Module 8
Network Security

Dr. Natarajan Meghanathan

Professor of Computer Science

Jackson State University, Jackson, MS 39232

E-mail: natarajan.meghanathan@jsums.edu

Module Topics

- 8.1   Classical Denial of Service (Spoofing-based) Attacks

- 8.2   Defense using Cryptography

- 8.3   IPSec

- 8.4   Firewalls and IDS

8.1 Classical Network Security Attacks

Dr. Natarajan Meghanathan
Professor of Computer Science
Jackson State University
E-mail: natarajan.meghanathan@jsums.edu

Reconnaissance

- **Port Scan**
- As a first step to prepare for an attack, an attacker runs a port scan tool (like nmap, netcat, etc) on the targeted machines and finds out some of the critical information about them:
  - Open ports and the services running on these ports
  - Operating system
  - Applications installed, etc
- The attacker then exploits the vulnerabilities of the operating system, services running on the open ports and the applications installed/running in the targeted machine to launch an attack.

- **IP Spoofing**
- An attacker can fake a message to come from an IP address that is not the actual sender.
- Port scanning can be done to identify both the victims (typically servers) as well as the hosts whose IP addresses can be spoofed and will serve as agents (zombies: compromised machines) for the attacker.

Distributed Denial-of-Service Attacks

- DDoS attacks involve breaking into hundreds or thousands of machines all over the Internet. Then, the attacker installs DDoS software on these burgled machines and controls them to launch coordinated attacks on victim sites.

- DDoS attacks typically exhaust bandwidth, router processing capacity and break network connectivity to the victims.

- <u>First stage</u>: Forming Zombies
  - The attacker uses any convenient attack (such as exploiting buffer overflow or tricking the victim to open and install unknown code from an email attachment) to plant a Trojan Horse on a target machine.

  - The attacker also installs software such as "rootkit" on these compromised machines. The rootkit helps to conceal the fact of the break-in, hide traces of subsequent malicious activities and also replaces the standard commands for displaying running processes with versions that fail to replace the attacker's processes. The compromised victim machine is referred to as a "zombie"

  - A zombie could also be used to break into some more machines, install the Trojan Horse and rootkits and convert them to be a zombie. A network of zombies is called a botnet.

Distributed Denial-of-Service Attacks

- <u>Second stage:</u> Launch the DDoS attack through the zombies.
  - The attacker choose a victim and sends a signal to all the zombies to launch the attack. Each zombie could launch a different type of attack on the victim.
  - A victim of the DDoS attack will thus have to counter multiple zombies, launching the same or different types of attacks.

1. Attacker plants Trojan horse in zombies

2. Zombies attack victim simultaneously on command
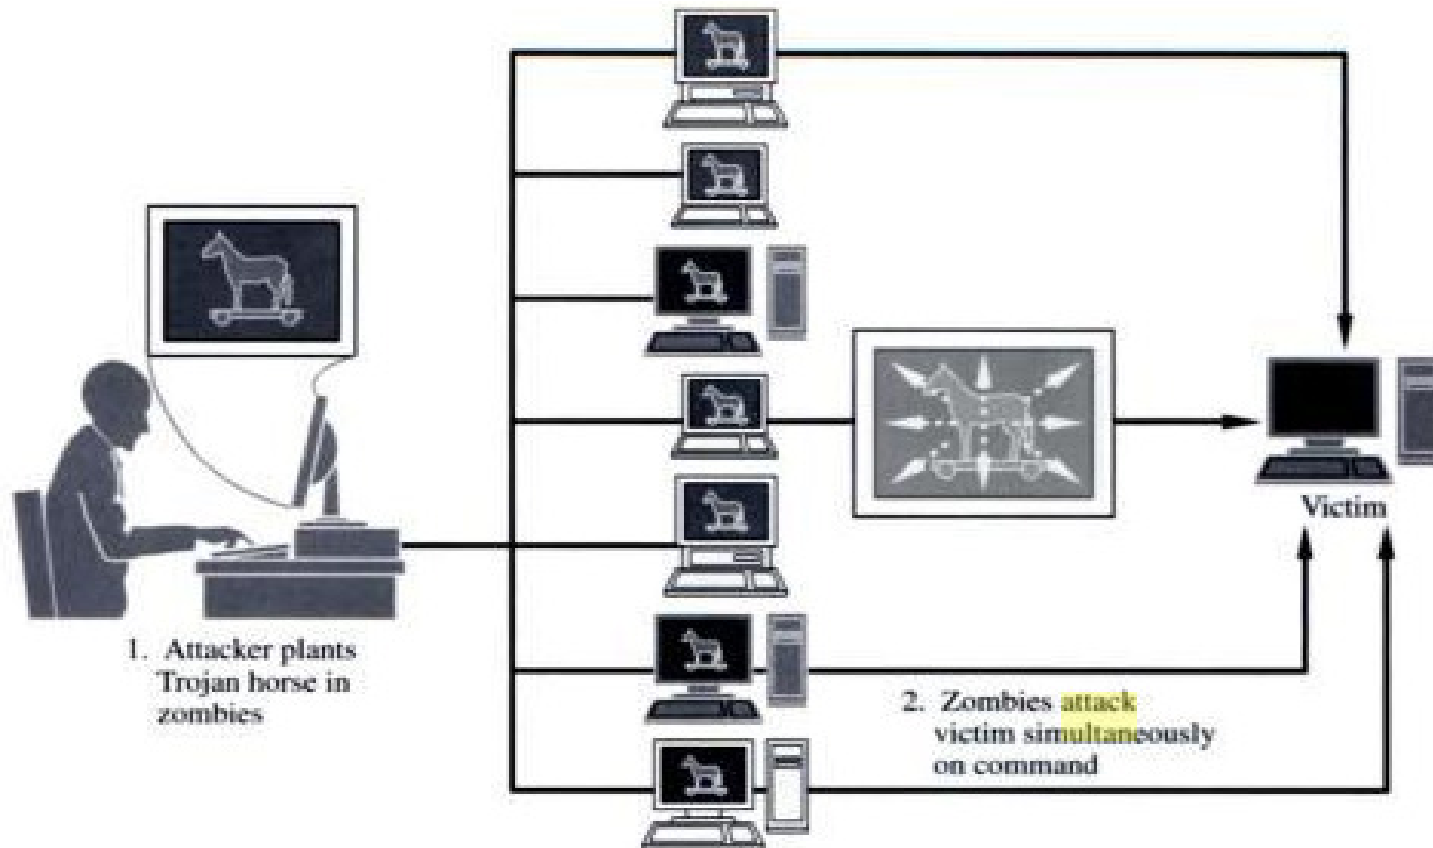
Victim

SYN Flood Attack

- An attacker initiates several SYN request messages that appear to come from different (spoofed) IP addresses to a targeted victim (typically a server).

- The victim allocates resources for each connection request, sends SYN+ACK packets and waits for the response from the initiators.

- The attacker does not respond with the ACK packets to complete the connection establishment process.

- The victim would then eventually end up denying connection requests from genuine clients, thus leading to a denial of service.

Smurf Attack

- The smurf attack, a kind of denial-of-service attack, floods a target system via spoofed broadcast ping Echo-Reply messages.
- A perpetrator sends a ping Echo-Request message (having a spoofed address of the intended victim) to the broadcast IP address of a network.
- All the hosts in that network on receiving the ping message, send a reply to the source of the ping, which is the victim machine.

