Security Services in a Modular Network Design

KEY POINT Security is an infrastructure service that increases the network's integrity by protecting network resources and users from internal and external threats.

Without a full understanding of the threats involved, network security deployments tend to be incorrectly configured, too focused on security devices, or lacking appropriate threat response options.

Security both in the Enterprise Campus (internal security) and at the Enterprise Edge (from external threats) is important. An enterprise should include several layers of protection so that a breach at one layer or in one network module does not mean that other layers or modules are also compromised; Cisco calls deploying layered security *defense-in-depth*.

Internal Security

Strongly protecting the internal Enterprise Campus by including security functions in each individual element is important for the following reasons:

- If the security established at the Enterprise Edge fails, an unprotected Enterprise Campus is vulnerable. Deploying several layers of security increases the protection of the Enterprise Campus, where the most strategic assets usually reside.
- Relying on physical security is not enough. For example, as a visitor to the organization, a potential attacker could gain physical access to devices in the Enterprise Campus.
- Often external access does not stop at the Enterprise Edge; some applications require at least indirect access to the Enterprise Campus resources. Strong security must protect access to these resources.

Figure 3-15 shows how internal security can be designed into the Cisco Enterprise Architecture.

Figure 3-15 *Designing Internal Security into the Network*

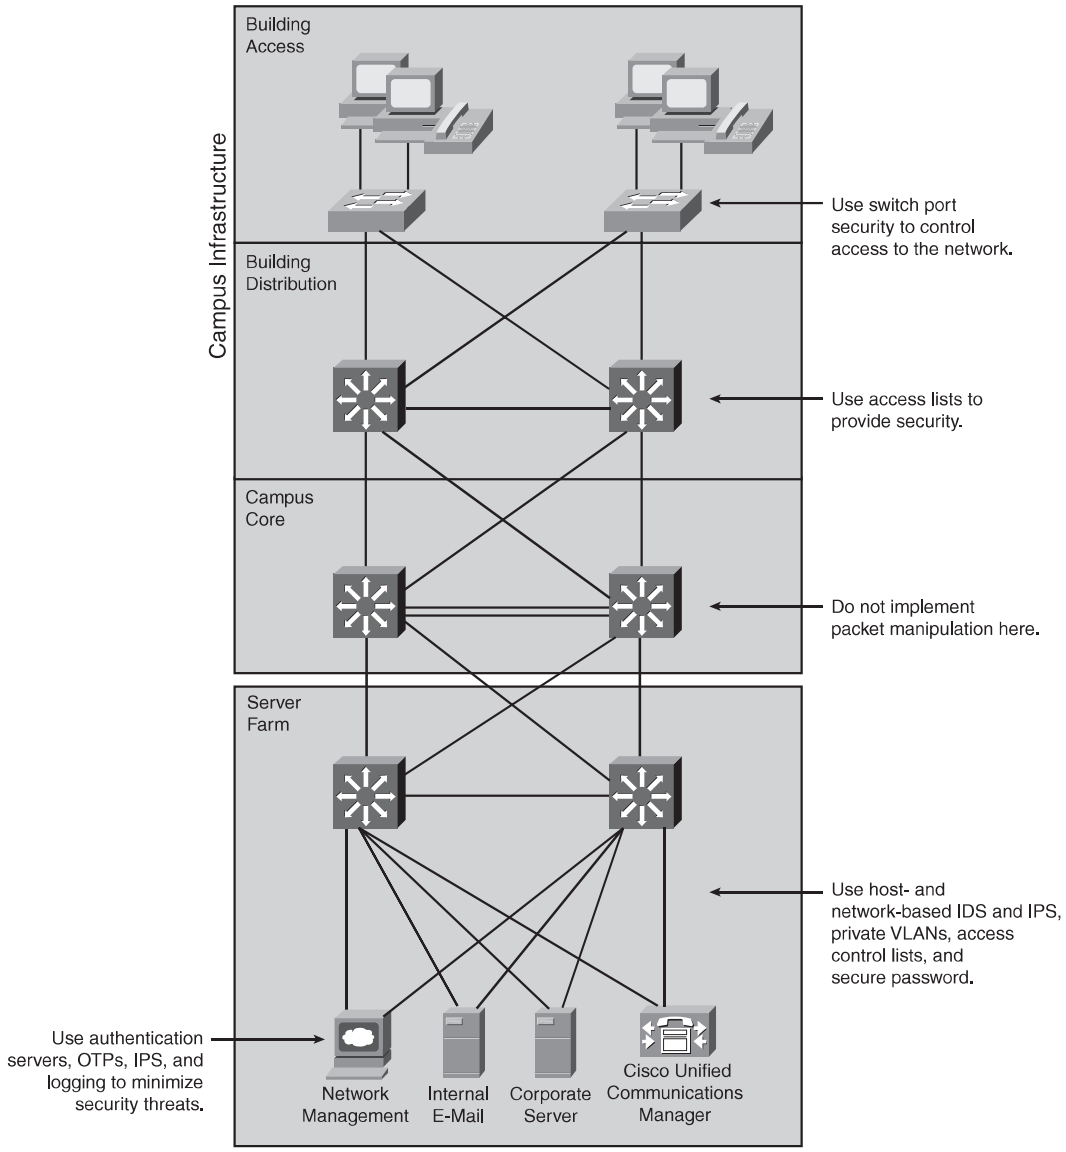

The following are some recommended security practices in each module:

