Contents

Topic 1: Scenario

A Stealthy Attack

Enterprise Network Intrusion Prevention Systems
CSEC 630 – Module 1

A Stealthy Attack

D&A Laboratories is a Fortune 500 company that provides laboratory services to the pharmaceutical industries. Despite the fact that D&A has an external firewall in place, a hacker, Bill West, managed to gain access to D&A's internal network.

Kate Simons, an information security analyst for D&A, quickly detected signs of Bill's attack using intrusion detection methods and tools. How did Bill manage to bypass D&A's firewall? Did Kate discover the attack in time to prevent any data loss?

Scenario

The names and companies used in the scenario and throughout the rest of the module are fictitious; any similarities to actual individuals or companies are coincidental.

The Sequence of Events

On a Thursday morning, D&A Laboratories was the target of a stealth attack by hacker Bill West. You will see how Bill infiltrated D&A's corporate network and how Kate Simons, an information security analyst for D&A, discovered the attack. Explore Bill's Web site to get a better understanding of him and his attack.

Profile of Bill West (age 26)

I'm based in Portland, Oregon. Currently between positions, I became interested in hacking when I figured out how to exploit a local coffee shop's "Pay for Wi-Fi" service. Using an open Domain Name System (DNS) connection, I was able to access the coffee shop's locked network.

Profile of Kate Simons (age 28)

Kate Simons has been working as an information security analyst for D&A Laboratories for two years. She has more than six years of experience in systems and network administration. At D&A, Kate analyzes malware, tracks botnets, and handles cyber attacks. As an information security analyst, Kate focuses on information assurance, network analysis, and security solutions for cyber threats.

She strongly believes in regularly reviewing the event logs because she knows that firewalls are only one form of defense, which can be bypassed by clever hackers.

Network Diagram

The diagram represents the internal network configuration of D&A Laboratories.

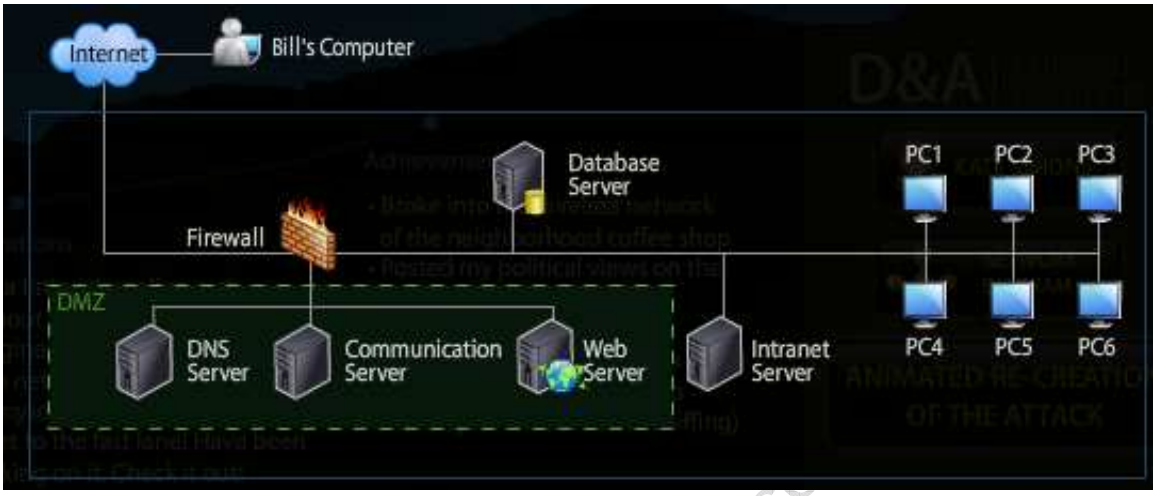

The network is comprised of a database server, intranet server, Web server, DNS server, and a communication server that serves 1,000 clients. The DMZ comprises the DNS server, Web server, and the communication server. The firewall connects the internal corporate network and the DMZ, which is accessible to the outside world, to the Internet.

Find out how Bill carried out the attack on D&A's network.

Step 1
Bill uses Scapy, a software program that can perform anonymized DNS queries.

Step 2
Bill decides to see how D&A's DNS server will respond if he queries the location of a nonexistent internal Web page. After a few attempts, Bill is able to break through the firewall.

Bill had supplied data about one of D&A's Web pages, so the query was able to fetch the DNS server's root certificate, allowing Bill to masquerade as a trusted party and an authoritative source for general information about the company's domain. In essence, the server was tricked into believing it was responding to a legitimate Web page request—when actually it was taking instructions from Bill.

Bill successfully managed to infiltrate D&A's network without any resistance from the firewall.

Step 3
Kate begins her routine analysis of the logs using Windows Event Viewer. She also inspects a network packet capture that she had previously set with a filter to separate the packets that have been allowed to enter and the packets that have been dropped.

She knows that firewalls by themselves could be an insufficient defense against sophisticated attacks and prefers to rely on regular analysis of logs and packet captures to check for intrusions.

Step 4
Kate reviews the records obtained from Syslog on server.DANDALABS.net. The original host names and MAC addresses have been hidden.

Syslog reports changes in MAC address for server.DANDALABS.net:

May 27 12:17:11 mail.DANDALABS.net /bsd: arp info overwritten for server.DANDALABS.net by **00:00:00:00:00:01** on fxp0
May 27 12:17:20 mail.DANDALABS.net /bsd: arp info overwritten for server.DANDALABS.net by **00:00:00:00:00:02** on fxp0
May 27 12:20:20 mail.DANDALABS.net /bsd: arp info overwritten for server.DANDALABS.net by **00:00:00:00:00:01** on fxp0
May 27 12:21:01 mail.DANDALABS.net /bsd: arp info overwritten for server.DANDALABS.net by **00:00:00:00:00:02** on fxp0
May 27 12:25:07 mail.DANDALABS.net /bsd: arp info overwritten for server.DANDALABS.net by **00:00:00:00:00:01** on fxp0
May 27 12:31:36 mail.DANDALABS.net /bsd: arp info overwritten for server.DANDALABS.net by **00:00:00:00:00:02** on fxp0
May 27 12:42:34 mail.DANDALABS.net /bsd: arp info overwritten for server.DANDALABS.net by **00:00:00:00:00:01** on fxp0

Kate discovers that the address for server.DANDALABS.net is being spoofed, or maybe another host on the network is attempting to use DANDALABS.net's IP address.

Step 5
It turns out that it was all a false positive, even though it looked like an attack because of the network traffic. Kate discovered that the anomaly picked up by the intrusion detection system (IDS) was caused by a load-balancing server that was not correctly configured.

Step 6
On analyzing the records, Kate finds that someone was trying to reconfigure the packets to use port 53 to gain access into the internal network.

> May 27 12:51:32 localhost portsentry[123]: attackalert: SYNC/Normal scan from host: 195.76.27.44/195.76.27.44 to TCP port: 53

She noticed that traffic used the TCP protocol with the source port of 65535. Also, during each particular scan, the attacker seemed to have used the same TCP sequence number and IP address for each packet.

Kate was aware that port 65535 is well-known as a port number used by hackers.

Reference: Zeltser, L. (2005). *Intrusion detection analysis: A case study*. Retrieved from http://zeltser.com/intrusion-detection-analysis/

Topic 2: Module Introduction

As Kate discovered, firewalls alone cannot completely secure a network. Even the addition of an intrusion detection and prevention system may not safeguard data. An industry best practice to address this problem is to implement another line of defense by segmenting sensitive data on the network. Implementing a layered defense approach is preferred, as this defense combines several tools and mechanisms to thwart potential intrusions.

This module starts with the elements of an Internet Protocol (IP) packet and how it can be spoofed. It then discusses types of firewalls, their benefits and limitations, and common techniques used to circumvent them. Finally, the module discusses approaches to detecting intrusions in the network by examining intrusion detection and prevention systems.

Topic 3: TCP/IP

TCP/IP Basics

The Transmission Control Protocol/Internet Protocol (TCP/IP) suite contains several protocols, including User Datagram Protocol (UDP) and Address Resolution Protocol (ARP). TCP/IP is the de facto standard used to communicate on the Internet.

Internet Protocol Addresses
When you access a Web site on the Internet, the IP address of your computer and other computers on your network is communicated to the server hosting the Web site. Likewise, the IP addresses of the computers that you access are also communicated to you. Computers store these IP addresses in the decimal notation in the form of a 32-bit address, which contains three parts—the network number, subnet number, and machine number or host address. For example, consider an IP address XXX.XXX.XXX.XXX, which is 4 octets long. Each octet is 8 bits, making a total of 256 decimal (with index at 0 and the maximum value going up to 255).

