Introduction to Computer Networks

CS640

Network Security

https://pages.cs.wisc.edu/~mgliu/CS640/F22/

Ming Liu
mgliu@cs.wisc.edu

Today

Last lecture

- What are the learned lessons on building network applications?

Today

- What is network security?
- How do networking attacks happen?
- How does the networking defense work?

Announcements

- Lab5 is due 12/14/2022, 11:59 PM
- Lab6 is due 12/19/2022, 11:59 PM
- Final exam: Dec 17, 2022 5:05 PM – 7:05 PM @Engineering Hall 1800

Problem

Computer networks are typically a shared resource, used by many applications with different interests

The Internet is particularly widely shared, being used by competing businesses, mutually antagonistic governments, and opportunistic criminals

Problem

Computer networks are typically a shared resource, used by many applications with different interests

> Unless security measures are taken, a network conversation or a distributed application may be compromised by an adversary

Th… ed by competing businesses, mutually antagonistic governments, and opportunistic criminals

Q: What is network security?

An Example

Suppose you are a customer using a credit card to order an item from a website

Network Security — Confidentiality

Suppose you are a customer using a credit card to order an item from a website

Network Security — Confidentiality

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary would eavesdrop on your network communication, reading your messages to obtain your credit card information

Network Security — Confidentiality

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary would eavesdrop on your network communication, reading your messages to obtain your credit card information
- What to do I: Encrypt messages so as to prevent an adversary from understanding the message content. A protocol that does so is said to provide **confidentiality**

Network Security — Confidentiality

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary would eavesdrop on your network communication, reading your messages to obtain your credit card information
- What to do I: Encrypt messages so as to prevent an adversary from understanding the message content. A protocol that does so is said to provide **confidentiality**
- What to do II: Solution I + conceal the quantity or destination of communication, which is called **traffic confidentiality**

Network Security — Data Integrity

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary who can't read the contents of your encrypted message might still be able to change a few bits in it, resulting in a valid order for a completely different item or perhaps 1000 units of the item

Network Security — Data Integrity

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary who can't read the contents of your encrypted message might still be able to change a few bits in it, resulting in a valid order for a completely different item or perhaps 1000 units of the item
- What to do: A protocol that detects such message tampering provides **data integrity**

Network Security — Originality

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary can alternatively transmit an extra copy of your message in a reply attack. It appears as though you had simply ordered another of the same data item you ordered the first time

Network Security — Originality

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary can alternatively transmit an extra copy of your message in a reply attack. It appears as though you had simply ordered another of the same data item you ordered the first time
- What to do: A protocol that detects replays provides **originality**

Network Security — Timeliness

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary can intercept your order, wait a while, and then transmit it – delay your order. Or an adversary could arrange for the item to arrive on your doorstep while you are away on vacation, when it can be easily snatched

Network Security — Timeliness

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary can intercept your order, wait a while, and then transmit it – delay your order. Or an adversary could arrange for the item to arrive on your doorstep while you are away on vacation, when it can be easily snatched
- What to do: A protocol that detects delaying tactics provides **timeliness**

Network Security — Authentication

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary can direct customer requests to a false website
  - URL -> incorrect IP address

Network Security — Authentication

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary can direct customer requests to a false website
  - URL –> incorrect IP address
- What to do: A protocol that ensures that you really are talking to whom you think you're talking provides **authentication**

Network Security — Access Control

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary can remotely access or even modify the website without authentication

Network Security — Access Control

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary can remotely access or even modify the website without authentication
- What to do: A protocol that enforces the rules regarding who is allowed to do what provides **access control**

Network Security — Availability

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary can overload the web server so that customers are unable to access the website

Network Security — Availability

Suppose you are a customer using a credit card to order an item from a website

- Threat: An adversary can overload the web server so that customers are unable to access the website
- What to do: A protocol that ensures a degrees of access provides **availability**

Network Security More

The Internet has notably been used as a means for deploying malicious code that exploits vulnerabilities in end-system

- Worms: pieces of self-replication code that spread over networks
- Viruses, which are spread by the transmission of “infected” files
- Infected machines can be arranged into botnets which can be used to inflict further harm, such as launching DoS attacks

Q: What is network security?

A: A concern to any entities in the network (including client and server)