- At the Building Access layer, access is controlled at the port level using the data link layer information. Some examples are filtering based on media access control addresses and IEEE 802.1X port authentication.

- The Building Distribution layer performs filtering to keep unnecessary traffic from the Campus Core. This packet filtering can be considered a security function because it does prevent some undesired access to other modules. Given that switches in the Building Distribution layer are typically multilayer switches (and are therefore Layer 3–aware), this is the first place on the data path in which filtering based on network layer information can be performed.
- The Campus Core layer is a high-speed switching backbone and should be designed to switch packets as quickly as possible; it should not perform any security functions, because doing so would slow down the switching of packets.
- The Server Farm module's primary goal is to provide application services to end users and devices. Enterprises often overlook the Server Farm module from a security perspective. Given the high degree of access that most employees have to these servers, they often become the primary goal of internally originated attacks. Simply relying on effective passwords does not provide a comprehensive attack mitigation strategy. Using host-based and network-based IPSs and IDSs, private VLANs, and access control provides a much more comprehensive attack response. For example, onboard IDS within the Server Farm's multilayer switches inspects traffic flows.

NOTE Private VLANs provide Layer 2 isolation between ports within the same broadcast domain.

- The Server Farm module typically includes network management systems to securely manage all devices and hosts within the enterprise architecture. For example, syslog provides important information on security violations and configuration changes by logging security-related events (authentication and so on). An authentication, authorization, and accounting (AAA) security server also works with a one-time password (OTP) server to provide a high level of security to all local and remote users. AAA and OTP authentication reduces the likelihood of a successful password attack.

IPS and IDS

IDSs act like an alarm system in the physical world. When an IDS detects something it considers an attack, it either takes corrective action or notifies a management system so that an administrator can take action.

HIDSs work by intercepting operating system and application calls on an individual host and can also operate via after-the-fact analysis of local log files. The former approach allows better attack prevention, and the latter approach is a more passive attack-response role.

Because of their specific role, HIDSs are often more effective at preventing specific attacks than NIDSs, which usually issue an alert only on discovering an attack. However, this specificity does not allow the perspective of the overall network; this is where NIDS excels.

Intrusion prevention solutions form a core element of a successful security solution because they detect and block attacks, including worms, network viruses, and other malware through inline intrusion prevention, innovative technology, and identification of malicious network activity.

Network-based IPS solutions protect the network by helping detect, classify, and stop threats, including worms, spyware or adware, network viruses, and application abuse. Host-based IPS solutions protect server and desktop computing systems by identifying threats and preventing malicious behavior.

This information was derived from the *SAFE Blueprint for Small, Midsize, and Remote-User Networks,* available at http://www.cisco.com/go/safe/, and the *Cisco Intrusion Prevention System Introduction,* available at http://www.cisco.com/en/US/products/sw/secursw/ps2113/index.html.

Authentication, Authorization, and Accounting

AAA is a crucial aspect of network security that should be considered during the network design. An AAA server handles the following:

- **Authentication—*Who?*** Authentication checks the user's identity, typically through a username and password combination.
- **Authorization—*What?*** After the user is authenticated, the AAA server dictates what activity the user is allowed to perform on the network.
- **Accounting—*When?*** The AAA server can record the length of the session, the services accessed during the session, and so forth.

The principles of strong authentication should be included in the user authentication. *Strong authentication* refers to the two-factor authentication method in which users are authenticated using two of the following factors:

- **Something you know**: Such as a password or personal identification number (PIN)
- **Something you have**: Such as an access card, bank card, or token
- **Something you are**: For example, some biometrics, such as a retina print or fingerprint
- **Something you do**: Such as your handwriting, including the style, pressure applied, and so forth

As an example, when accessing an automated teller machine, strong authentication is enforced because a bank card (something you have) and a PIN (something you know) are used.

Tokens are key-chain-sized devices that show OTPs, one at a time, in a predefined order. The OTP is displayed on the token's small LCD, typically for 1 minute, before the next password in the sequence appears. The token is synchronized with a token server, which has the same predefined list of passcodes for that one user. Therefore, at any given time, only one valid password exists between the server and a token.

