PC & Network Security

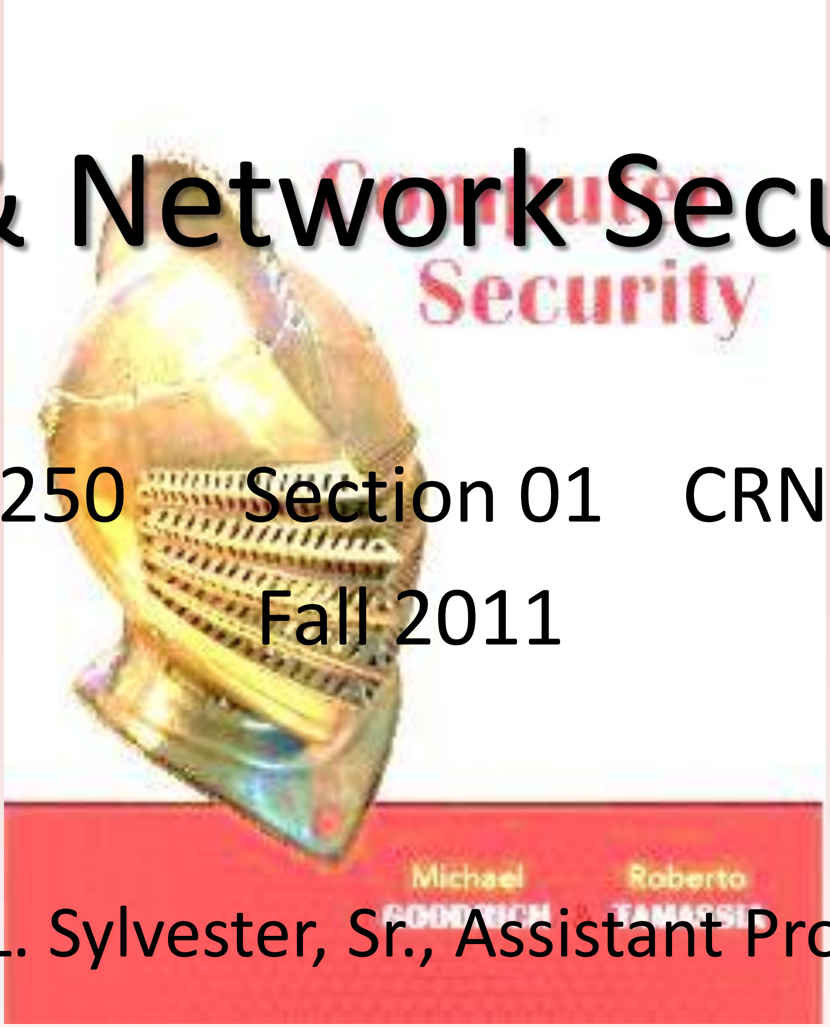

CNET – 250   Section 01   CRN: 19234

Fall 2011

David L. Sylvester, Sr., Assistant Professor

Chapter 5

Network Security I

Network Security Concepts

The internet was originally developed during Cold War as a communication network that could survive military attacks. For this reason, rather than basing communication on switched paths that connect communication parties, the internet was designed so that communication occurs through sequences of data packets. A data packet is a finite-length set of bits, which is divided into two parts: a header, which specifies where the packet is going and contains various overhead and bookkeeping details, and a payload, which is the actual information that is being communicated.

So if entities wish to communicate using the Internet, they must chop their messages into packets, attach a header on the front of each one, and then have those packets find their way through the internet to reach their respective destinations.

Network Topology

A network’s topology is its connection structure. The computers in a network are host nodes that can be sources and destinations of messages, and the routers in the network are communication nodes through which messages flow. The physical connections between nodes define the channels through which messages travel, so that packets move by being passed from one node to the next in order to get from their source node to their destination node.

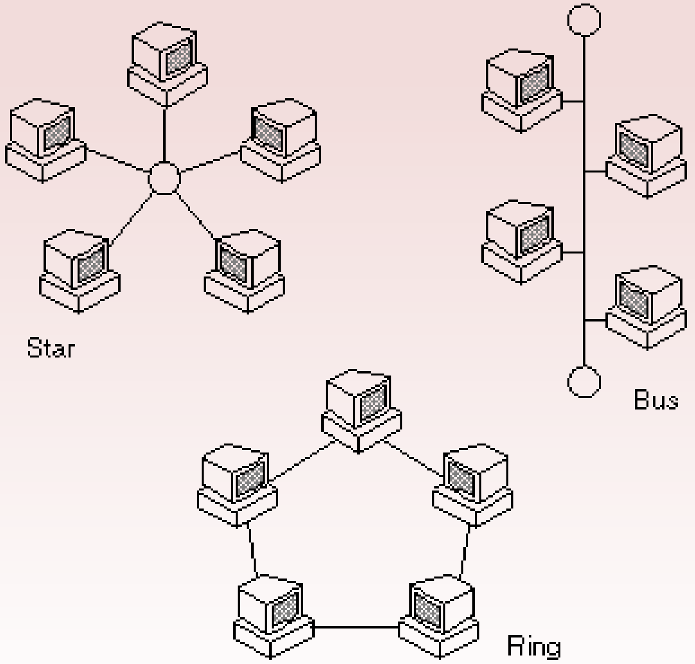

Star - each station is directly connected to a common central node.

Ring - the network consists of a set of repeaters joined by point-to-point links in a closed loop.

Bus - all stations attach, through appropriate hardware interfacing known as a tap, directly to a linear transmission medium, or bus.

Networks are characterized based on the distance between their nodes.

•Local-Area Network (LAN)	(room, floor, building)
•Wide-Area Network (WAN)	(buildings, cities, countries)
•Metropolitan-Area Network (MAN)	(building within cities)
•Small-Area Network (SAN)	(several feet –in home)

Internet Protocol Layers

The architect of the internet is modeled conceptually as being partitioned into layers, which collectively are called the internet protocol stack. Each layer provides a set of services and functionality guarantees for higher layers and, to the extent possible, each layer does not depend on details from higher levels. Likewise, the interface each layer provides to higher levels is designed to provide only the essential information from this layer that is needed by the higher levels – lower level details are hidden from the higher levels.

Five Standard Layers for Internet Communication

1. Physical Layer
2. Link Layer
3. Network Layer
4. Transport Layer
5. Application Layer

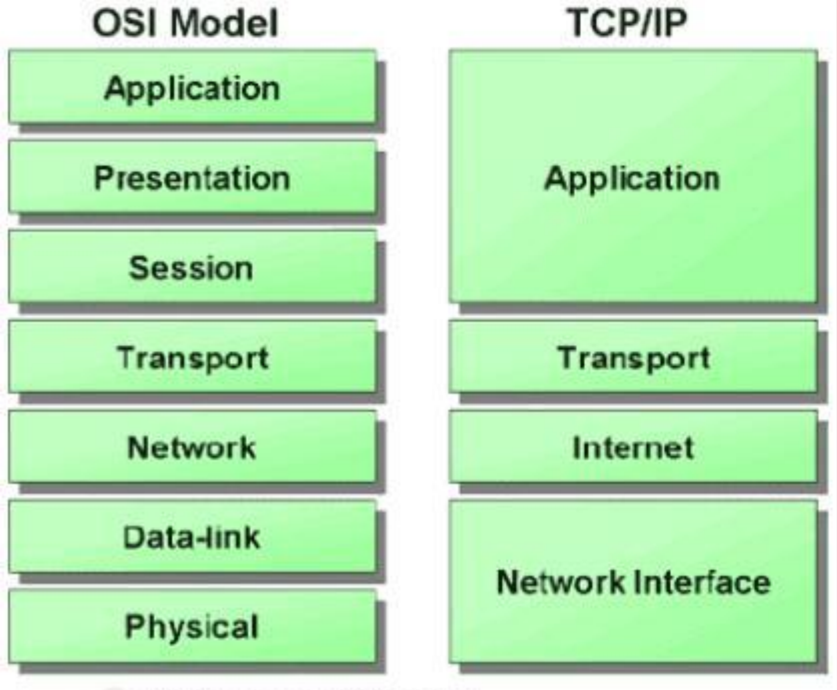

TCP/IP and the OSI model

Physical Layer. This layer has the task of moving the actual bits between the nodes of the network, on a best effort basis. For example, this level deals with details related to whether connections are done with copper wires, coaxial cables, optical-fiber cables or wireless radio. The abstraction it provides to the next higher level is an ability to transmit bits between a pair of network nodes.

Link layer. This layer has the task of transferring data between a pair of network nodes or between nodes in a local-area network and to detect errors that occur at the physical layer.

The link layer deals with the logical aspect of sending information across network links and how to find good routing paths in a local-area network. It include such protocols as Ethernet, which is used to route packets between computers sharing a common connection. Link layers use MAC (media access control) addresses consisting of link layers of 48-bit addresses.

The MAC address is a unique value associated with a network adapter. MAC addresses are also known as hardware addresses or physical addresses. They uniquely identify an adapter on a LAN.

MAC addresses are 12-digit hexadecimal numbers (48 bits in length). By convention, MAC addresses are usually written in one of the following two formats:

MM:MM:MM:SS:SS:SS

or

MM-MM-MM-SS-SS-SS

The first half of a MAC address contains the ID number of the adapter manufacturer. These IDs are regulated by an Internet standards body. The second half of a MAC address represents the serial number assigned to the adapter by the manufacturer. In the example, 00:A0:C9:14:C8:29. The prefix 00A0C9 indicates that the manufacturer is Intel Corporation.

