Contents

Topic 1: Analogy

The Postal Service and TCP/IP

Advanced TCP/IP and Introduction to Penetration Testing
CSEC 640 – Module 1

The Postal Service and TCP/IP
The analogy of the postal service is commonly used to introduce the topic of data transmission on the Internet. This module develops and applies this analogy to introduce some key concepts.

Keep this analogy in mind as you progress through the module to uncover subtle similarities between the workings of the postal service and the transmission of data over the Internet.

Slide 1
Run an Internet search for the term Transmission Control Protocol/Internet Protocol (TCP/IP), and you will surely see an analogy comparing TCP/IP with the postal service. The following example will help you understand the basis of this popular analogy.

Someone writes a letter, addresses the envelope, and drops it in a mailbox. When the letter arrives, the recipient focuses on the contents of the letter, not how it reached the destination. The sender and the recipient are communicating directly, oblivious to the services and protocols that facilitated the communication. A number of factors are responsible for ensuring that the recipient receives the letter on time and in good condition.

Similarly, with TCP/IP, the recipient could be sitting at a computer in Hong Kong and viewing a Web page from a server in Africa simply by typing a URL in the Web browser. Entering the URL sets off a series of functions that enable the computer to retrieve and display the Web page.

This comparison is the essence of the analogy between the postal service and TCP/IP—the protocol that enables exchange of information over the Internet.

Slide 2
Taking this analogy a step further carries us to some of the main topics covered in this module, such as the Internet Control Message Protocol (ICMP), IP fragmentation, and the Traceroute tool.

For instance, if the sender writes the wrong address on a letter, the letter will be returned with an “address unknown” notification. Similarly, operating systems use ICMP to send error messages and alerts when data is not transmitted successfully.

Slide 3
The postal service sends multiple delivery trucks from one city to another because it would be impossible for one truck to carry all the mail.

IP fragmentation and TCP flow control serve the same purpose of sending data in bite-sized chunks and controlling how many of these “chunks” are sent at one time.

Slide 4

A letter sent overseas passes through multiple post offices before it reaches its intended recipient. The sender’s local post office, the overseas receiver’s local post office, and each post office on the way puts a stamp, or postmark, on the envelope. This postmark shows when each post office dispatched the letter on to the next one.

Consider the Traceroute tool in this context. This tool records the route between two devices on different networks and serves to check connectivity between two devices that may be physically separated, even across continents.

Topic 2: Module Introduction

This module discusses key protocols, tools, and concepts related to TCP/IP, including:

- Internet Control Message Protocol (ICMP)
- IP fragmentation
- TCP flow control mechanisms (such as TCP Window Size and TCP Sliding Window)
- Traceroute

It is important to examine these protocols and tools in order to understand the structure of the TCP/IP packet and gain insight into how the TCP/IP protocol works. This module develops a foundation for upcoming modules that contain in-depth discussions on the security implications of these concepts. This module also briefly discusses the basic steps and procedures in penetration testing.

Topic 3: Advanced TCP/IP

Internet Control Message Protocol

ICMP is one of the most widely deployed protocols of the Internet Protocol Suite. It is used mainly by operating systems or network devices to send error messages or perform simple queries.

The diagnostic tool ping is used to send a simple request to determine whether a host is alive. This request is based on the ICMP communication protocol.

An ICMP request is used to perform a ping against a remote host to verify connectivity. If the remote host is alive, it responds with an ICMP reply. You may run the following command in a Windows command prompt: **ping –n 1 www.umuc.edu**

Note that –n 1 means only one ICMP request will be issued. The number of ping requests can be changed by using the –n switch when executing the ping command. Without the –n switch, ping will issue four ICMP requests by default. One of the purposes of the Windows ping command issuing four ping requests is to give the issuer an average response time.

Since ICMP lacks a TCP header, it does not support reliable data transmission.

Topic 3: Advanced TCP/IP

Comparing TCP and ICMP

ICMP facilitates a simple means of communication between hosts, or a host and a router, to alert them about an error or issue. The unique characteristics of this protocol are:

- Service ports do not have to be activated or listening. Every operating system can respond to an ICMP echo request (ping).
- ICMP supports broadcast traffic.

The main distinguishing factors between TCP and ICMP are as follows:

- ICMP does not have a TCP header. Instead, it has an ICMP header.
- The TCP header has a sequence number and a port number that enables reliable transmission of data and services. However, the main purpose of ICMP is to perform simple queries and report errors—reliable transmission of data is not one of its requirements.

Topic 3: Advanced TCP/IP

The ICMP Packet Structure

The diagram shows the ICMP message header format.

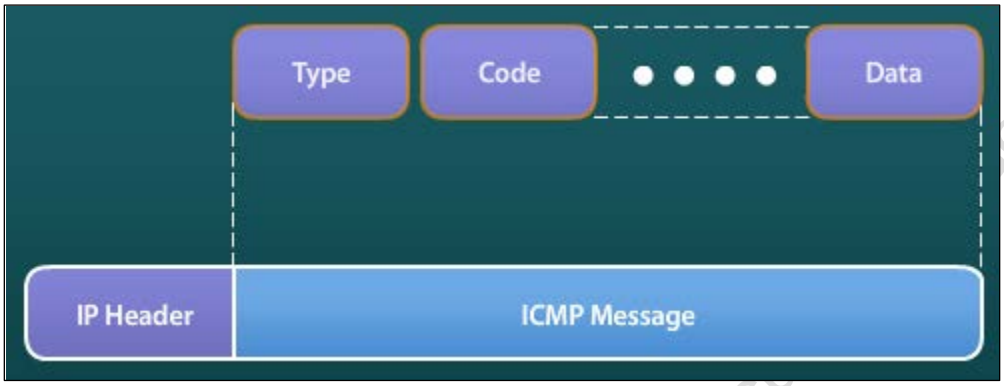

Type
The Type field defines the type of ICMP message.

Code
The Code field determines the subtype of the ICMP message.

Data
In error messages, the Data field includes information that helps identify the original packet that caused the error. In query messages, the Data field carries extra information based on the type of the query.

