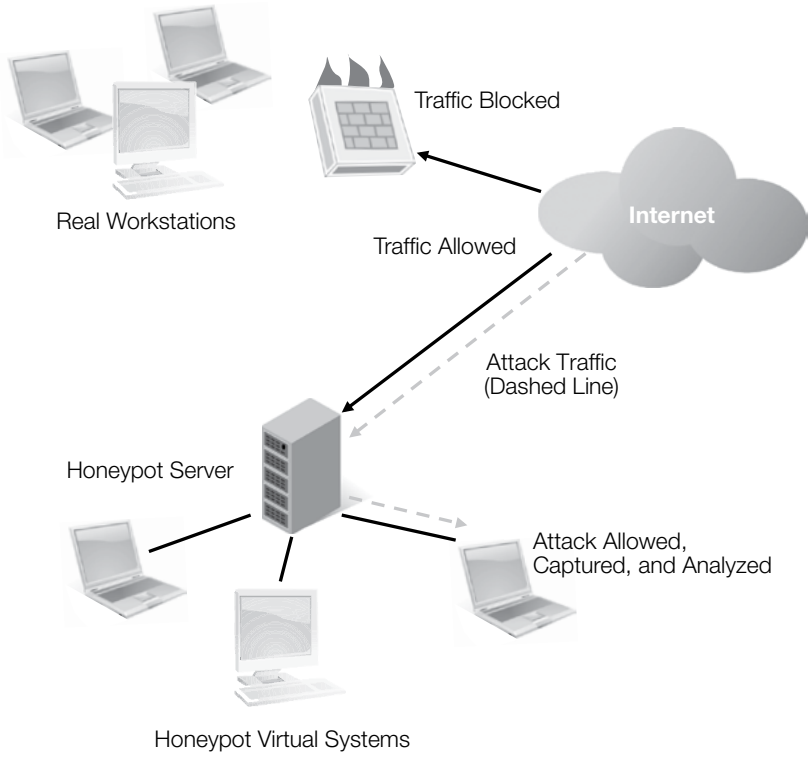

FIGURE 12-5
A honeypot captures and allows analysis of attack traffic.

Defenders who want to take a more active role in their defense can use honeypots, honeynets, and darknets to provide insight into attacks and who the attackers might be. Defenders can use the information they gain from honeypots, honeynets, and darknets to build firewall rules, to improve the defenses of the systems they deploy, and to understand what new attacks may be used on their real networks and devices.

The Honeynet project provides prebuilt and configured honeypots and honeynet systems. You can read more about this at *https://www.honeynet.org*.

Active Response

In traditional information security scenarios, an **active response**, or "hacking the hacker" is normally considered to be a bad idea. Civilians and corporations are typically prohibited from attacking others under local, national, or international law. Cyberwarfare opens a new possibility for defenders who may be authorized or even expected to use active responses as part of their computer network defense responsibilities.

Active responses are simply computer network attack capabilities linked to computer network defense. When an attack occurs, or when one is known or suspected, defenders may strike at the systems or networks that are attacking or will attack. This can allow the defenders to weather an attack relatively safely if their own systems are not attacked, or if their actions blunt the attack.

Active responses require a careful balance by defenders, however, as many attacks come from systems or networks shared by civilians. Although international law has not solidified around cyberwarfare activities, attacks against civilian infrastructure can have unexpected and undesired effects.

CHAPTER SUMMARY

This chapter discussed computer network defense in cyberwar. U.S. civilian agencies and the military use the concept of mission assurance, or the ability to provide available and usable networks to the organization, as a key element in network defense. Critical to mission assurance are, first, a strong network security design that supports the classification of data and second, systems based on assigning roles and rights to users and objects, such as programs and computers.

Network defenders must also be familiar with a variety of network security devices commonly used in secure networks. These include border defenses like firewalls, intrusion detection and prevention systems, and security configurations deployed on network routers and switches. They also include systems designed to gather network logs and security events, bait systems designed to gather information about attacks, and a wide range of other tools. The U.S. military envisions a secure network that uses these tools and detection capabilities to react to threats and provide an active response to attacks by reconfiguring itself to stop or limit them as they occur.

KEY CONCEPTS AND TERMS

Access control list (ACL)
Active defense
Active response
Application-aware firewalls
Darknet
Firewalls
Honeynets
Honeypot
Information assurance
Intrusion detection systems (IDSs)
Intrusion prevention systems (IPSs)
Local area networks (LANs)
Mission assurance
Mission Assurance Categories (MAC)
Network address translation (NAT)
Packet filter firewalls
Routers
Rulesets
Secure Sockets Layer (SSL)
Security Information and Event Management (SIEM)
Self-defending network
Stateful packet inspection firewalls
Transport Layer Security (TLS)
Wide area networks (WANs)