Finding your MAC Address of Network Adapter

1. From **Start** menu, search field, type **cmd** and press **Enter**.
2. Type in **ipconfig /all**

The MAC address is the six pairs of hexadecimal numbers in the field labeled "Physical Address."

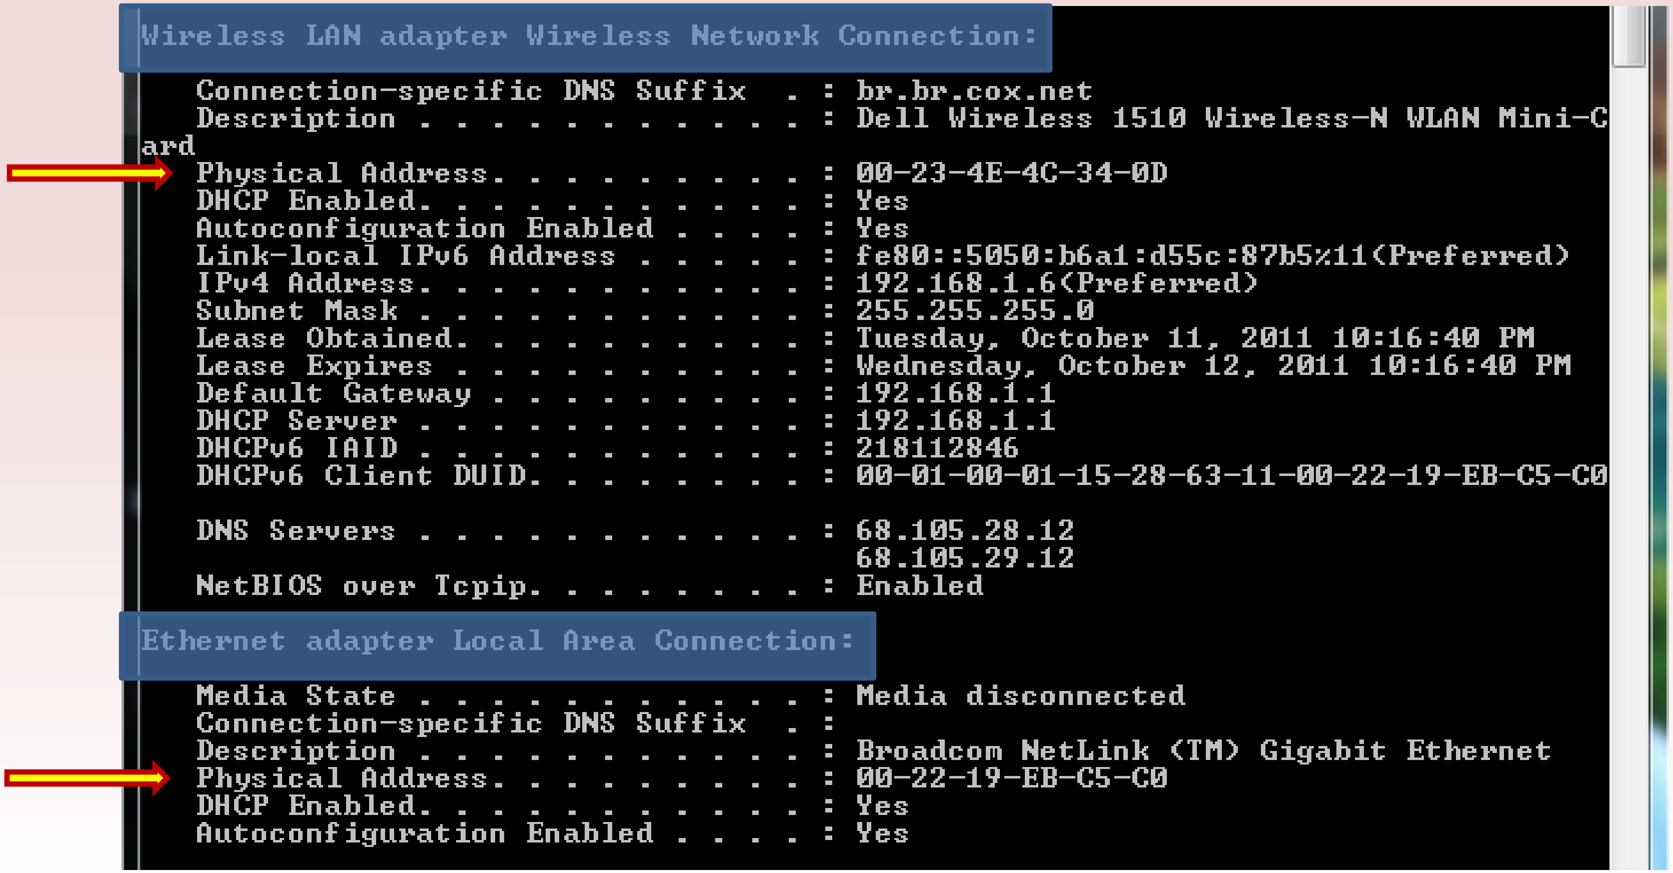

```
Wireless LAN adapter Wireless Network Connection:

   Connection-specific DNS Suffix  . : br.br.cox.net
   Description . . . . . . . . . . . : Dell Wireless 1510 Wireless-N WLAN Mini-C
ard
   Physical Address. . . . . . . . . : 00-23-4E-4C-34-0D
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   Link-local IPv6 Address . . . . . : fe80::5050:b6a1:d55c:87b5%11(Preferred)
   IPv4 Address. . . . . . . . . . . : 192.168.1.6(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : Tuesday, October 11, 2011 10:16:40 PM
   Lease Expires . . . . . . . . . . : Wednesday, October 12, 2011 10:16:40 PM
   Default Gateway . . . . . . . . . : 192.168.1.1
   DHCP Server . . . . . . . . . . . : 192.168.1.1
   DHCPv6 IAID . . . . . . . . . . . : 218112846
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-15-28-63-11-00-22-19-EB-C5-C0

   DNS Servers . . . . . . . . . . . : 68.105.28.12
                                       68.105.29.12
   NetBIOS over Tcpip. . . . . . . . : Enabled

Ethernet adapter Local Area Connection:

   Media State . . . . . . . . . . . : Media disconnected
   Connection-specific DNS Suffix  . :
   Description . . . . . . . . . . . : Broadcom NetLink (TM) Gigabit Ethernet
   Physical Address. . . . . . . . . : 00-22-19-EB-C5-C0
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
```

Network Layer. This layer also known as the internet layer for the internet, is to provide for the moving of packets between any two hosts, on the best effort basis. It provides a way of individually addressing each host using a numerical label called its IP address. The main protocol provided by this layer is the Internet protocol (IP). This protocol is subdivided into version (IPv4), which uses 32bit IP addresses, and a version 6 (IPv6), which uses 128bit IP addresses.

Best efforts means that there are no guarantees that any given packet will be delivered. Thus, reliability checks must be done by a higher layer.

Example IP Address (IPv4)

Decimal:

238 . 17 . 159 . 4

Binary:

11101110 00010001 10011111 00000100

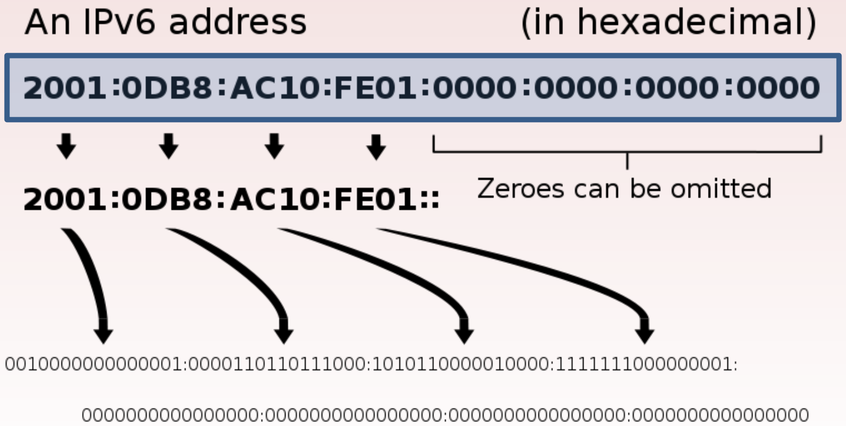

Transport Layer. This layer supports communication and connections between applications, based on IP addresses and ports, which are 16-bit address for application-level protocols to use. The transport layer provides a protocol called Transmission Control Protocol (TCP), which enables a virtual connection between a client and server and guarantees delivery of all packets in an ordered fashion and a protocol User Datagram Protocol (UDP), which assumes no prior setup and delivers packets as quickly as possible, but with no delivery guarantee.

Application Layer. This layer provides protocols that support useful functions on the internet, based on the services provided by the transport layer. Protocols include: Hypertext Transfer Protocol (HTTP), which use TCP and supports web browsing, DNS, which uses UDP and supports the use of useful names for hosts instead of IP addresses, Simple Mail Transfer Protocol (SMTP) and Internet Message Access Protocol (IMAP), which uses TCP and supports electronic mail, Secure Sockets Layer (SSL), which uses TCP and supports encrypted connections, and Voice Over Internet Protocol (VoIP), which uses UDP and supports internet telephone messaging.