An organization can use some of the bits in the host segment of the address to identify a specific subnetwork, or subnet, which is a separate part of an organization's network. A subnet can be used to divide local area network (LAN) segments across multiple physical locations, allowing an organization to connect to the Internet through one external-facing IP address.

The table illustrates the structure of an IP address.

IP Address	150.4.13.9			
IP Component	**Network Address – Octet 1**	**Network Address – Octet 2**	**Subnet**	**Host Address**
	150	4	13	9
Value	$2^7 + 2^4 + 2^2 + 2^1$	2^2	$2^3 + 2^2 + 2^0$	$2^3 + 2^0$
Binary	10010110	00000100	00001101	00001001

The original IP address standard was IP version 4 (IPv4), which is still used for most Internet addresses. However, because of the need for more addresses, a shift toward the more robust standard IP version 6 (IPv6) has begun.

Here is some information related to IPv4.

Larger networks need to accommodate more IP addresses. Private IP address spaces were originally created to delay IPv4 address exhaustion, but they are also a feature of the next generation IP, IPv6. Private IP addresses are created for specific purposes, such as when globally routable addresses are not mandatory and when IP addresses are not available for the intended network applications. Computers in homes, offices, and enterprise LANs primarily use these addresses.

This table lists the reserved address space for private networks.

Address Range	Reason
10.0.0.0 - 10.255.255.255	Private Network
127.0.0.1	Local Loopback
169.254.0.0 - 169.254.255.255	Private Network
172.16.0.0 - 172.31.255.255	Private Network
192.168.0.0 - 192.168.255.255	Private Network
240.0.0.0 - 255.255.255.255	Private Network

TCP Segment

TCP uses IP as an underlying protocol. While TCP allows an application to transmit data in an unstructured sequence of bytes, IP requires that the data be transparently packaged as discrete messages. TCP messages are also known as segments, indicating that each segment is a portion of the overall data stream.

TCP segments are extremely flexible and can perform more than one function. They are designed to carry control information and data at the same time, reducing the number of segments sent.

Here is an example: TCP does not need to send separate acknowledgments because every TCP message includes an acknowledgment. Similarly, a message can combine a request to terminate a connection with data being sent to another device.

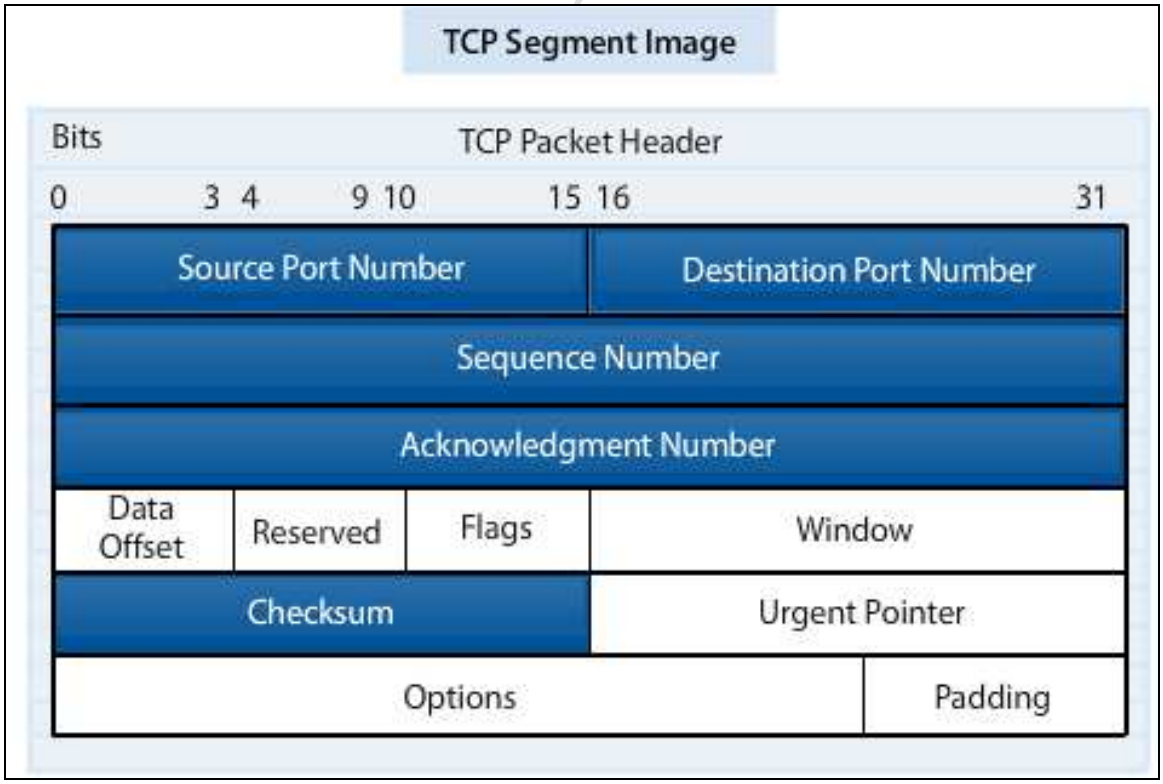

This table provides a description of the key fields in the TCP Header.

Field	Size	Function
Source port	16 bits	Represents the TCP port used for connection by the segment creator
Destination port	16 bits	Represents the TCP port used for connection by the segment recipient
Sequence number	32 bit	Uniquely identifies the TCP segment
Acknowledgment number	32 bit	Indicates the sequence number that the segment creator expects to receive next
Checksum	NA	Contains the checksum value of the segment

Ports

Ports are unique identifiers for applications transmitting data on the systems. They identify the source and destination of a connection and transmit information to an application or service. Ports are numeric identifiers, ranging from 1 to 65,535.

An application reserves a 16-bit unsigned port on either end of a TCP connection. The receiving application associates TCP data packets with a specific TCP connection through its sockets. Sockets include the source host address, source port, destination host address, and destination port.

A server, therefore, can cater to several clients as long as the client establishes simultaneous connections to a single destination port from different source ports.

Some well-known registered port numbers and their corresponding services are displayed in the table.

Port	Application
TCP port 80	HTTP (Web)
TCP port 20 and 21	FTP (Port 20 is the data channel while port 21 is the control channel)
TCP port 22	SSH
TCP port 23	Telnet
TCP port 53	DNS
TCP port 161	SNMP

Topic 3: TCP/IP

TCP/IP Working

Three-Way Handshake

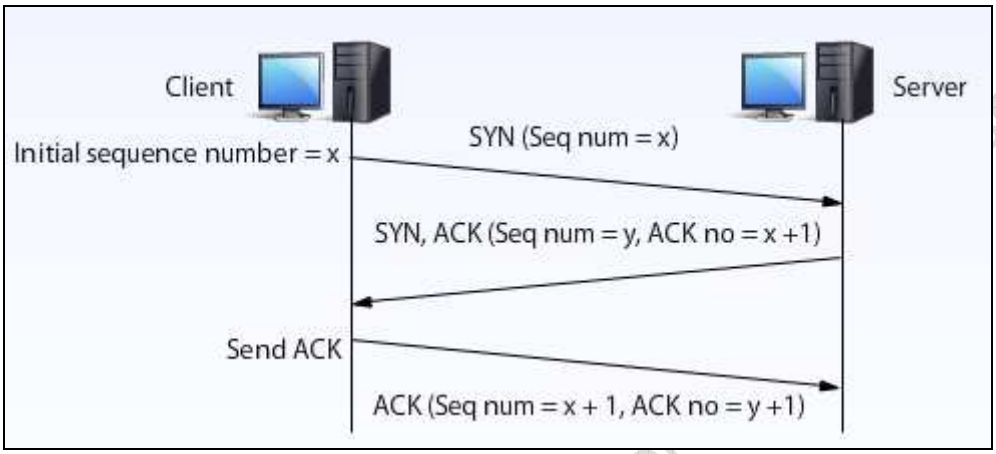

To establish a reliable connection between a client and a server, the three-way handshake occurs before any data exchanges. The three-way handshake consists of three steps shown below.

Step 1: Connection Request
The client chooses a random number x, called an initial sequence number (ISN) and initiates a connection by sending a packet with an SYN bit set and the ISN "x."