Topic 3: Advanced TCP/IP

ICMP Queries

ICMP uses two types of queries:

- Echo request or ICMP request
- Echo reply or ICMP reply

Echo requests and echo replies are specifically used to test the IP connection between two hosts. Assume that Host A, the sender, wants to see if it can reach Host B, the target. Host A sends an ICMP echo request to Host B. If Host B is alive and listening, it responds by sending an ICMP echo reply message. When Host A receives this reply message, it knows that it can successfully communicate with Host B.

ICMP includes a pair of echo request and reply messages specifically for testing an IP connection between two hosts. Suppose Host A (sender) wants to see if it can reach Host B (target). Host A sends an ICMP echo request to Host B. If Host B is alive and listening, it responds to Host A (original sender) with an ICMP echo reply message. When Host A receives this reply message, it knows that it is able to communicate successfully with Host B.

ICMP Request

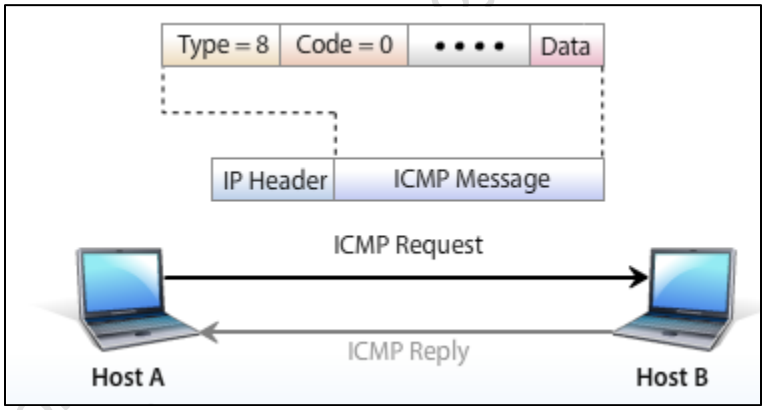

ICMP Reply

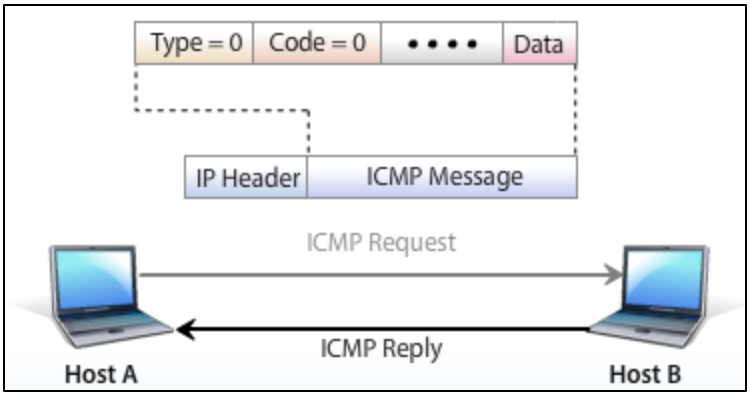

Topic 3: Advanced TCP/IP

ICMP Error Messages

When a host or a device has an error, it uses ICMP to report the error condition. A few important ICMP error message types include types 3, 5, and 11.

Error Type 3

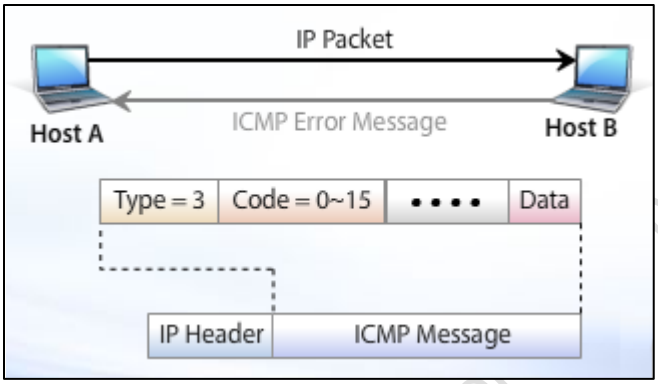

Error Type	Code	Error Description
3	0~15	**Destination Unreachable:** The IP packet cannot reach the destination. The value in the Code field explains the reason for the error.

Destination Unreachable Error Subtypes

Code	Description	Reason for the Error Message
0	Network Unreachable	No routing table entry is available for destination network. Typically, routers send this error message.
1	Host Unreachable	If an IP packet reaches a router in the network a destination host is attached to, and the host is not responding, a “Host Unreachable” message is sent back.
2	Protocol Unreachable	A “Protocol Unreachable” message is generated if the designated transport protocol within a datagram (packet) cannot be supported in the transport layer of the final destination.
3	Port Unreachable	A “Port Unreachable” message is generated if the designated protocol is unable to inform the host (sender) of the incoming message.

Error Type 5

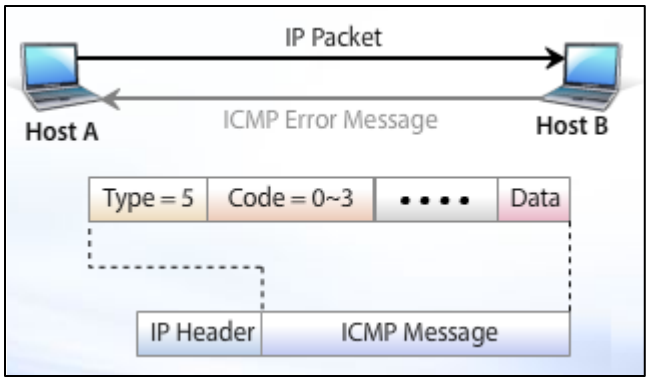

Error Type	Code	Error Description
5	0~3	**Redirect:** An alternative route for the packet is designated. The value in the Code field explains the reason for the route change.