The Open Systems Interconnection (OSI) model differs slightly from the TCP/IP model. Rather the having five layers, it has seven layers, as the application layer is divided into a strict application layer for host application-to-network processes; also a presentation layer for data representation; and a session layer for interhost communication.

The TCP/IP model will be use in order to focus on the internet security issues. With the TPC/IP model, packets are sent with metadata stored in locations called headers and sometimes in the final portion of the packet which would be known as the footer. The actual data being transmitted is called the payload. For all but the topmost layer, the payload stores a packet of the layer immediately above. This nesting of packets is called encapsulation.

Each packet from a higher layer becomes the data for the next lower-layer packet, with headers added to the beginning, and for frames, a footer added at the end.

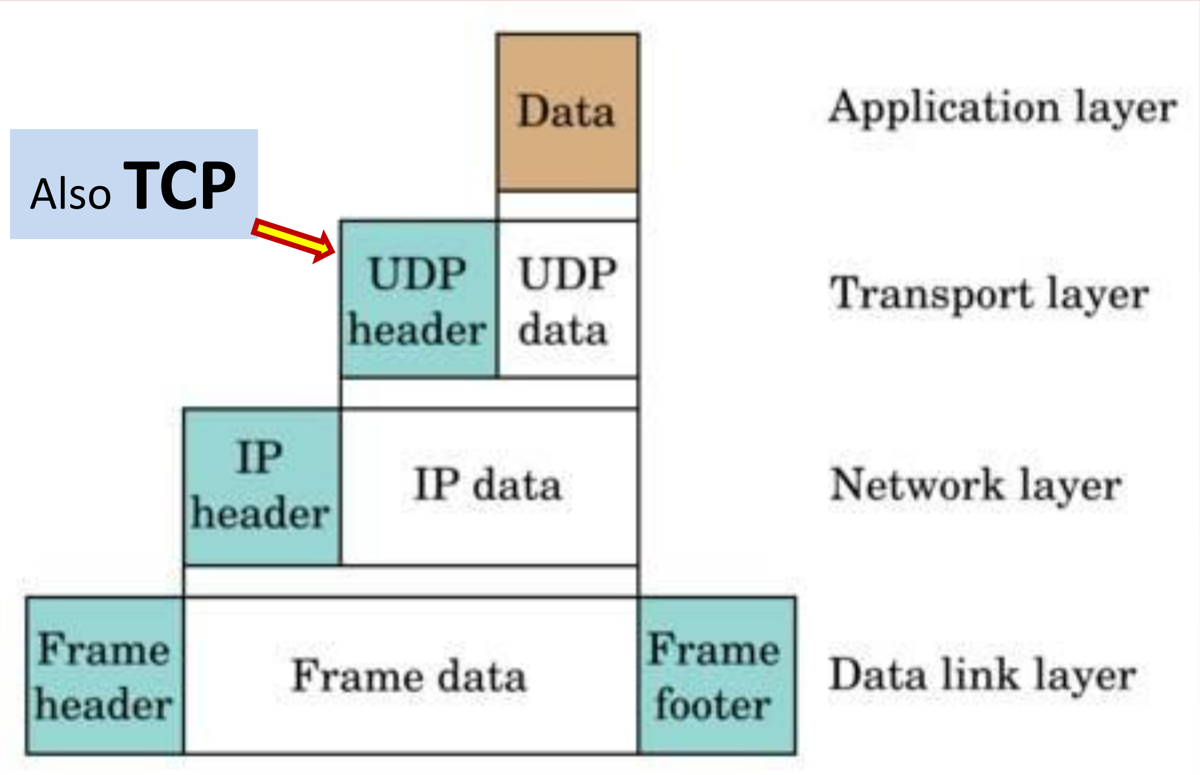

Also **TCP**

Using the Internet Protocol Suite. The Internet Protocol stack provides a useful set of functions and abstractions that make the internet possible. These functions were designed during the time when the internet was not used for any malicious intent, so may safeguard were not put in place. The layered model of the internet protocol suite, helps designers to build software that uses appropriate services and provides the right service guarantees.

Web server transmission to client's web browser.

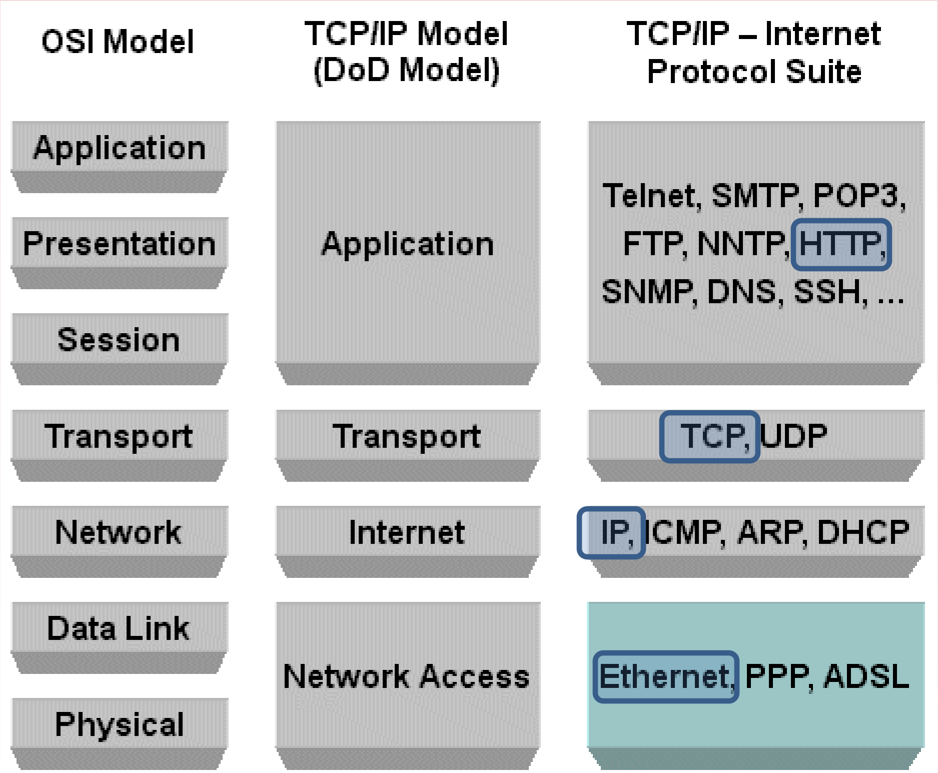

OSI Model	TCP/IP Model (DoD Model)	TCP/IP – Internet Protocol Suite
Application	Application	Telnet, SMTP, POP3, FTP, NNTP, HTTP, SNMP, DNS, SSH, …
Presentation		
Session		
Transport	Transport	TCP, UDP
Network	Internet	IP, ICMP, ARP, DHCP
Data Link	Network Access	Ethernet, PPP, ADSL
Physical		

Network Security Issues

Due to the widespread use of computers and networks, the internet has become a huge benefit to society. Computer networks are vulnerable to numbers of attacks where hackers are searching for personal information.

How Networking Impacts Computer Security Goals.

In reference to the CIA concept (*Confidentiality*)

- No confidentiality required at any layer.
- Does not encrypt data (header or the data)

If encryption is needed, the (Hypertext Transfer Protocol with Secure Sockets Layer (HTTPS) protocol could be use at the applications level.

Remember, SSL is the protocol primarily developed with secure, safe Internet transactions in mind.

In reference to the CIA concept (*Integrity*)

- Simple checksums or hash sums are used to validate data and header content.
- These checksums are good for validating small amounts of data in the headers or data.
- Again, if true data integrity is required, is should be done at application layer and with alternate protocols at lower levels.

In reference to the CIA concept (*Availability*)

- A large volume of data requests can cause a web server to become unavailable.
  - These requests can be from legitimate users having a sudden interest in that web site
  - These requests can also be attacks coming from many compromised hosts intending to shut down the web site.

In reference to the AAA concept (*Assurance*)

- By default, packets are allowed to travel between any host (source or destination) node in a network.
  - If permissions and policies are added to the network to control data flow, it must be done explicitly, which can be done using firewalls.

In reference to the AAA concept (*Authenticity*)

- Standard internet protocols do not have a place to put digital signatures in headers and footers.
- There is no notation of user identities; data is exchanged between computers not people.
  - To allow for signatures, it must be done explicitly at the application layer with an alternate protocol.

In reference to the AAA concept (*Anonymity*)

- The problem of anonymity is solved based on the fact that by default there is no identity of users on the internet.
  - This is can be a good thing for human rights workers, but bad if it lets an identity thief steal credit card information without getting caught.

The Link Layer

Most modern operating systems include a TCP/IP implementation and allow programs to interact with the internet Protocol stack via a simple interface.

Ethernet
One of the most popular ways to transmit internet traffic is Ethernet, which refers to both the physical medium used (normally cable) as well as the linked layer protocol standardized as IEEE 802.3.

Frames are transmitted through the Ethernet cable using electrical pulses. The data can be received by other machines that are logically connected to that cable on the same LAN. The portion of the LAN that has the same logical connection is called a network segment. A collision occurs when two machines on the same network segment each transmit a frame at the same time. When a collision occurs, these frames are discarded and retransmitted.

Dealing with Collisions. In the event of a collision, each of the transmitting machines wait a random amount of time, and then retransmits, in the hope of avoiding a second collision. This process is repeated each time a collision occurs. The Ethernet protocol is designed to ensure that eventually every machine on the network segment successfully transmits its frame. Packet collision on a LAN can cause a slowdown if there are a large number of machines that are locally connected to each other.