Step 2: Connection Acknowledgment
As soon as the server receives the SYN packet from the client, it takes the following actions:

1. The server chooses a random number, y. The number y is used for its own TCP sequence number.
2. The server responds with an SYN, ACK packet with the sequence number "y" and acknowledgment number (ACK No) "x + 1."

Note that this SYN, ACK packet has both SYN and ACK flags set to 1.

Step 3: Data Transmission
The client sends an ACK packet with the sequence number "x + 1" and the acknowledgment number (ACK No) "y + 1."

Note that the ACK packet has the ACK flag set to 1. At this point, the connection is established between the client and server, and the client is able to send data through the connection.

Retransmission

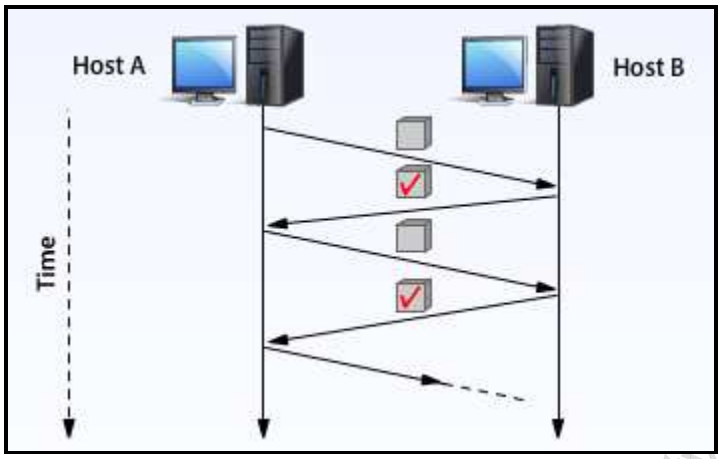

In addition to ensuring that the requested data reaches the client, TCP generates repeat requests for data that might be dropped by IP during transmission. This is possible because TCP:

- Assigns unique sequence numbers to segments based on the order in which they are received
- Ensures that the sender gets notified when the client receives the requested data

The steps through which TCP ensures retransmission are:

1. The sender packages data into a TCP segment.
2. The segment is labeled with the next sequence number.
3. The segment is handed to IP for delivery across the Internet.
4. The receiver receives the segment and makes sure it is not corrupted.
5. The receiver sends an acknowledgment segment back to the sender indicating that it received the specific segment.

Topic 3: TCP/IP

Internet Core Protocols

Software applications transmit data through the application layer. TCP and UDP operate at the transport layer. There are several other protocols that also provide specialized services, including:

1. UDP
2. DNS
3. HTTP
4. HTTPS
5. FTP
6. SMTP
7. Telnet
8. Secure Shell

Here is information on each of these.

1. **User Datagram Protocol**
   UDP is a stateless connection. This means that if a data packet in the sequence is dropped, the sender will not resend the lost packet. Further, if the data packet is delivered to the recipient, no acknowledgment is received. This process is in contrast to TCP, which is very reliable.

   UDP, however, is preferred for transmitting very small packets of data quickly, especially when it is not critical to track the delivery of all the packets, such as in the case of Voice over IP and streaming videos, which are time-sensitive and can tolerate some data loss.

   This image depicts the structure of the UDP header.

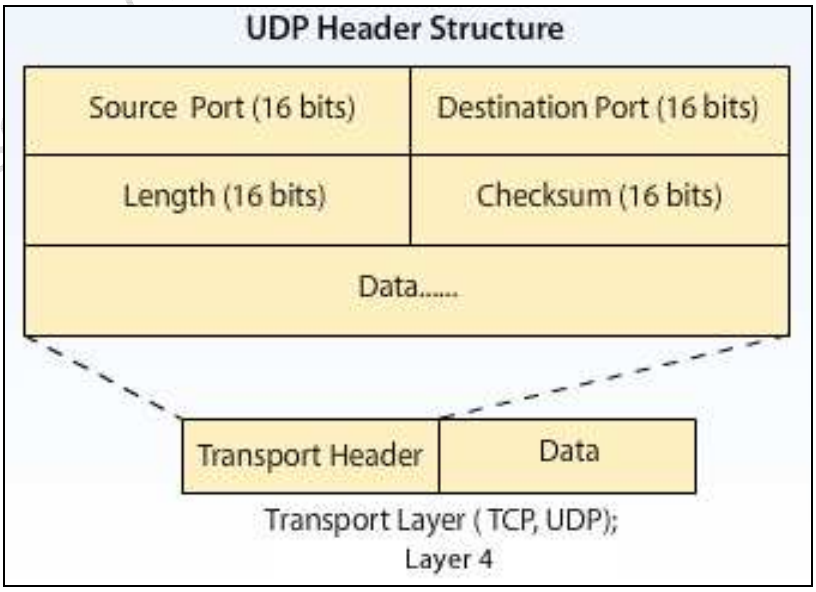

2. **Domain Name Service**
   Domain Name Service (DNS) is a protocol used to convert human-readable domain names into computer-readable IP addresses. For example, humans would recognize www.umuc.edu as a valid Web address, but computer systems would need the address in an understandable format such as 131.171.9.150.

   Domains are important to the DNS system. Some important domains include .com, .net, .org, and .edu.

3. **Hypertext Transfer Protocol**
   Hypertext Transfer Protocol (HTTP) transfers content over the Internet to clients. HTTP servers use TCP port 80 to communicate with client applications. HTTP uses eight commands to transmit pages. For example, GET and POST are methods for providing Web form data to the server. GET can also be used to download content from the Internet.

4. **Secure Hypertext Transfer Protocol**
   Secure Hypertext Transfer Protocol (HTTPS) supports secure transmission of confidential information, such as credit card and Social Security numbers, over the Internet. It works in the same manner as HTTP, except that:
   - All the data being transmitted is encrypted using asymmetric key encryption.
   - It uses TCP port 443.

   HTTPS guarantees that the data being exchanged between a client and the server cannot be intercepted by any unrelated entity.

5. **File Transfer Protocol**
   File Transfer Protocol (FTP) enables users to efficiently exchange files over the Internet. However, because FTP does not encrypt data, it is not considered a secure protocol. Like HTTP, it can be intercepted by eavesdropping. For transmission of confidential files, Secure FTP (SFTP) and Secure Copy Protocol (SCP) are more secure than FTP.

   Although FTP can operate on any TCP port, it is typically assigned port 21.

6. Simple Mail Transfer Protocol

Simple Mail Transfer Protocol (SMTP) is most commonly used for electronic mail exchange between servers. It works over TCP port 25.

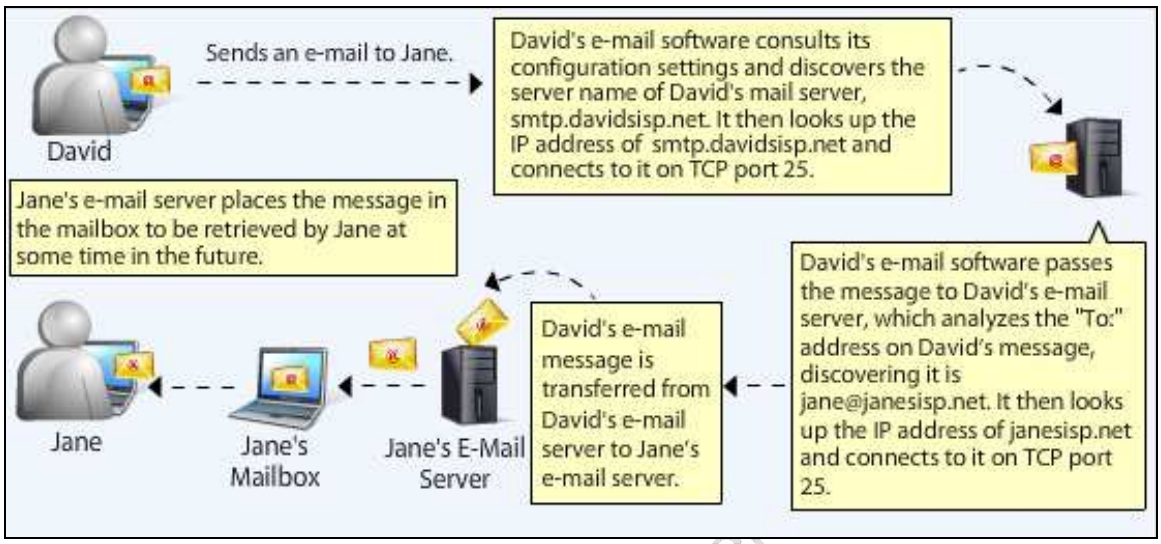

Any of the following methods may be used to download messages from the server to the recipient's mailbox:

- Post Office Protocol (POP), which deletes messages from the server after downloading them to the recipient's mailbox
- Internet Message Access Protocol (IMAP), which retains messages on the server even after they have been downloaded in the recipient's mailbox

7. Telnet

Telnet provides a means to deliver information from the client to the command shell. It runs over TCP. Telnet servers typically listen on port 23.