This information was derived from Cisco Press's *Campus Network Design Fundamentals* by Diane Teare and Catherine Paquet, 2006.

External Threats

When designing security in an enterprise network, the Enterprise Edge is the first line of defense at which potential outside attacks can be stopped. The Enterprise Edge is like a wall with small doors and strong guards that efficiently control any access. The following four attack methods are commonly used in attempts to compromise the integrity of the enterprise network from the outside:

- **IP spoofing**: An IP spoofing attack occurs when a hacker uses a trusted computer to launch an attack from inside or outside the network. The hacker uses either an IP address that is in the range of a network's trusted IP addresses or a trusted external IP address that provides access to specified resources on the network. IP spoofing attacks often lead to other types of attacks. For example, a hacker might launch a denial of service (DoS) attack using spoofed source addresses to hide his identity.
- **Password attacks**: Using a packet sniffer to determine usernames and passwords is a simple password attack; however, the term *password attack* usually refers to repeated brute-force attempts to identify username and password information. Trojan horse programs are another method that can be used to determine this information. A hacker might also use IP spoofing as a first step in a system attack by violating a trust relationship based on source IP addresses. First, however, the system would have to be configured to bypass password authentication so that only a username is required.
- **DoS attacks**: DoS attacks focus on making a service unavailable for normal use and are typically accomplished by exhausting some resource limitation on the network or within an operating system or application.
- **Application layer attacks**: Application layer attacks typically exploit well-known weaknesses in common software programs to gain access to a computer.

DoS Attacks

DoS attacks are different from most other attacks because they are not generally targeted at gaining access to a network or its information. Rather, these attacks focus on making a service unavailable for normal use. They are typically accomplished by exhausting some resource limitation on the network or within an operating system or application.

When involving specific network server applications, such as a web server or an FTP server, these attacks focus on acquiring and keeping open all the available connections supported by that server, thereby effectively locking out valid users of the server or service. DoS attacks are also implemented using common Internet protocols, such as TCP and Internet Control Message Protocol (ICMP).

Rather than exploiting a software bug or security hole, most DoS attacks exploit a weakness in the overall architecture of the system being attacked. However, some attacks compromise a network's performance by flooding the network with undesired and often useless network packets and by providing false information about the status of network resources. This type of attack is often the most difficult to prevent, because it requires coordinating with the upstream network provider. If traffic meant to consume the available bandwidth is not stopped there, denying it at the point of entry into your network does little good, because the available bandwidth has already been consumed. When this type of attack is launched from many different systems at the same time, it is often referred to as a *distributed denial of service attack*.

This information was derived from the *SAFE Blueprint for Small, Midsize, and Remote-User Networks*, available at http://www.cisco.com/go/safe/.

Application Layer Attacks

Hackers perform application layer attacks using several different methods. One of the most common methods is exploiting well-known weaknesses in software commonly found on servers, such as SMTP, HTTP, and FTP. By exploiting these weaknesses, hackers gain access to a computer with the permissions of the account that runs the application—usually a privileged system-level account. These application layer attacks are often widely publicized in an effort to allow administrators to rectify the problem with a patch. Unfortunately, many hackers also subscribe to these same informative mailing lists and therefore learn about the attack at the same time (if they have not discovered it already).

The primary problem with application-layer attacks is that they often use ports that are allowed through a firewall. For example, a hacker who executes a known vulnerability against a web server often uses TCP port 80 in the attack. A firewall needs to allow access on that port because the web server serves pages to users using port 80. From a firewall's perspective, the attack appears as merely standard port 80 traffic.

This information was derived from the *SAFE Blueprint for Small, Midsize, and Remote-User Networks*, available at http://www.cisco.com/go/safe/.

Figure 3-16 shows these four attack methods and how they relate to the Enterprise Edge modules.