Hubs and Switches.

The simplest way to connect machines in a local-area network is to use an Ethernet hub, which is a device that logically connects multiple devices together, allowing them to act as a single networks segment. Hubs typically forward all frames to all attached devices, doing nothing to separate each attached device. Hubs may generate large amounts of unnecessary traffic, since each frame is duplicated and broadcast to all the machines connected to that network segment.

A better way of connect machines in a small local area network is to use switches – namely, an Ethernet switch.

When devices are first connected to an Ethernet switch, it acts like a hub, sending out frames to all connected machines. Over time, however, a switch learns the addresses of the machines that are connected to its various ports. Having attained the address information, a switch will then only forward each frame it receives, along the cable it knows is connect to the destination for that frame.

The selectivity that comes from a switch learning the address of the machine it connects reduces the possibility of collisions and increases the effective speed of the network (bandwidth). Switches also reduce the risks of network eavesdropping, since network frames forwarded by a switch are less likely to be seen by machines that are not destinations.

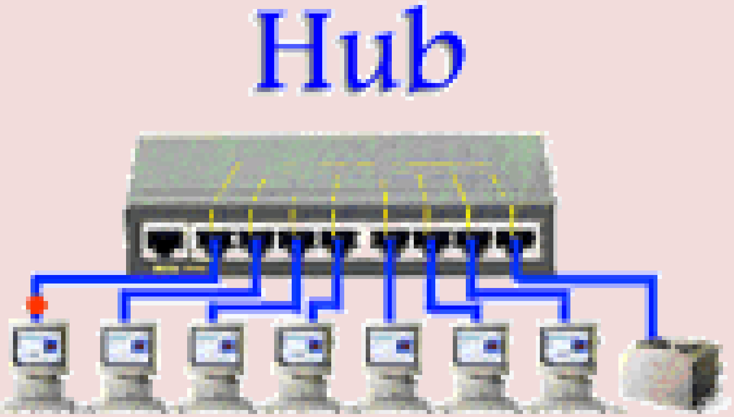

A hub copies and transmits traffic to all attached devices.

Looking at this 8 port hub: If a packet is sent to a port of the hub, this packet is also broadcast to the other 7 ports. Only the computer which this packet was meant for will retain the packet. The rest will just discard the packet. This is clearly not as efficient as a switch whereby the packet goes directly to the port it was specifically meant for.

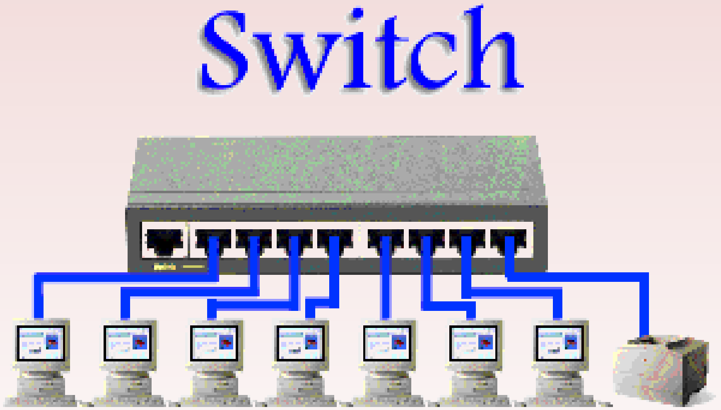

A switch only transmits frames to the appropriate destination device.

A switch is actually a multi-port hardware device that receives incoming data at one port and then forwards the data to another port. Basically a Layer 2 switch is able to send a packet directly to the host computer. It works by broadcasting an ARP request to all the machines connected to the switch. In return, the machines will respond to the switch by sending their MAC addresses and subsequently, the switch will be aware of the machines that are hooked up to its individual ports.

Media Access Control (MAC) Addresses

MAC addresses are used at the link layer to facilitate the routing of frames to the correct destination. Remember, switches learn the location of network devices from their MAC addresses and they forward frames to the appropriate segments. Each **Ethernet frame** consists of the source and destination MAC addresses, a checksum confirming data integrity, and a payload section, which contains data from higher layers, such as the IP layer.

Bits	Field	
0 - 55	Preamble (7bytes)	Header
56 – 63	Start-of-Frame Delimiter (1byte)	
64 – 111	MAC Destination (6 bytes)	
112 – 159	MAC Source (6 bytes)	
160 – 175	Ethertype/Length (2 bytes)	
176 – 543+	Payload (46 – 1500 bytes)	Payload
543+ - 575+	CRC-32 Checksum (4 bytes)	Footer
575+ - 671+	Interframe Gap (12 bytes) (minimum idle period between transmission)	

EtherType	Protocol
0x0800	Internet Protocol, Version 4 (IPv4)
0x0806	Address Resolution Protocol (ARP)
0x0842	Wake-on-LAN Magic Packet, as used by ether-wake and Sleep Proxy Service
0x1337	SYN-3 heartbeat protocol (SYNdog)
0x22F3	IETF TRILL Protocol
0x6003	DECnet Phase IV
0x8035	Reverse Address Resolution Protocol (RARP)
0x809B	AppleTalk (Ethertalk)
0x80F3	AppleTalk Address Resolution Protocol (AARP)
0x8100	VLAN-tagged frame (IEEE 802.1Q)
0x8137	Novell IPX (alt)
0x8138	Novell
0x8204	QNX Qnet
0x86DD	Internet Protocol, Version 6 (IPv6)
0x8808	MAC Control
0x8809	Slow Protocols (IEEE 802.3)
0x8819	CobraNet
0x8847	MPLS unicast
0x8848	MPLS multicast
0x8863	PPPoE Discovery Stage
0x8864	PPPoE Session Stage
0x886F	Microsoft NLB heartbeat [3]
0x8870	Jumbo Frames
0x887B	HomePlug 1.0 MME
0x888E	EAP over LAN (IEEE 802.1X)
0x8892	PROFINET Protocol
0x889A	HyperSCSI (SCSI over Ethernet)
0x88A2	ATA over Ethernet
0x88A4	EtherCAT Protocol
0x88A8	Provider Bridging (IEEE 802.1ad)
0x88AB	Ethernet Powerlink
0x88CC	LLDP
0x88CD	SERCOS III
0x88D8	Circuit Emulation Services over Ethernet (MEF-8)
0x88E1	HomePlug AV MME
0x88E5	MAC security (IEEE 802.1AE)
0x88F7	Precision Time Protocol (IEEE 1588)
0x8902	IEEE 802.1ag Connectivity Fault Management (CFM) Protocol / ITU-T Recommendation Y.1731 (OAM)
0x8906	Fibre Channel over Ethernet
0x8914	FCoE Initialization Protocol
0x9000	Configuration Test Protocol (Loop)[4]
0x9100	Q-in-Q
0xCAFE	Veritas Low Latency Transport (LLT)[

ARP Spoofing

The Address Resolution Protocol (ARP) is a link-layer protocol that provides services to the network layer. ARP is used to find a host's hardware address given its network layer address. It is used to determine the MAC address associated with a given IP address. Unfortunately, there is a man-in-the-middle attack against this protocol, called ARP spoofing.

How ARP Works. When a source machine wants to send a packet to a destination machine on the local-area-network, at the network layer, the source machine knows the destination IP address. Since the sending of the packet is delegated to the link layer, the source machine needs to identify the MAC address of the destination machine. In the ARP protocol, the resolution of IP addresses into MAC addresses is accomplished by means of a broadcast message that queries all the network interfaces on the LAN, so that the appropriate destination can respond.

ARP Spoofing

An attacker, simply sends an ARP reply to a target (Alice), who is associates the IP address of the LAN gateway (Bob), with EVE's MAC address. Eve also sends an ARP reply to Bob associating Alice's IP address with Eve's MAC address. After this ARP cache poisoning has taken place, Bob thinks Alice's IP address is associate with Eve's MAC address and Alice thinks Bob's IP address is associated with Eve's MAC address. Thus, all traffic between Alice and Bob, the gateway to the internet, is routed through Eve.

Eve has control over traffic between gateway.

- Can passively observe traffic
- Sniff passwords/information
- Alter data between Bob and Alice
- Deny severe to Alice

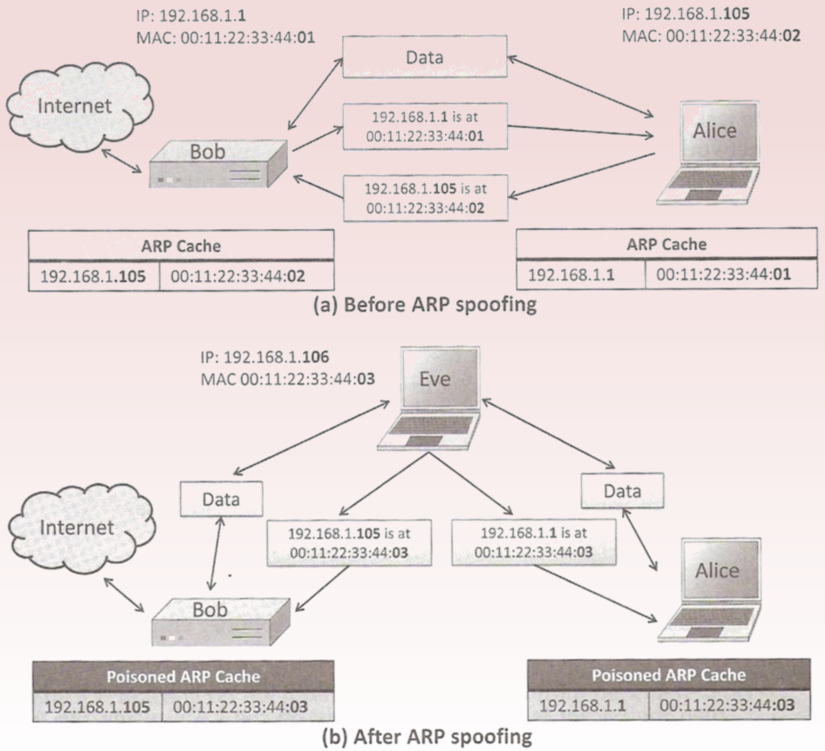

(a) Before ARP spoofing

(b) After ARP spoofing

The Network Layer

The task of the network layer is to move packets between any two hosts in a network, on a best effort basis. It relies on the services provided by the link layer to do this.