Victim

Attacker

Attacker sends broadcast ECHO request to network, with victim's return address

All network hosts reply to victim

Victim is saturated with ECHO replies from entire network

- <u>Solution:</u> After the incidence of several Smurf attacks, routers in the Internet were configured not to forward packets having a broadcast IP address as the destination address. Hosts are also configured not to respond to ping requests that were sent to them as a broadcast message.
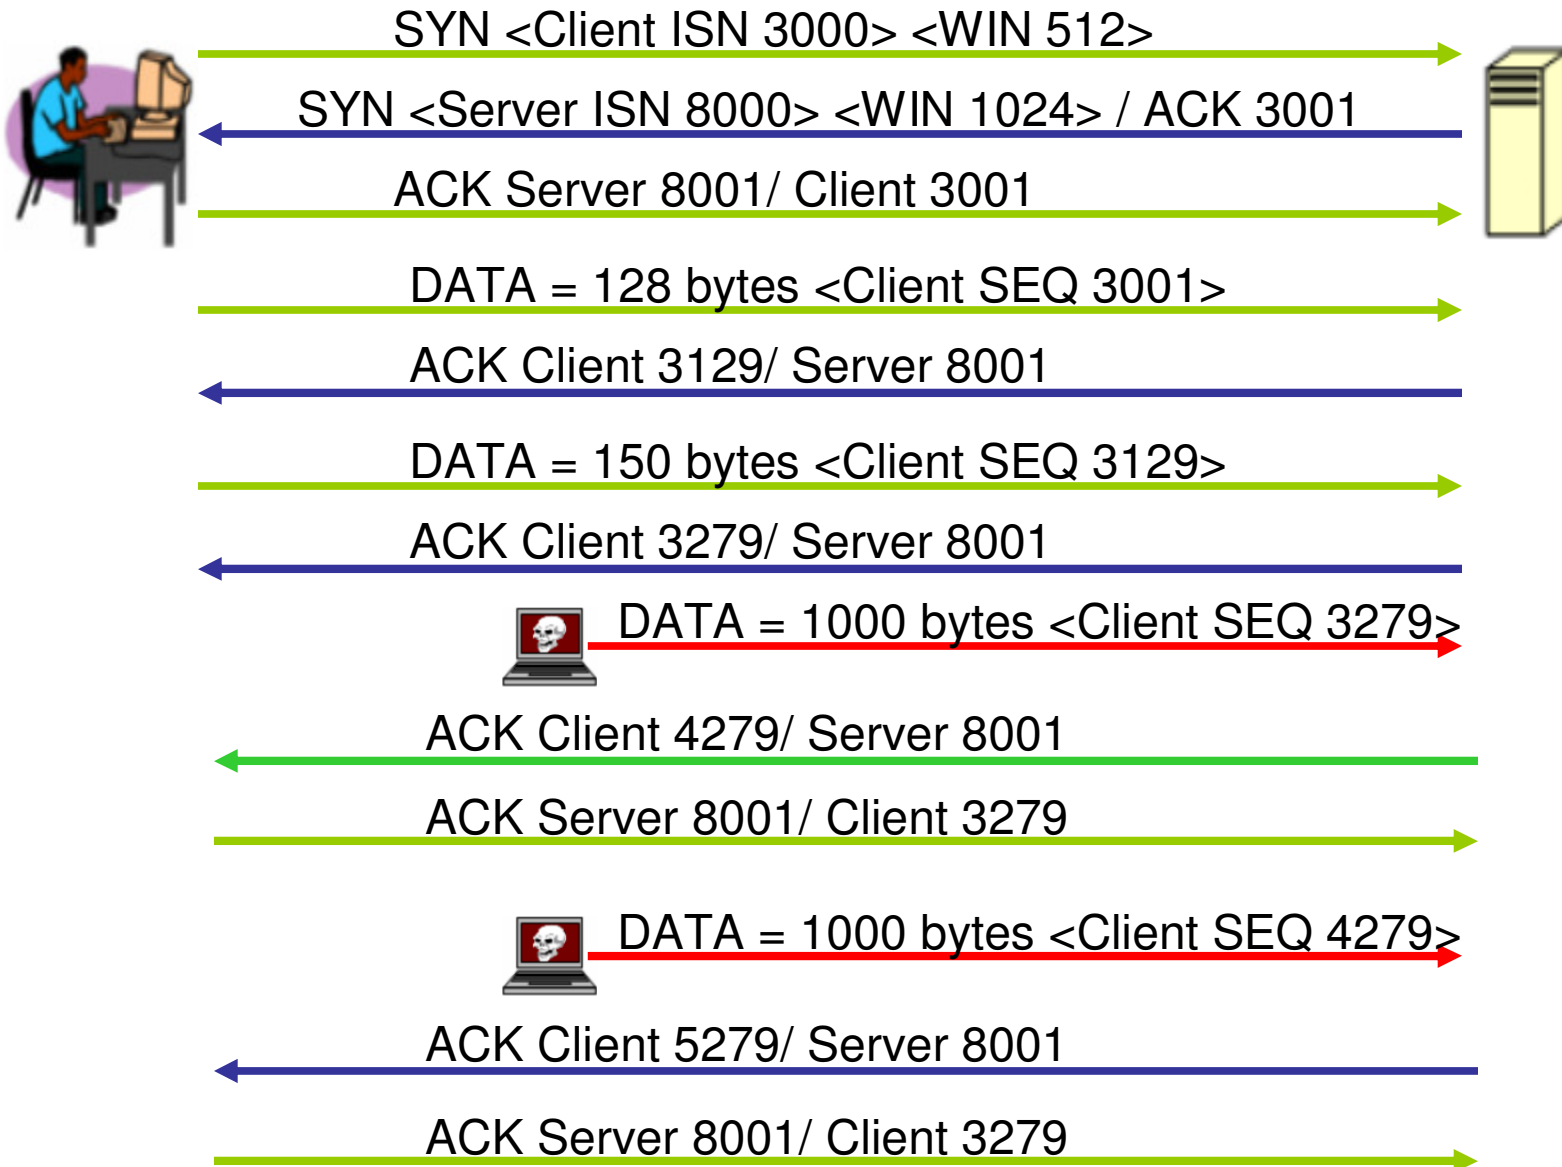
Echo-Chargen Attack

- The Chargen (Character Generator) service is an internet-layer protocol and is intended for testing and measurement purposes.

- A host may connect to a server that supports the Chargen protocol, on either TCP or UDP port 19.

- Upon opening a TCP connection, the server starts sending arbitrary characters to the connecting host and continues until the host closes the connection.

- In the UDP version of the protocol, the server sends an UDP packet containing a random number (between 0 and 512) of characters every time it receives an UDP packet from the connecting host.

- Attack Scenario: The attacker sends a packet to machine A's UDP port 19 with machine B as the forged source address and UDP port 7 (Echo protocol) as the source port.
  - The attacker spoofs a conversation between the two services and redirects the output of each service to the other, creating a rapidly expanding spiral of traffic.
  - Eventually, the attack begins to consume memory and processor power at the targeted devices A and B, causing them to become non-responsive to user commands.

TCP Session Hijacking

- Session hijacking is the act of taking over an already established TCP session and injecting your own packets into that stream so that your commands are processed as the authentic owner of the session.
- A TCP session is merely identified by the quadruple: Client IP address, Client Port number, Server IP address and Server Port number.
- Any packets that reach either machine with the above identifiers are assumed to be part of the existing session.
- Hence, if an attacker can spoof these items, they can pass TCP packets to the client or server and have those packets processed as coming from the other machine.
- Two steps: Desynchronize the session and Inject own commands
- Desynchronizing the session: Predict the sequence number to be used by a client (or server) and use that sequence number before the client (or server) gets a chance to.
  - How to do it? Use Local Session Hijacking or Blind Session Hijacking
  - <u>Local Session Hijacking:</u> If we have access to the network and can sniff the TCP session, we can tell the next expected sequence number from the ACK packets exchanged
  - <u>Blind Session Hijacking:</u> If we do not have the ability to sniff the TCP session between the client and server, then we have try all options and guess the expected sequence number.

Desynchronizing a Session

SYN <Client ISN 3000> <WIN 512>

SYN <Server ISN 8000> <WIN 1024> / ACK 3001

ACK Server 8001/ Client 3001

DATA = 128 bytes <Client SEQ 3001>

ACK Client 3129/ Server 8001

DATA = 150 bytes <Client SEQ 3129>

ACK Client 3279/ Server 8001

DATA = 1000 bytes <Client SEQ 3279>

ACK Client 4279/ Server 8001

ACK Server 8001/ Client 3279

DATA = 1000 bytes <Client SEQ 4279>

ACK Client 5279/ Server 8001

ACK Server 8001/ Client 3279

TCP ACK Storm

ACK Client 5279/ Server 8001

ACK Server 8001/ Client 3279

ACK Client 5279/ Server 8001

ACK Server 8001/ Client 3279

ACK Client 5279/ Server 8001

ACK Server 8001/ Client 3279

DATA = 1000 bytes <Client SEQ 5279>

ACK Client 6279/ Server 8001

ACK Server 8001/ Client 3279

ACK Client 6279/ Server 8001

ACK Server 8001/ Client 3279

ACK Client 6279/ Server 8001

ACK Server 8001/ Client 3279

TCP Session Hijacking

- When the attacker successfully hijacks the TCP session and injects its own data (and spoofs the server as if the data is coming from the original client), the server will acknowledge the receipt of the data by sending to the original client an ACK packet.

- This ACK packet will most likely contain a sequence number that the client is not expecting, so when the client receives this packet, it will try to resynchronize the TCP Session with the server by sending it an ACK packet with the sequence number that it is expecting.

- This ACK packet will in turn contain a sequence number that the server is not expecting and so the server will resend its last ACK packet.

- This cycle will go on and on, and this rapid passing back and forth of the ACK packets creates the TCP ACK Storm.

- The attacker could keep injecting more and more data, the size of the ACK storm increases and can quickly degrade the network performance.

- The original client will have to eventually get exhausted after a certain number of resynchronization attempts and close the connection with the server.

8.2: Defense using Cryptography

Dr. Natarajan Meghanathan

Professor of Computer Science

Jackson State University
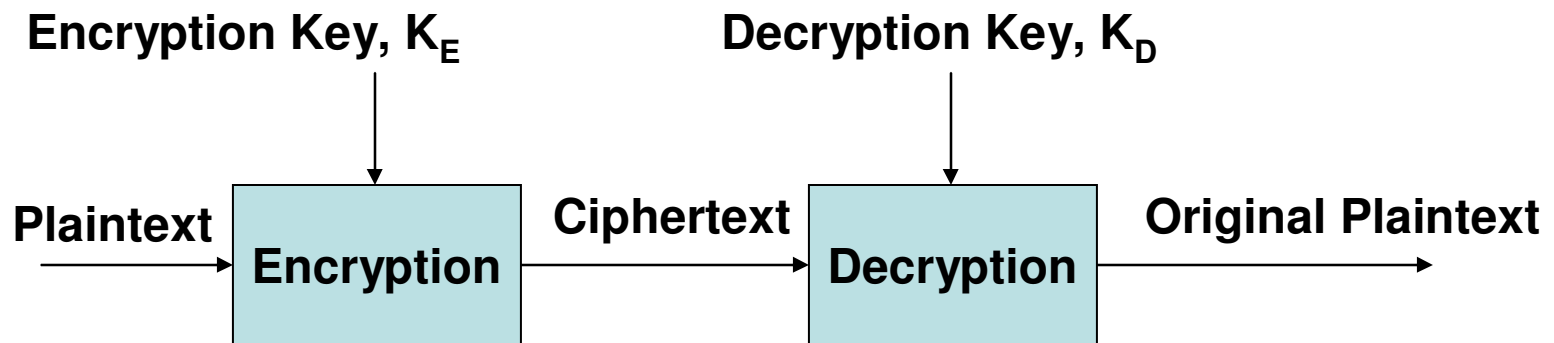
E-mail: natarajan.meghanathan@jsums.edu