Error Type 11

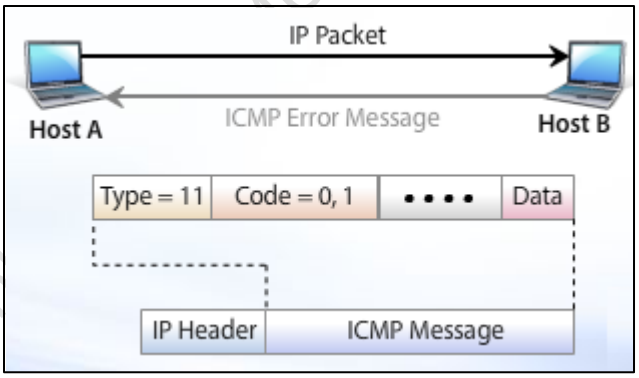

Error Type	Code	Error Description
11	0,1	**Time Exceeded:** This error message is sent when the code value in the Time-to-Live field (or TTL field) of the IP packet becomes 0, or there is a timeout for reassembly of packets.

Topic 3: Advanced TCP/IP

IP Fragmentation

An Internet Protocol (IP) packet is fragmented into a number of pieces if its size exceeds the limit imposed by data link layer protocols such as Ethernet, Fiber Distributed Data Interface (FDDI), and so on. This limit is called the maximum transmission unit (MTU). The theoretical maximum size of the IP datagram is 65,536 bytes, but the data link layer protocol generally imposes a limit that is much smaller. The various data link layer protocols have different MTU values. For instance, the MTU values for Ethernet and FDDI are 1,500 bytes and 4,352 bytes, respectively.

If an IP packet follows a route that includes a network with an MTU value and the packet size exceeds this value, the packet is fragmented into smaller pieces as shown in the diagram. An IP packet can be fragmented at a sender or at intermediate routers. Reassembly of the original IP packet only occurs at a destination host.

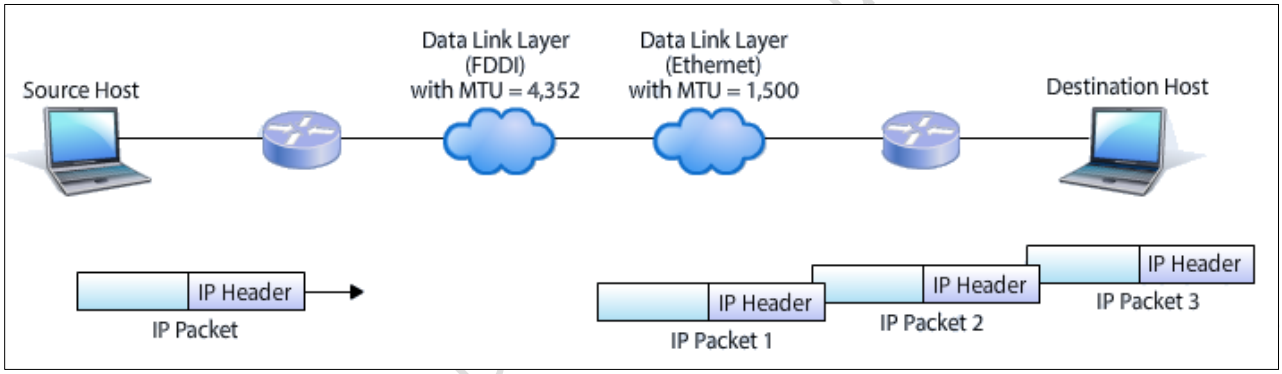

Topic 3: Advanced TCP/IP

IP Header Fields Used in IP Fragmentation

Explore the IP header fields that are used for IP fragmentation and then perform some simple activities to test your understanding.

Introduction

The IP header fields highlighted in this diagram are the ones that come into play when an IP packet is fragmented.

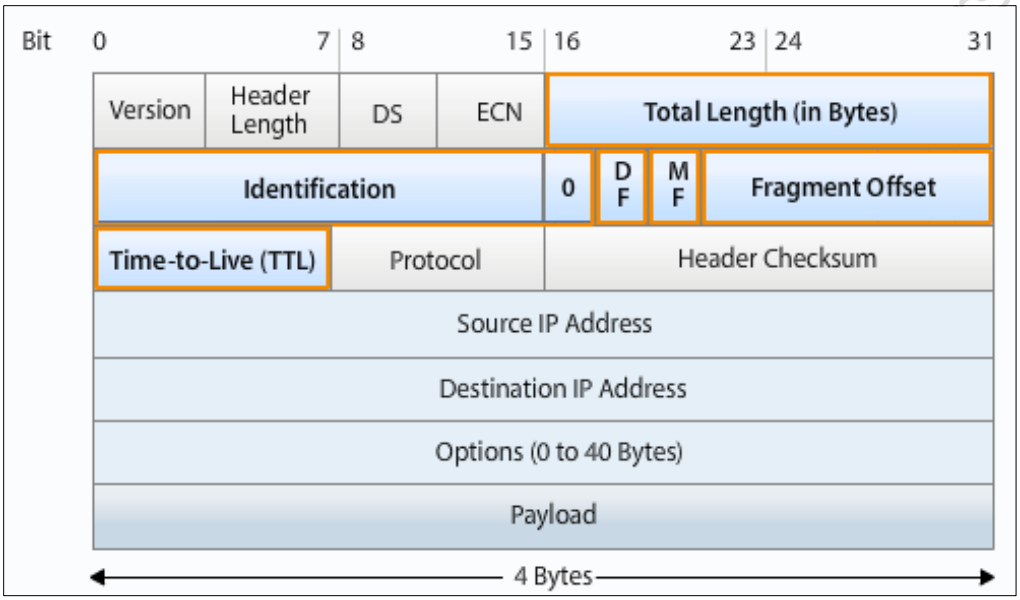

Reference: Information Sciences Institute. (1981). *Internet protocol: DARPA Internet program protocol specification*. Retrieved from http://www.rfc-editor.org/rfc/pdfrfc/rfc791.txt.pdf

Total Length

After fragmenting, the Total Length field indicates the length of each fragment. This field indicates the total size of a fragment, including the header and data.

Identification

The Identification field is a 16-bit field. When an IP packet is fragmented, the value in the Identification field is the same in all fragments.