IP

The internet protocol (IP) is the network-level protocol that performs a best effort to route a data packet a source node to a destination node in the internet. In IP, every node is given a unique numerical address which is a 32-bit number under version 4 (IPv4) and a 128 bit number under version 6 (IPv6). Both the source and destination of any transmission are specified by an IP address.

- If the packet is addressed to a machine on the same LAN as the host, then the packet is transmitted directly on the LAN, using the ARP protocol to determine the MAC address of the destination machine.

- If the packet is addressed to a machine that is not on the LAN, then the packet is transmitted to a specially designated machine on the LAN called a gateway, which will handle then next step of the routing. The ARP protocol is used to determine the MAC address of the gateway.

The Structure of the Internet. Routers are designed to be very fast. For each packet received, the router performs one of three possible actions.

1. **Drop** – if the packet is expired.
2. **Deliver** – if the destination is a machine on one of the LANs to which the router is connected.
3. **Forward** – If the destination of the packet does not belong to the LANs of the route, then the packet is forwarded to a neighboring router.

There are two primary protocols that determine how the next hops are encoded in internet routing tables:

- » Open Shortest Path First (OSPF)
- » Border Gateway Protocol (BGP)

OSPF determines how packets are routed within a autonomous system and is based on a policy that packets should travel along shortest paths.

BGP, on the other hand, determines how packets are routed between autonomous systems and is based on policies dictated by contractual agreements between different ASs. (*Routes may not be the shortest path.*)

An **Autonomous System** (**AS**) is a collection of connected Internet Protocol (IP) routing prefixes under the control of one or more network operators that presents a common, clearly defined routing policy to the Internet.

IPv4 Packet

Bit Offset	0-3	4-7	8-15	16-18	19-31
0	Version	Header Length	Service Type	Total Length	
32	Identification			Flags	Fragment Offset
64	Time to Live		Protocol	Header Checksum	
96	Source Address				
128	Destination Address				
160	(Options)				
160+	Data Data Data Data Data Data Data Data Data				

Header

Payload

The internet is divided into autonomous systems, so routing tables have to be able to direct traffic to clusters of nodes, not just and individual destination.

To accomplish this, the IP addressing scheme takes into account the fact that networks are partitioned into logical groupings known as subnetworks or subnets. IP address can be written into two portions, a network portion that denotes an IP prefix used by all machines on a particular network, and a host portion which identifies a particular network device.

Subnet mask are used to define the address range of a particular network. Ranges of IP address are based on the size of the organization.

Class	Leading bits	Size of *network number* bit field	Size of *rest* bit field	Number of networks	Addresses per network	Start address	End address
Class A	0	8	24	128 (2^7)	16,777,216 (2^{24})	0.0.0.0	127.255.255.255
Class B	10	16	16	16,384 (2^{14})	65,536 (2^{16})	128.0.0.0	191.255.255.255
Class C	110	24	8	2,097,152 (2^{21})	256 (2^8)	192.0.0.0	223.255.255.255
Class D (multicast)	1110	not defined	not defined	not defined	not defined	224.0.0.0	239.255.255.255
Class E (reserved)	1111	not defined	not defined	not defined	not defined	240.0.0.0	255.255.255.255

Internet Control Message Protocol

The Internet Control Message Protocol (ICMP) is a network-layer protocol that is used by hosts to perform a number of basic testing and error notification tasks. ICMP is primarily used for network diagnostics tasks, such as determining if a host is alive and finding the path followed by a packet. ICMP packets carry various types of messages, including the following:

- Echo request, which asks the destination machine to acknowledge the receipt of the packet.
- Echo response, which acknowledges the receipt of a packet in reply on an echo request.
- Time exceeded, is an error notification that a signifying the a packet has expired, that is , its Time-to-Live (TTL) is zero.
- Destination unreachable, is an error notification that the packet could not be delivered.

Ping is a utility that uses the ICMP protocol to verify whether or not a particular host is receiving packets. Ping sends and ICMP echo request message to the destination host, which in turn replies with an ICMP echo response message.

IP Spoofing

Given that the source address is never checked on nearly every operating system that provides an interface by which it can make network connections with arbitrary IP header information, spoofing an IP address is a simple matter of specifying the desired IP in the source field of an IP packet data structure before transmitting that data to the network. The modification of the IP address in this case (modifying the sender's IP in the source of the packet header), is called IP spoofing.

If an attacker sends an IP packet with a spoofed source address, then he will not receive any response from the destination server. In fact, with a spoofed source IP address on an outbound packet, the machine with the spoofed IP address will receive any response from the destination server, not he attacker. Therefore, if the attacker is using IP spoofing on his outbound packets, he must either not care about any responses from these packets or he has some other way of receiving responses.