Types of Encryption

- <u>Symmetric encryption</u>: The same key performs, both encryption and decryption.
  - $P = D(K, E(K, P))$

Key → **Plaintext** → **Encryption** → **Ciphertext** → **Decryption** → **Original Plaintext**

- <u>Asymmetric encryption</u>: distinct, very different keys, one for encryption and the other for decryption only

Encryption Key, K_E   **Decryption Key, K_D**

Plaintext → **Encryption** → **Ciphertext** → **Decryption** → **Original Plaintext**

# Public-Key Encryption

- Let K_{PRIV} and K_{PUB} be the private key and public key of a user. Then,
  - $P = D(K_{PRIV}, E(K_{PUB}, P))$
  - $P = D(K_{PUB}, E(K_{PRIV}, P))$

- <u>Exchange of Secret Message using Asymmetric Encryption</u>
  - Let K_{PUB-S}, K_{PRI-S} denote the public and private keys of Sender S. Similarly, let K_{PUB-R} and K_{PRI-R} be the public and private key of Receiver R. Let M be the secret message to be sent from S to R.
  - S sends to R the following:
    - $E (K_{PUB-R} E(K_{PRI-S}, M))$
  - The inner encryption guarantees that the secret message M came from S and the outer encryption guarantees that only the receiver R could open the outer encryption of the message and get access to the inner encryption.

Use of Public-Key Encryption to Provide Integrity and Authentication

- $M \parallel E_{Pri-S}(Hash(M))$

Use of Public-Key Encryption to Provide Confidentiality

- $E_{Secret-Key}(M) \parallel E_{Pub-R}(Secret-key)$

Source Side

PU_R

Destination Side

$K_S \longrightarrow$ EP

M

K_S

EC $\longrightarrow \parallel$

Dec $\longrightarrow$ M

K_S

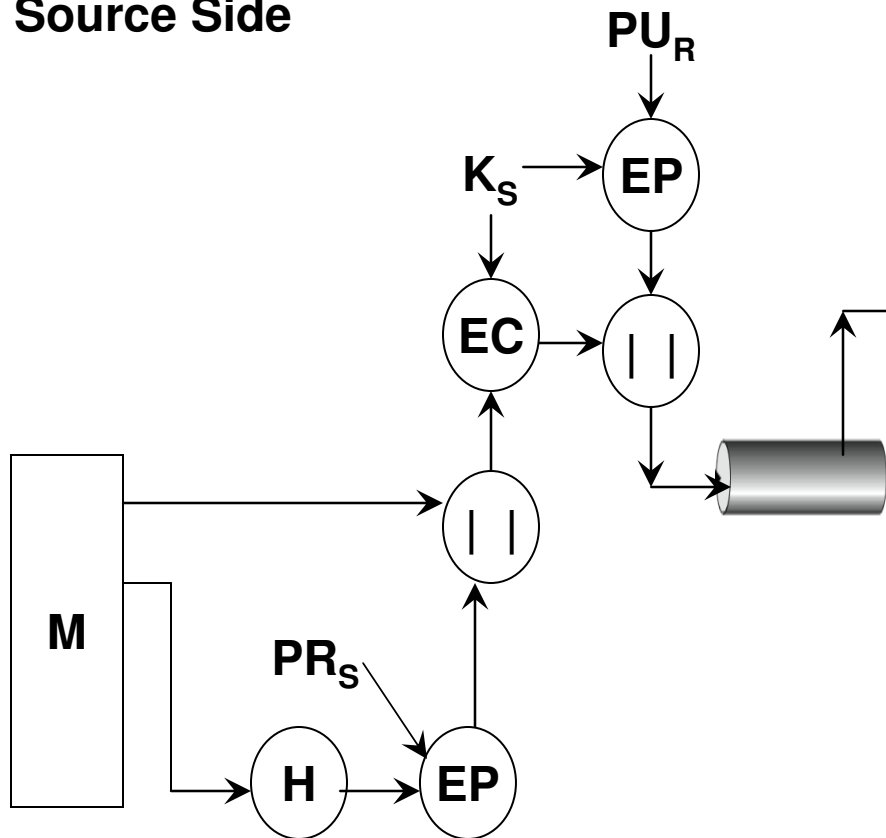Dec

$E(PU_R, K_S)$

PR_R

Note that the secret key is generated for each message for better security and also to avoid the need for key distribution.
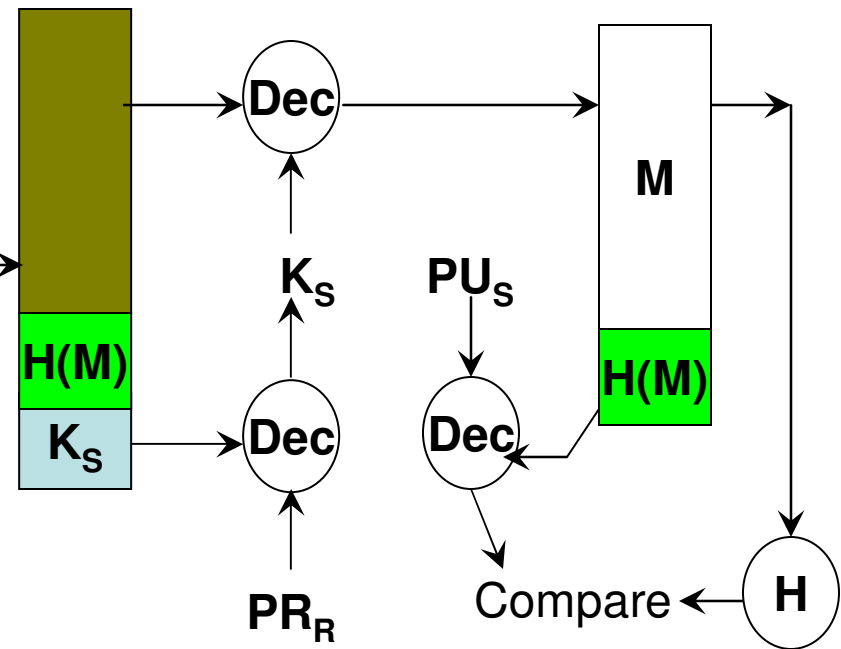
Use of Public-Key Encryption to Provide Confidentiality, Integrity and Authentication