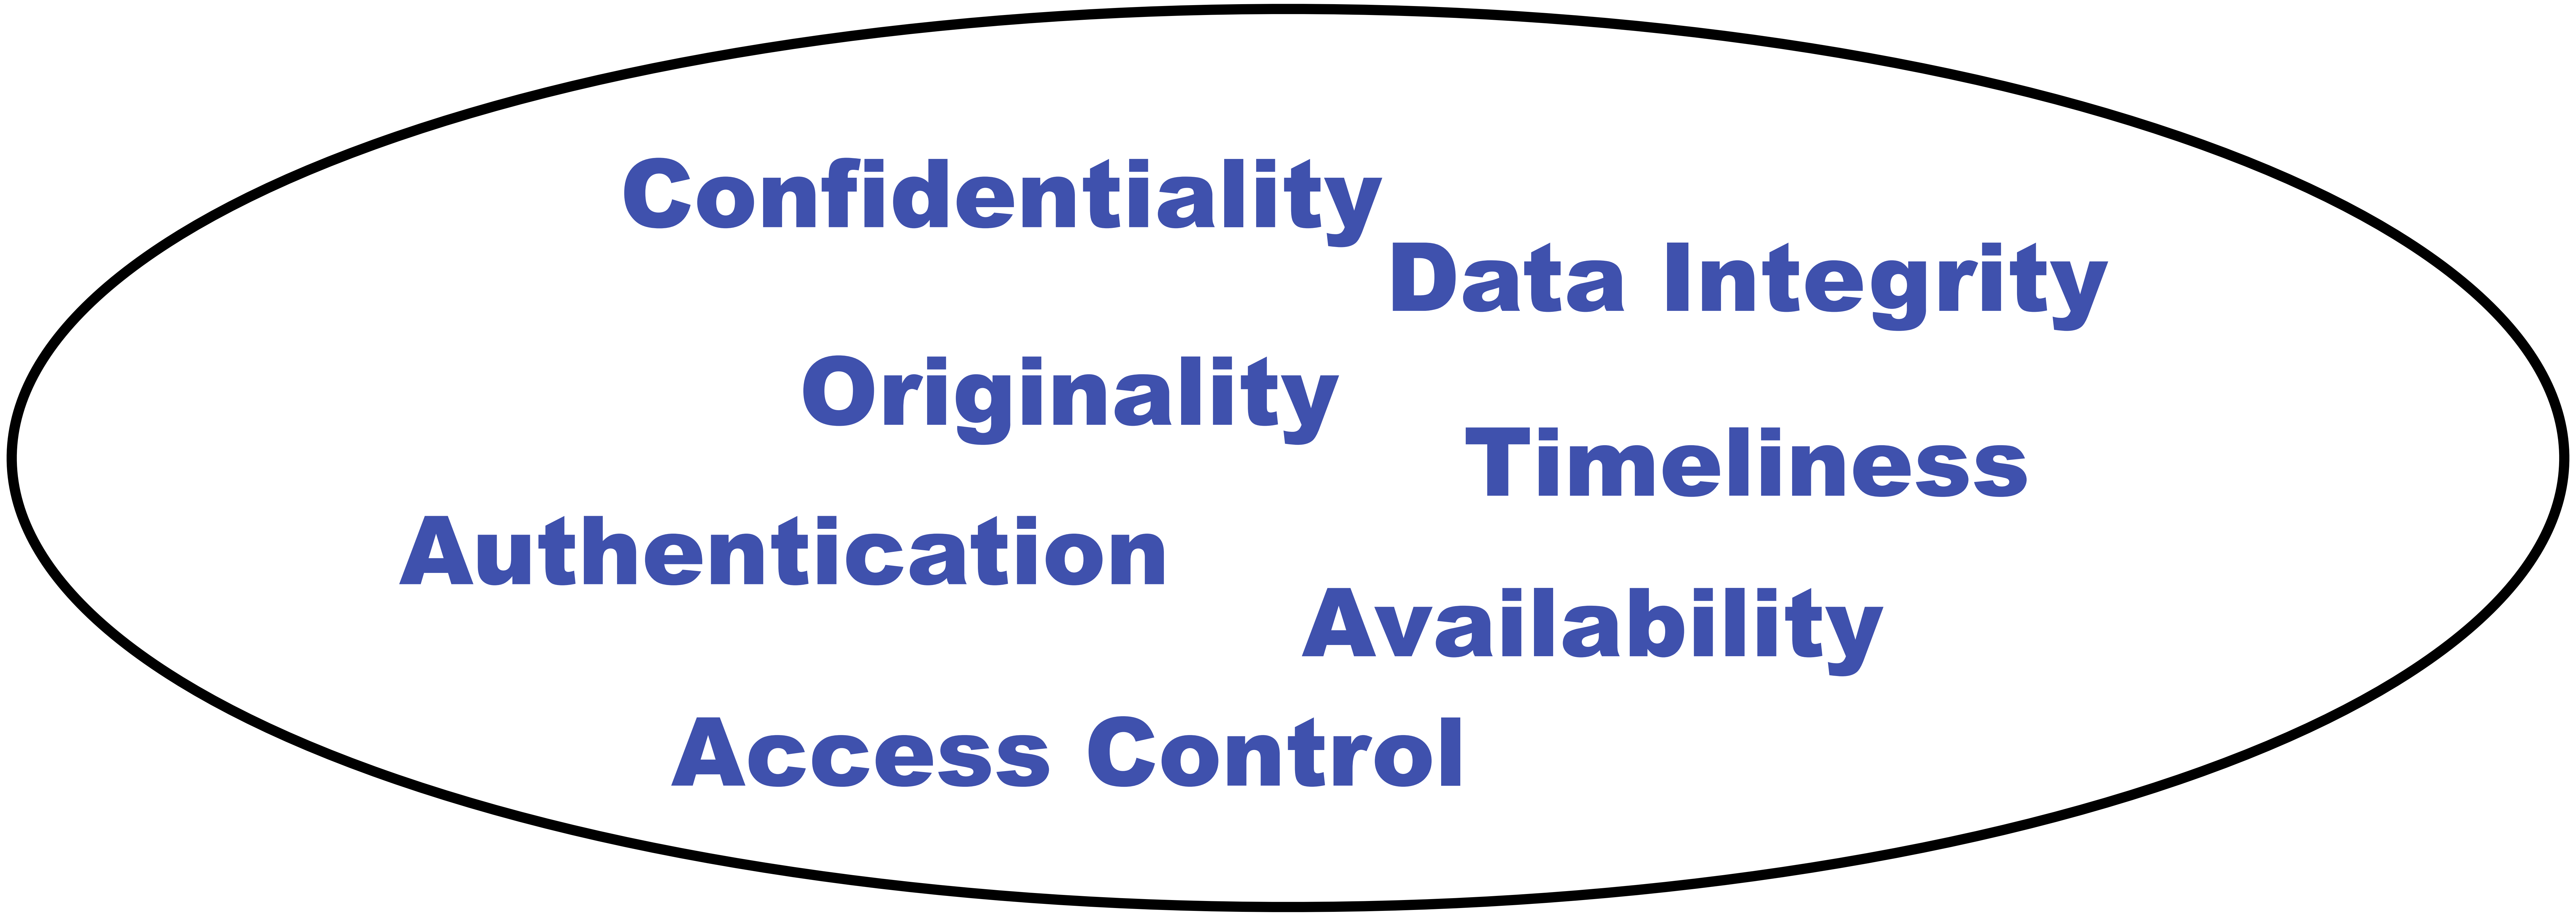

Q: How do networking attacks happen?