DF (Do Not Fragment)

When the DF bit is set to 1, it signifies that the IP packet cannot be fragmented. When it is set to 0, the IP packet may be fragmented.

MF (More Fragments, or Many Fragments)

When a network device such as a router fragments an IP packet, it sets the MF bit value to 1 for all fragments except the last one. Based on this value, a destination host keeps expecting fragments until it encounters a fragment with MF = 0.

Fragment Offset
The Fragment Offset field solves the problem of sequencing fragments by indicating to the destination device where each particular fragment should be placed in the overall message. This field specifies the fragment's position within the original datagram, measured in 8-byte units. Thus, every fragment except the last one contains data in multiples of 8 bytes.

Time-to-Live (TTL)
The TTL field specifies how long the packet is allowed to live on the network in terms of router hops. Every time a packet passes through a router, the TTL value in the IP header is reduced by one. If the TTL value drops to zero, the packet is assumed to have exceeded its lifetime and is discarded.

Activity 1: Identifying Fragments

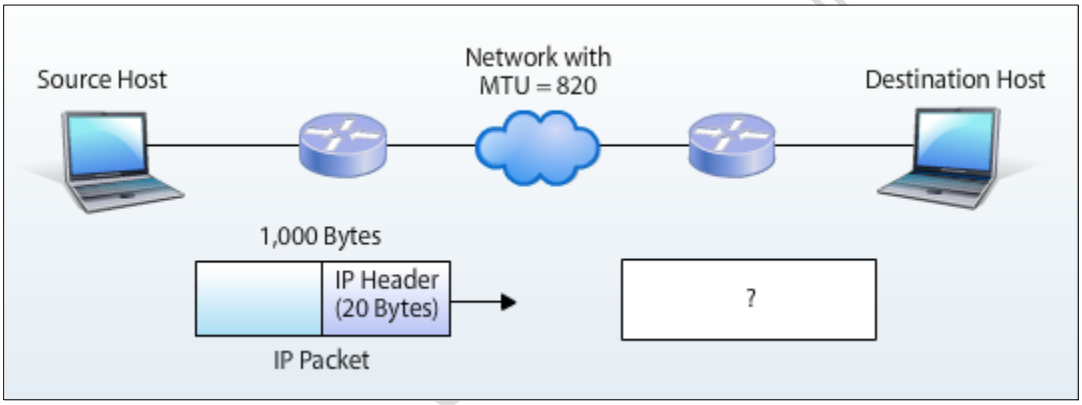

Question: As shown in the diagram, the source host tries to send a payload size of 1,000 bytes to the destination host across the network with an MTU of 820 bytes. Assume that the size of the IP header is 20 bytes. Into how many fragments will the original packet be divided?
a. Three
b. Two
c. Five
d. Four

Correct answer: Option b

Feedback:
The total size of the IP packet is 1,020 bytes, i.e., a payload of 1,000 bytes + 20 bytes, which is the size of the IP header. This packet must be fragmented when it travels across the network with an MTU of 820 bytes, since 1,020 is greater than 820. The number of IP fragments can be obtained by dividing the total size of the IP packet by the MTU, or 820 bytes in this case. Dividing 1,020 by 820 yields a result of 1.243, which signifies that there will be two fragments.

Activity 2: Identifying Field Values

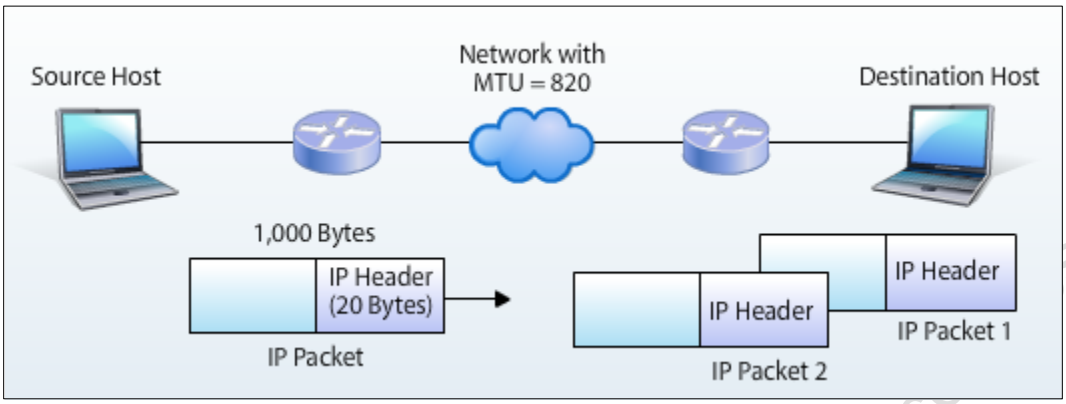

Question 1: Consider the diagram shown here. How much data would be transmitted in IP Packet 1, the first packet to reach the destination host?
a. 780 bytes
b. 860 bytes
c. 820 bytes
d. 800 bytes

Correct answer: Option d

Feedback:
The source host sends an IP packet with a payload size of 1,000 bytes. The total length of the IP packet is 1,020 bytes, i.e., 1,000 bytes + 20 bytes for the IP header. Given the MTU of 820 bytes, this 1,020-byte packet must be fragmented when it travels across the network.

In IP Packet 1, the IP header uses 20 bytes of the MTU, while the rest is used up by the data. Since the MTU is 820 bytes, the data transmitted in IP Packet 1 is 800 bytes (820 - 20). The remaining data to be transmitted in the second fragment, or IP packet 2, is 200 bytes (1,000-800).

Question 2: What value would the Fragment Offset field contain for IP Packet 2, the second packet to reach the destination host?
a. 0
b. 800
c. 100
d. 1

Correct answer: Option c

Feedback:
The Fragment Offset field specifies the fragment's position within the original datagram, measured in 8-byte units. The fragment offset for the first fragment is 0. Since the payload size of the first fragment is 800 bytes, the fragment offset for the second packet is 800 divided by 8, which is 100; therefore, the payload of the second fragment starts with the 801st byte.