$$E_{\text{Secret-Key}}(M \parallel E_{\text{Pri-S}}(\text{Hash } (M))) \parallel E_{\text{Pub-R}}(\text{Secret-key})$$
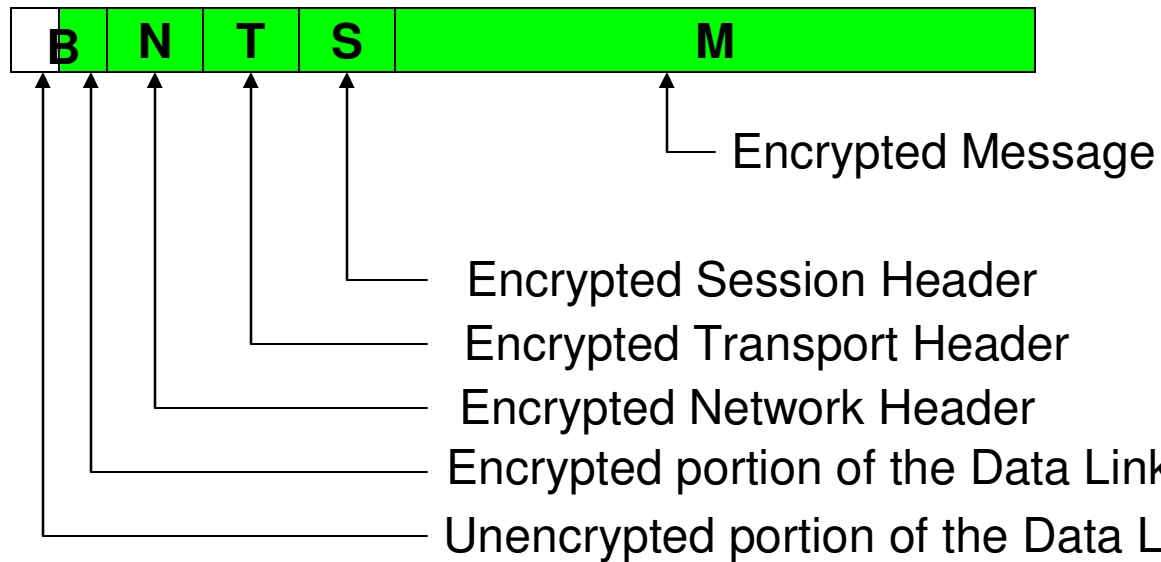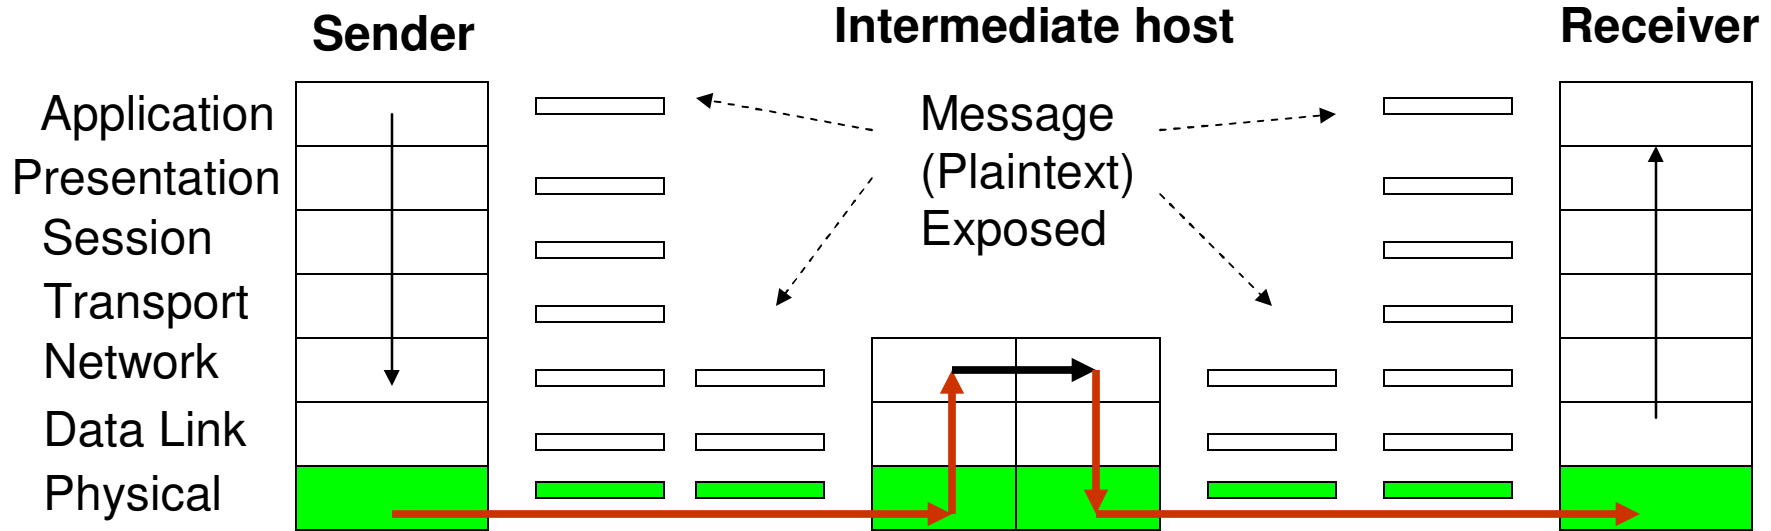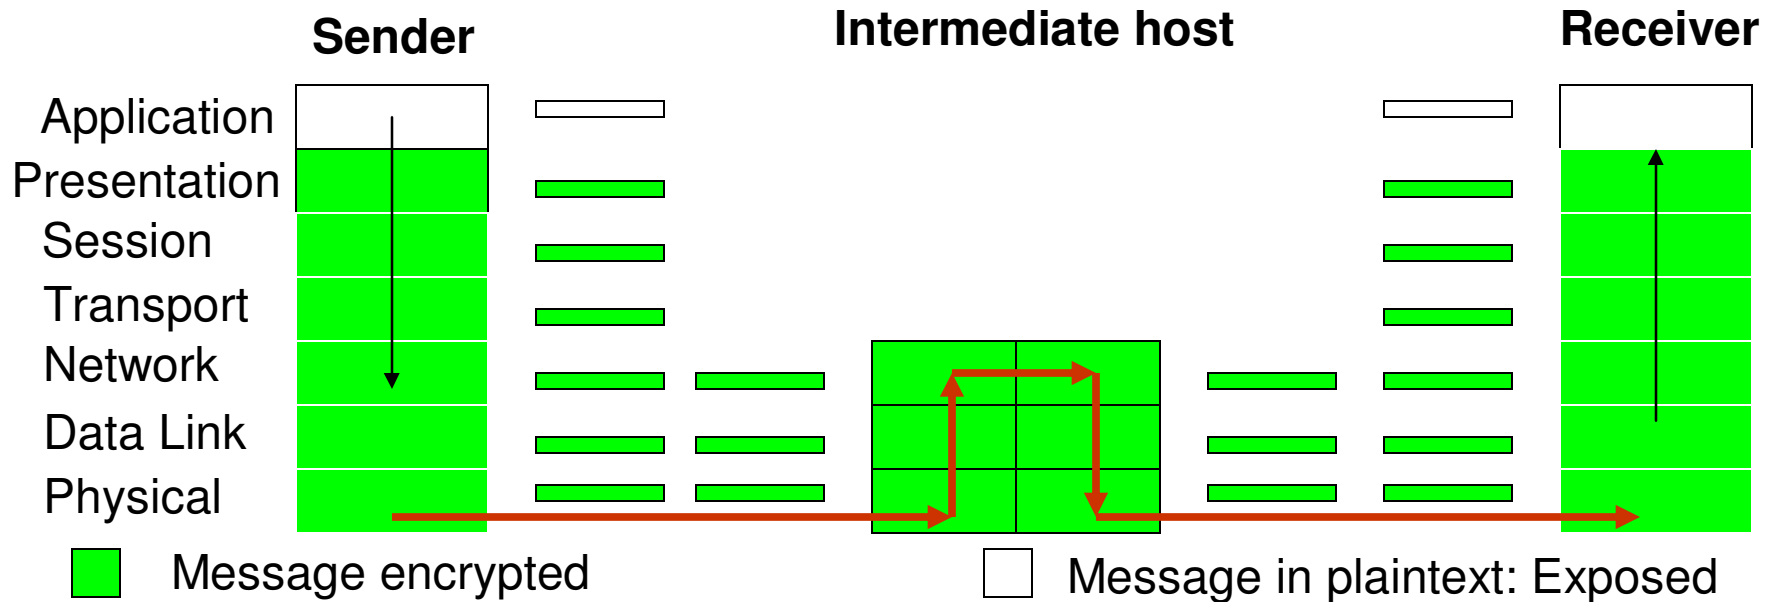
Link Encryption vs. End-to-End Encryption

- Link Encryption: Encrypt a packet at every network on its way from the source network to the destination network.
  - Encryption handled at the network routers
  - Data is exposed at the intermediate routers
  - Every network on the source-destination path needs to support encryption; otherwise no use
  - Encryption keys are to be maintained for every pair of network routers.

- End-to-End Encryption: The application user encrypts the data
  - No need for the networks to support encryption
  - Data is not exposed at intermediate routers
  - Encryption keys need to be maintained on per-user basis

Link Encryption

Message encrypted

Message in plaintext: Exposed

Sender	Intermediate host	Receiver

Application
Presentation
Session
Transport
Network
Data Link
Physical

Message (Plaintext) Exposed

B	N	T	S	M

Encrypted Message

Encrypted Session Header

Encrypted Transport Header

Encrypted Network Header

Encrypted portion of the Data Link Header

Unencrypted portion of the Data Link Header

End-to-End Encryption

B	N	T	S	M

Unencrypted headers

Sender   **Intermediate host**   **Receiver**

Application
Presentation
Session
Transport
Network
Data Link
Physical

Message encrypted   Message in plaintext: Exposed
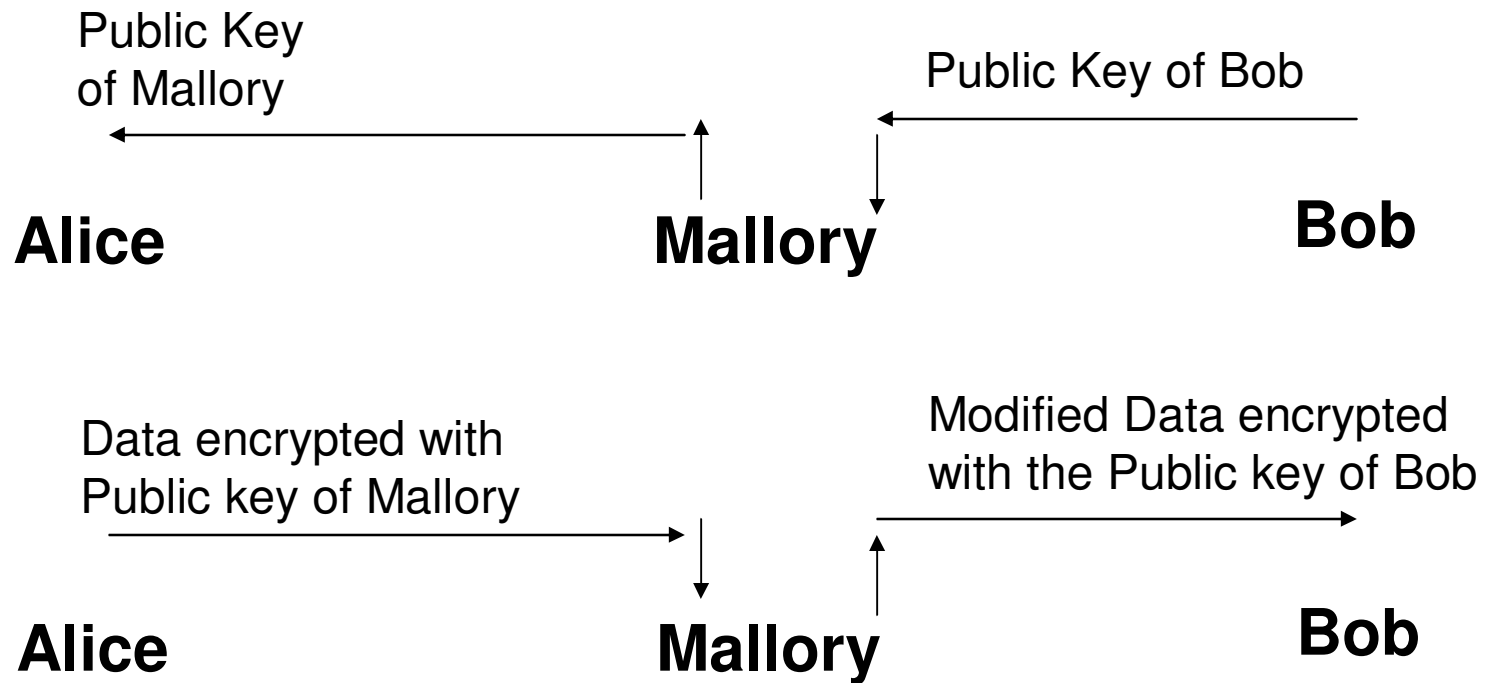
Man-in-the-Middle Attack

- Man-in-the-middle (MITM) attack is an attack in which an attacker is able to read, insert and modify at will, messages between two parties without either party knowing that the link between them has been compromised.