How IP Spoofing Works

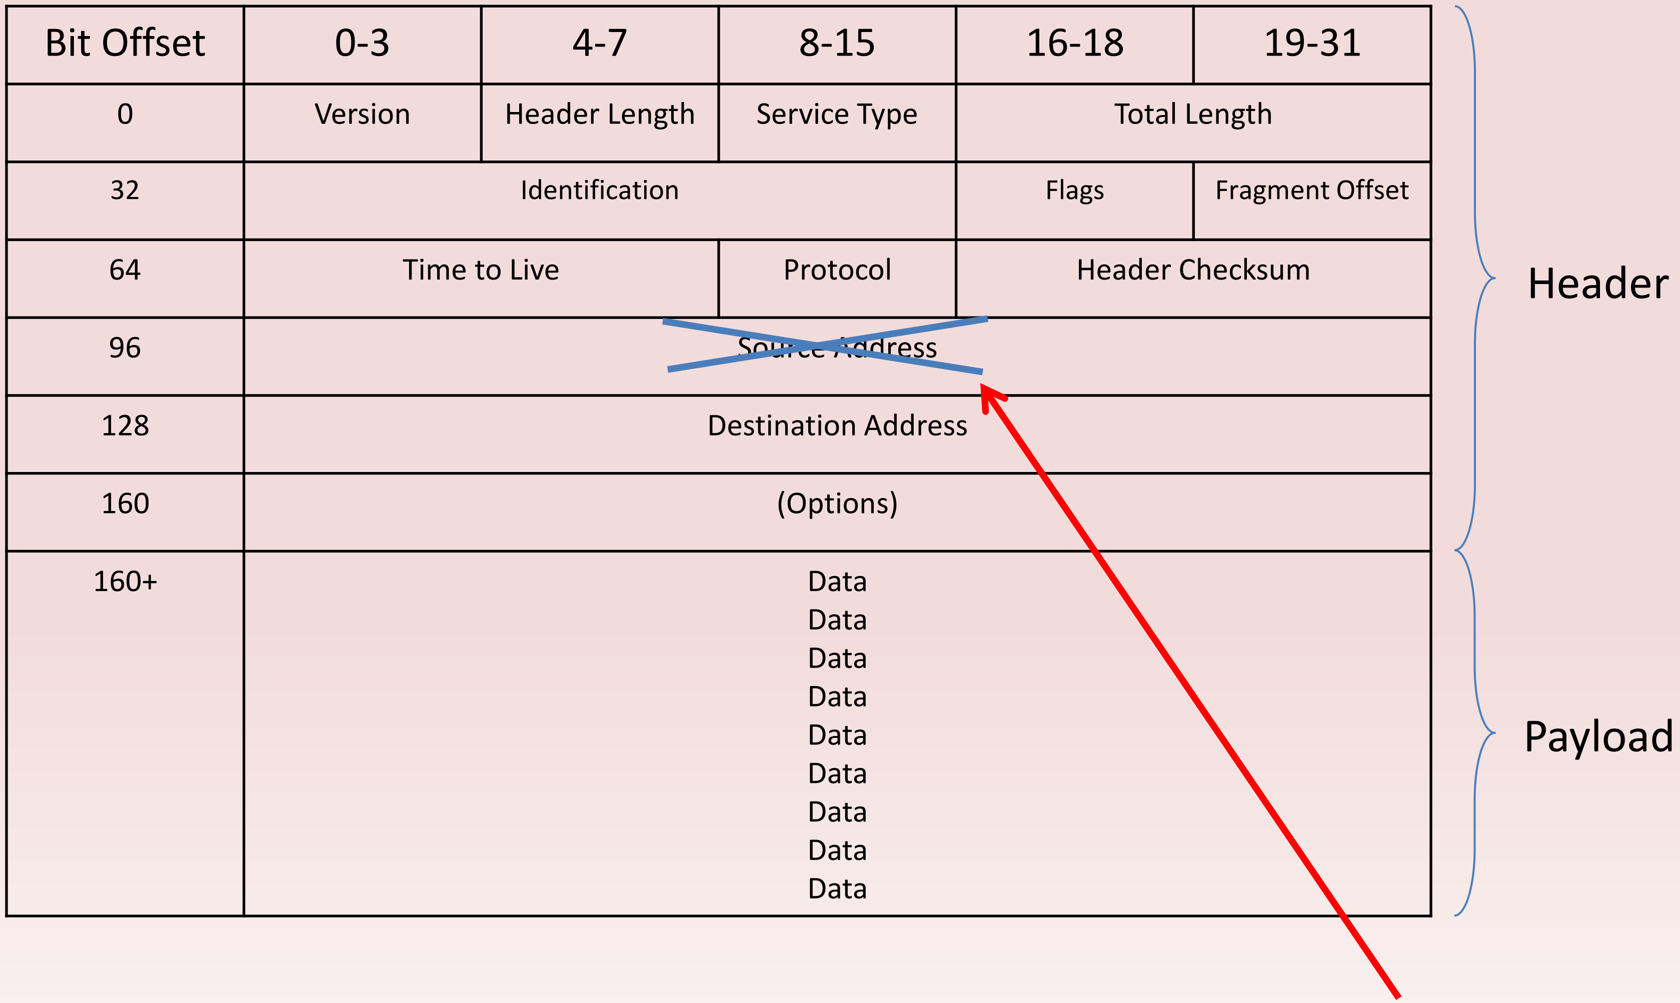

Bit Offset	0-3	4-7	8-15	16-18	19-31
0	Version	Header Length	Service Type	Total Length	
32	Identification			Flags	Fragment Offset
64	Time to Live		Protocol	Header Checksum	
96	~~Source Address~~				
128	Destination Address				
160	(Options)				
160+	Data Data Data Data Data Data Data Data Data				

Header

Payload

Overwrite source address with a different IP address

Dealing with IP Spoofing. IP spoofing cannot be prevented, but there are a number of ways of dealing with IP spoofing.

1. Using border routers, which can span across two or more subnetworks, that can be configured to block packets from outside their administrator domain that have source addresses from inside that domain.
2. By implementing traceback techniques (tracing the path of a packet back to its actual source address.) *If source is not matched, then requests can be made to various autonomous systems along this path to block packets from that location.*

Packet Sniffing

Packet sniffing allows for eavesdropping, due to the fact that payload data of IP packets are not encrypted. Packet sniffing is the process of listening in on the traffic in a network that is intended for the network. (Wireless or Wired) If a network interface is operating in promiscuous mode, it allows attackers to examine data transmitted over a particular network segment.

Defense Against Packet Sniffing.

To reduce the impact of packet sniffing, encryption mechanisms should be utilized in higher-level protocols to prevent attackers from recovering sensitive data.

Ex: Web Traffic

Starts at the application level with HTTP,
Encapsulated a TCP packet at the transport level,
An IP packet in the network layer,
And Ethernet or WiFi at the link layer frame.

But if the HTTPS protocol is used, encryption would be employed at the application layer.

The Transport Layer

The transport layer builds on top of the network layer, which supports communication between machines (host and destination nodes) to provide communication between processes. This is done by viewing the IP address of each machine. (*Each machine having a unique IP address.*)

Two primary protocols operate at the transport layer for the internet:

1. Transmission Control Protocol (TCP)
   - Reliable delivery of data (intact and in order)
2. User Datagram Protocol (UDP)
   - Unreliable delivery of data (Best effort)

Transmission Control Protocol
TCP is a critical protocol for the internet, since it takes the IP Protocol which routes packets between machines in a best effort fashion, and creates a protocol that can guarantee transmission of a stream of bits between two virtual ports.

A TCP session starts out by establishing a communication connection between the sender and receiver. Once connection has been created, the parties can then communicate over the established channel.

TCP ensures reliable transmission by using a sequence number that is initialized during its transmission.

TCP also manages the amount of data that can be sent by one party while avoiding overwhelming the processing resources of the receiver or the bandwidth of the network. This is known as flow control, which is managed by a technique known as a sliding window protocol.
Sliding Window Protocol

- Packets are sent but received by the destination node by the frame in sequence (frame increments)
- When frame seize, and the destination is still waiting on a frame, the destination node adjusts its stored acknowledge number, shifting the sliding window of sequence numbers
- The sender sets a timer on acknowledgment response. If no response, data is assumed to be lost and is retransmitted.

Congestion Control. TCP tackles the network problem of congestion control. Congestion control is not implemented into TCP packets, but rather is based on information gathered by keeping track of acknowledgments for previously sent data and the time required for certain operations.

In TCP, connection sessions are maintained beyond the life of a single packet. The states of the connection sessions are:

1. Open
2. Exchange data and Acknowledgement
3. Close

TCP Connections. TCP uses a three-way handshake to establish a reliable connection stream between two parties.

1. SYN flag where packets to the desired destination. *SYN is short for synchronization.*
2. SYN-ACK flag is sent by the server, indicating that server is wish to accept the connection. *The SYN flag is also sent.*
3. ACK packet to indicate a successful connection.

User Datagram Protocol

The UDP protocol does not make a guarantee about the order or correctness of its packet delivery. It has no initial handshake to establish a connection, but rather allows parties to send messages, known as datagrams, immediately. If a sender wants to communicate via UDP, it need only use a socket (defined with respect to a port on a receiver.

While UDP features a 16 bit checksum to verify the integrity of each individual packet, there is no sequence number scheme, so transmissions can arrive out of order or may not arrive at all. It is assumed that checking for missing packets in a sequence of datagrams is left to applications processing these packets. As a result, UDP can be much faster than TCP, which often requires retransmission and delaying of packets.

UDP is often used in time-sensitive applications and TCP is used for applications where data order and data integrity is important.

Network Address Translator (NAT)

When people add computers, printer, and other network devices to their home networks, they typically do not buy new IP addresses and setup the new addresses directly on the internet. Instead, they use network address translation (NAT), which allows all the machines on a local-area-network to share a single public IP address. This public IP address represents the point of contact with the internet for the entire LAN, while machines on the network have private addresses that are only accessible from within the network.

Since NAT allows an entire network to be assigned a single public IP address, widespread use of NAT has significantly delayed the inevitable exhaustion of the IPv4 address space. The private IP address are in the form:

192.168.x.x

172.16.x.x thru 172.31.x.x

10.x.x.x

Home Network Using NAT Router

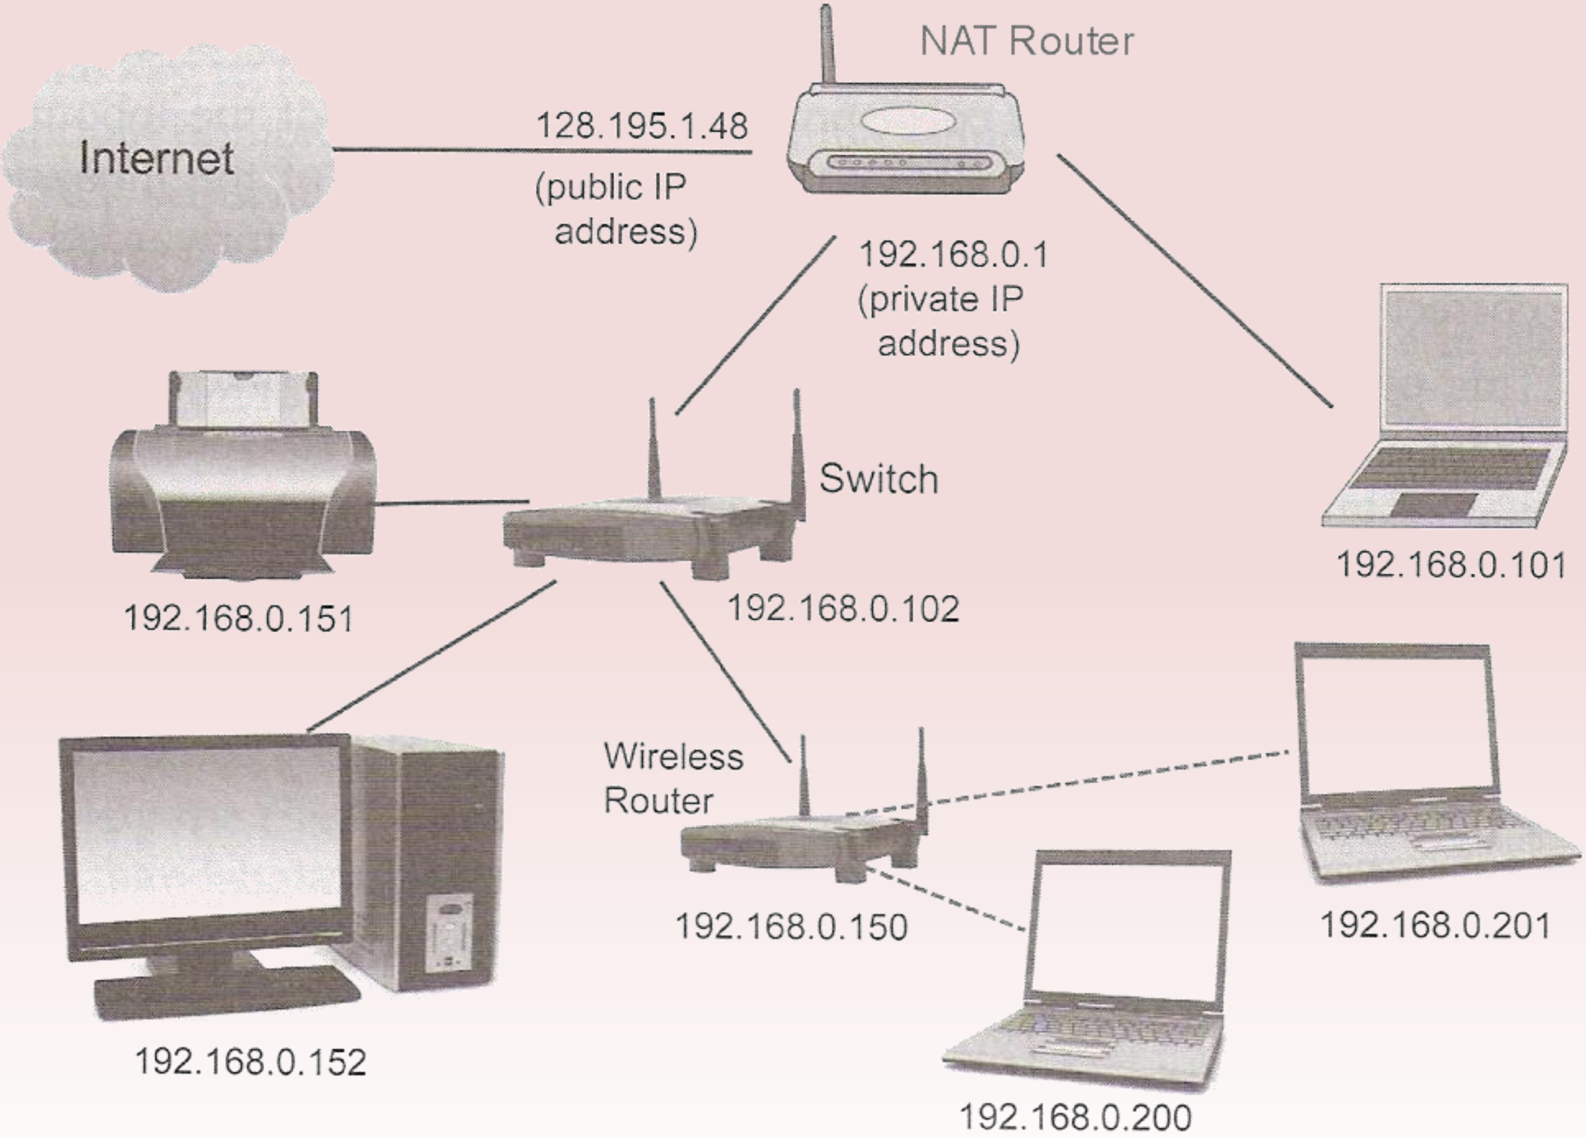

How NAT Works. To translate between private and public IP addresses, the NAT router maintains a lookup table that contains entries of the following form:

(*private source IP, private source port, destination IP, public source port*)

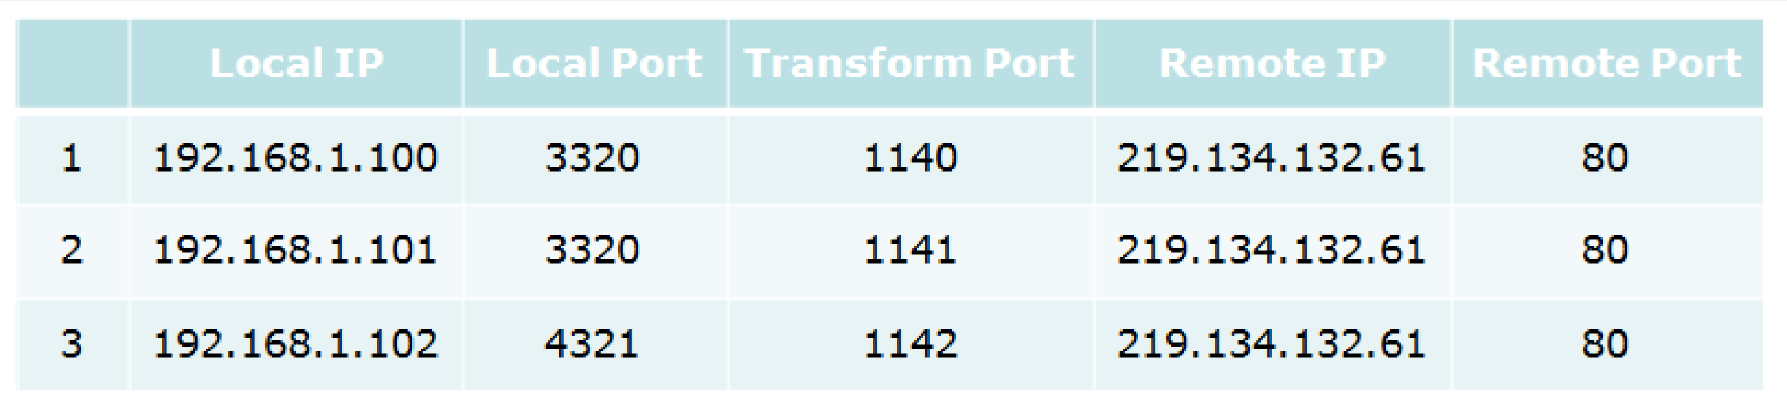

	Local IP	Local Port	Transform Port	Remote IP	Remote Port
1	192.168.1.100	3320	1140	219.134.132.61	80
2	192.168.1.101	3320	1141	219.134.132.61	80
3	192.168.1.102	4321	1142	219.134.132.61	80

A NAT router dynamically rewrites the header of all inbound and outbound TCP and UDP packets. When a machine on the internal network attempts to sent a packet to an external IP address, the NAT router creates a new entry in the lookup table associated with the source machine's private IP address and the internal source port of the transmitted packet.

The router rewrites the source IP address to be that of the NAT device’s public IP, opens a new public source port, and rewrites the IP header’s source port field to contain the newly opened port. This public port and destination IP address are recorded alongside the private source IP and private internal port in the NAT device’s lookup table.

After receiving a response, the NAT router checks its lookup table for any entries whose public source port corresponds to the destination port of the inbound packet and whose destination IP address corresponds to the source IP of the inbound packet. Finally, the NAT router rewrites the IP headers of the inbound packet according to the lookup table, so that the packet is forwarded to the correct private IP address and private port.

TCP Session Hijacking

Early TCP stacks implemented sequence numbers by using a single counter that was incremented by 1 with each transmission. Without using any randomness, it was trivial to predict the next sequence number, which is the key to a TCP sequence prediction attack.

Modern TCP stack implementations are pseudo-random number generators to determine sequence numbers, which makes a TCP sequence prediction attack more difficult, but not impossible.

Example of a TCP Sequence Prediction Attack:

1. The attacker launches a denial-of-service attack against the client victim to prevent that client from interfering with the attack.
2. The attacker sends a SYN packet to the target server, spoofing the source IP address to be that of the client victim.
3. After waiting a short period of time for the server to send a reply to the client, the attacker concludes the TCP handshake by sending an ACK packet with the sequence number set to a prediction of the next expected number, again spoofing the source IP to be that of the client victim.
4. The attacker can now send requests to the server as if he is the client victim.

Blind Injection. Makes it possible to inject a packet containing a command that creates a connection back to the attacker. The TCP sequence prediction attack only allows one-way communication.

ACK Storms. Is a side-effect of blind injection that causes a client and server to become out-of-synchronization with respect to sequence numbers, since the server got a synchronized message that the client never sent.

Complete Session Hijacking. When an attacker is on the same network segment as the target server and/or client, an attacker can completely hijack and existing TCP session. This attack is possible because an attacker can use packet sniffing to see the sequence numbers of the packets used to establish the session. Given this information, an attacker can inject a packet with a highly probable sequence number (and a well-chosen attack command) to the server using a spoofed source IP address impersonating the client.

Countermeasures. Countermeasures to TCP session hijacking attacks involve the use of encryption and authentication, either at the network layer, such as using IPsec, or at the application layer, such as using application-layer protocols that encrypt entire sessions. Also, web sites should avoid creating session that begin with secure authentication measures but subsequently switch over to unencrypted exchanges. These sessions trade off efficiency for security.

Denial of Service Attacks

Any attack that is designed to cause a machine or piece of software to be unavailable and unable to perform its basic functionality is known as a denial-of-service (DOS) attack. This includes any situation that causes a server to not function properly, but most often refers to deliberate attempts to exceed the maximum available bandwidth of a server.

Internet Control Message Protocol (ICMP) Attacks

There are two simple DOS attacks that exploits ICMP. They are:

1. Ping Flood Attack
2. Smurf Attack

<u>Ping Flood Attack</u>. The ping utility sends an ICMP echo request to a host, which in turn replies with an ICMP echo response. Normally, ping is used as a simple way to see if a host is working properly, but in a ping flood, a powerful machine can perform a DOS attack on a weaker machine.

To carry out the ping flood attack, a powerful machine sends a massive amount of echo requests to a single victim server. If the attacker can create many more ping requests that the victim can process, and the victim has enough network bandwidth to receive all these requests, then the victim server will be overwhelmed with the traffic and start to drop legitimate connections.