Network-related Attacks

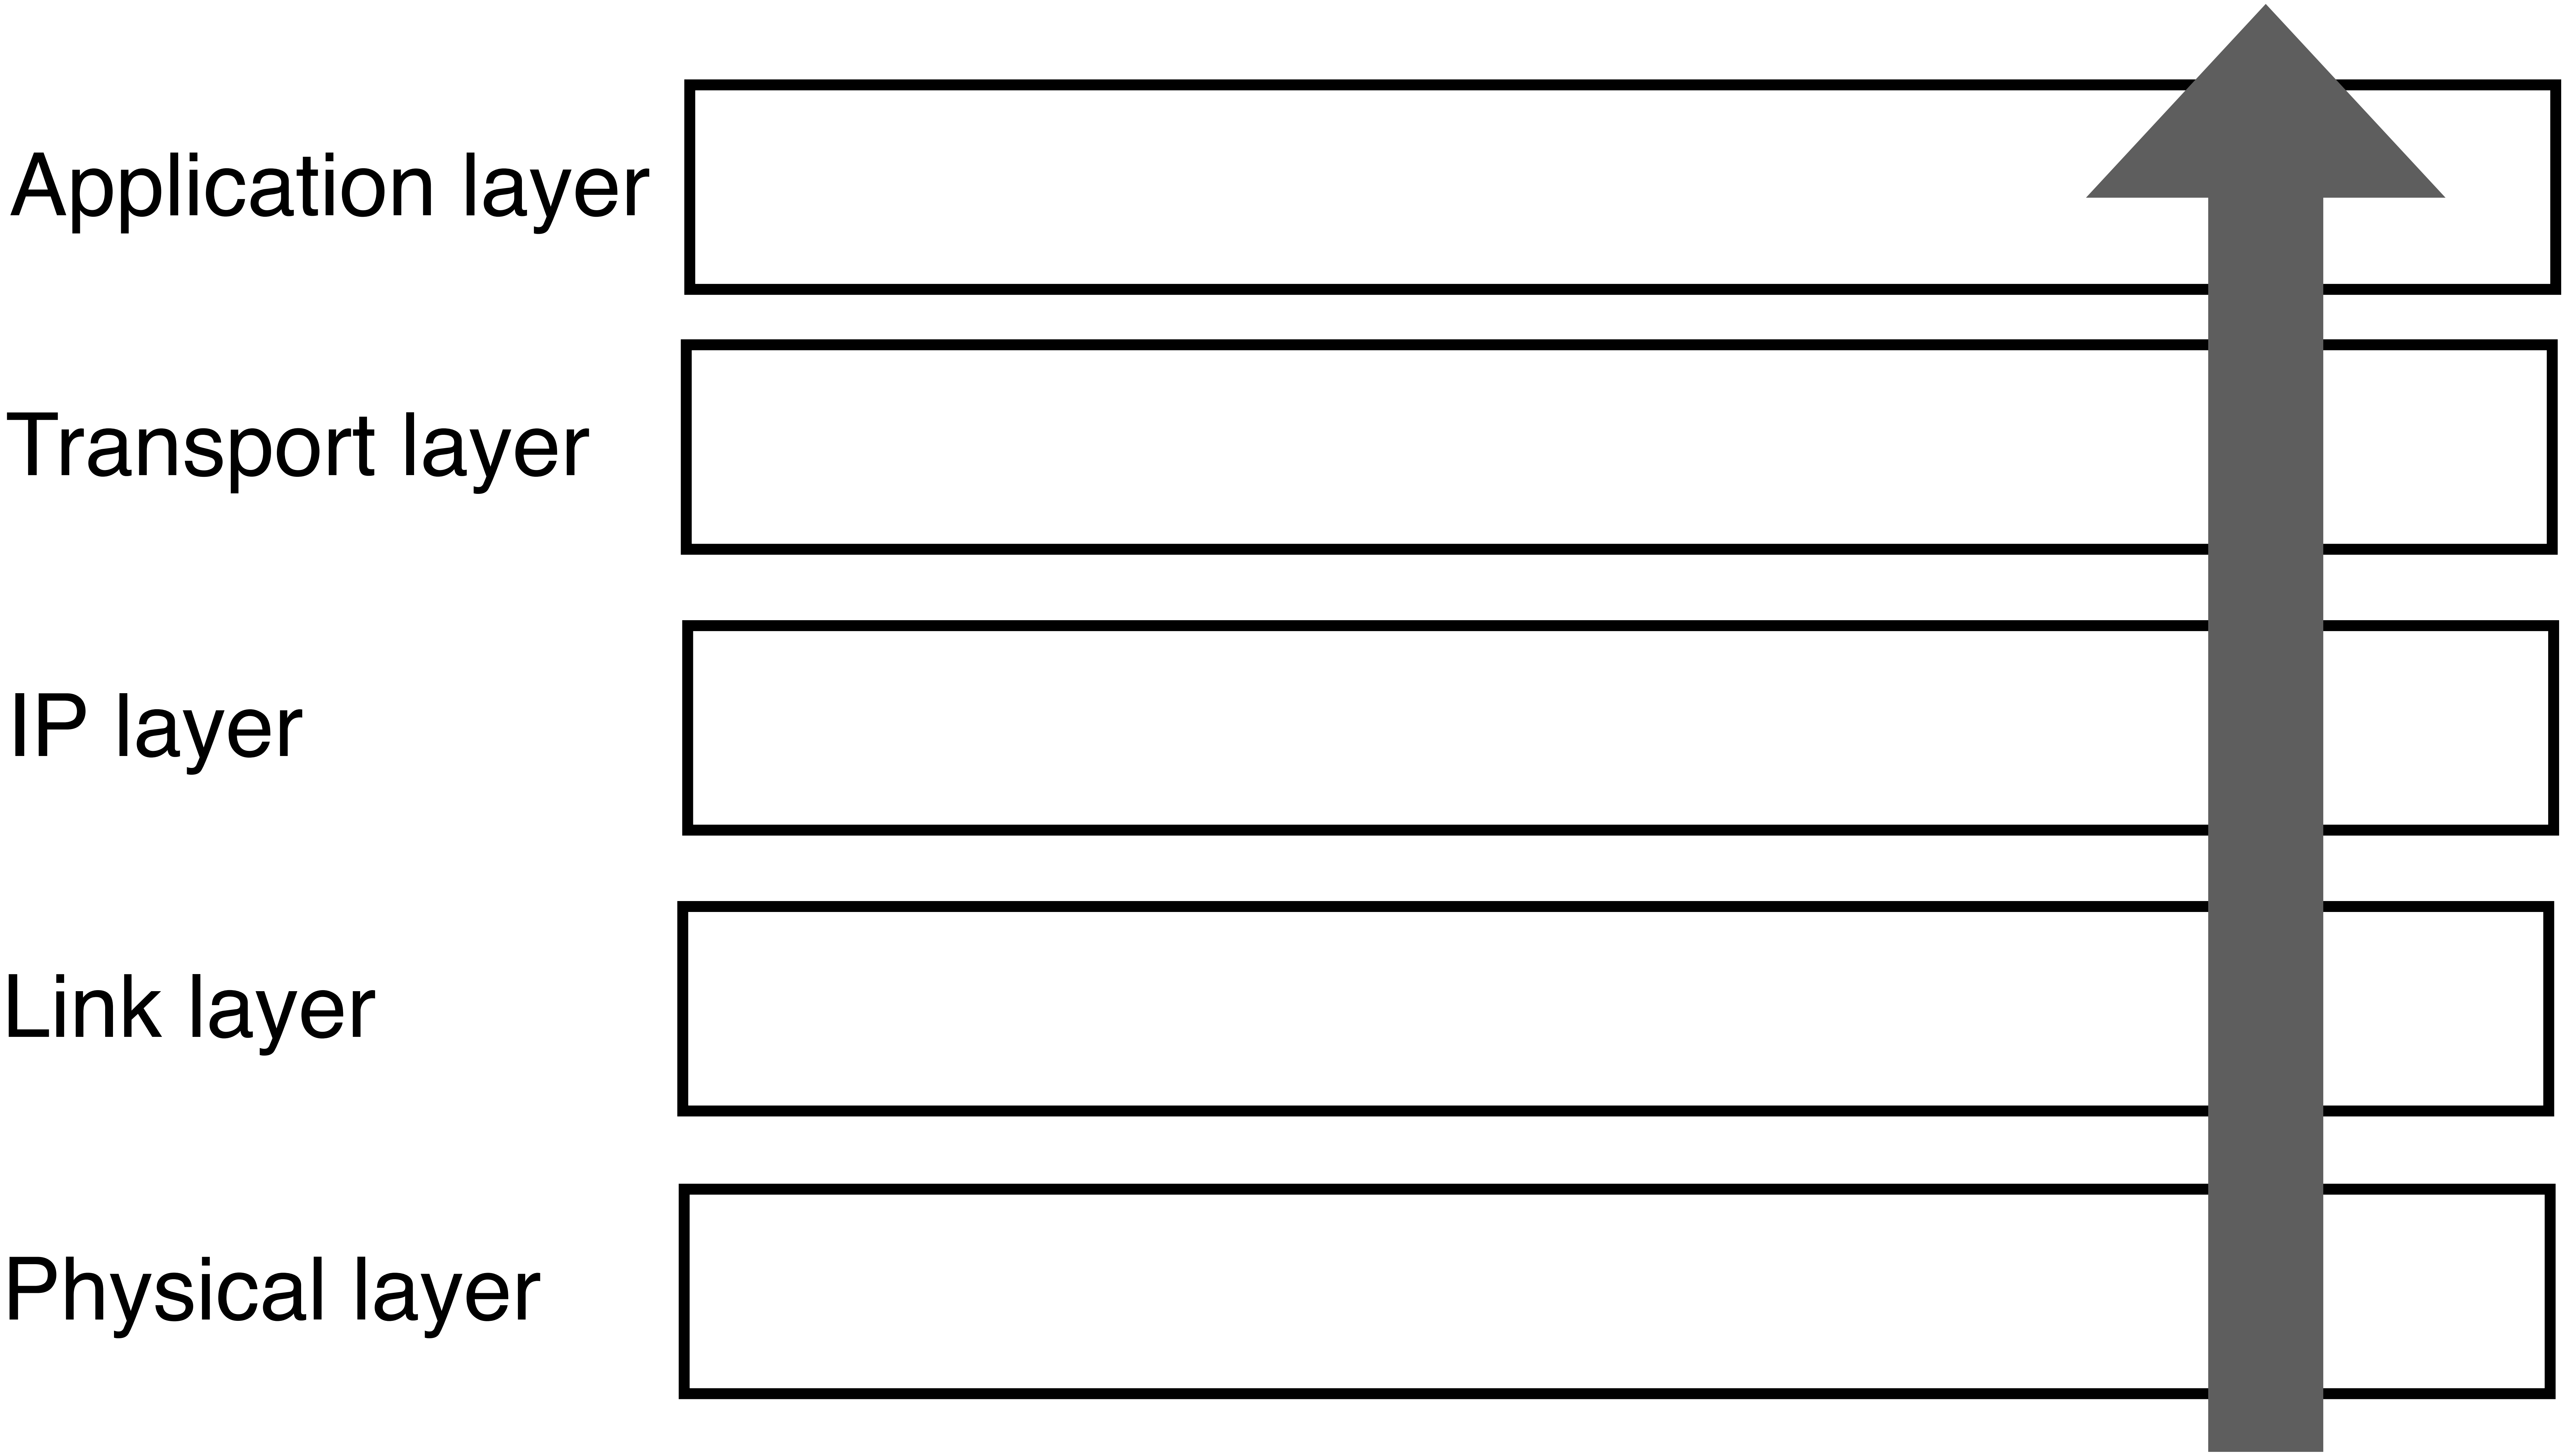

Physical Layer Attacks

If an attacker gains physical control of a device, he/ she can control a device’s behavior