However, Telnet is rarely used now because it is an insecure protocol, comparable to a TCP connection that provides no data encryption and requires no authentication. Thus, it is vulnerable to eavesdropping attacks in which the attacker would know everything that the user typed, including the user name and password.

8. Secure Shell

Secure Shell (SSH) can be thought of as a secure version of Telnet. However, SSH is resistant to attacks by eavesdroppers. Although SSH provides security through data encryption, the encryption is only as strong as the passwords and encryption keys that users choose.

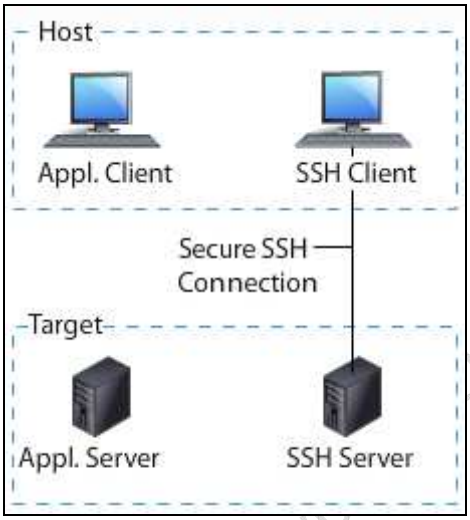

Topic 3: TCP/IP

Three-Way Handshake: Walk-Through

Consider a Web browser with the IP address 192.168.195.1 initiating a three-way handshake to connect a Web server with the IP address 192.168.195.130 using the packet analyzer Wireshark. Three packets (SYN, SYN-ACK, and ACK) are captured during the three-way handshake.

Step 1: Establishing a Connection

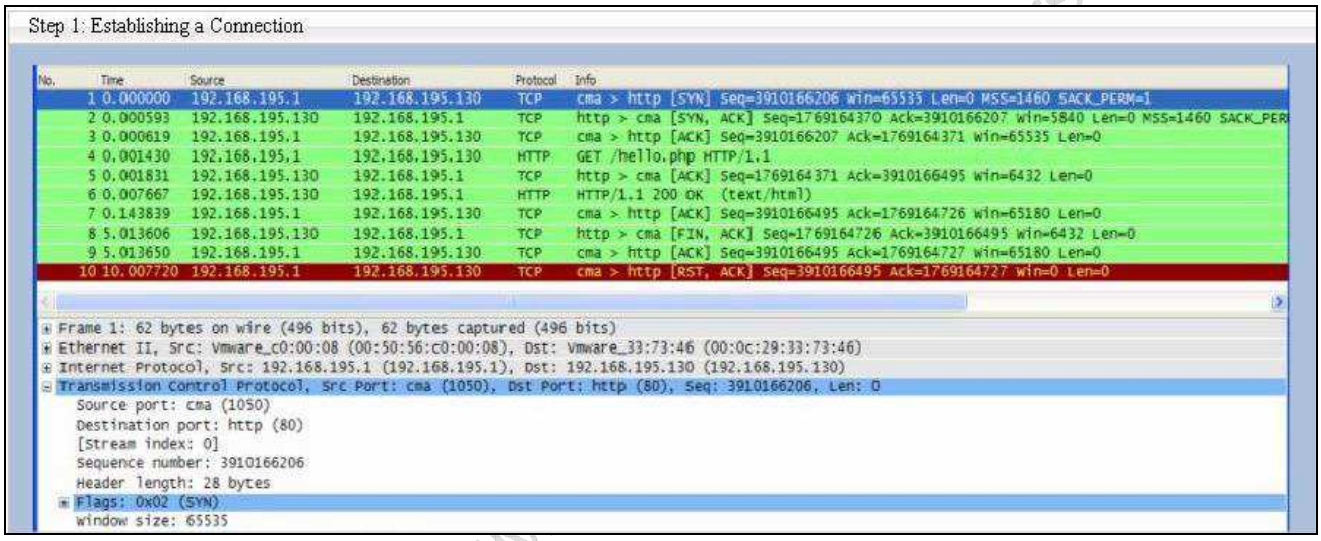

Reference: Wireshark product screenshot reprinted with permission from the Wireshark Foundation.

In this step, the client initiates a connection by sending an SYN packet to the server. The SYN packet contains the initial sequence number (ISN) of the packet. For the connection to be established, the sequence numbers of the client and the server need to be synchronized.

The details of this packet are highlighted in blue in the image. The sequence number in this case is 3910166206.

Step 2: SYN-ACK

In this step, the server acknowledges the client's synchronization request. This acknowledgment is proof that it is a response to the SYN sent by the client. Concurrently, the server sends its request, along with its sequence number, to the client for synchronization.

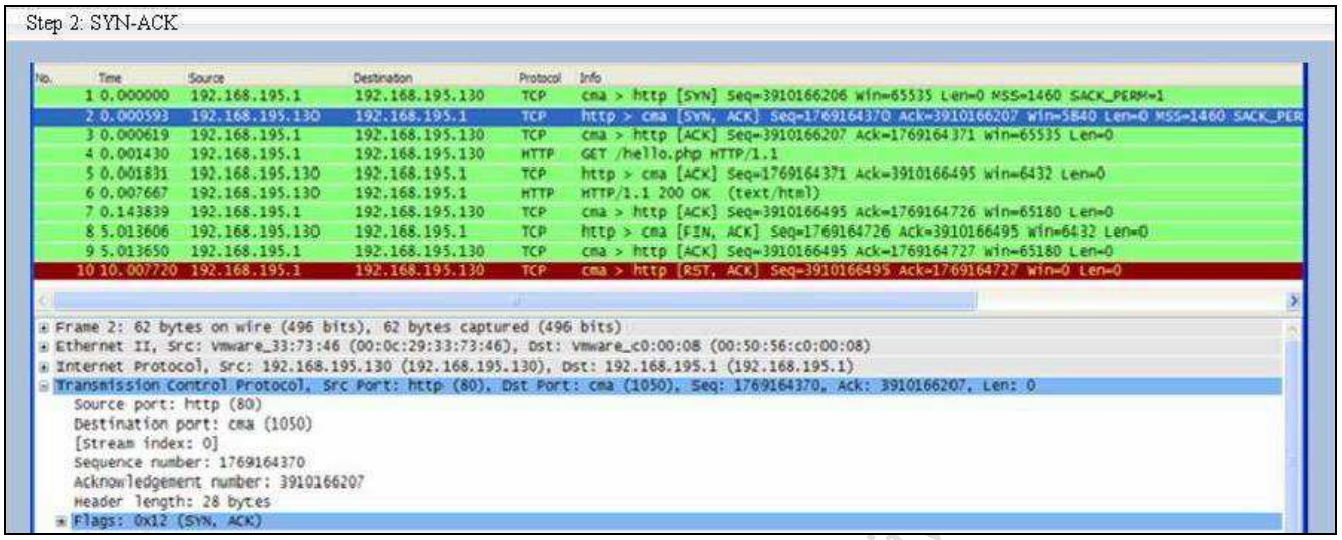

Reference: Wireshark product screenshot reprinted with permission from the Wireshark Foundation.

The server uses the sequence number sent by the client, adds one to it, and sends it back to the client as the acknowledgment number. In this case, the acknowledgment number is 3910166207.

