International Jurisdiction Issues

A unique legal challenge for prosecuting information security offenses deals with jurisdictional issues. For example, an attacker in one country could launch an attack from a computer in another country that targets a computer in yet another country. The international boundaries that were virtually crossed could pose significant challenges to litigators.

Fortunately, governments are beginning to collaborate on such investigations and prosecutions. For example, organizations that share law enforcement information between countries include G8, Interpol, and the European Union.

Understanding the Methods of Network Attacks

You might have noticed that this book has thus far referred to computer criminals as "attackers" rather than "hackers." This wording is intentional, because not all hackers have malicious intent, even though the term "hacker" often has a negative connotation. In this section, you will gain additional insight into the mind-set and characteristics of various hackers.

Additionally, you will be introduced to a variety of methods that attackers can use to infiltrate a computing system. To help mitigate such attacks, Cisco recommends the Defense in Depth design philosophy, which also is covered in this section, in addition to a collection of best practices for defending your network.

Vulnerabilities

A *vulnerability* in an information system is a weakness that an attacker might leverage to gain unauthorized access to the system or its data. In some cases, after a vulnerability is discovered, attackers write a program intended to take advantage of the vulnerability. This type of malicious program is called an *exploit*.

However, even if a system has a vulnerability, the likelihood that someone will use that vulnerability to cause damage varies. This likelihood is called *risk*. For example, a data center might be vulnerable to a fire breaking out in the building. However, if the data center has advanced fire suppression systems and hot standby backups at another physical location, the risk to the data is minimal.

When you make plans to address vulnerabilities, consider the varied types of vulnerabilities. For example, consider the following broad categories of vulnerabilities:

- Physical vulnerabilities, such as fire, earthquake, or tornado
- Weaknesses in a system's design
- Weaknesses in the protocol(s) used by a system

- Weaknesses in the code executed by a system
- Suboptimal configuration of system parameters
- Malicious software (for example, a virus)
- Human vulnerabilities (whether intentional or unintentional)

For example, consider human vulnerabilities. Because most attacks against information systems are launched from people on the “inside,” controls should be set up to prevent the intentional or unintentional misuse of information systems.

Social engineering is an example of unintentional misuse. To illustrate this concept, consider a situation in which an outside attacker calls a receptionist. The attacker pretends to be a member of the company’s IT department, and he convinces the receptionist to tell him her username and password. The attacker then can use those credentials to log into the network.

To prevent a single inside user from accidentally or purposefully launching an attack, some organizations require that two users enter their credentials before a specific act can be carried out, much like two keys being required to launch a missile.

Also, many employees are concerned with accomplishing a particular task. If stringent security procedures seem to stand in their way, the employees might circumnavigate any safeguards to, in their minds, be more productive. Therefore, user education is a critical component of any organizational security policy.

Potential Attackers

Another element of defending your data is identifying potential attackers who might want to steal or manipulate that data. For example, a company might need to protect its data from corporate competitors, terrorists, employees, and hackers, to name just a few.

The term “hacker” is often used very generically to describe attackers. However, not all hackers have malicious intent.

Table 1-5 lists various types of “hackers.”

Key Topic

Table 1-5 *Types of Hackers*

Type of "Hacker"	Description
White hat hacker	A white hat hacker has the skills to break into computer systems and do damage. However, he uses his skills to help organizations. For example, a white hat hacker might work for a company to test the security of its network.
Black hat hacker	A black hat hacker, also known as a "cracker," uses his skills for unethical reasons (for example, to steal funds).
Gray hat hacker	A gray hat hacker can be thought of as a white hat hacker who occasionally strays and acts unethically. For example, a gray hat hacker might be employed as a legitimate network security tester. However, in the course of his ethical duties, he finds an opportunity for personal gain and acts unethically to obtain that personal gain.
Phreaker	A phreaker is a hacker of a telecommunications system. For example, a phreaker known as "Captain Crunch" used a toy whistle he found in a box of Captain Crunch cereal (which generated a 2600-Hz tone) to trick phone systems into letting him place free long distance calls. Convincing a telecommunications carrier to permit free long distance calls in this manner is an example of "phreaking."
Script kiddy	A script kiddy is a user who lacks the skills of a typical hacker. Rather, he downloads hacking utilities and uses those utilities to launch attacks, rather than writing his own programs.
Hacktivist	A hacktivist is a hacker with political motivations, such as someone who defaces the website of a political candidate.
Computer security hacker	A computer security hacker is knowledgeable about the technical aspects of computer and network security systems. For example, this person might attempt to attack a system protected by an IPS by fragmenting malicious traffic in a way that would go undetected by the IPS.
Academic hacker	An academic hacker typically is an employee or student at an institution of higher education. The academic hacker uses the institution's computing resources to write "clever" programs. Typically, these hackers use their real names (unlike the pseudonyms often used by computer security hackers), and they tend to focus on open-standards-based software and operating systems (for example, Linux).
Hobby hacker	A hobby hacker tends to focus on home computing. He might modify existing hardware or software to, for example, use software without a legitimate license. For example, code that "unlocks" an Apple iPhone might be the work of a hobby hacker.

As shown in Table 1-5, "hackers" come in many flavors, which leads to the question, "What motivates a hacker?" Some hackers might work for governments to try to gather intelligence from other governments. Some attackers seek financial gain through their attacks. Other hackers simply enjoy the challenge of compromising a protected information system.