Question 3: What value would the DF bit contain for both IP fragments?
a. 0
b. 1

Correct answer: Option a

Feedback:
The DF (Do Not Fragment) bits in both packets are set to 0. When the DF bit is set to 1, it signifies that the IP packet cannot be fragmented.

Question 4: What value would the MF bit contain for the IP Packet 1?
a. 1
b. 0

Correct answer: Option a

Feedback:
The MF (Many Fragments) bit value is set to 1 for all fragments except the last one to reach the destination host, which has its MF bit value set to 0. Since IP Packet 1 is the first packet to reach the destination host, its MF bit value would be 1.

Topic 3: Advanced TCP/IP

TCP Flow Control Mechanisms

TCP flow control serves a very important function—it regulates the flow of data between devices and prevents the receiver from being overwhelmed by the amount of data it receives. TCP flow control is achieved using the TCP Window Size and TCP Sliding Window mechanisms.

- **TCP Flow Control**
  The goal of TCP flow control is to ensure that the receiver of data is not overwhelmed by too much data from the sender. The data receiver controls the rate of transmission by sending a control message to the sender.

  It is essential to have a TCP flow control mechanism in a network environment in which network devices or hosts process received data at different rates. For instance, if a high-capacity Web server transmits data to a smartphone that is capable of receiving data at a much slower rate, the smartphone must regulate data transmission to prevent being overloaded.

- **TCP Window Size**
  A TCP header has a window size field. When the receiver of data sends acknowledgment packets to the sender, TCP uses the window size field to tell the sender how many bytes it may transmit. This window size specifies the number of bytes, starting with the acknowledgment number, that the receiving host's TCP layer is currently prepared to receive. The window size field is included in every TCP segment starting with the three-way handshake.

  If the receiver receives too much data from the sender, the receiver reduces its acceptable window size. If more data can be handled, the receiver increases its acceptable window size. Note that both sender and receiver specify their window sizes, since TCP is a full duplex service. In a full duplex service, the sender specifies the TCP sequence number and acknowledgment number in the TCP header and sends the TCP packet to the receiver, which does the same in turn when sending a TCP packet to the sender.

- **TCP Sliding Window**
  A mechanism called the TCP Sliding Window is used to speed data flow in a busy network. Without the TCP sliding window mechanism, a sender must wait for an acknowledgment from the receiver before it can transmit the next data segment, as seen in the TCP Window Size mechanism. This waiting period could result in a delay in transmitting data, especially in a busy network.

Topic 3: Advanced TCP/IP

Exploring TCP Window Size

TCP uses the window size field to specify the number of bytes, starting with the acknowledgment number, that the receiving host's TCP layer is currently prepared to receive. On receiving a TCP packet, the receiver calculates the acknowledgment number as a sum of the sequence number (in the packet received) and the number of bytes it receives.

Use the step list below to sequentially build a diagram that depicts the operation of the window size mechanism.

TCP Window Size Mechanism

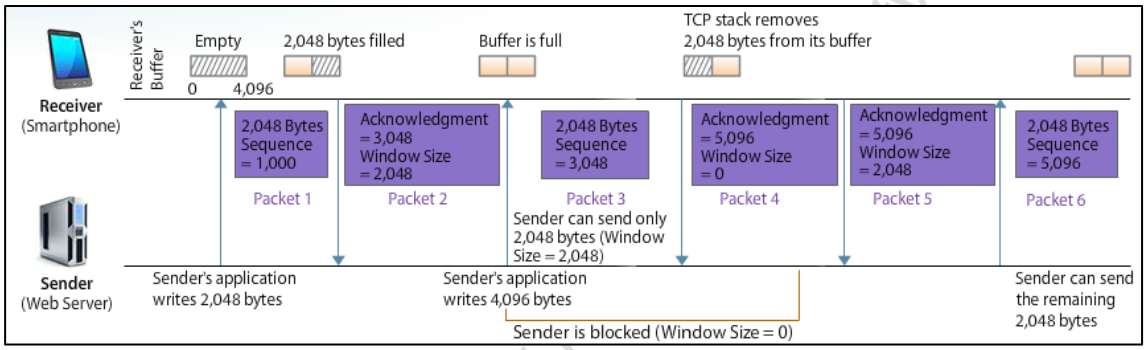

Step 1
Initially, the size of the receiver's buffer is 4,096 bytes and the buffer is empty. The sender sends 2,048 bytes of data to the receiver. The data size is 2,048 bytes and the sequence number is 1,000 (as seen in Packet 1).

Step 2
The receiver receives the data and puts it into its buffer. At this point, 2,048 bytes are available in the receiver's buffer. Then, the receiver sends an acknowledgment to the sender with a window size of 2,048 bytes, (4,096 – 2,048 = 2,048 bytes) and acknowledgment of 3,048 bytes (1,000 + 2,048 = 3,048), as seen in Packet 2.

Step 3
The sender wants to transmit 4,096 bytes of data. However, it can send only 2,048 bytes of data because of the previous window size of 2,048 bytes sent by the receiver. Therefore, the sender sends 2,048 bytes of data with sequence = 3,048, as seen in Packet 3. The receiver's buffer is now full upon receiving the 2,048 bytes of data.

Step 4
The receiver sends an acknowledgment with window size = 0 and acknowledgment = 5,096 (based on Packet 3, we can calculate the acknowledgment as 3,048 + 2,048 = 5,096). The sender is blocked from sending any more data and waits for another acknowledgment with a window size greater than 0. In the meantime, the receiver consumes half the data in its buffer. Then, since half of its buffer is available, the receiver sends acknowledgment of 5,096 and a window size of 2,048 in Packet 5.

Step 5
Upon receiving the acknowledgment, the sender starts to transmit 2,048 bytes of data with the sequence number of 5,096 (from Packet 5).

Topic 3: Advanced TCP/IP

TCP Sliding Window: An Example

Sliding window is a flow control mechanism that uses the receiver's window size. The sender computes its usable window, which is how much data it can immediately send without receiving any acknowledgment from the receiver.