Step 3: Data Transmission

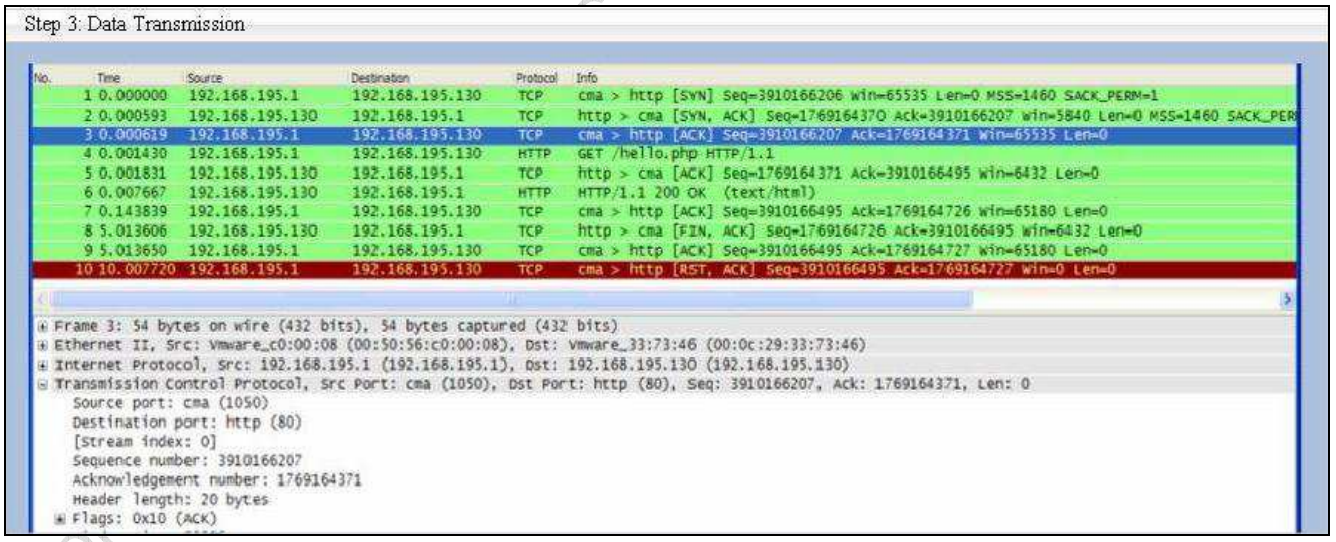

Reference: Wireshark product screenshot reprinted with permission from the Wireshark Foundation.

In this step, the client acknowledges the server's request for synchronization. The client generates the acknowledgment number in the same manner as the server in the previous step.

This step completes the process of the three-way handshake, which establishes a reliable connection.

Topic 4: Firewalls

Purpose and Types of Firewalls

Introduction

A firewall provides a virtual barrier between a network's internal resources and the "untrusted" outside world. It works with the help of very specific rules applied to routers. These rules ensure that no unwanted traffic gets into the system while allowing incoming traffic in response to DNS queries originating from within the network. The rules also ensure that data that should not leave the network (for example, information that a virus might be trying to transmit out from a system) is contained within the system. As a general principle, the routing mechanism determines the destination of the data. Once the destination is determined, the firewall decides whether the data has the required permission to go to the specified destination.

The main types of firewalls include:

1. Stateful packet inspection firewalls
2. Application-layer firewalls
3. Network address translation firewalls
4. Port address translation firewalls

Here is more information about each type of firewall.

1. **Stateful Packet Inspection Firewalls**
   One of the ways to prevent unauthorized entry of data into a network is to inspect each data packet individually to see whether the intended recipient of the packet actually sent a request for it. Stateful packet inspection (SPI) firewalls use this approach to detect network intrusions.

2. **Application-Layer Firewalls**
   A more efficient and reliable approach to stateful packet inspection is to scan data packets not only at the network layer, but also at the application layer. Application-layer firewalls are built to detect deviations in the normal behavior of application-layer protocols, such as HTTP or SMTP.

   For example, an application-layer protocol for an interactive Web site would be familiar with the usual number of page requests on the site. Thus, it would detect—and identify as abnormal traffic—page requests that exceed the normal number by a large margin. Such detection is not possible with stateful packet inspection, which relies merely on inspecting data packets at the network layer.

3. **NAT and PAT Firewalls**
   Network address translation (NAT) and port address translation (PAT) firewalls work by hiding the internal network addresses used within an organization's LAN from the outside world. By adding this defense measure to the traditional SPI firewall, NAT and PAT firewalls provide greater security.

Network Address Translation Firewall
Network address translation (NAT) firewalls combine two different technologies, an NAT component and a traditional SPI firewall. The role of the NAT component in the firewall is to temporarily replace the contents of the SRC ADDRESS field in the IP address with a public address, thereby "hiding" the system's internal IP address from the external network. This hiding may be done either dynamically, in which public addresses from a global pool are assigned when required by systems, or in a static manner, in which each system has a fixed public address assigned to it.

Port Address Translation Firewall
Port address translation (PAT) is commonly used in cases where there is a need for multiple systems to use the same IP address at the same time. A PAT firewall involves the use of PAT technology to remap and hide the port numbers of the private systems. The hidden ports help make the systems more difficult to attack.

Here is how the PAT firewall works:
1. Systems on the private side of the PAT firewall generate a TCP segment.
2. PAT changes the headers of the TCP segment and changes the source IP address to a global one.
3. PAT sends the TCP segment outside the network to establish a three-way handshake with the server.
4. The server responds to the TCP segment, returning data, which is again received by PAT.
5. PAT forwards the data to the appropriate client.

Thus, at no point does PAT allow the client systems' IP addresses to be exposed to the external network, thereby ensuring that the systems remain protected.

NAT/PAT firewalls automatically log exceptions and deviations from preset rules. These logs are useful for discovery and analysis of security incidents. These firewalls can also be configured to generate alerts to bring intrusions to the attention of the network administration.

Although safer than the traditional SPI firewalls, NAT and PAT firewalls are not impervious to attacks unless coupled with other modern defenses. For example, an NAT/PAT firewall would be ineffective in protecting the systems in the private network from attacks by malicious software launched from an employee's e-mail because the attack comes directly from within the secure network. Similarly, infected systems can bypass the firewall and spread the infection.

Topic 4: Firewalls

Activity: Firewalls

Here are some questions that will help you check your knowledge about firewalls.

Question 1: What is the primary purpose of a firewall?
a. To prevent the spread of malware, such as viruses, worms, and Trojans
b. To scan for possible intrusions and to take protective action when they are observed
c. To secure data while it is being transmitted over the Internet
d. To protect a single computer or an entire network from unwanted network traffic

Correct answer: Option d

Feedback:
The primary purpose of a firewall is to protect a single computer or an entire network from unwanted network traffic.

Question 2: What does an application-layer firewall do?
a. It prevents users from using applications that require lots of bandwidth, such as Internet radio and video.
b. It watches application-layer traffic such as HTTP and SMTP for signs of abnormal activity, which could indicate that an attack is under way.
c. It examines data at the network and transport layer, watching for suspicious network traffic.
d. It uses complex mathematics to protect data at the application layer after it has left the local network.

Correct answer: Option b

Feedback:
An application-layer firewall watches application-layer traffic such as HTTP and SMTP for signs of abnormal activity, which could indicate that an attack is under way.

Question 3: What is the advantage of using a stateful packet inspection firewall?
a. The ability to monitor the network for abnormal activity and to take steps to block a potential attack from spreading
b. The ability to securely and invisibly bridge two physically separate networks using the Internet
c. The ability to examine application-layer traffic, such as HTTP or SMTP, and monitor that traffic for abnormal activity
d. The ability to track the status of each connection and block invalid packets from reaching their intended target

Correct answer: Option d

Feedback:
A stateful packet inspection firewall has the ability to track the status of each connection and block invalid packets from reaching their intended target.

Question 4: What is the difference between an NAT firewall and a PAT firewall?

a. A PAT firewall does not prevent unsolicited incoming traffic from reaching its intended target; an NAT firewall prevents this type of traffic from passing.
b. A PAT firewall is capable of examining encrypted traffic; an NAT firewall examines packets for malicious content.
c. NAT firewalls allow a single IP address to be shared by multiple computers simultaneously; a PAT firewall allows an IP address to be used by only one computer at a time.
d. PAT firewalls allow a single IP address to be shared by multiple computers simultaneously; an NAT firewall allows an IP address to be used by only one computer at a time.

Correct answer: Option d

Feedback:
PAT firewalls allow a single IP address to be shared by multiple computers simultaneously. An NAT firewall, on the other hand, allows an IP address to be used by only one computer at a time.

Topic 4: Firewalls

Attacks from Outside the Organization

Attackers come up with newer and more sophisticated attacks every day. Most of these attacks work by fooling the internal networks into believing that the attacker is a trusted party. After they are mistakenly identified as trusted parties, the attackers gain access to internal network data, including IP addresses of the systems.

Here are the steps that are carried out in an external attack.