- The attacker must be able to observe and intercept messages going between the two victims.

- Example: (MITM attack on public-key cryptography)
  - Suppose Alice wishes to communicate with Bob.
  - Mallory wants to eavesdrop their conversation or also possibly deliver a false message to Bob.
  - First, Alice must ask Bob for his public key.
  - If Bob sends his public key to Alice, but Mallory is able to intercept it, a MITM attack can begin.
  - Mallory sends a forged message to Alice that claims to have come from Bob, but contains Mallory's public key
  - Alice believes the public key received to be that of Bob's. So, Alice encrypts the message she wishes to send to Bob using the public key received and transmits on the link to Bob.
  - Mallory could now intercept the message, decrypt it with his private key and get the actual contents of the message.
  - Mallory now again encrypts the message (could be even altered too) with Bob's public key and transmits the message to Bob.
  - Bob on receiving the message, decrypts the message with his private key and reads the contents of the message assuming it came from Alice

Man-in-the-Middle Attack

Public Key
of Mallory

Public Key of Bob

Alice   **Mallory**   **Bob**

Data encrypted with
Public key of Mallory

Modified Data encrypted
with the Public key of Bob

**Alice**   **Mallory**   **Bob**
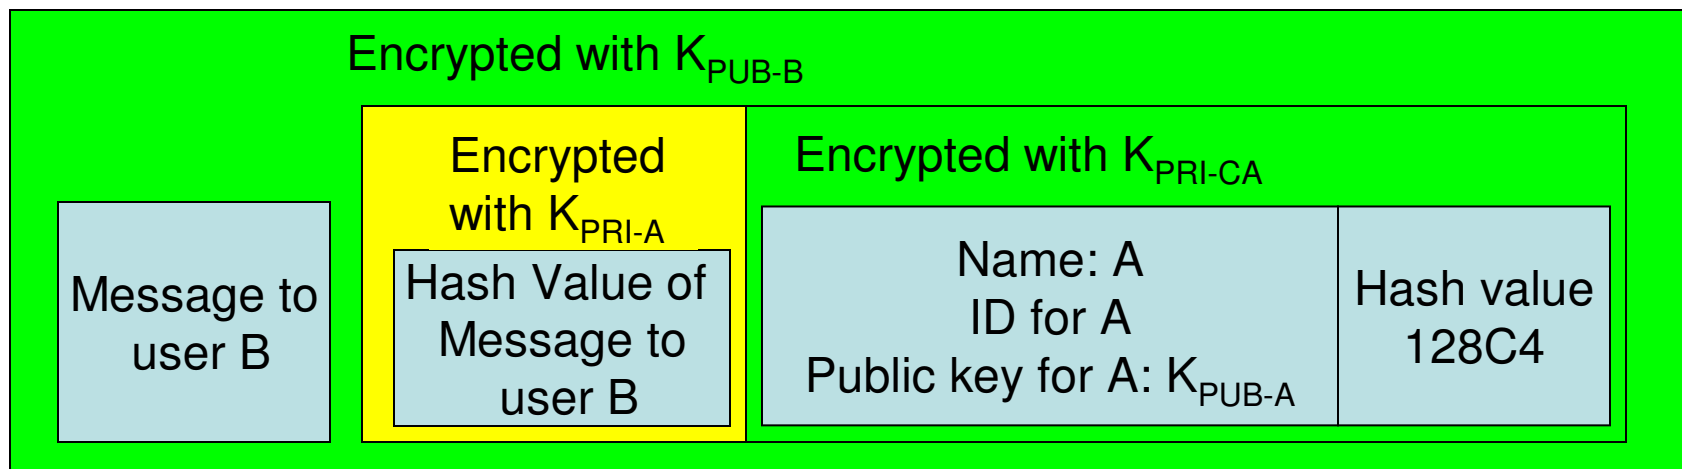
Public-Key Certificates

- Each of us adopt a "trust threshold" – a degree to which we are willing to believe an unidentified individual.

- We will use the concept of "vouching for" by a third party as the basis of trust in settings where two parties do not know about each other.

- Certification Authority (CA): Is an entity that issues digital certificates that contain a public key and the identity of the owner.

- The CA attests that the public key contained in the digital certificate belongs to the person (CA is a sort of digital notary).

Certificates

Digital Certificate for the Public Key of A

Encrypted with K_{PRI-CA}

Name: A ID for A Public key for A: K_{PUB-A}	Hash value 128C4

User A sending to user B

Encrypted with K_{PUB-B}

Message to user B	**Encrypted with K_{PRI-A}** Hash Value of Message to user B	**Encrypted with K_{PRI-CA}** Name: A ID for A Public key for A: K_{PUB-A}	Hash value 128C4

Note: The certificates are created and formatted based on the X.509 standard, which outlines the necessary fields of a certificate and the possible values that can be inserted into these fields. The latest X.509 version is v.3.

8.3: IPSec

Dr. Natarajan Meghanathan
Professor of Computer Science
Jackson State University
E-mail: natarajan.meghanathan@jsums.edu

IPSec

- Before two machines send the messages using their IP addresses, they have to establish an IPSec SA (Security Association)
- IPSec SA
    - The two hosts A and B exchange their public-key certificates (that has their IP address and public-key certified).
    - All further communication are encrypted with the public key of the receiver (so that it can be decrypted only by the receiver with its private key)
    - The two hosts A and B negotiate on the encryption and keyed-hashing algorithms to use for confidentiality and integrity + authentication respectively.
    - The two hosts establish a session key (for integrity and authentication) using the public-key encryption based Diffie-Hellman key exchange mechanism.
    - Using the session key, the two hosts can then establish a secret key for confidentiality in communication.
    - IPSec SA is unidirectional: If machines A and B want to send messages back and forth, they have to establish an IPSec SA in each direction.
    - An IPSec SA from A to B is said to be outbound at A and inbound at B.
    - An IPSec SA from A to B is identified by the tuple <SPI, IPaddress of A> where SPI is the Security Parameter Index value, locally unique at A. The combination of the SPI with the IP address of the host makes the tuple globally unique.

IPSec

- The IPSec header is inserted in between the IP header and transport layer header. There is no need for support from any higher layers.

IP Header	IPSec Header	TCP/ UDP Header	Data

- IPSec headers:
  - <u>Authentication Header (AH)</u>: used for integrity + authentication
  - <u>Encapsulated Security Payload Header (ESP)</u>: used for confidentiality, integrity + authentication.
- IPSec Modes:
  - <u>Transport Mode:</u> When IPSec SA is directly established between the two end hosts. Message is secure all the way from the source host to the destination host
  - <u>Tunnel Mode:</u> When IPSec SA is established between the gateway routers of the two end hosts. Message is not secure in the source and destination networks. Need to use IP-in-IP encapsulation to encapsulate the IP datagram with the IP addresses of the two ultimate end hosts.