When TCP has a large file to transmit, it breaks it into equal pieces. Each piece of data is called a TCP segment, and the size of each piece of data is called the maximum segment size, or MSS. In this example, we assume that the server, which is sending data, has an MSS of 1,000 bytes, and the client, or receiver, has a window size of 4,000 bytes.

Working of Sliding Window Mechanism

Step 1

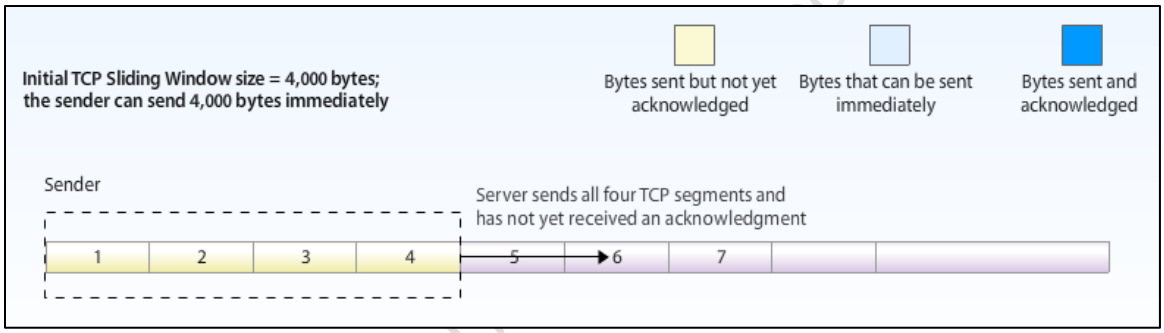

Step 2

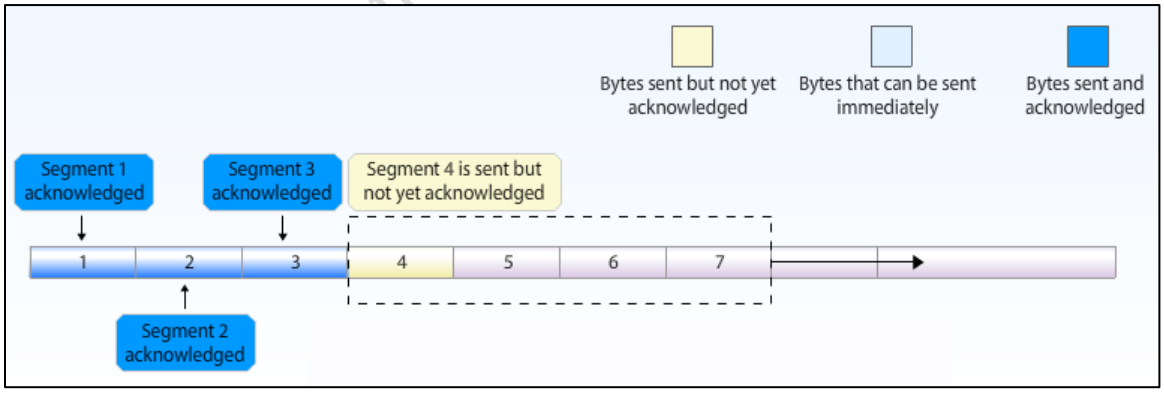

Topic 3: Advanced TCP/IP

Introduction to Traceroute

What Is Traceroute?

Traceroute is a tool that records the route between two devices on different networks. This tool is typically used to check the connectivity between any two devices. The Traceroute output indicates the possible path, or route, taken to reach the destination host.

On UNIX and Linux operating systems, the Traceroute tool uses User Datagram Protocol (UDP) packets with a destination port number starting at 33434—a random high port number. On the other hand, the Windows-based Traceroute tool uses ICMP echo requests, known as ping packets.

How Does Traceroute Work?

Consider an example to understand how Traceroute works. Assume that a network has Host A, Router A, Router B, and Host B, as shown here. At Host A, a user executes a Traceroute command: **Traceroute Host B**.

Step 1

Traceroute sends an ICMP packet with TTL = 1. When Router A receives the packet, it decreases the TTL value by 1 and sends an ICMP error message to the effect that the packet exceeded its time to live in transit.

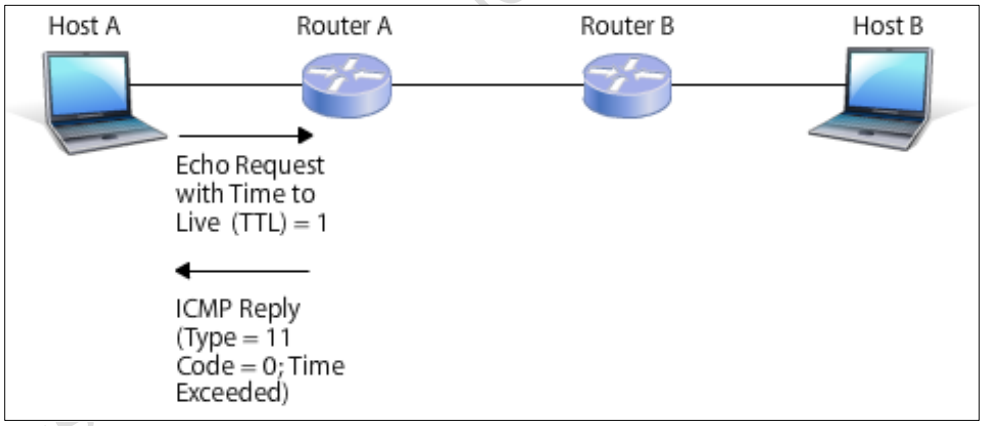

Step 2

Host A repeats this process until it receives an ICMP echo reply message. Note that each time a router receives an ICMP request, it decreases the TTL value by 1.