- Wiretapping
- Hardare modification

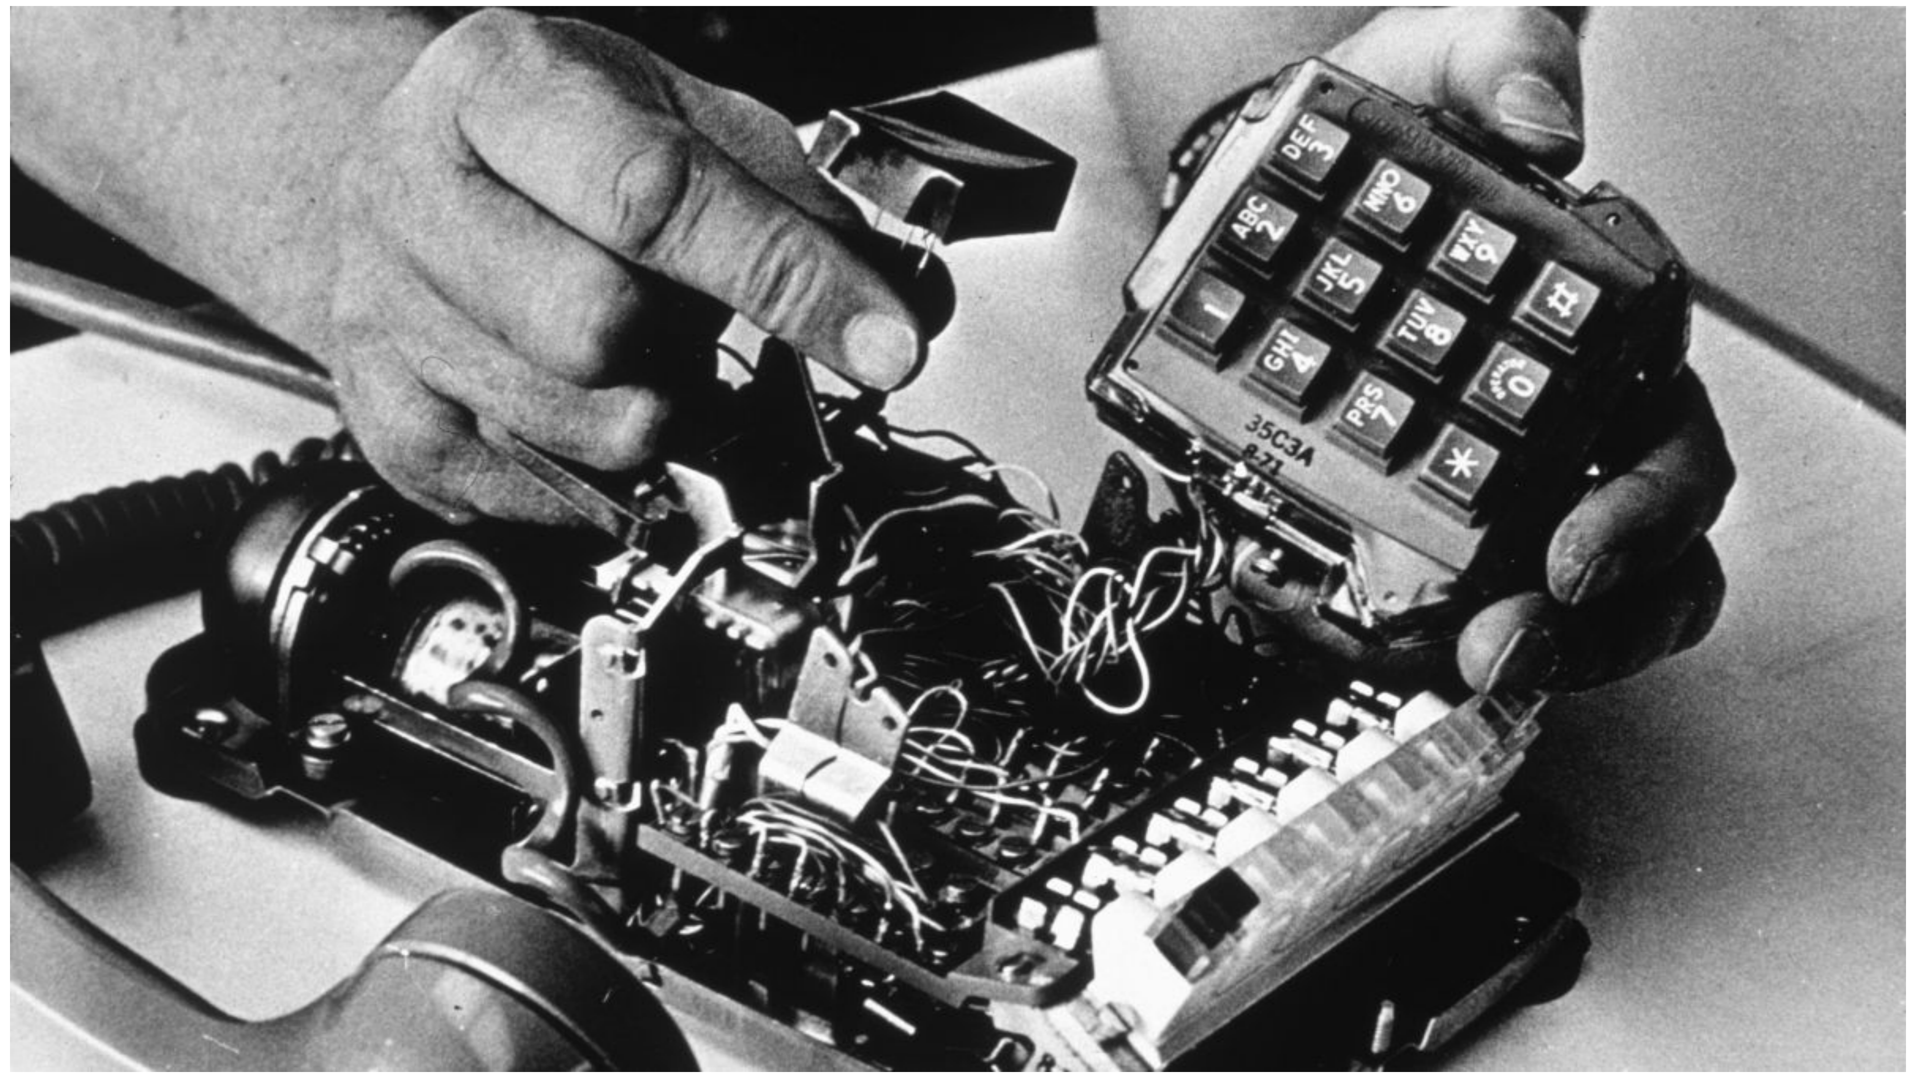

Link Layer Attacks — ARP Spoofing

ARP is a stateless protocol. When a host gets an ARP reply, no matter whether it sends an ARP request or not, it accepts the ARP entry and updates the cache

Link Layer Attacks — ARP Spoofing

ARP is a stateless protocol. When a host gets an ARP reply, no matter whether it sends an ARP request or not, it accepts the ARP entry and updates the cache