Typical IPSec Scenarios

Transport Mode

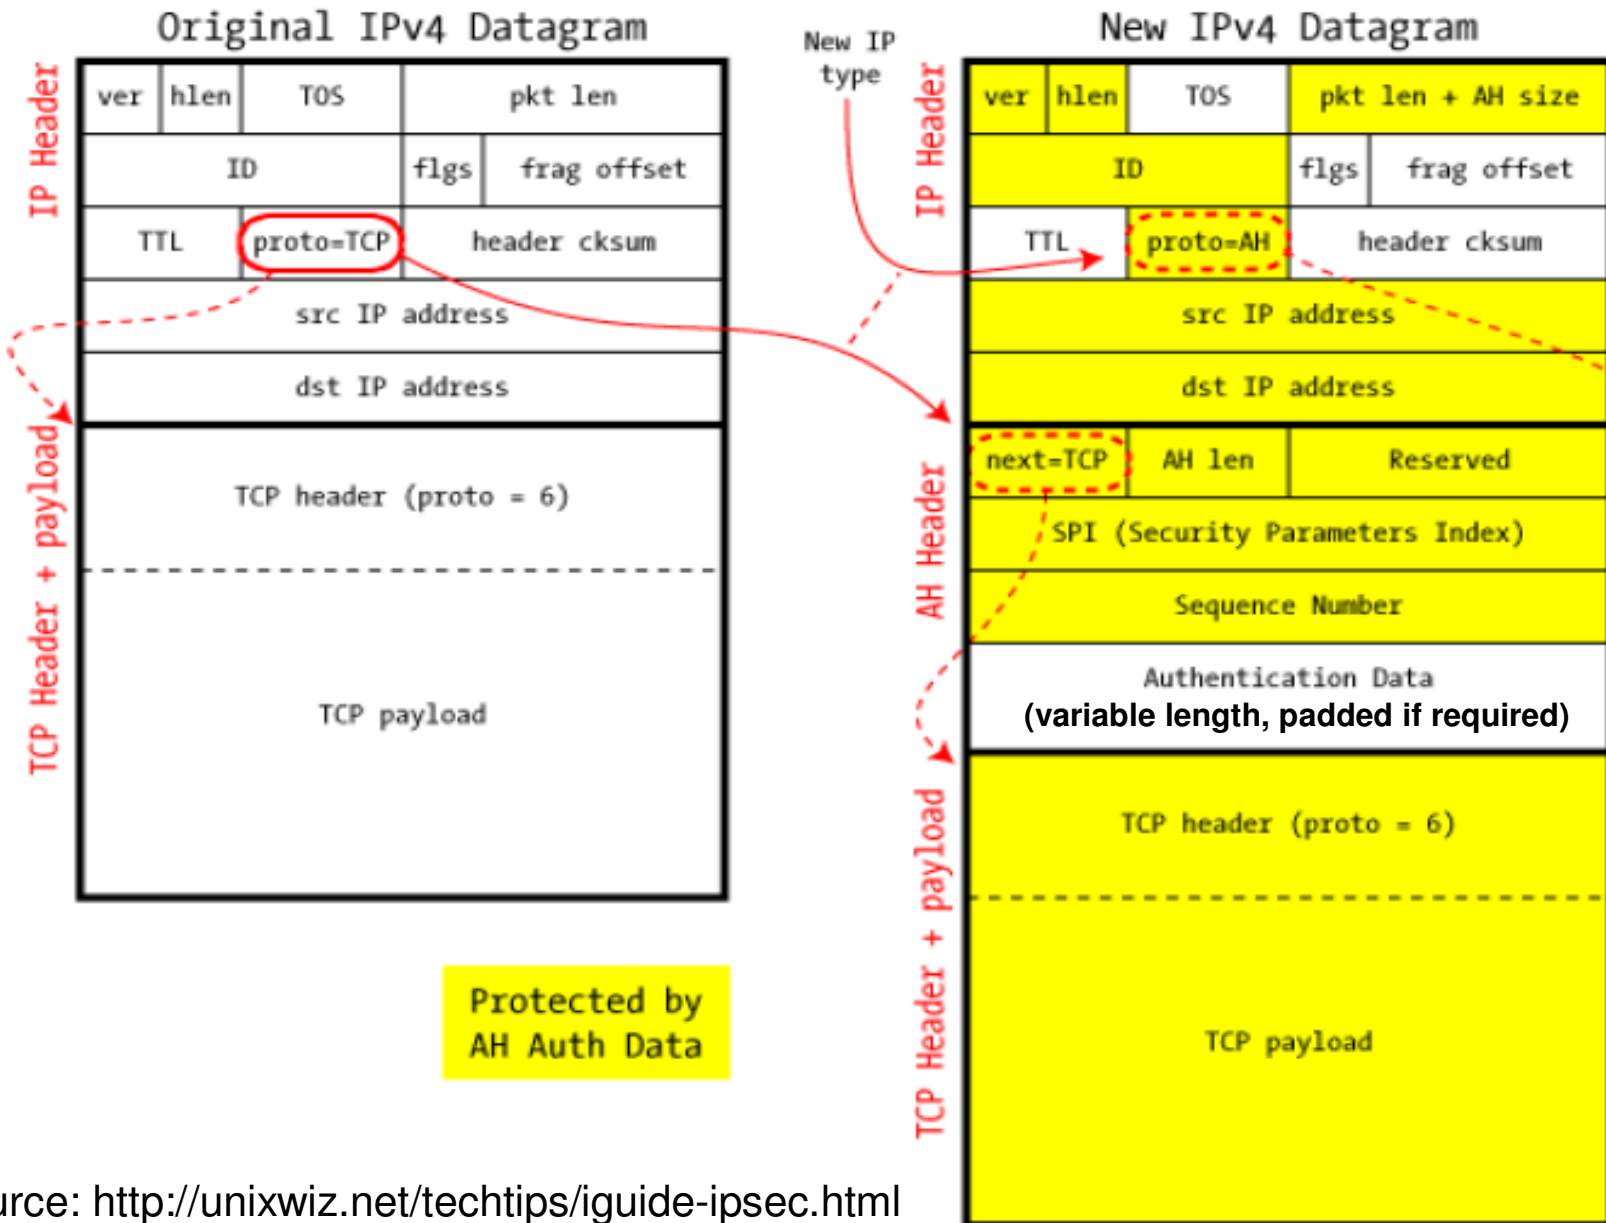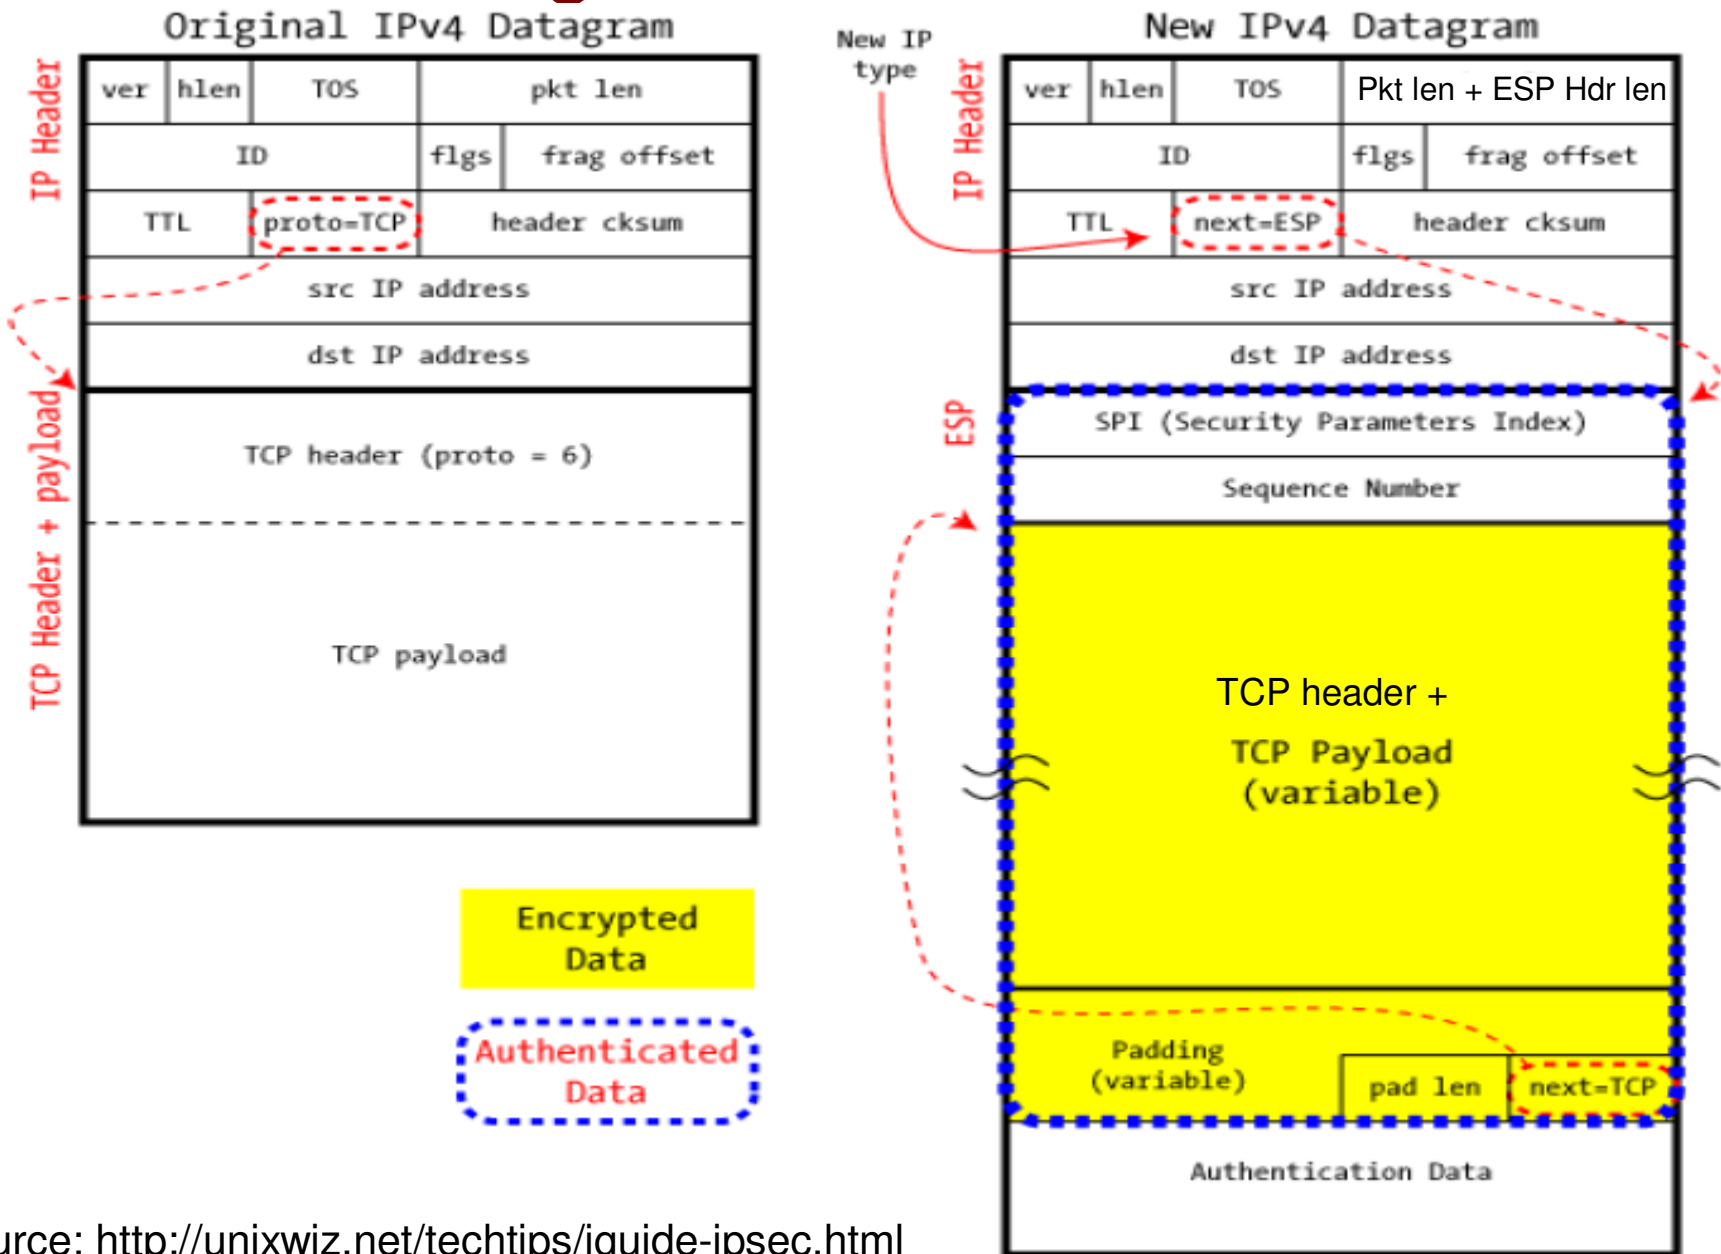| IP Header IP$_S$ IP$_D$	IPSec Header SA: S -- D	TCP Header	Data