This book details several specific attacks that an attacker can launch. However, at this point, you should be familiar with five broad categories of attacks:

Key Topic

- **Passive**: A passive attack is difficult to detect, because the attacker isn't actively sending traffic (malicious or otherwise). An example of a passive attack is an attacker capturing packets from the network and attempting to decrypt them (if the traffic was encrypted originally).
- **Active**: An active attack is easier to detect, because the attacker is actively sending traffic that can be detected. An attacker might launch an active attack in an attempt to access classified information or to modify data on a system.
- **Close-in**: A close-in attack, as the name implies, occurs when the attacker is in close physical proximity with the target system. For example, an attacker can bypass password protection on some routers, switches, and servers if he gains physical access to those devices.
- **Insider**: An insider attack occurs when legitimate network users leverage their credentials and knowledge of the network in a malicious fashion.
- **Distribution**: Distribution attacks intentionally introduce "back doors" to hardware or software systems at the point of manufacture. After these systems have been distributed to a variety of customers, the attacker can use his knowledge of the implanted back door to, for example, access protected data, manipulate data, or make the target system unusable by legitimate users.

The Mind-set of a Hacker

Hackers can use a variety of tools and techniques to "hack" into a system (that is, gain unauthorized access to a system). Although these methods vary, the following steps illustrate one example of a hacker's methodical process for hacking into a system:

Step 1 Learn more about the system by performing reconnaissance. In this step, also known as "footprinting," the hacker learns all he can about the system. For example, he might learn the target company's domain names and the range of IP addresses it uses. He might perform a port scan to see what ports are open on a target system.

Step 2 Identify applications on the system, as well as the system's operating system. Hackers can use various tools to attempt to connect to a system, and the prompt they receive (for example, an FTP login prompt or a default web page) could provide insight into the system's operating system. Also, the previously mentioned port scan can help identify applications running on a system.

Step 3 Gain access to the system. Social engineering is one of the more popular ways to obtain login credentials. For example, public DNS records provide contact information for a company's domain name. A hacker might be able to use this information to convince the domain administrator to reveal information about the system. For example, the hacker could pretend to be a representative of the service provider or a government agency. This approach is called *pretexting*.

Step 4 Log in with obtained user credentials, and escalate the hacker's privileges. For example, a hacker could introduce a Trojan horse (a piece of software that appears to be a legitimate application but that also performs some unseen malicious function) to escalate his privileges.

Step 5 Gather additional usernames and passwords. With appropriate privileges, hackers can run utilities to create reports of usernames and/or passwords.

Step 6 Configure a "back door." Accessing a system via a regular username/password might not be how a hacker wants to repeatedly gain access to a system. Passwords can expire, and logins can be logged. Therefore, hackers might install a back door, which is a method of gaining access to a system that bypasses normal security measures.

Step 7 Use the system. After a hacker gains control of a system, he might gather protected information from that system. Alternatively, he might manipulate the system's data or use the system to launch attacks against other systems with which the system might have an established trust relationship.

Defense in Depth

Because a security solution is only as strong as its weakest link, network administrators are challenged to implement a security solution that protects a complex network. As a result, rather than deploying a single security solution, Cisco recommends multiple, overlapping solutions. These overlapping solutions target different aspects of security, such as securing against insider attacks and securing against technical attacks. These solutions should also be subjected to routine testing and evaluation. Security solutions should also overlap in a way that eliminates any single point of failure.

Defense in Depth is a design philosophy that achieves this layered security approach. The layers of security present in a Defense in Depth deployment should provide redundancy for one another while offering a variety of defense strategies for protecting multiple aspects of a network. Any single points of failure in a security solution should be eliminated, and weak links in the security solution should be strengthened.

The Defense in Depth design philosophy includes recommendations such as the following:

Key Topic

- Defend multiple attack targets in the network.
    - — Protect the network infrastructure.
    - — Protect strategic computing resources, such as via a Host-based Intrusion Prevention System (HIPS).
- Create overlapping defenses. For example, include both Intrusion Detection System (IDS) and IPS protections.
- Let the value of a protected resource dictate the strength of the security mechanism. For example, deploy more resources to protect a network boundary as opposed to the resources deployed to protect an end-user workstation.
- Use strong encryption technologies, such as AES (as opposed to DES) or Public Key Infrastructure (PKI) solutions.

Consider the sample Defense in Depth topology shown in Figure 1-2. Notice the two e-mail servers—external and internal. The external e-mail server acts as an e-mail relay to the internal e-mail server. Therefore, an attacker attempting to exploit an e-mail vulnerability would have to compromise both e-mail servers to affect the internal corporate e-mail.

Also notice the use of a Network-based Intrusion Detection System (NIDS), a Network Intrusion Prevention System (NIPS), and a Host-based Intrusion Prevention System (HIPS). All three of these mitigation strategies look for malicious traffic and can alert or drop such traffic. However, these strategies are deployed at different locations in the network to protect different areas of the network. This overlapping yet diversified protection is an example of the Defense in Depth design philosophy.

However, if all security solutions in a network were configured and managed by a single management station, this management station could be a single point of failure. Therefore, if an attacker compromised the management station, he could defeat other security measures.