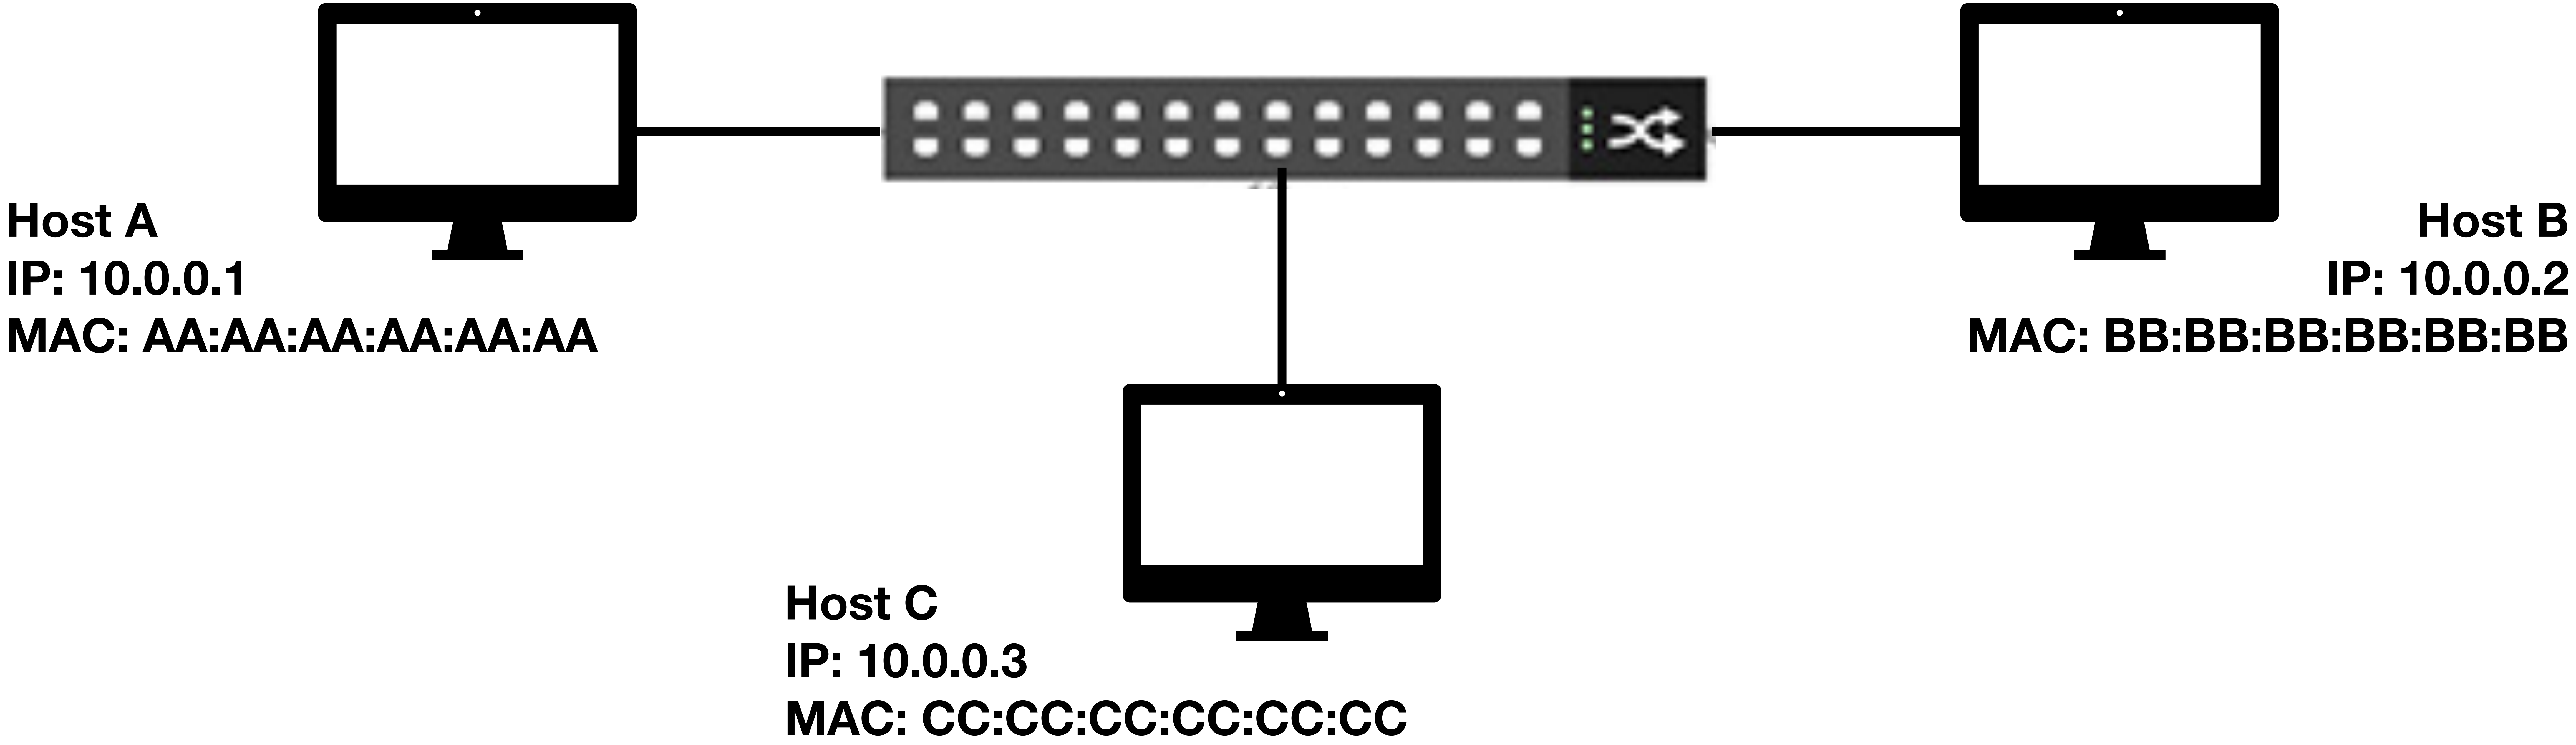

Link Layer Attacks — ARP Spoofing

ARP is a stateless protocol. When a host gets an ARP reply, no matter whether it sends an ARP request or not, it accepts the ARP entry and updates the cache