Tunnel Mode

IP Header IP$_{R1}$ IP$_{R3}$	IPSec Header SA: R1 – R3	IP Header IP$_S$ IP$_D$	TCP Header	Data

IP4 Datagram with Authentication Header

Original IPv4 Datagram

IP Header			
ver	hlen	TOS	pkt len
ID		flgs	frag offset
TTL	proto=TCP		header cksum
src IP address			
dst IP address			

TCP Header + payload:
TCP header (proto = 6)
TCP payload

New IP type

New IPv4 Datagram

IP Header			
ver	hlen	TOS	pkt len + AH size
ID		flgs	frag offset
TTL	proto=AH		header cksum
src IP address			
dst IP address			

AH Header:
next=TCP | AH len | Reserved
SPI (Security Parameters Index)
Sequence Number
Authentication Data
(variable length, padded if required)

TCP Header + payload:
TCP header (proto = 6)
TCP payload

Protected by AH Auth Data

IP4 Datagram with ESP Header

Original IPv4 Datagram

ver	hlen	TOS	pkt len

IP Header

ver	hlen	TOS	pkt len
ID		flgs	frag offset
TTL	proto=TCP	header cksum	
src IP address			
dst IP address			

TCP Header + payload

TCP header (proto = 6)

TCP payload

Encrypted Data

Authenticated Data

New IPv4 Datagram

New IP type

IP Header

ver	hlen	TOS	Pkt len + ESP Hdr len
ID		flgs	frag offset
TTL	next=ESP	header cksum	
src IP address			
dst IP address			

ESP

SPI (Security Parameters Index)

Sequence Number

TCP header + TCP Payload (variable)

Padding (variable)	pad len	next=TCP

Authentication Data

Source: http://unixwiz.net/techtips/iguide-ipsec.html

8.4 Firewalls and IDS

Dr. Natarajan Meghanathan
Professor of Computer Science
Jackson State University
E-mail: natarajan.meghanathan@jsums.edu

Firewalls

- A firewall is a software running on a dedicated host computer on which no other application is run.
  - To prevent someone from changing the firewall rules by exploiting the vulnerabilities of the other applications that may be run on the host.

- A firewall is as good as it is configured with, depending on the needs of the admin.

- Allowable or non-allowable traffic are typically identified with source/destination IP/network addresses and ports.

- **Filtering:** Egress filtering (filter outgoing traffic); Ingress filtering (filter incoming traffic)

- **Typical Firewall designs:**
  - **Default-deny approach (white-list):** Have a list of allowable traffic and block the rest.
  - **Default-allow approach (black-list):** Have a list of non-allowable traffic and allow the rest.
  - A good design needs to have a hybrid of these two approaches

Packet Filters

- A packet filter firewall is a stateless firewall that looks at only the packet headers to decide whether or not to drop a packet.
  - Stateless: Does not keep track of the decisions taken on any packet.
    - The decision taken on a packet is independent of the decision taken on the preceding packets.
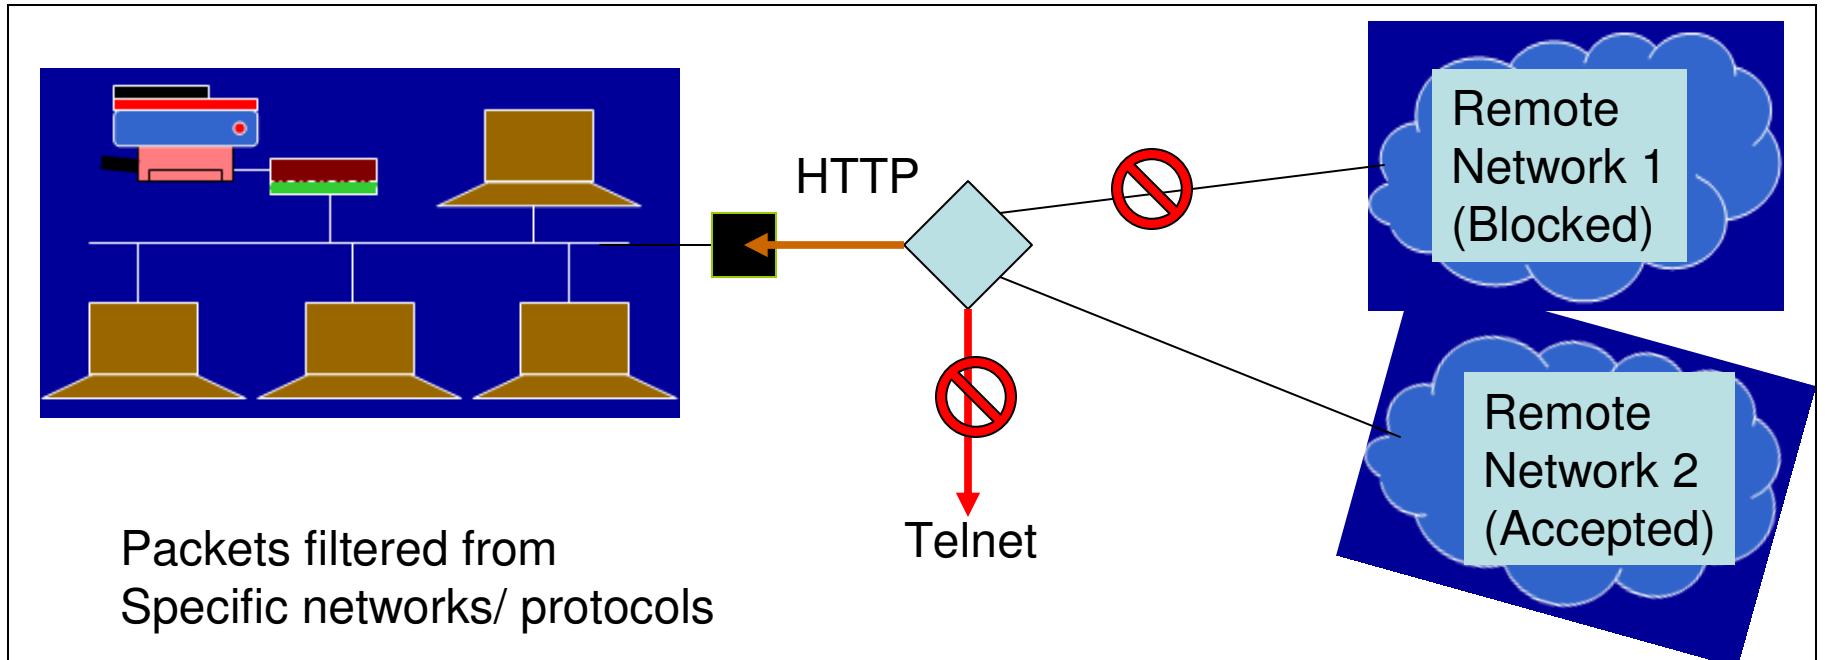
End-to-End Connection ←

Application
Transport
Network
Link Layer
Physical Layer

→ **End-to-End Connection**

Layers supported by Packet Filter and Stateful Firewalls

The code for packet filters will become lengthy as we want to block traffic belonging to specific networks, IP addresses and transport layer protocols.

Need efficient filtering algorithms

Packet Filters

Packets filtered from
Specific networks/ protocols

HTTP

Telnet

Remote
Network 1
(Blocked)

Remote
Network 2
(Accepted)

100.50.25.x

Subnet 100.50.25.x

Screening
Router

Incoming Packets with
source address spoofed
to belong to Subnet
100.50.25.x are filtered

Attacks Detected by Packet Filters

- **IP Spoofing Attacks:** Have the packet filter configured not to let in packets having a source address that corresponds to the internal network.
  - For example, the attacker has spoofed the source IP address to be the IP address of a machine belonging to the network being protected by the firewall.