The Smurf Attack. This attack takes advantage of the broadcast feature that many networks uses, by which a user can send a packet that is received by every IP address on the network. The smurf attacks exploits this feature by sending ICMP packets with a source address set to the target and with a destination address set to the broadcast address of a network.

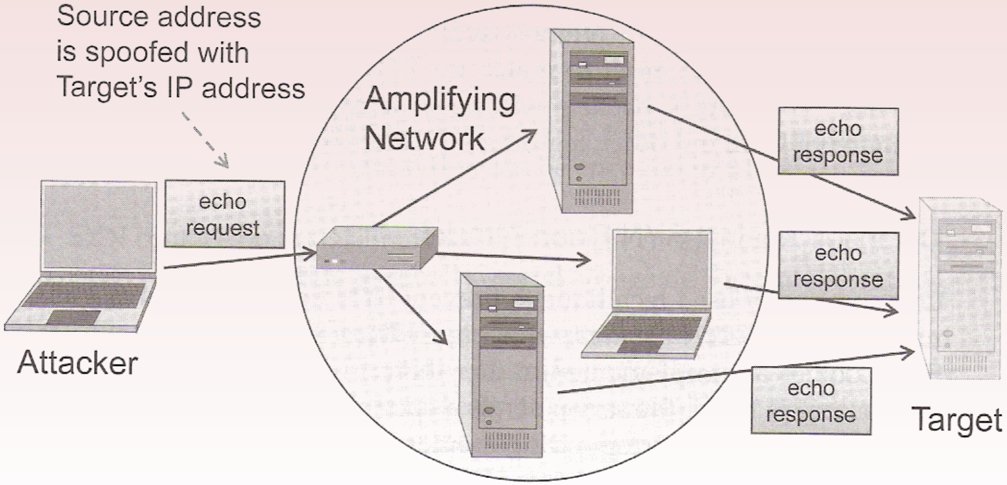

To prevent smurf attacks, administrators should configure hosts and routers on their networks to ignore broadcast request. In addition, routers should be configured to avoid forwarding packets directly to broadcast address.

SYN Flood Attacks

This is another type of denial-of-service attack. In the SYN flood attack an attacker sends a large number of SYN packets to the server, ignores the SYN/ACK replies, and never sends the expected ACK packets. An attacker initiating this attack will probably use random spoofed source addresses in the SYN packets, so that the SYN/ACK replies are sent to random IP addresses. If an attacker sends a large amount of SYN packets with no corresponding ACK packets, the server's memory will fill up with sequence numbers that it is remembering in order to match up TCP sessions with expected ACK packets.

One commonly used technique to prevent SYN flooding is to use SYN cookies, where the server encodes information in the TCP sequence number.

TCP Sequence Number

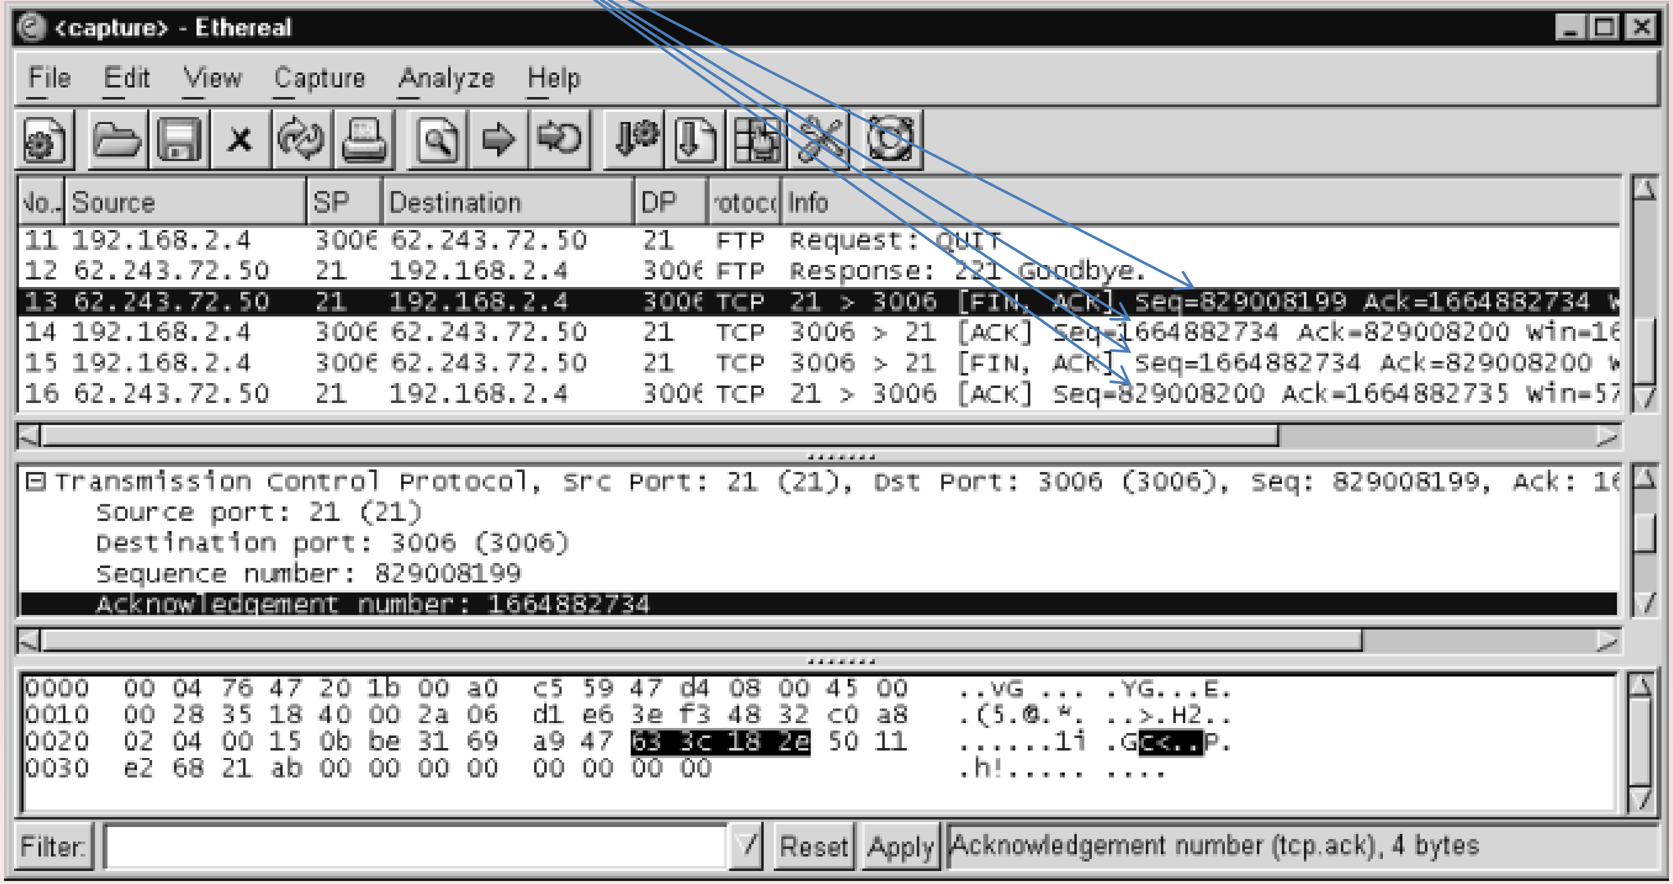

```
No. Source         SP   Destination    DP   rotoco Info
11  192.168.2.4    3006 62.243.72.50   21   FTP  Request: QUIT
12  62.243.72.50   21   192.168.2.4    3006 FTP  Response: 221 Goodbye.
13  62.243.72.50   21   192.168.2.4    3006 TCP  21 > 3006 [FIN, ACK] Seq=829008199 Ack=1664882734 W
14  192.168.2.4    3006 62.243.72.50   21   TCP  3006 > 21 [ACK] Seq=1664882734 Ack=829008200 Win=16
15  192.168.2.4    3006 62.243.72.50   21   TCP  3006 > 21 [FIN, ACK] Seq=1664882734 Ack=829008200 W
16  62.243.72.50   21   192.168.2.4    3006 TCP  21 > 3006 [ACK] Seq=829008200 Ack=1664882735 Win=57
```

```
⊟ Transmission Control Protocol, Src Port: 21 (21), Dst Port: 3006 (3006), Seq: 829008199, Ack: 16
    Source port: 21 (21)
    Destination port: 3006 (3006)
    Sequence number: 829008199
    Acknowledgement number: 1664882734
```

```
0000   00 04 76 47 20 1b 00 a0  c5 59 47 d4 08 00 45 00   ..vG ... .YG...E.
0010   00 28 35 18 40 00 2a 06  d1 e6 3e f3 48 32 c0 a8   .(5.@.*. ..>.H2..
0020   02 04 00 15 0b be 31 69  a9 47 63 3c 18 2e 50 11   ......1i .Gc<..P.
0030   e2 68 21 ab 00 00 00 00  00 00 00 00               .h!..... ....
```

Acknowledgement number (tcp.ack), 4 bytes

IP Traceback

Early IP traceback techniques relied on logging each packet forwarded by each router. While this approach maybe effective, it places significant space requirements on routers. A commonly proposed alternative relies on a technique known as packet marking. In this approach, routers probabilistically or deterministically mark forwarded packets with information related to the path that packet has taken up to that point.