Figure 1-2 *Defense in Depth*

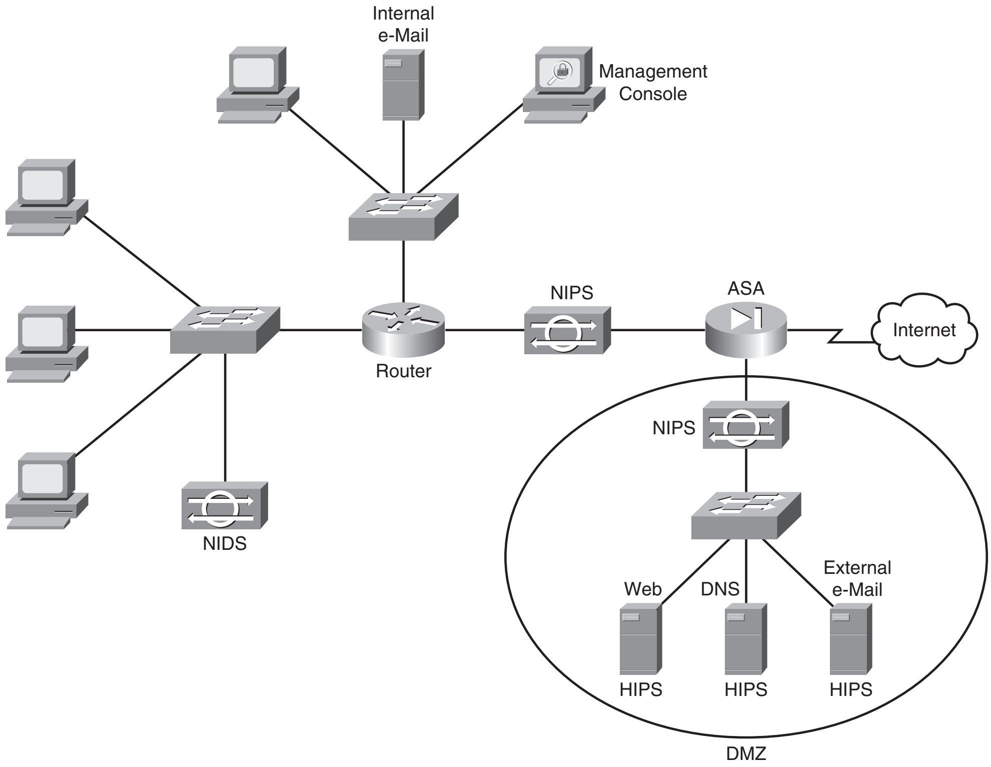

In the "Potential Attackers" section you read about five classes of attacks; Table 1-6 provides examples of overlapping defenses for each of these classes.

Key Topic

Table 1-6 *Defending Against Different Classes of Attacks*

Attack Class	Primary Layer of Defense	Secondary Layer of Defense
Passive	Encryption	Applications with integrated security
Active	Firewall at the network edge	HIPS
Insider	Protecting against unauthorized physical access	Authentication
Close-in	Protecting against unauthorized physical access	Video monitoring systems
Distribution	Secured software distribution system	Real-time software integrity checking

Understanding IP Spoofing

Attackers can launch a variety of attacks by initiating an IP spoofing attack. An IP spoofing attack causes an attacker's IP address to appear to be a trusted IP address. For example, if an attacker convinces a host that he is a trusted client, he might gain privileged access to a host. The attacker could also capture traffic, which might include credentials such as usernames and passwords. As another example, you might be familiar with denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks. The perpetrators of such attacks might use IP spoofing to help conceal their identities.

To understand how an IP spoofing attack is possible, consider the operation of IP and TCP. At Layer 3, the attacker can easily modify his packets to make the source IP address appear to be a "trusted" IP address. However, TCP, operating at Layer 4, can be more of a challenge.

From your early studies of TCP, you might recall that a TCP session is established using a three-way handshake:

Key Topic

1. The originator sends a SYN segment to the destination, along with a sequence number.
2. The destination sends an acknowledgment (an ACK) of the originator's sequence number along with the destination's own sequence number (a SYN).
3. The originator sends an ACK segment to acknowledge the destination's sequence number, after which the TCP communication channel is open between the originator and destination.

Figure 1-3 illustrates the TCP three-way handshake process.

Figure 1-3 *TCP Three-Way Handshake*

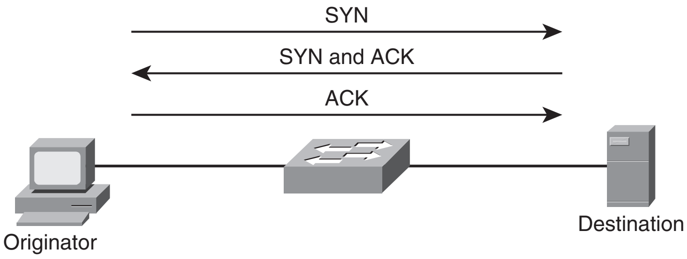

For an attacker to "hijack" a session being set up between a legitimate originator and a destination, the attacker needs to know the TCP sequence numbers used in the TCP segments. If the attacker successfully guesses or predicts the correct TCP sequence numbers, he can send a properly constructed ACK segment to the destination. If the attacker's ACK segment reaches the destination before the originator's ACK segment does, the attacker becomes trusted by the destination, as illustrated in Figure 1-4.

Figure 1-4 *IP Spoofing*

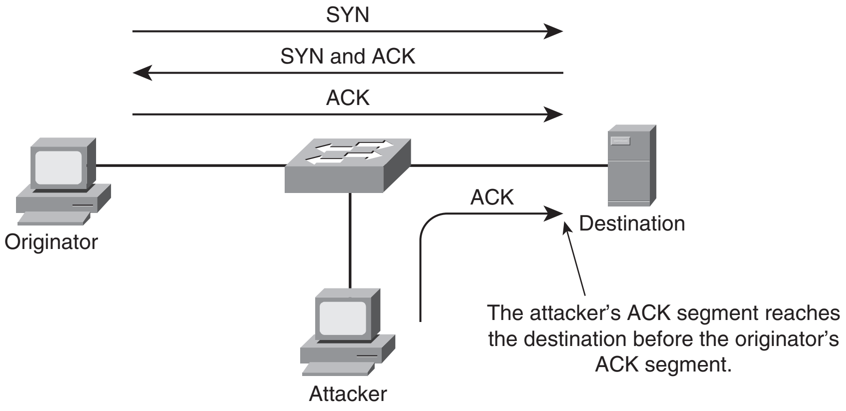
SYN
SYN and ACK
ACK
Originator
ACK
Destination
Attacker
The attacker's ACK segment reaches the destination before the originator's ACK segment.

How an attacker guesses or predicts correct TCP sequence numbers depends on the type of IP spoofing attack being launched. Table 1-7 describes two categories of IP spoofing attacks.