Step 1
A network's internal systems are very likely to interact freely with at least a few external systems:

- Employees' personal computers that use remote desktop to access the corporate network
- Networks of third-party companies that support administrative and outsourcing functions
- Networks of the other branch offices

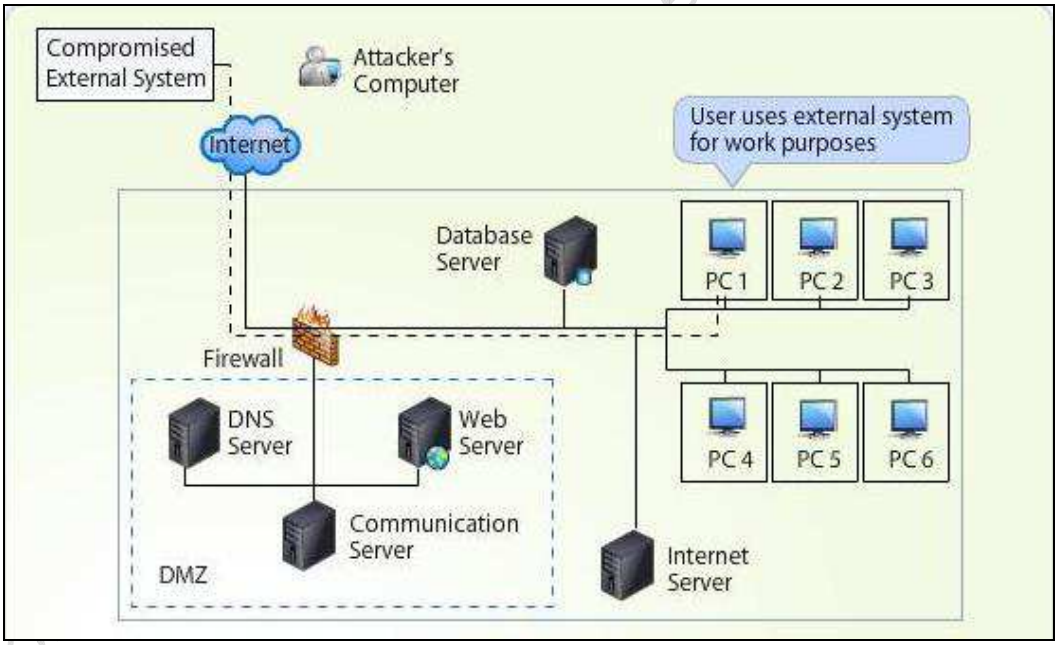

Step 2
When a user from the internal network accesses the trusted (now compromised) external system, a "cookie," which provides information about the connection, is placed on the Web server. This cookie will be used to bypass authentication.

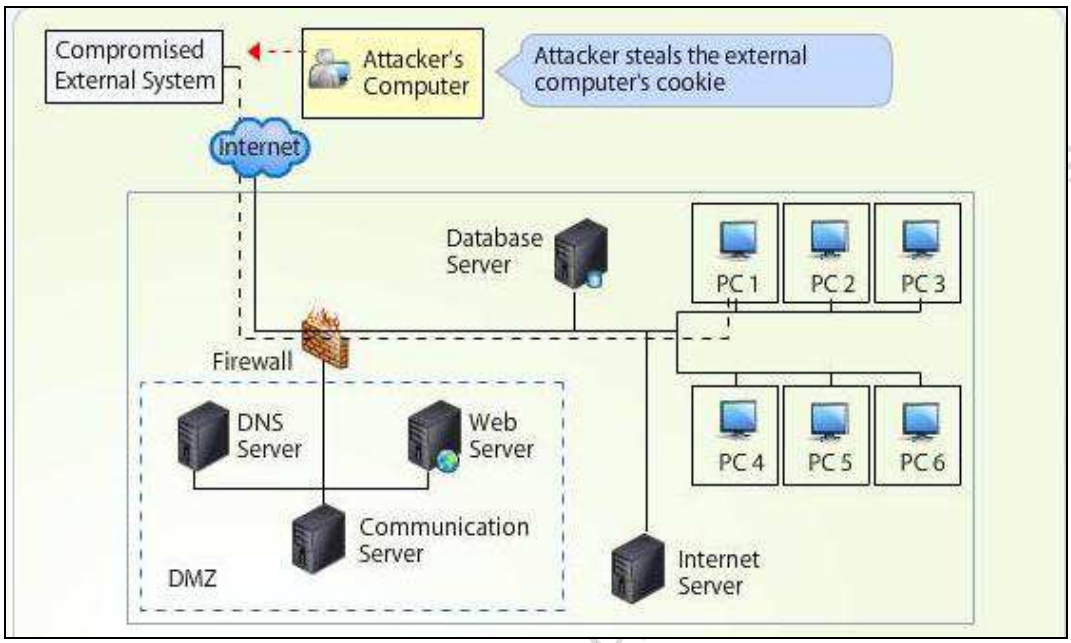

Step 3
The attacker connects to the user's computer, using a protocol that enables multiple users to share the same computer and to connect remotely to the computer.

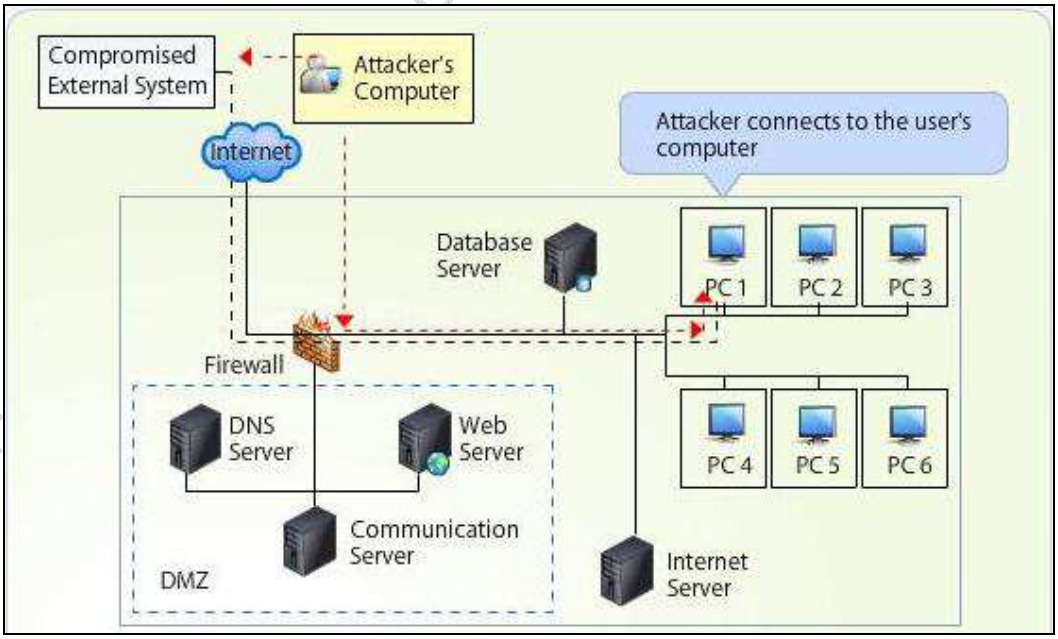

Step 4
The attacker accesses the Web browser on the user's computer and then connects the user's computer to a fake Web server instead of the corporate server.

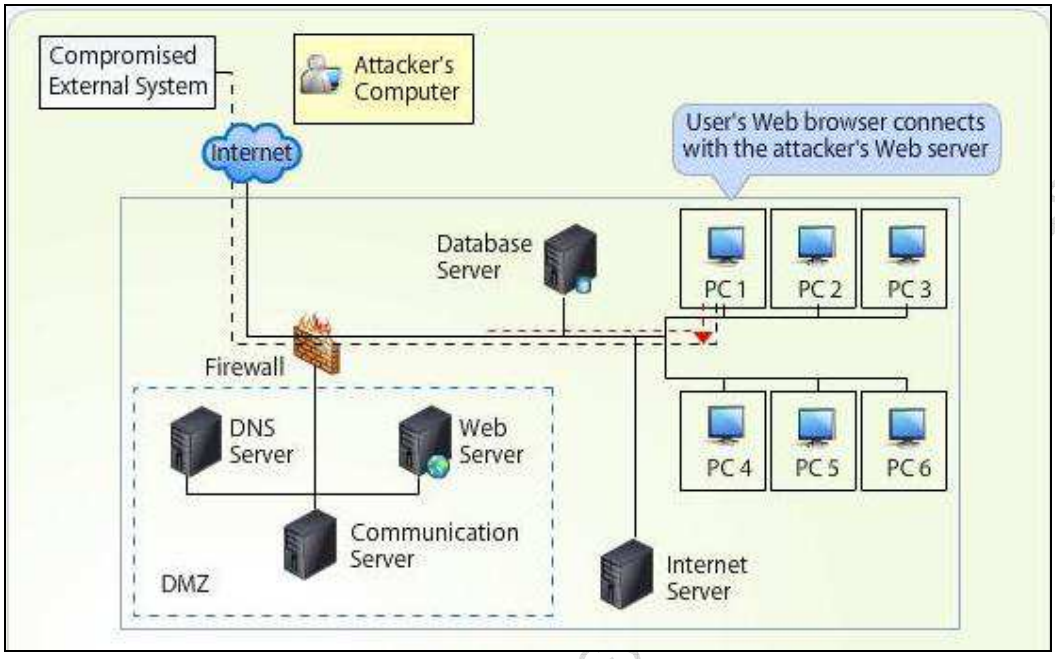

Step 5
The attacker inserts malicious content into the Web page that the user's computer is trying to access.