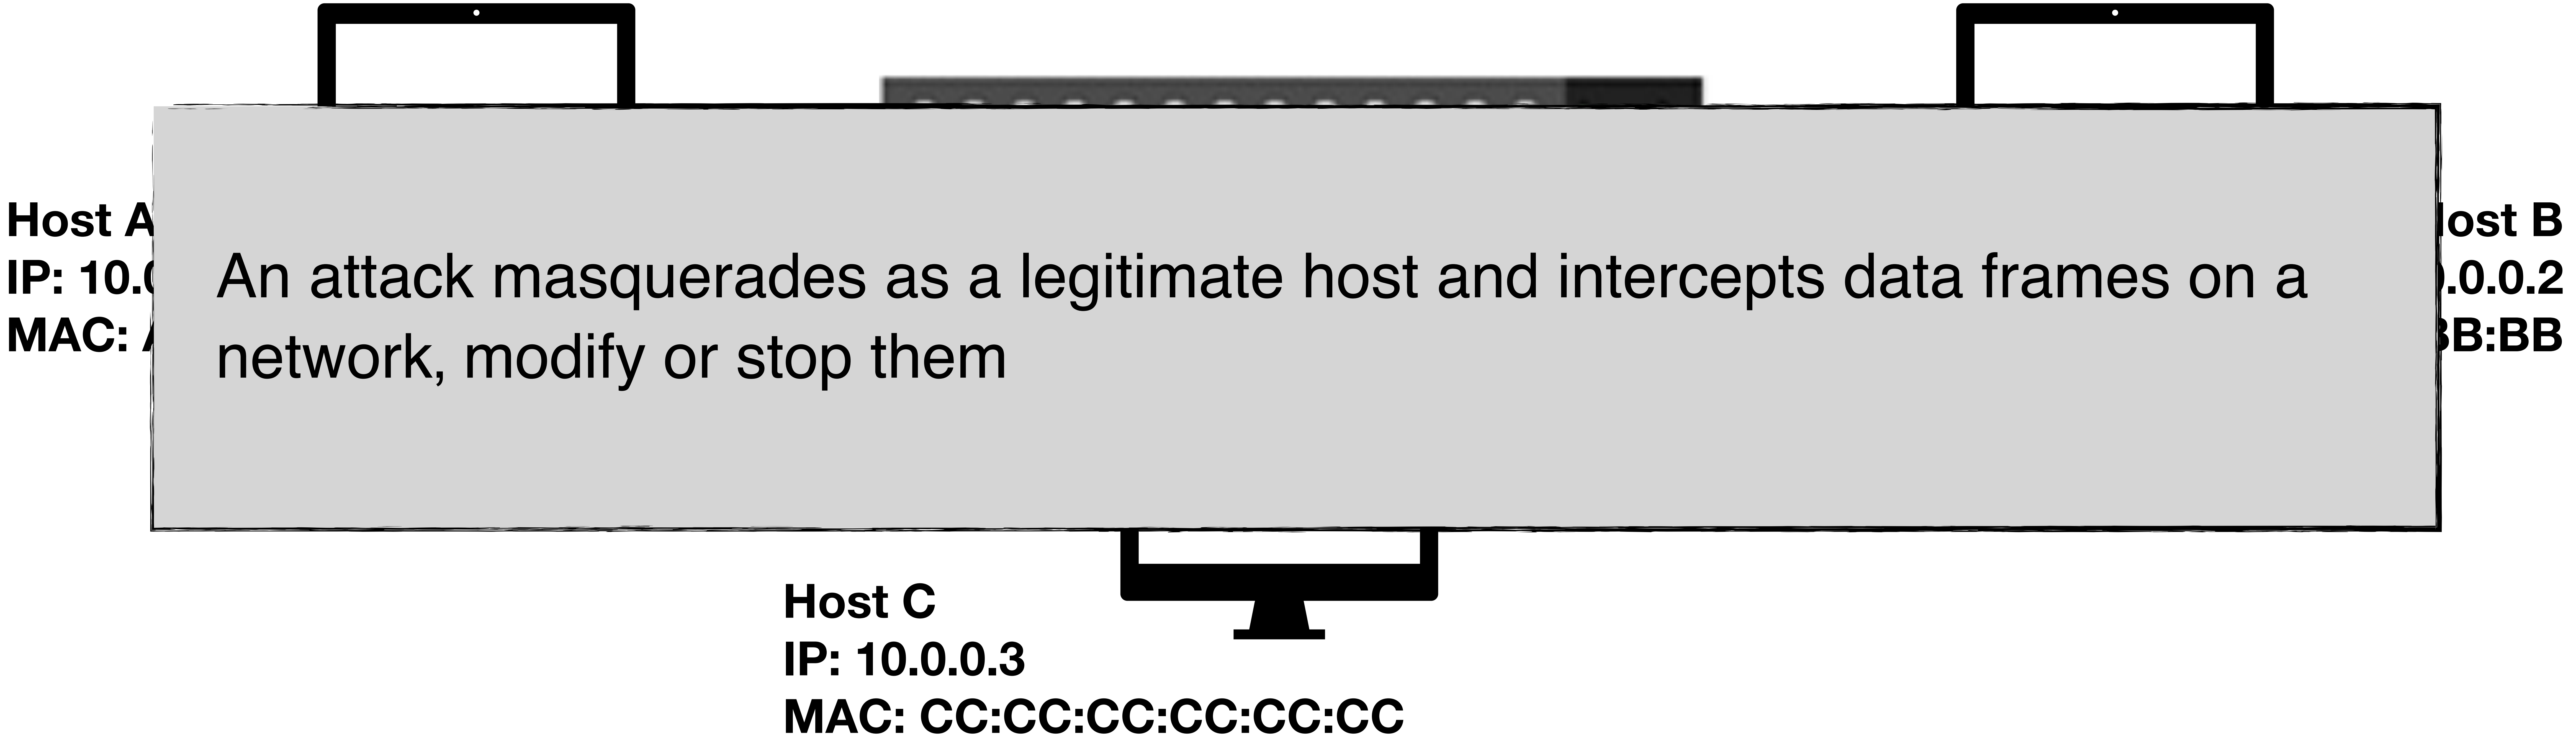

An attack masquerades as a legitimate host and intercepts data frames on a network, modify or stop them

Link Layer Attacks — MAC Flooding

An Ethernet switch has a fixed-size TCAM table to store MAC addresses, etc.

Link Layer Attacks — MAC Flooding

An Ethernet switch has a fixed-size TCAM table to store MAC addresses, etc.