Key Topic

Table 1-7 *Types of IP Spoofing Attacks*

Type of Attack	Description
Nonblind spoofing	Nonblind spoofing occurs when the attacker and the destination are on the same subnet. By being on the same subnet, the attacker might be able to use a packet-capture utility to glean sequence numbers.
Blind spoofing	Blind spoofing occurs when the attacker is not on the same subnet as the destination. Therefore, obtaining correct TCP sequence numbers is more difficult. However, using techniques such as *IP source routing* (described next), an attacker can accurately determine those sequence numbers.

Launching a Remote IP Spoofing Attack with IP Source Routing

If an attacker uses a feature known as IP source routing, he can specify a complete routing path to be taken by two endpoints. Consider Figure 1-5. The attacker is on a different subnet than the destination host. However, the attacker sends an IP packet with a source route specified in the IP header, which causes the destination host to send traffic back to the spoofed IP address via the route specified. This approach can overcome the previously described challenge that an attacker might have when launching a remote IP spoofing (blind spoofing) attack.

Figure 1-5 *IP Source Routing*

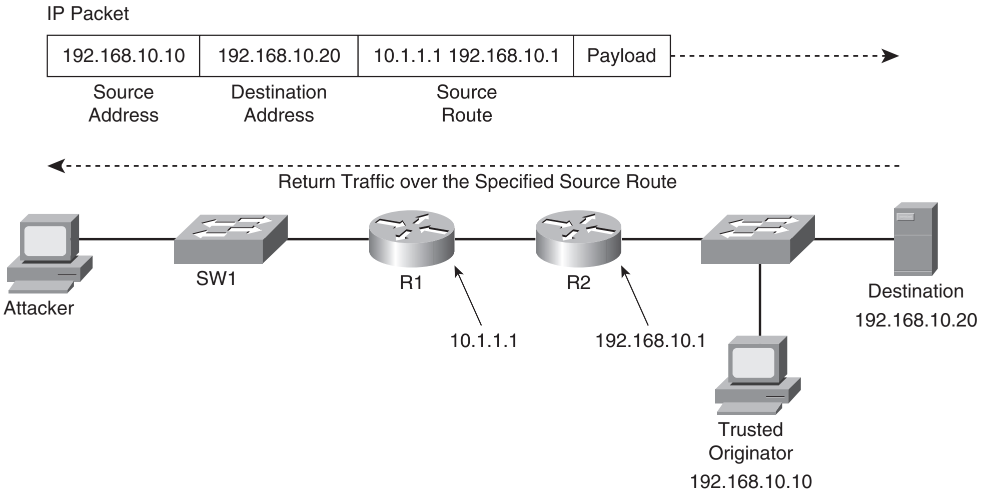

Source routing has two variations:

- **Loose**: The attacker specifies a list of IP addresses through which a packet must travel. However, the packet could also travel through additional routers that interconnect IP addresses specified in the list.

- **Strict**: The IP addresses in the list specified by the attacker are the only IP addresses through which a packet is allowed to travel.

Key Topic

Launching a Local IP Spoofing Attack Using a Man-in-the-Middle Attack

If an attacker is on the same subnet as the target system, he might launch a man-in-the-middle attack. In one variant of a man-in-the-middle attack, the attacker convinces systems to send frames via the attacker's PC. For example, the attacker could send a series of gratuitous ARP (GARP) frames to systems. These GARP frames might claim that the attacker's Layer 2 MAC address was the MAC address of the next-hop router. The attacker could then capture traffic and forward it to the legitimate next-hop router. As a result, the end user might not notice anything suspicious.

Another variant of a man-in-the-middle attack is when the attacker connects a hub to a network segment that carries the traffic the attacker wants to capture, as shown in Figure 1-6. Alternatively, an attacker could connect to a Switch Port Analyzer (SPAN) port on a Catalyst switch, which makes copies of specified traffic and forwards them to the configured SPAN port. The attack could then use a packet-capture utility to capture traffic traveling between end systems. If the captured traffic is in plain text, the attacker might be able to obtain confidential information, such as usernames and passwords.

Figure 1-6 *Man-in-the-Middle Attack*

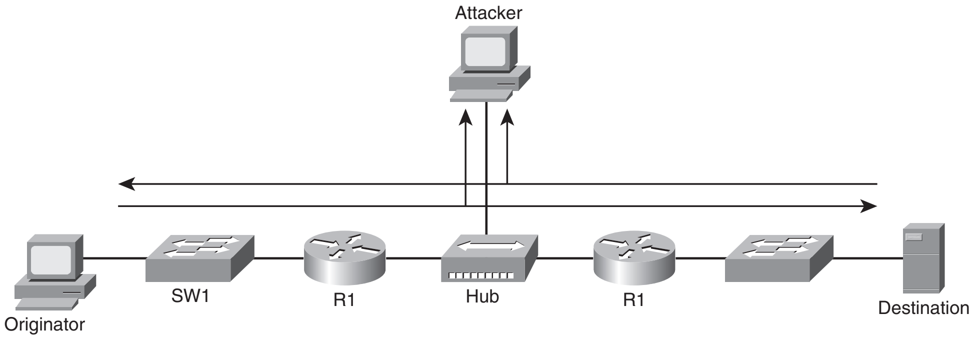

Protecting Against an IP Spoofing Attack

The following approaches can be used to mitigate IP spoofing attacks:

- Use access control lists (ACL) on router interfaces. As traffic comes into a router from an outside network, an ACL could be used to deny any outside traffic claiming to be addressed with IP addressing used internally on the local network. Conversely, ACLs should be used to prevent traffic leaving the local network from participating in a DDoS attack. Therefore, an ACL could deny any traffic leaving the local network that claimed to have a source address that was different from the internal network's IP address space.

- Encrypt traffic between devices (for example, between two routers, or between an end system and a router) via an IPsec tunnel. In Figure 1-7, notice that the topology is now protected with an IPsec tunnel. Even though the attacker can still capture packets via his rogue hub, the captured packets are unreadable, because the traffic is encrypted inside the IPsec tunnel.