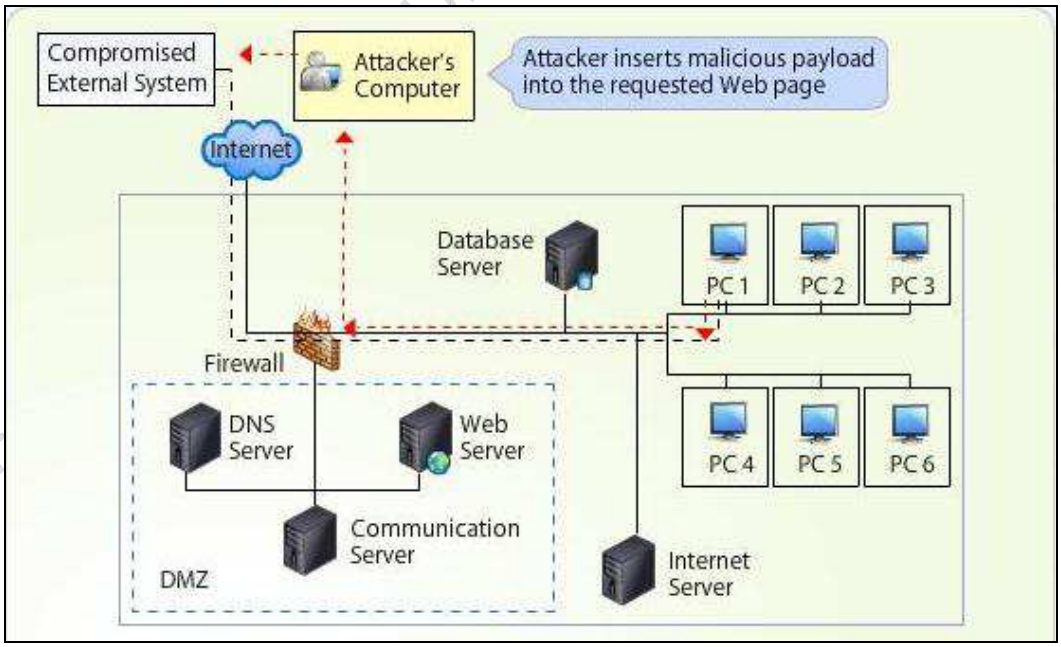

Step 6
The attacker's malicious content—inserted on the Web page—gets executed on the user's computer.

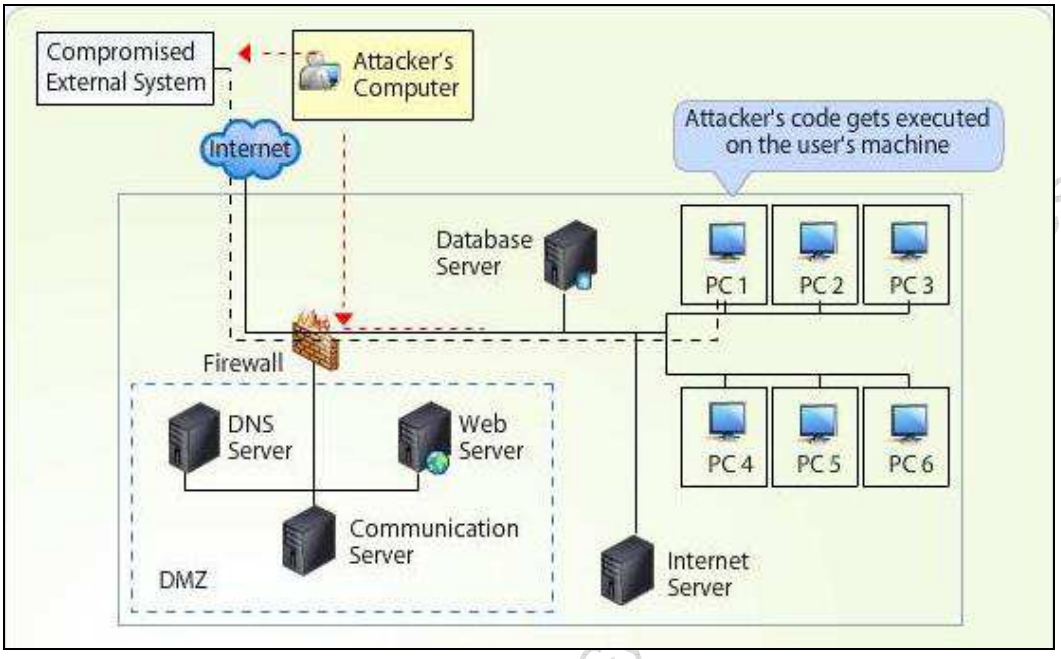

Topic 5: Intrusion Detection Systems

Detection Strategies

Intrusion detection systems (IDSs) protect host computers and networks by alerting the network administrator about illegal probes and attacks.

Here is information related to the common approaches used to detect intrusions.

1. **Signature- (or Rule-) Based Anomaly Detection**
   The approach works on the principle that every component of a network's traffic can be identified based on its signature. Thus, IDSs using signature detection are set up so they respond in different ways to different signatures. For example, an IDS for a private Web site server could be set up so that it generates alerts for incoming traffic that uses any port other than port 80 while allowing requests that are directed to port 80.

   An obvious disadvantage of this approach is that the system, which has been set up for certain signatures, will recognize only those. Attackers, whose signatures are unknown to the IDS, will easily bypass detection. To make sure the IDS can detect the relevant signatures, it is important to keep the IDS updated with the latest attackers' signatures.

   On the other hand, the system could also generate "false positive" alerts when it comes across signatures that are harmless but closely resemble attackers' signatures. These false positives can be avoided by tweaking the sensitivity of the IDS so that it overlooks all but those signatures that match the potential attackers very closely.

   Snort: Signature-Based IDS

   Snort is a network IDS that can detect various attacks and probes based on the signatures (or rules) set up by the information security administrator. It comprises the following components:

   - **Packet decoder:** It processes each captured packet to identify and isolate protocol headers at the data link, network, transport, and application layers.
   - **Detection engine:** It analyzes each packet using rules defined for this configuration of Snort by the administrator.
   - **Logger:** It creates a log including each packet that matches a rule, if specified. The administrator can then use the log file for later analysis.
   - **Alerter:** It can be sent for each detected packet to a file, a socket, or a database.

2. Anomaly-Based Detection

This approach, which is commonly used in network IDSs, is based on the fact that any deviation from a network's usual behavior implies that something in the system is abnormal. The parameters commonly used to quantify a network's behavior include interval timer, resource utilization, etc.

This approach is foolproof compared to signature detection, because it "learns" the normal behavior pattern of a network system over time and then uses this information to highlight abnormalities without the need for continual updates.

After an initial period of learning, anomaly detection-based IDSs become effective at providing alerts, even against completely unknown attacks—unlike signature detection systems—because they're fine-tuned to identify the slightest abnormal behavior in the network.

Because of the way it works, this approach is more likely to generate false positive alerts than the signature-based approach.

Anomaly-Based Detection: Example

In the D&A attack, Kate Simons used anomaly-based detection to identify the hacker. To narrow her search, Kate analyzed the audit logs and finally managed to trace the attacker's IP address. Here are the steps that she followed.

Step 1: Assigning Unique Identifiers

Kate analyzed the data in the Snort alerts as well as in the Snort scan reports. These two sources of data provided different perspectives on the same event. However, because of the thousands of records, Kate decided to focus on the most active hosts. Information from those hosts was parsed with custom-written Perl scripts and saved into database files. Each alert had its own identifier, which was used to index the files. Kate then used the scripts to initialize the environment and parse Snort alert files and scan detection files.

Step 2: Correlating Information

Next, Kate used scripts to track information by source and destination host, by alert name, by activity, and by date of activity. The scripts also provided information on a value from these fields that Kate used during the analysis.