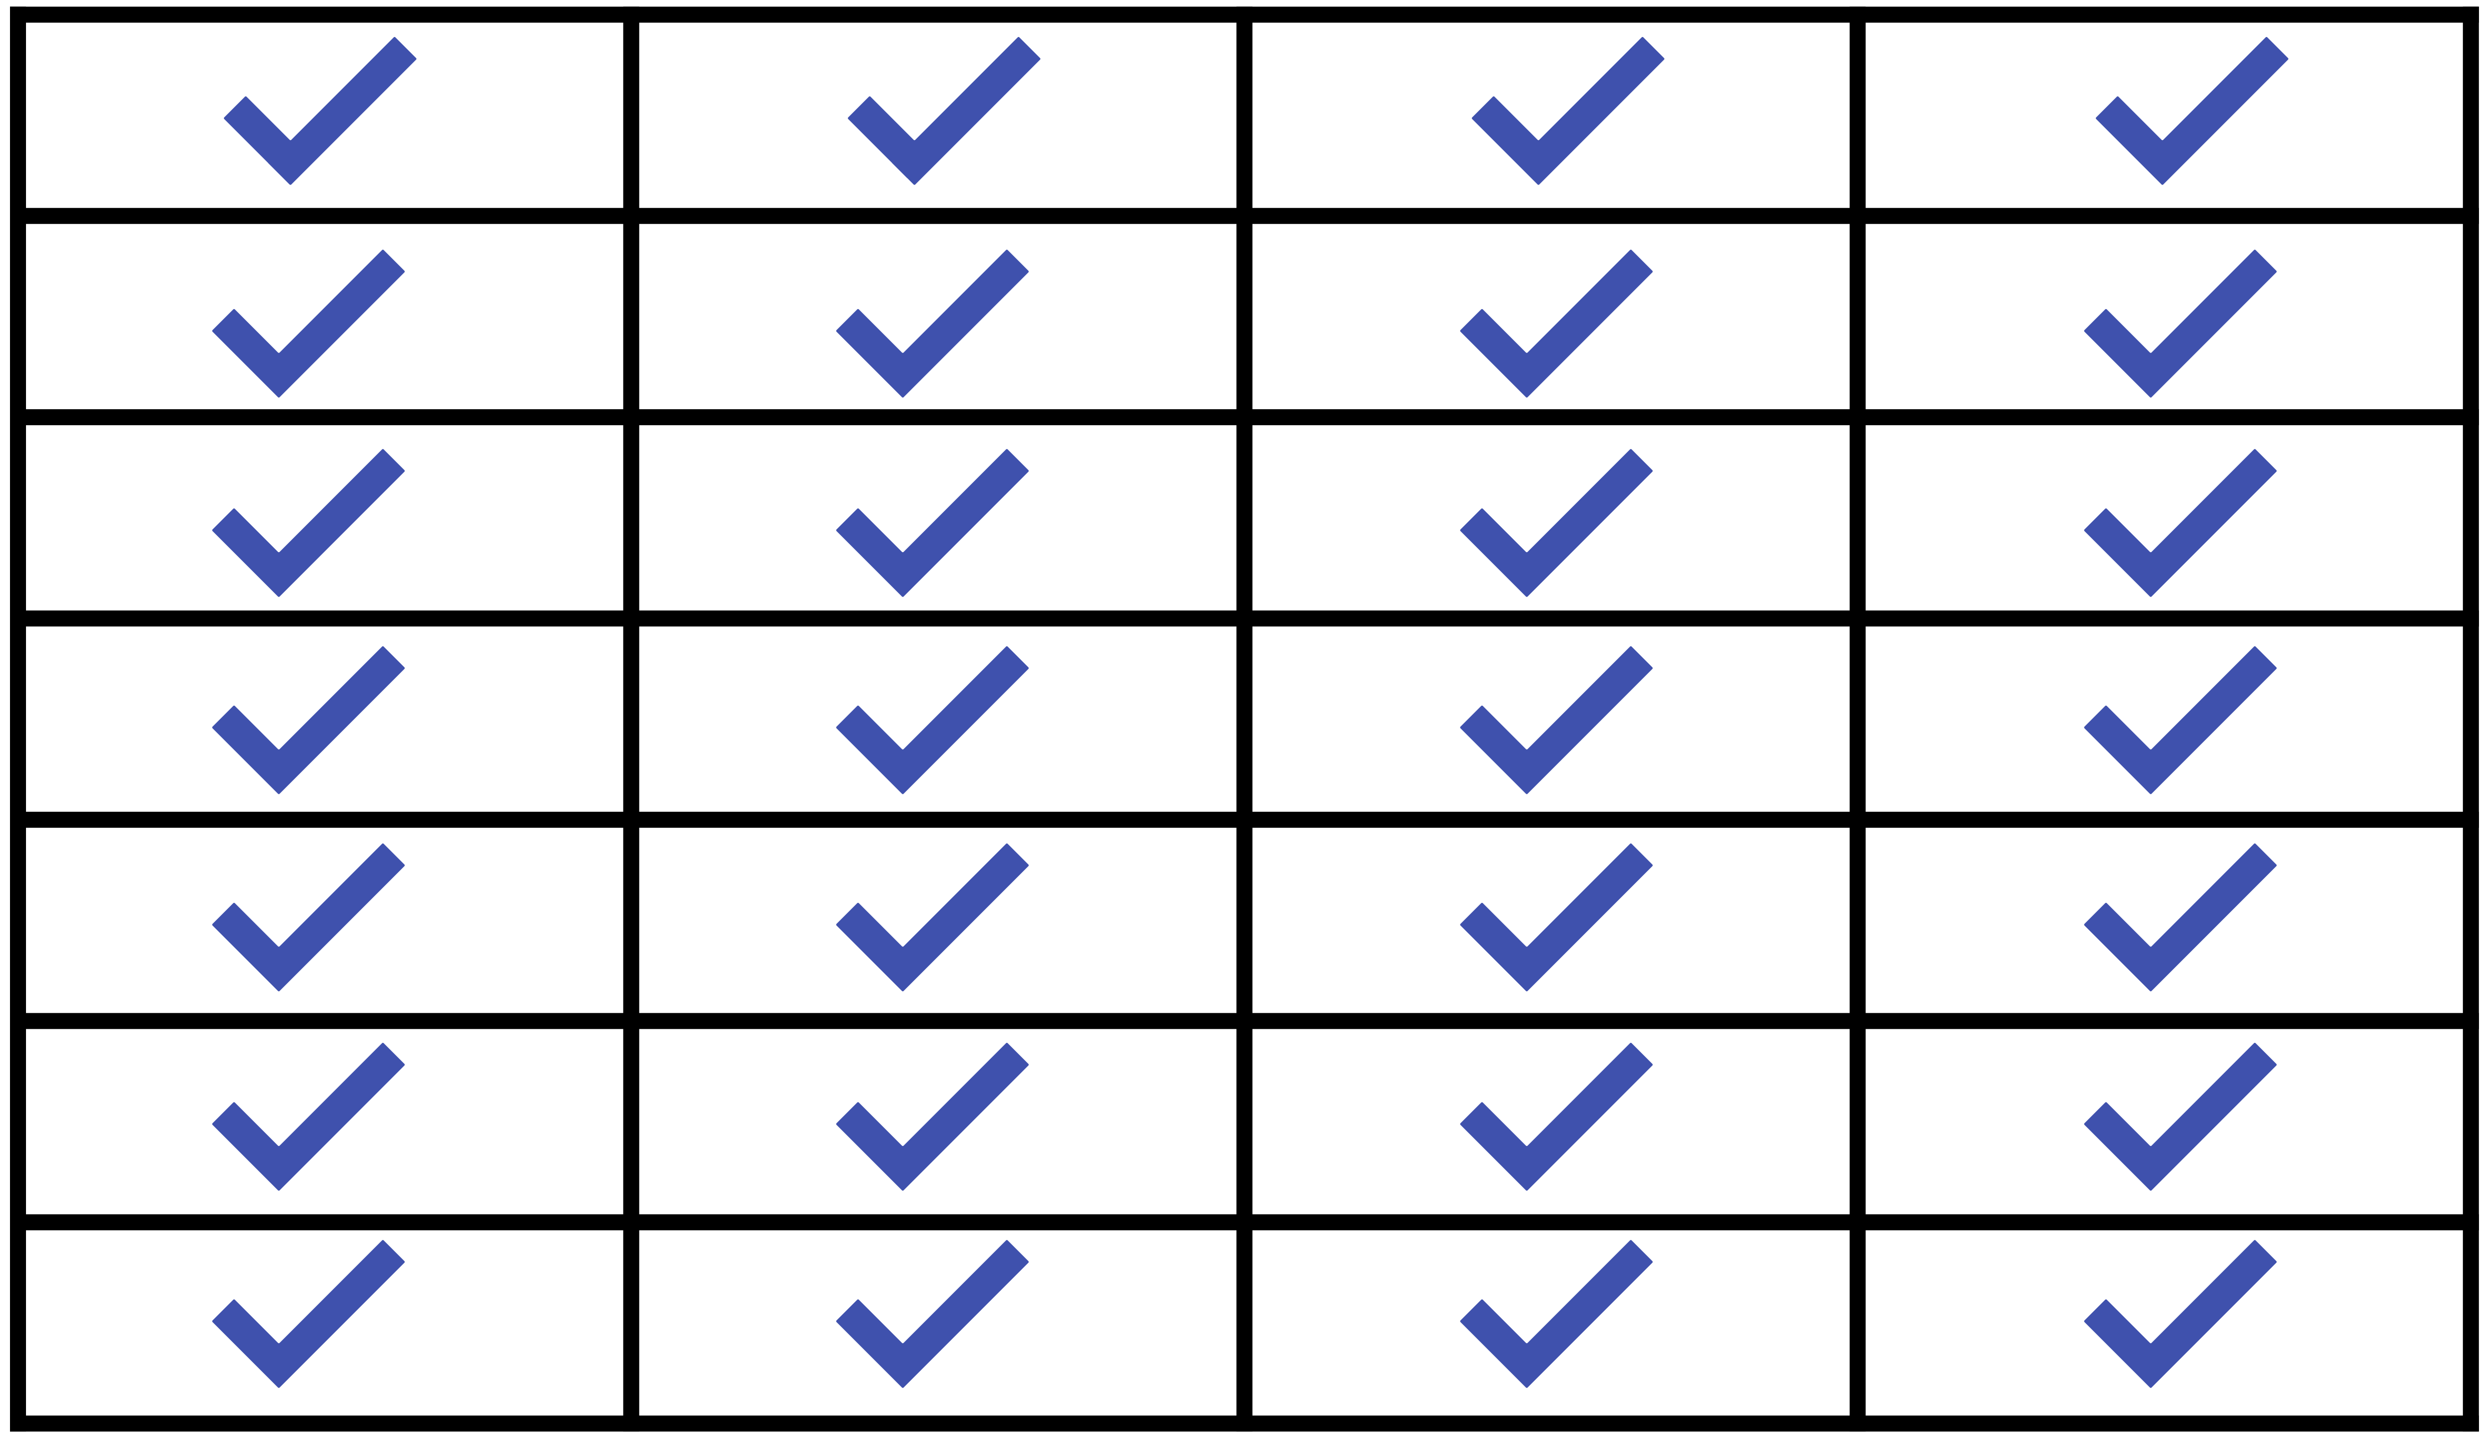

Link Layer Attacks — MAC Flooding

An Ethernet switch has a fixed-size TCAM table to store MAC addresses, etc.