Figure 1-7 *Protecting Traffic in a Tunnel*

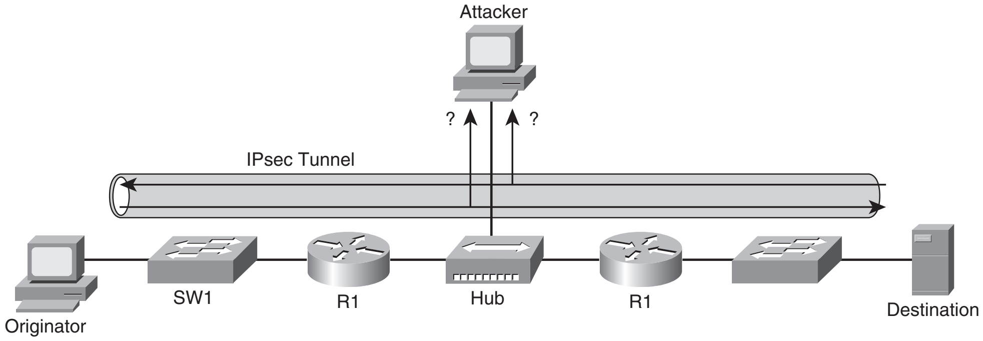

- Use cryptographic authentication. If the parties involved in a conversation are authenticated, potential man-in-the-middle attackers can be thwarted. Potential attackers will not be successfully authenticated by the other party in the conversation.

Understanding Confidentiality Attacks

A confidentiality attack (see Figure 1-8) attempts to make "confidential" data (such as personnel records, usernames, passwords, credit card numbers, and e-mails) viewable by an attacker. Because an attacker often makes a copy of the data, rather than trying to manipulate the data or crash a system, confidentiality attacks often go undetected. Even if auditing software to track file access were in place, if no one suspected an issue, the audit trail might never be examined.

Figure 1-8 *Confidentiality Attack*

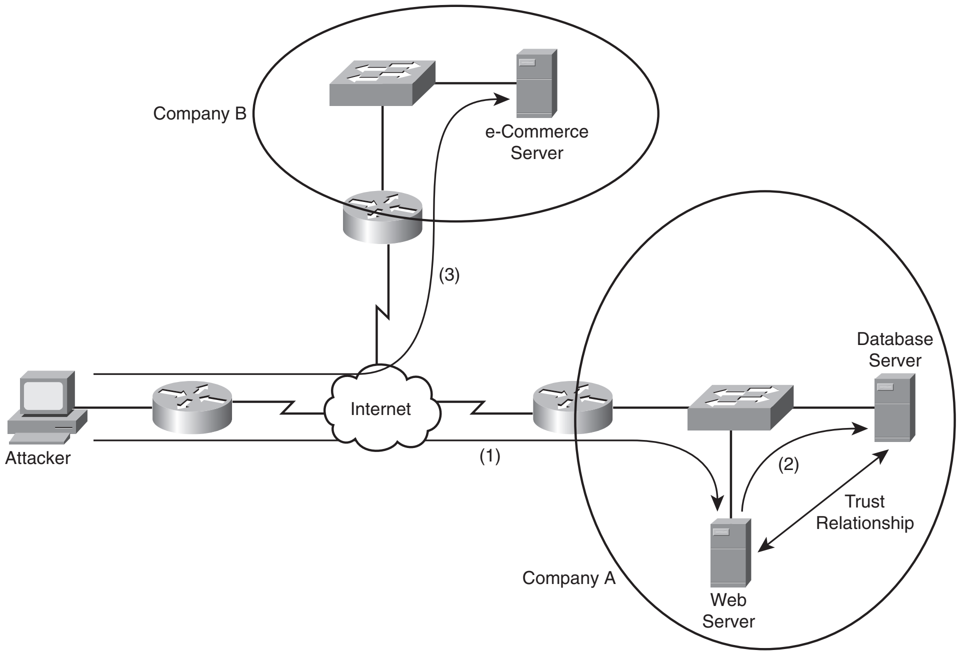

In Figure 1-8, a web server and a database server have a mutual trust relationship. The database server houses confidential customer information, such as credit card information. As a result, Company A decides to protect the database server (for example, patching known software vulnerabilities) better than the web server. However, the attacker leverages the trust relationship between the two servers to obtain customer credit card information and then make a purchase from Company B using the stolen information. The procedure is as follows:

Step 1 The attacker exploits a vulnerability in Company A’s web server and gains control of that server.

Step 2 The attacker uses the trust relationship between the web server and the database server to obtain customer credit card information from the database server.

Step 3 The attacker uses the stolen credit card information to make a purchase from Company B.

Table 1-8 identifies several methods that attackers might use in a confidentiality attack.

Key Topic

Table 1-8 *Confidentiality Attack Strategies*

Tactic	Description
Packet capture	A packet-capture utility (such as Wireshark, available at http://www.wireshark.org) can capture packets visible by a PC’s network interface card (NIC) by placing the NIC in promiscuous mode. Some protocols (for example, Telnet and HTTP) are sent in plain text. Therefore, an attacker can read these types of captured packets, perhaps allowing him to see confidential information.
Ping sweep and port scan	A confidentiality attack might begin with a scan of network resources, to identify attack targets on a network. A ping sweep could be used to ping a series of IP addresses. Ping replies might indicate to an attacker that network resources can be reached at those IP addresses. As soon as a collection of IP addresses is identified, the attacker might scan a range of UDP and/or TCP ports to see what services are available on the host at the specified IP addresses. Also, port scans often help attackers identify the operating system running on the target system.
Dumpster diving	Because many companies throw away confidential information, without proper shredding, some attackers rummage through company dumpsters in hopes of discovering information that could be used to compromise network resources.
Electromagnetic interference (EMI) interception	Because data is often transmitted over wire (for example, unshielded twisted-pair), attackers can sometimes copy information traveling over the wire by intercepting the EMI being emitted by the transmission medium. These EMI emissions are sometimes called “emanations.”