Figure 3-16 *External Threats*

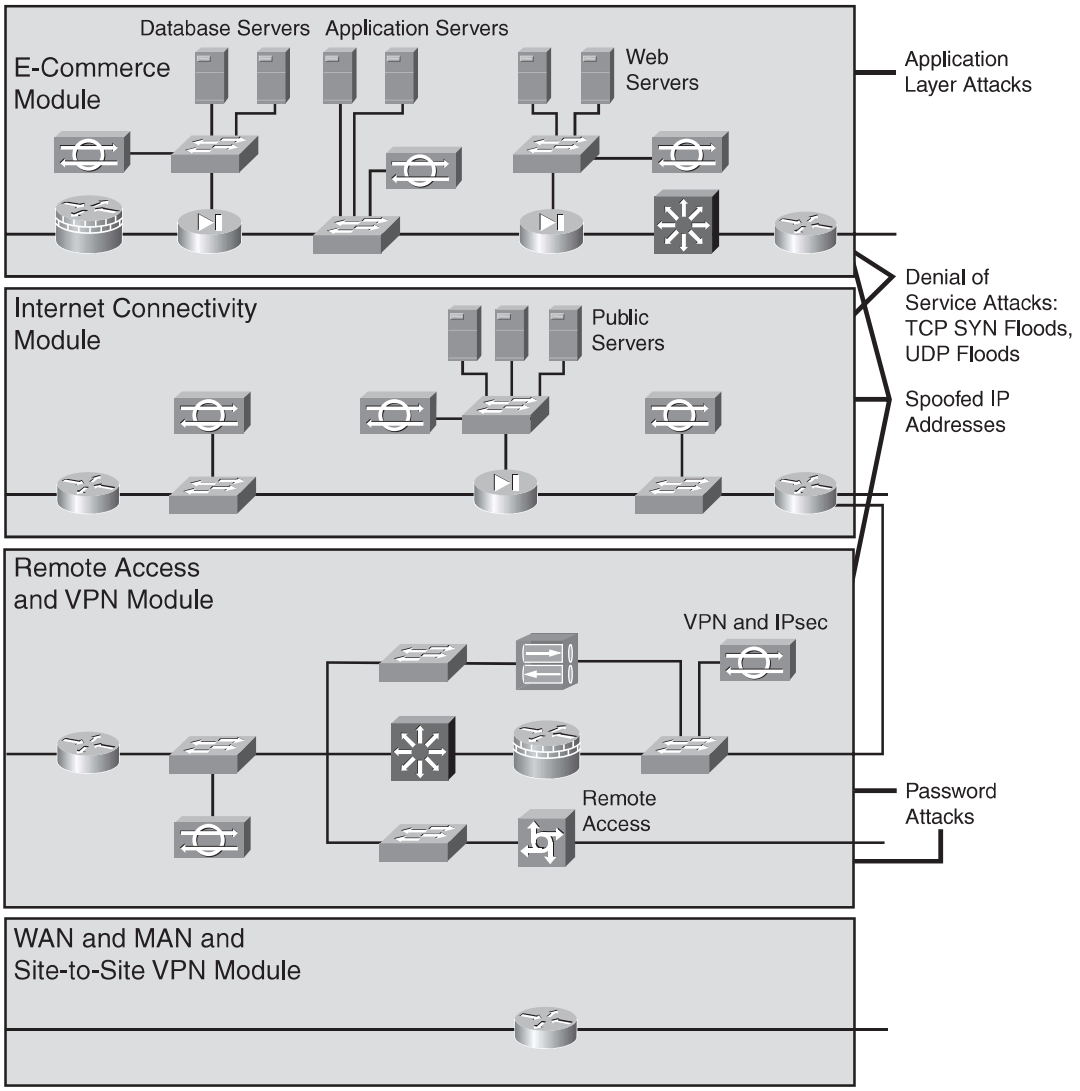

Because of the complexity of network applications, access control must be extremely granular and flexible yet still provide strong security. Tight borders between outside and inside cannot be defined, because interactions are continuously taking place between the Enterprise Edge and

Enterprise Campus. The ease of use of the network applications and resources must be balanced against the security measures imposed on the network users.

NOTE Chapter 10, "Evaluating Security Solutions for the Network," covers security in the network in more detail.

High-Availability Services in a Modular Network Design

Most enterprise networks carry mission-critical information. Organizations that run such networks are usually interested in protecting the integrity of that information. Along with security, these organizations expect the internetworking platforms to offer a sufficient level of resilience.

This section introduces another network infrastructure service: high availability. To ensure adequate connectivity for mission-critical applications, high availability is an essential component of an enterprise environment.

Designing High Availability into a Network

Redundant network designs duplicate network links and devices, eliminating single points of failure on the network. The goal is to duplicate components whose failure could disable critical applications.

Because redundancy is expensive to deploy and maintain, redundant topologies should be implemented with care. Redundancy adds complexity to the network topology and to network addressing and routing. The level of redundancy should meet the organization's availability and affordability requirements.

KEY POINT Before selecting redundant design solutions, analyze the business and technical goals and constraints to establish the required availability and affordability.

Critical applications, systems, internetworking devices, and links must be identified. Analyze the risk tolerance and the consequences of *not* implementing redundancy, and ensure that you consider the trade-offs of redundancy versus cost and simplicity versus complexity. Duplicate any component whose failure could disable critical applications.

Redundancy is not provided by simply duplicating all links. Unless all devices are completely fault-tolerant, redundant links should terminate at different devices; otherwise, devices that are not fault-tolerant become single points of failure.

KEY POINT Because many other modules access the Server Farm and Campus Core modules, they typically require higher availability than other modules.