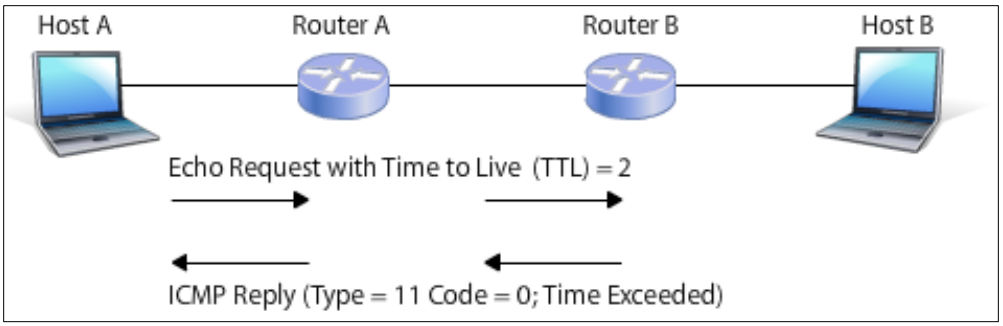

Step 3
Host A repeats this process until it receives an ICMP echo reply message. Note that each time a router receives an ICMP request, it decreases the TTL value by 1.

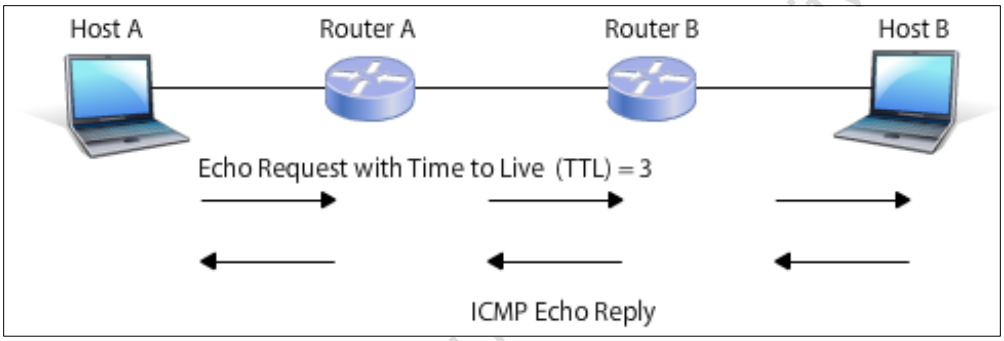

Activity: Test Your Understanding

Packet 1

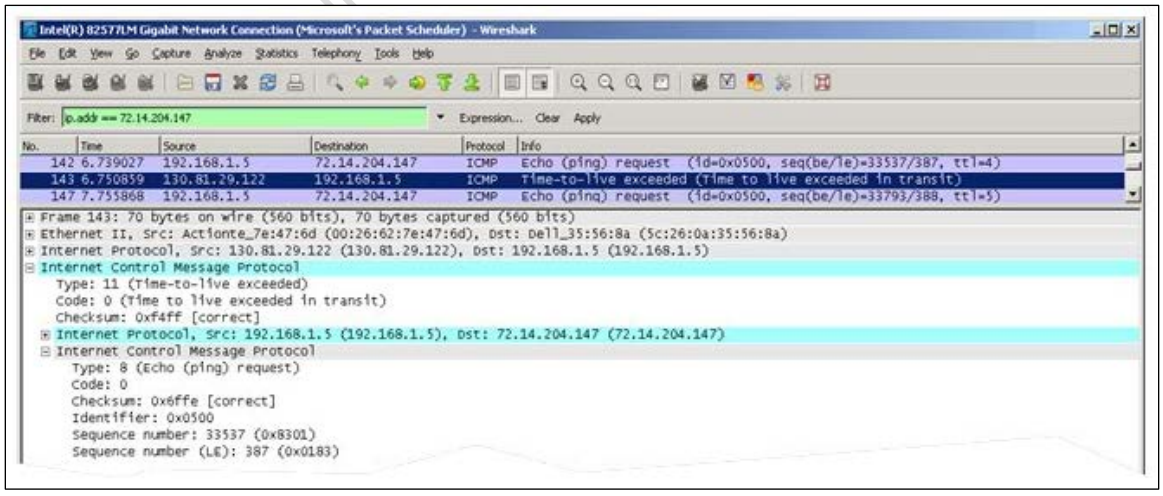

Reference: This screenshot has been included with permission from the Wireshark Foundation.

Packet 2

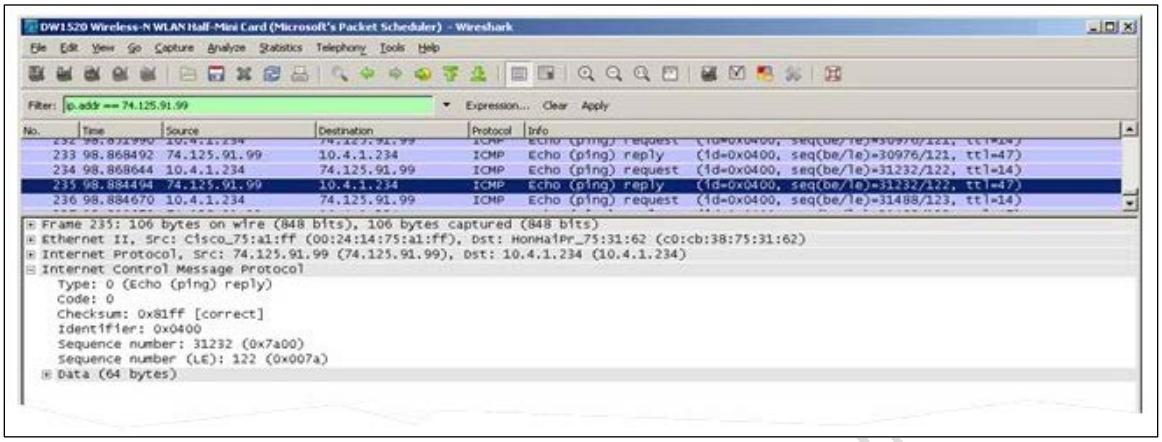

Reference: This screenshot has been included with permission from the Wireshark Foundation.