Table 1-8 *Confidentiality Attack Strategies (Continued)*

Key Topic

Tactic	Description
Wiretapping	If an attacker gains physical access to a wiring closet, he might physically tap into telephone cabling to eavesdrop on telephone conversations. Or he might insert a shared media hub inline with a network cable. This would let him connect to the hub and receive copies of packets flowing through the network cable.
Social engineering	Attackers sometimes use social techniques (which often leverage people's desire to be helpful) to obtain confidential information. For example, an attacker might pose as a member of the IT department and ask a company employee for her login credentials "for the IT staff to test the connection."
Sending information over overt channels	An attacker might send or receive confidential information over a network using an overt channel. An example of using an overt channel is tunneling one protocol inside another (for example, sending instant messaging traffic via HTTP). *Steganography* is another example of sending information over an overt channel. An example of steganography is sending a digital image made up of millions of pixels, with "secret" information encoded in specific pixels. Only the sender and receiver know which pixels represent the encoded information.
Sending information over covert channels	An attacker might send or receive confidential information over a network using a covert channel, which can communicate information as a series of codes and/or events. For example, binary data could be represented by sending a series of pings to a destination. A single ping within a certain period of time could represent a binary 0, and two pings within that same time period could represent a binary 1.

Understanding Integrity Attacks

Integrity attacks attempt to alter data (that is, compromise its integrity). Figure 1-9 shows an example of an integrity attack.

Figure 1-9 *Integrity Attack*

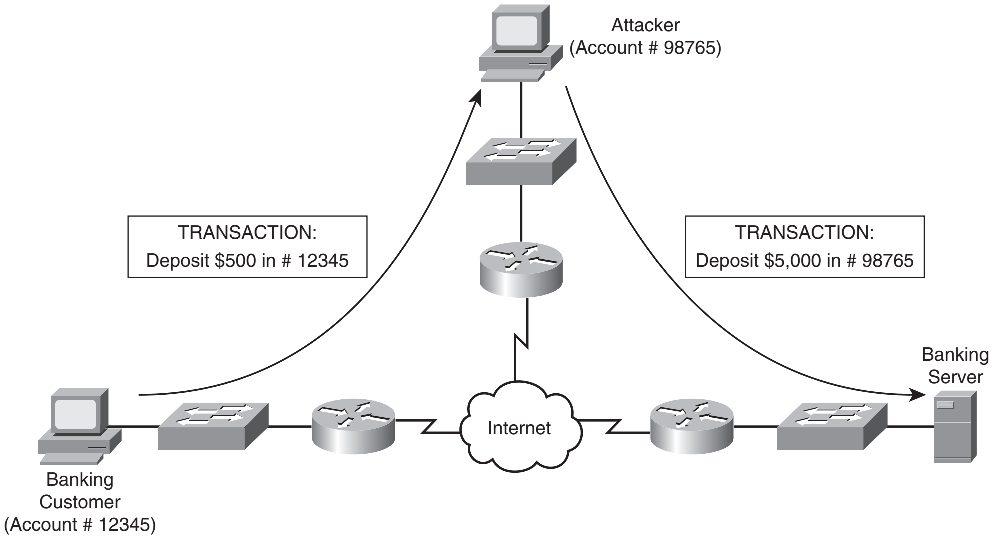

Traffic diverted to attacker due to a man-in-the-middle attack.

In the figure, an attacker has launched a man-in-the-middle attack (as previously described). This attack causes data flowing between the banking customer and the banking server to be sent via the attacker's computer. The attacker then can not only intercept but also manipulate the data. In the figure, notice that the banking customer attempts to deposit $500 into her account. However, the attacker intercepts and changes the details of the transaction, such that the instruction to the banking server is to deposit $5,000 in the attacker's account.

The following list describes methods that attackers might leverage to conduct an integrity attack:

Key Topic

- **Salami attack**: This is a collection of small attacks that result in a larger attack when combined. For example, if an attacker had a collection of stolen credit card numbers, he could withdraw small amounts of money from each credit card (possibly unnoticed by the credit card holders). Although each withdrawal is small, they add up to a significant sum for the attacker.
- **Data diddling**: The process of data diddling changes data before it is stored in a computing system. Malicious code in an input application or virus could perform data diddling. For example, a virus, Trojan horse, or worm could be written to intercept keyboard input. It would display the appropriate characters on-screen so that the user would not see a problem. However, manipulated characters would be entered into a database application or sent over a network.

- **Trust relationship exploitation**: Different devices in a network might have a trust relationship between themselves. For example, a certain host might be trusted to communicate through a firewall using specific ports, while other hosts are denied passage through the firewall using those same ports. If an attacker could compromise the host that had a trust relationship with the firewall, the attacker could use the compromised host to pass normally denied data through a firewall. Another example of a trust relationship is a web server and a database server mutually trusting one another. In that case, if the attacker gained control of the web server, he might be able to leverage that trust relationship to compromise the database server.

- **Password attack**: A password attack, as the name suggests, attempts to determine a user's password. As soon as the attacker gains the username and password credentials, he can attempt to log into a system as that user, and therefore inherit that user's set of permissions. Various approaches are available for determining passwords:

  — **Trojan horse**: A program that appears to be a useful application captures a user's password and then makes it available to the attacker.

  — **Packet capture**: A packet-capture utility can capture packets seen on a PC's NIC. Therefore, if the PC can see a copy of a plain-text password being sent over a link, the packet-capture utility can be used to glean the password.

  — **Keylogger**: A keylogger is a program that runs in the background of a computer, logging the user's keystrokes. After a user enters a password, it is stored in the log created by the keylogger. An attacker then can retrieve the log of keystrokes to determine the user's password.