- **Source routing attacks:** where source specifies the route that a packet should take to bypass security measures, should discard all source routed packets

- **Tiny fragment attacks:** intruder uses the IP fragmentation option to create extremely small fragments and force the TCP header information into fewer separate fragments to circumvent filtering rules needing full header info; can enforce minimum fragment size to include full header.

Stateful Inspection Firewalls

- Stateful firewalls (also called circuit firewalls) examine the headers of each packet with regards to their placement within the packet series belonging to a specific session/connection and maintain records of all connections passing through the firewall.

- **Attacks and actions prevented with Stateful Inspection Firewalls**
- **Session Hijacking Attack**: Stateful firewalls can remember the sequence numbers expected on both sides as part of a TCP session and can block attempts to hijack the session, when an intruder sends several TCP segments with different sequence numbers (trial-and-error).

- **SYN Flood Attack**: Stateful firewalls can remember the number of connection requests that have been let through for an IP address/TCP port and block connection requests beyond a threshold.
  - Also, do not let more than a certain number of simultaneous TCP connections to originate per (source) IP address.

- **Bandwidth Exhaustion**: Do not let more than a specific amount of data to be transferred per day from the inside network to any outside IP address.

- **ICMP-based Attacks:** The volume of ICMP packets (like Echo Reply/Request, Destination not reachable, etc) that are transmitted in and out of the networks should be within a threshold.

Application Firewall

- The packet filter and stateful firewall look at only the packet headers. The application proxy firewall scans through the entire packet (including the application data) and makes sure if it could be forwarded in/out.

- An application firewall protecting an internal network of clients from being attacked by an external server/user is called a <u>Proxy Firewall</u>.
  - Example: An application firewall that protects an internal network of desktop/ office machines from users attempting to connect after office hours.

- An application firewall protecting an internal network of servers from being attacked by an external client is called a <u>Reverse Proxy Firewall</u>
  - Example: A reverse proxy firewall hosted to protect a sales network (comprising of various servers – database server, file server, etc) monitors every incoming packet to make sure it does not have any malicious scripts to cause any command injection attacks (XSS, XSRF or SQL-injection) or buffer overflow attacks.

Personal Firewalls

- <u>Motivation:</u> Home users, individual workers, and small businesses use cable modems or DSL connections with unlimited, always-on access.

- These people need a firewall, but a separate firewall computer to protect a single workstation can seem too complex and expensive.

- A workstation could be vulnerable to malicious code or malicious active agents (ActiveX controls or Java applets), leakage of personal data stored in the workstation, and vulnerability scans (like nmap) to identify potential weaknesses.

- A personal firewall is an application program that runs on a workstation to screen traffic on the workstation and block unwanted traffic leaving or entering the workstation to the network to which it is connected.

- A user could configure the personal firewall to accept traffic only from certain sites, and not from specific sites, and to generate logs of activities happened in the past

- A personal firewall could be also configured with a virus scanner which would be then automatically invoked to scan any incoming data to the workstation.

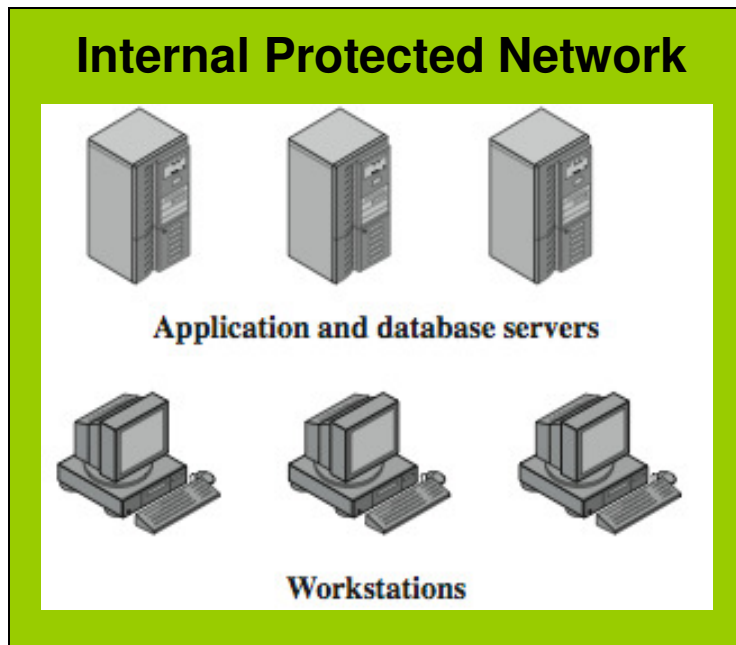# What Firewalls Can and Cannot Block

- <u>Firewalls cannot alone secure an environment.</u>
- A firewall can control only what comes into the network or what goes out of the network.
- Firewalls cannot protect from internal threats (through disgruntled employees).
- Firewalls cannot protect against malware imported via laptop, PDA, or portable storage device infected outside the network, then attached and used internally.
- A firewall is often a single point of failure for a network.
  - A more layered approach like a screening router (packet filter), followed by a stateful firewall, a proxy firewall, followed by a personal firewall may be more helpful. The number of packets processed by a firewall in the layered approach would decrease as we progress from the packet filter to the personal firewall.
- Firewalls must be frequently configured and updated to take into account the changes in the internal and external environment and based on the review of the firewall activity reports that may indicate intrusion attempts.

# Demilitarized Zone (DMZ) Networks

- A DMZ network (also called perimeter network) is a subnet that contains an organization's services that are exposed to a larger untrusted network (like the Internet).

- In other words, the DMZ comprises of hosts that provide services to users outside the internal LANs, such as e-mail, web, DNS servers.
- Because of the higher chances of these hosts being compromised, they are placed into their own sub-network in order to protect the rest of the network if an intruder were to succeed in attacking them.

- Thus, a DMZ network adds an additional layer of security to an organization's LAN – an external attacker only has access to the hosts in the DMZ and not to any other internal networks.

- Hosts in the DMZ provide services to both the internal and external networks – an external ("front-end") firewall monitors the traffic between the DMZ network and the external Internet; while, an internal ("back-end") firewall monitors the traffic between the DMZ hosts and the internal network clients.

DMZ Networks

Servers that are more accessed from inside and/or store much of confidential information

Internal Protected Network

Application and database servers

Workstations

Front-end Firewall

Internet

Boundary router

Servers that are more accessed from outside and/or do not store much of confidential information

LAN Switch

LAN Switch

Back-end Firewall

Email server

Web server(s)

DNS server

DMZ Network

Intrusion Detection Systems (IDS)

- An IDS to the networking world is like a burglar alarm to the physical world.
- The main purpose of an IDS is to identify suspicious or malicious activity, note activities that deviate from normal behavior, catalog and classify the activity, and, if possible, respond to the activity.
- Typically run in promiscuous mode.

- <u>False Positives</u>: An IDS classifies an activity/traffic as malicious, when it is actually not.
- <u>False Negatives</u>: An IDS classifies an activity/traffic as not malicious, when it is actually malicious.

Logical Depiction of IDS Components

Source: Figure 13.2 from Conklin and White – Principles of Computer Security, 2nd Edition

Categories of IDS

- **Based on Signatures/Rule sets**
  - <u>Signature-based IDS:</u> Relies on pre-defined attack and traffic patterns (called signatures); Cannot detect new attacks whose pattern does not match with the existing signatures
    - Likely to have more false negatives than false positives
  - <u>Anomaly-based IDS:</u> Starts with a minimal set of rules; anything that deviates from the rules is classified as suspicious; the network admin decides how to appropriately classify the suspicious activity and updates the rule set
    - Likely to have more false positives than false negatives.
- **Based on Placement and Coverage**
  - <u>Network-based IDS (NIDS):</u> Placed next to the firewall; looks at all the incoming and outgoing traffic; a few NIDSs are sufficient
    - Cannot handle encrypted traffic and detect attacks that originate within the hosts in the network
  - <u>Host-based IDS (HIDS):</u> One HIDS per host; could detect attacks targeted at specific hosts (through changes in the log files/critical files at the host); more like a personal firewall/anti-virus scanner

Intrusion Prevention Systems (IPS)

- While an IDS is typically passive (just analyzing the traffic and reporting the attack), an IPS will react to an attack traffic.

- The reactions of an IPS could range from something simple, such as sending a TCP reset message to interrupt a potential attack and disconnect a session, to something complex, such as dynamically modifying firewall rules to reject all traffic from specific source IP addresses for the next few hours or days.

- When configured with the private keys of the servers in the internal network, IPSs would be able to decrypt the SSH connection establishment messages between a client and server and extract the session keys that would be used during the complete session. This gives an added advantage for the IDS/IPS to handle encrypted traffic.

Honeypot

- A honeypot is a trap to detect, deflect or in some manner counteract attempts at unauthorized use of information systems.
- A honeypot is usually a computer, and sometimes data or an unused IP address space that appears to be part of a network but which is actually isolated, unprotected and monitored, and which seems to contain information or a resource that would be of value to attackers.
- Honeypots have no production value and hence should not see any legitimate traffic or activity. Whatever they capture can be surmised as malicious or unauthorized.
- A honeynet is a network of honeypots. A honeynet is used for monitoring a larger and/or more diverse network in which one honeypot may not be sufficient.
- A honeypot/ honeynet is more of a preventative approach of detecting potential attackers existing in the Internet who may target the organization network in the near future.
- Honeypots could be used to fake as open relays to attract spam emails and determine the source e-mail address and destination e-mail addresses used by the spammers.
  - An open relay is an e-mail server that allows anyone on the Internet to send email through it.
  - Once they find an open relay, spammers keep sending the spam email to the open relay and expect it to spread the spam.
- Note that no ordinary e-mail will come to a honeypot. All it receives could be categorized as spam.