Your Traceroute program has just received the two packets shown here. Analyze the enlarged view of each image and answer the following question:

Question: Which of these packets originated from the final destination in the network path rather than from a network node in the path, such as a router, to the final destination?

a. Packet 1
b. Packet 2

Correct answer: Option b

Feedback:
Packet 1 is from one of the routers in the path. This conclusion is evident from its Type field showing a value of 11. This value reveals that the packet contains an error message—that the packet exceeded its time to live in transit. Therefore, Packet 1 must have originated from one of the network nodes (routers) in the path.

On the other hand, the Type field in the ICMP header for Packet 2 shows a value of 0. Therefore, this packet originated from the final destination in the network path.

Topic 4: Introduction to Penetration Testing

What Is Penetration Testing?

Penetration testing is the practice of arranging for a trusted third-party company to attempt to compromise the computer network or digital resources of an organization in order to assess the organization's security.

Penetration testing:

- Provides a list of all unpatched computers
- Covers all relevant attack vectors
- Documents all activities performed
- Tests the system as a whole, including existing defense mechanisms

Topic 4: Introduction to Penetration Testing

Steps in Penetration Testing

The main difference between penetration testing and real attacks is permission. An organization permits penetration testers to test its computing resources by drawing up a well-defined contract that predetermines the scope and level of testing. Penetration testers usually perform four steps to perform an attack and gain as much access as possible.

Step 1: Performing Reconnaissance
The first step involves obtaining as much information as possible on the target network. Testers usually do this by:

- Downloading the target's Web site for offline analysis, using tools such as Wget and Teleport Pro
- Using technical tools such as NSLookup and Dig to uncover information on the hosts that are active on the target network
- Identifying publicly accessible services such as e-mail and Web servers

Step 2: Scanning and Enumeration
Scanning involves attempting to connect to a target system to observe a response. Many tools, such as Nmap, are available for scanning. Enumeration is used to gather in-depth information about the target systems, such as open shares, operating systems running, and user accounts.

A few key terms used in enumeration are:

- **Fingerprinting:** Fingerprinting is the process of discovering the underlying operating systems on the target network.
- **Footprinting:** Footprinting goes a step beyond fingerprinting. Using footprinting, testers can obtain the following details regarding the target system:
  - Host names
  - IP addresses
  - Active running ports and services
  - Operating systems
  - A network diagram

Step 3: Gaining Access
Gaining access is the most important and lengthy step in penetration testing. At this stage, a penetration tester moves from scanning and enumerating to actually performing an attack. The attack must adhere to the scope of the contract that exists between the tester and the organization's management. Otherwise, the attack could be considered malicious and potentially illegal. This step may involve almost any approach to gain access, such as Web server and Domain Name System (DNS) attacks, denial of service attacks, or e-mail attacks including spam and Trojans.

Step 4: Reporting Problems to Management
Upon completion of penetration testing, the testing team reports its findings to the organization's management. Penetration testing helps the organization in the following ways:

- It makes system administrators and technical staff aware of how their network can be compromised and allows them to accurately estimate budget increases for security technology upgrades.
- It aids managerial decision making with respect to purchasing new security devices and software.

Topic 5: Summary

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

- The Internet Control Message Protocol (ICMP) is one of the most widely deployed protocols of the Internet Protocol Suite. The main distinguishing factors between TCP and ICMP include:
  - ICMP does not have a TCP header. Instead, it has an ICMP header.
  - The TCP header has a sequence number and a port number. This header enables reliable transmissions of data and services. However, the main purpose of ICMP is to perform simple queries and report errors.

- A few important ICMP error types include error type 3 (Destination Unreachable), error type 5 (Redirect), and error type 11 (Time Exceeded).

- The Internet Protocol (IP) implements datagram fragmentation to fragment a packet if its size exceeds the limit imposed by a data link layer protocol. This limit is called the maximum transmission unit (MTU).

- TCP flow control regulates the flow of data between devices and prevents the receiver from being overwhelmed by the amount of data it receives. TCP flow control is achieved using the TCP Window Size and TCP Sliding Window mechanisms.

- Traceroute is a tool that records the route between two devices on different networks. This tool is typically used to check the connectivity between any two devices.

- Penetration testing is the practice of arranging for a trusted third-party company to attempt to compromise the computer network or digital resources of an organization in order to assess the organization's security.

Glossary

Term	Definition
Acknowledgment Number	In a TCP connection, an acknowledgment number is used by a receiver to inform the sending host that transmitted data has been received successfully.
ICMP	The Internet Control Message Protocol (ICMP) is one of the most widely deployed protocols of the Internet Protocol Suite. It is mainly used by operating systems or network devices to send error messages or perform simple queries.
IP Fragmentation	The Internet Protocol (IP) implements datagram fragmentation to fragment a packet if its size exceeds the limit imposed by the data link layer protocol. This limit is called the maximum transmission unit (MTU).
Penetration Testing	Penetration testing is the practice of arranging for a trusted third-party company to attempt to compromise the computer network or digital resources of an organization in order to assess the organization's security.
Ping	Ping is a diagnostic tool that sends a simple request to determine if a host on a network is alive. This request is based on the ICMP communication protocol.
Sequence Number	In a TCP connection, a host machine maintains a 32-bit (TCP) sequence number to keep track of how much data it has sent.
TCP/IP	Transmission Control Protocol/Internet Protocol (TCP/IP) is a protocol used to control, facilitate, and govern the interconnection of computer systems, as well as the connection of these systems to the Internet.
TCP Sliding Window	A mechanism called the TCP Sliding Window is used to speed up data flow in a (busy) network. Without the TCP Sliding Window mechanism, a sender must wait for an acknowledgment from a receiver before it can transmit the next data segment.
TCP Window Size	A TCP header has a window size field. In acknowledgments, TCP uses the window size field to tell the sender how many bytes it may transmit. This window size specifies the number of bytes, starting with the acknowledgment number, that the receiving host's TCP layer is currently prepared to receive.
Traceroute	Traceroute is a tool that records the route between two devices on different networks. This tool is typically used to check the connectivity between any two devices.