  — **Brute force**: A brute-force password attack tries all possible password combinations until a match is made. For example, the brute-force attack might start with the letter a and go through to the letter z. Then the letters aa through zz are attempted, until a password is determined. Therefore, using a mixture of uppercase and lowercase letters in passwords, in addition to special characters and numbers, can help mitigate a brute-force attack.

  — **Dictionary attack**: A dictionary attack is similar to a brute-force attack, in that multiple password guesses are attempted. However, the dictionary attack is based on a dictionary of commonly used words, rather than the brute-force method of trying all possible combinations. Picking a password that is not a common word can help mitigate a dictionary attack.

- **Botnet**: A software “robot” typically is thought of as an application on a machine that can be controlled remotely (for example, a Trojan horse or a back door in a system). If a collection of computers is infected with such software robots, called “bots,” this collection of computers (each of which is called a “zombie”) is known as a “botnet.” Because of the potentially large size of a botnet, it might compromise the integrity of a large amount of data.
- **Hijacking a session**: Earlier in this chapter, you read about how an attacker could hijack a TCP session (for example, by completing the third step in the three-way TCP handshake process between an authorized client and a protected server). If an attacker successfully hijacked a session of an authorized device, he might be able to maliciously manipulate data on the protected server.

Understanding Availability Attacks

Availability attacks attempt to limit a system’s accessibility and usability. For example, if an attacker could consume the processor or memory resources on a target system, that system would be unavailable to legitimate users.

Availability attacks vary widely, from consuming the resources of a target system to doing physical damage to that system. Attackers might employ the following availability attacks:

Key Topic

- **Denial of service (DoS)**: An attacker can launch a DoS attack on a system by sending the target system a flood of data or requests that consume the target system’s resources. Alternatively, some operating systems and applications might crash when they receive specific strings of improperly formatted data, and the attacker could leverage such operating system and/or application vulnerabilities to render a system or application inoperable. The attacker often uses IP spoofing to conceal his identity when launching a DoS attack, as shown in Figure 1-10.

Figure 1-10 *Denial-of-Service Attack*

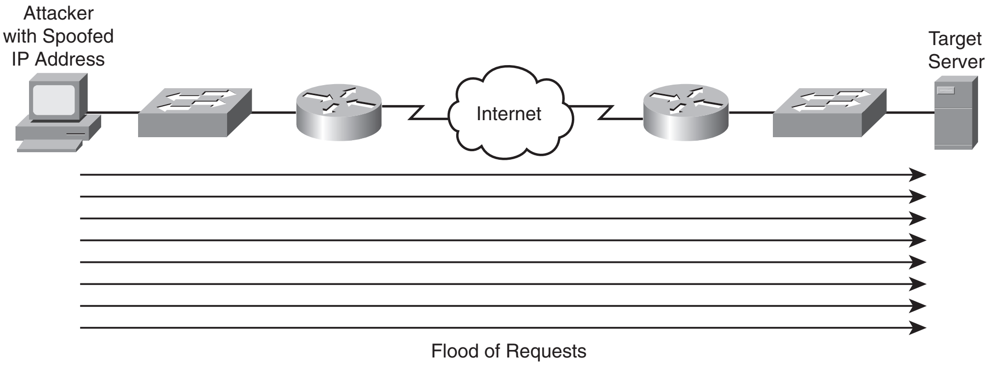

- **Distributed denial of service (DDoS)**: DDoS attacks can increase the amount of traffic flooded to a target system. Specifically, the attacker compromises multiple systems. The attacker can instruct those compromised systems, called "zombies," to simultaneously launch a DDoS attack against a target system.

- **TCP SYN flood**: Earlier in this chapter you reviewed the three-way TCP handshake process. One variant of a DoS attack is for an attacker to initiate multiple TCP sessions by sending SYN segments but never completing the three-way handshake. As illustrated in Figure 1-11, the attack can send multiple SYN segments to a target system, with false source IP addresses in the header of the SYN segment. Because many servers limit the number of TCP sessions they can have open simultaneously, a SYN flood can render a target system incapable of opening a TCP session with a legitimate user.

Figure 1-11 *TCP SYN Flood Attack*

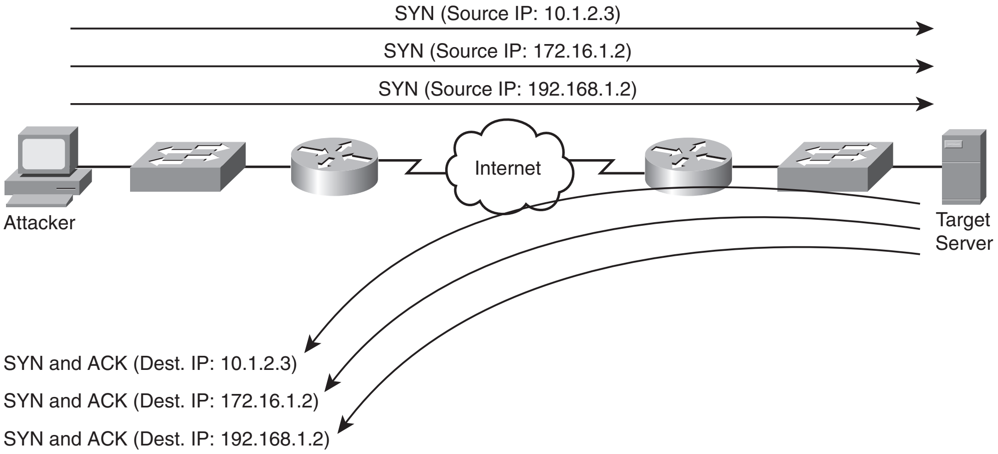
SYN (Source IP: 10.1.2.3)
SYN (Source IP: 172.16.1.2)
SYN (Source IP: 192.168.1.2)
Internet
Attacker
Target Server
SYN and ACK (Dest. IP: 10.1.2.3)
SYN and ACK (Dest. IP: 172.16.1.2)
SYN and ACK (Dest. IP: 192.168.1.2)