Step 3: Analyzing Information

Kate then inserted the files into a spreadsheet to sort the information. She noted suspicious activity with the help of a mapping program called The Brain.

Reference: Zeltser, L. (2005). *Intrusion detection analysis: A case study*. Retrieved from http://zeltser.com/intrusion-detection-analysis/

Topic 5: Intrusion Detection Systems

Shallow and Deep Packet Inspection

Early intrusion detection systems used shallow packet inspection (SPI) to examine only certain parts of a data packet such as the source, destination, and connection state.

Deep packet inspection (DPI), on the other hand, examines a data packet closely and in its entirety. Even the optional fields in a data packet are examined. Because this kind of inspection consumes a considerable amount of memory and processing capacity, deep packet inspection is limited to cases where a detailed examination is absolutely essential, for example, protocol analysis. To be able to perform protocol analysis, administrators must first enable deep packet inspection.

Topic 5: Intrusion Detection Systems

Honeypots

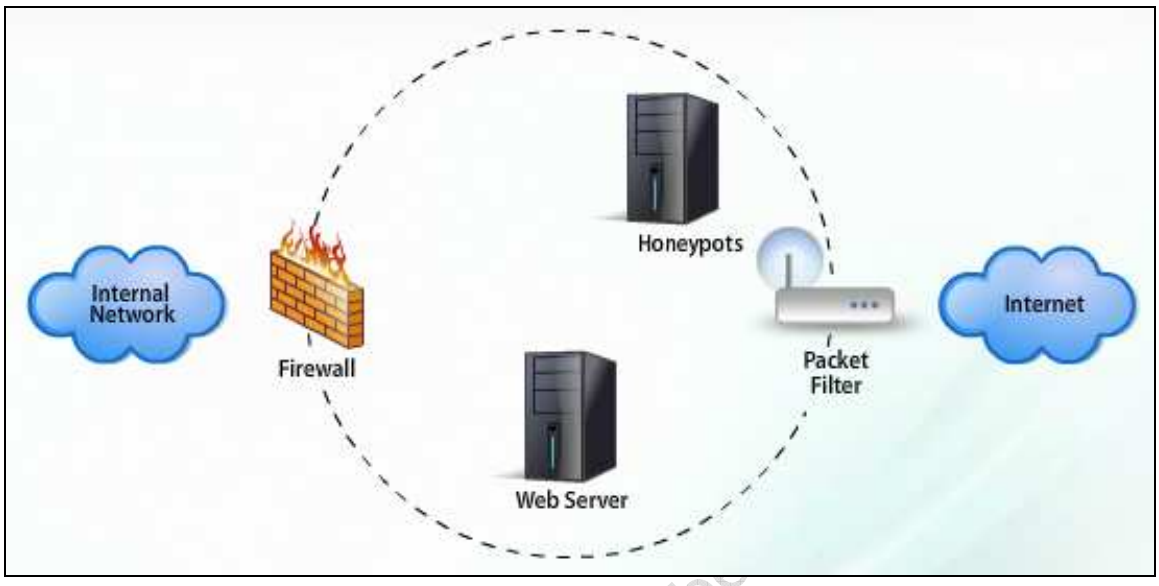

If they have sufficient time and resources, information security analysts try to outsmart attackers by leaving computers or networks as “honeypots” or “honeynets” open for attack. The idea is to lure attackers into attempting intrusions while closely monitoring their activities in order to obtain critical information that will help to identify the fingerprints of potential intruders. This information can be used to make the actual network more resistant to potential attacks, and, possibly, even bait attackers so they can be caught and reported to law enforcement.

Obviously, a honeypot serves its purpose only as long as it is not identified by an attacker. If identified, the honeypot can be easily used to mislead the information security personnel.

Further, it is important to completely segregate the honeypot from the actual network to make sure it does not become an inadvertent launching pad for attacks on the system.

Topic 6: Summary

We have come to the end of Module 1. The key concepts covered in this module are listed below.

- Firewalls alone do not provide an organization with a sufficient defense.
- The Transmission Control Protocol/Internet Protocol (TCP/IP) suite contains several protocols that can be used for various purposes.
- TCP allows an application to transmit data in an unstructured sequence of bytes. TCP segments are extremely flexible and can carry both control information and data at the same time.
- There are two approaches to intrusion detection—signature (rule) based detection and anomaly-based detection—that can be used to analyze and identify intrusions.
- Audit records provide vital information about inbound and outbound traffic in the network, and analyzing these records is an essential way to guard against attacks.
- Honeypots or honeynets are networks that are left open for attack with the intention of luring attackers into attempting an intrusion. By monitoring the honeypot, critical information related to potential intruders can be obtained. Many times, attackers are even baited and reported.

Glossary

Term	Definition
Application-Layer Firewalls	Application-layer firewalls are built to detect deviations in the normal behavior of application-layer protocols, such as HTTP or SMTP.
ARP Cache Poisoning	An attacker uses ARP cache poisoning to gather the information exchanged between computers.
Domain Name Service	The domain name service (DNS) translates human-readable Internet addresses such as www.xyz.com into their Internet Protocol (IP) addresses equivalent.
DNS Poisoning	Domain name poisoning or spoofing takes place when DNS information is replaced with other information. Poisoning can take place through a denial of service (DoS) attack or a zone transfer.
DoS or DDoS	A denial of service (DoS) or distributed denial of service (DDoS) attack floods a target site with large volumes of traffic, which consume the system's resources, slowing it down and rendering service unavailable. In doing do, a DoS or DDoS denies access to legitimate users.
File Transfer Protocol	File Transfer Protocol (FTP) is an application protocol that uses the TCP/IP protocol (or the Internet) to transfer files between computers.
Firewall	A firewall is the hardware or software that prevents unauthorized users from accessing a computer or a network.
Hypertext Transfer Protocol	Hypertext Transfer Protocol (HTTP) transmits Web pages to clients.
Secure Hypertext Transfer Protocol	Secure Hypertext Transfer Protocol (HTTPS) supports secure transmission of confidential information, such as credit card and Social Security numbers, over the Internet by using the SSL protocol in conjunction with HTTP.
IDS	An intrusion detection system (IDS) detects malicious activities on the network and reports them to the system administrator.
Internet Control Message Protocol	The Internet Control Message Protocol (ICMP) integrates with the Internet Protocol (IP). It reports error, control, and informational messages between a host and a gateway.
IPS	An intrusion prevention system (IPS) is an extension of an IDS, but unlike an IDS, which only detects the attack, the IPS prevents the malicious attacks.
IP Address	Internet Protocol (IP) address is a numeric representation used to identify devices within a computer network or Web sites on the Internet. There are two IP address versions that currently exist—IPv4 and IPv6.

Term	Definition
Network Address Translation	Network address translation (NAT) facilitates communication between two networks by translating an IP address used in one network (designated as the inside network) to an IP address used in another network (designated as the outside network).
Network Address Translation (NAT) Firewall	Network address translation (NAT) firewalls include an NAT component, which temporarily replaces the contents of the SRC ADDRESS field in the IP address with a public address, thereby "hiding" the system's internal IP address from the external network.
Port	A port is a hardware circuitry used to link one device with another.
Port Address Translation	Port address translation (PAT) conserves IP addresses by enabling several devices on a local area network (LAN) to map to a single public IP address.
Port Address Translation (PAT) Firewall	A PAT firewall involves the use of PAT technology to remap and hide the port numbers of the private systems. The hidden ports help make the systems more difficult to attack.
Secure Mail Transfer Protocol	Secure Mail Transfer Protocol (SMTP) is mostly used for electronic mail exchange between servers and clients.
Secure Shell	Secure Shell (SSH) is a data exchange protocol that allows data to be exchanged using a secure channel between two network devices.
Socket	A socket is a combination of the host IP address, the network interface, and the port number assigned to it.
Script Kiddie	"Script kiddie" is a derogatory term used to refer to an amateur hacker.
Stateful Packet Inspection Firewalls	Stateful packet inspection firewalls detect intrusions and prevent unauthorized entry of data into a network by inspecting each data packet individually to see whether the intended recipient of the packet actually sent a request for it.
Telnet	Telnet enables remote use and supervision of systems. Network administrators monitor and control systems remotely using Telnet.
User Datagram Protocol	User Datagram Protocol (UDP) is a network protocol that allows computers to communicate over the Internet. UDP packets are smaller than TCP packets because they do not contain data fields to perform integrity checks to ensure that all packets arrived.