- **ICMP attacks**: Many networks permit the use of ICMP traffic (for example, ping traffic), because pings can be useful for network troubleshooting. However, attackers can use ICMP for DoS attacks. One ICMP DoS attack variant called "the ping of death" uses ICMP packets that are too big. Another variant sends ICMP traffic as a series of fragments in an attempt to overflow the fragment reassembly buffers on the target device. Also, a "Smurf attack" can use ICMP traffic directed to a subnet to flood a target system with ping replies, as shown in Figure 1-12. Notice in the figure that the attacker sends a ping to the subnet broadcast address of 172.16.0.0/16. This collection of pings instructs devices on that subnet to send their ping replies to the target system at IP address 10.2.2.2, thus flooding the target system's bandwidth and processing resources.

NOTE For illustrative purposes, Figure 1-12 shows only three systems in the subnet being used for the Smurf attack. However, realize that thousands of systems could potentially be involved and send ping replies to the target system.

Figure 1-12 *Smurf Attack*

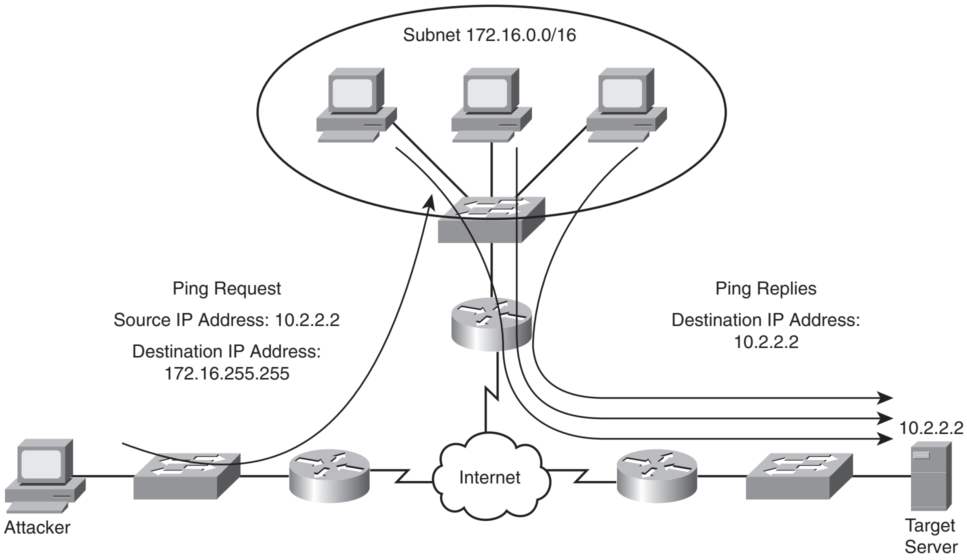
Subnet 172.16.0.0/16
Ping Request
Source IP Address: 10.2.2.2
Destination IP Address: 172.16.255.255
Ping Replies
Destination IP Address: 10.2.2.2
10.2.2.2
Internet
Attacker
Target Server

- **Electrical disturbances**: At a physical level, an attacker could launch an availability attack by interrupting or interfering with the electrical service available to a system. For example, if an attacker gained physical access to a data center's electrical system, he might be able to cause a variety of electrical disturbances:
  - **Power spike**: Excess power for a brief period of time
  - **Electrical surge**: Excess power for an extended period of time
  - **Power fault**: A brief electrical outage
  - **Blackout**: An extended electrical outage
  - **Power sag**: A brief reduction in power
  - **Brownout**: An extended reduction in power

To combat such electrical threats, Cisco recommends that you install uninterruptible power supplies (UPS) and generator backups for strategic devices in your network. Also, you should routinely test the UPS and generator backups.

- **Attacks on a system's physical environment**: Attackers could also intentionally damage computing equipment by influencing the equipment's physical environment. For example, attackers could attempt to manipulate such environmental factors as the following:

  — **Temperature**: Because computing equipment generates heat (for example, in data centers or server farms), if an attacker interferes with the operation of the air conditioning system, the computing equipment could overheat.

  — **Humidity**: Because computing equipment is intolerant of moisture, an attacker could, over time, cause physical damage to computing equipment by creating a high level of humidity in the computing environment.

  — **Gas**: Because gas can often be flammable, if an attacker injects gas into a computing environment, small sparks in that environment could cause a fire.

  Consider the following recommendations to mitigate such environmental threats:

  — Computing facilities should be locked (and inaccessible via a dropped ceiling, a raised floor, or any other way other than a monitored point of access).

  — Access should require access credentials (for example, via a card swipe or a fingerprint scan).

  — Access points should be visually monitored (for example, via local security personnel or remotely via a camera system).

  — Climate control systems should maintain temperature and humidity and send alerts if specified temperature and humidity thresholds are exceeded.

  — The fire detection and suppression systems should be designed not to damage electronic equipment.

Best-Practice Recommendations

You now have a fundamental understanding of threats targeting network and computing environments. Cisco recommends the following best practices to help harden the security of your network:

Key Topic

- Routinely apply patches to operating systems and applications.
- Disable unneeded services and ports on hosts.
- Require strong passwords, and enable password expiration.
- Protect the physical access to computing and networking equipment.
- Enforce secure programming practices, such as limiting valid characters that can be entered into an application's dialog box.
- Regularly back up data, and routinely verify the integrity of the backups.
- Train users on good security practices, and educate them about social engineering tactics.
- Use strong encryption for sensitive data.
- Defend against technical attacks by deploying hardware- and software-based security systems (for example, firewalls, IPS sensors, and antivirus software).
- Create a documented security policy for company-wide use.