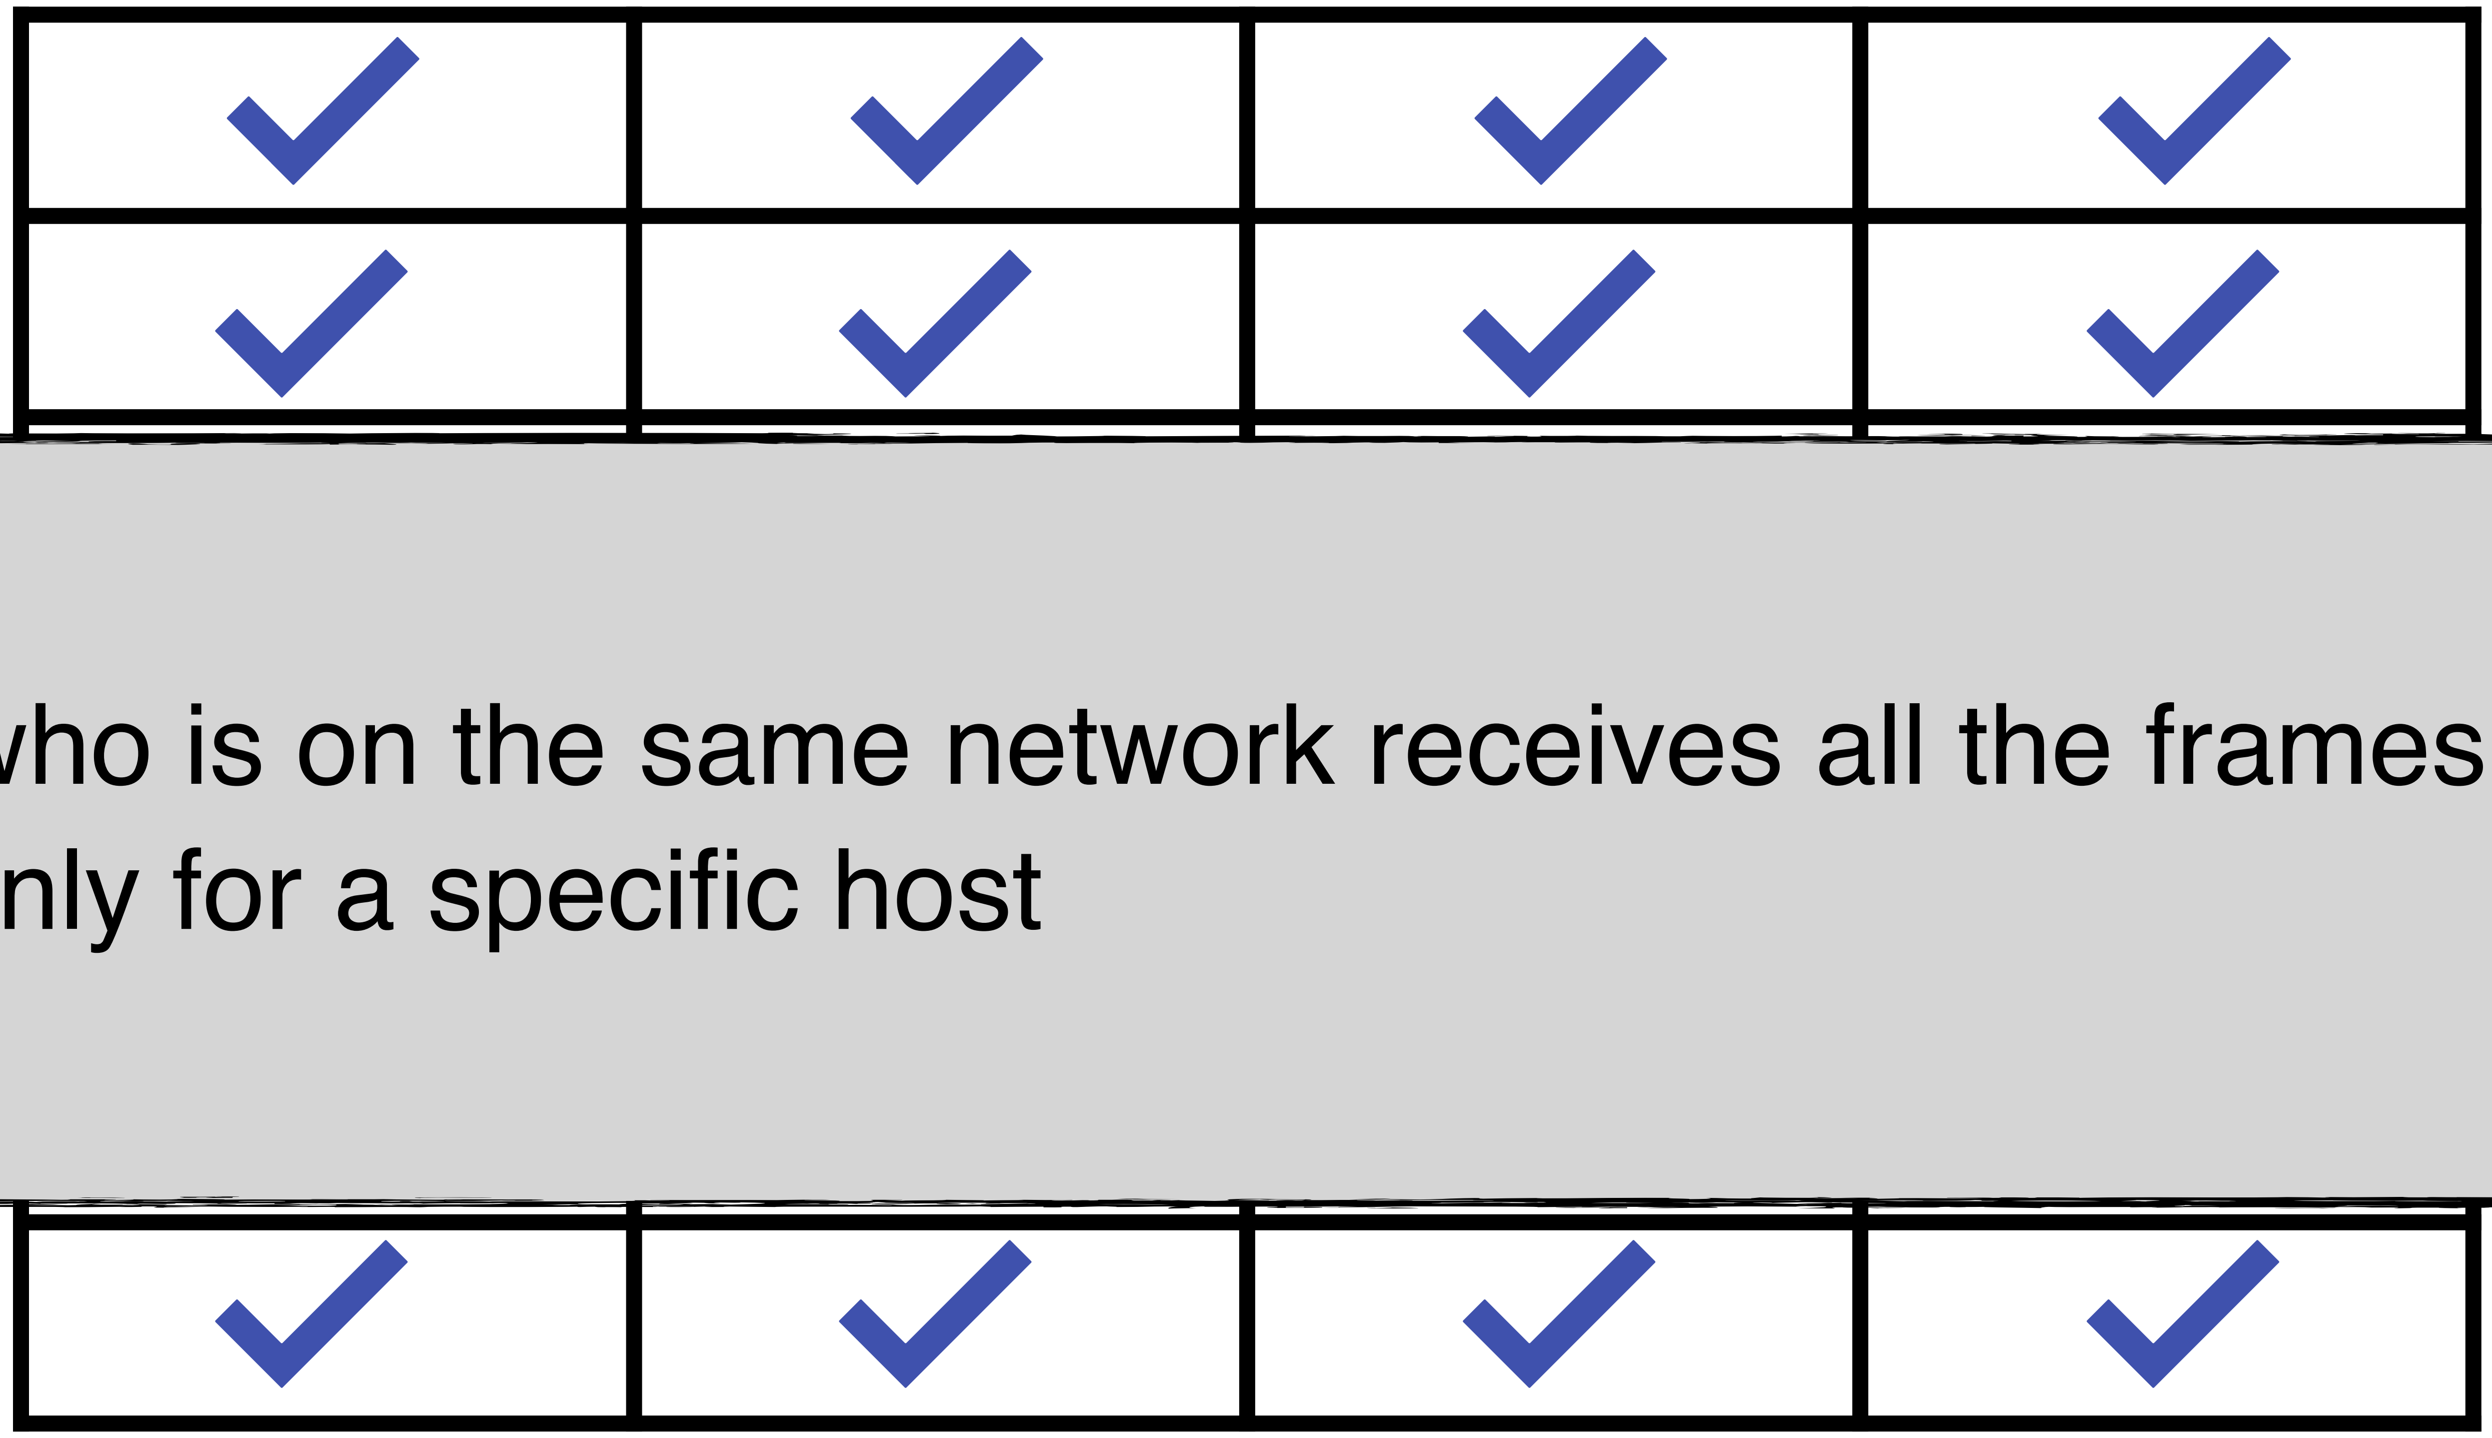

An attack who is on the same network receives all the frames which were designed only for a specific host

Link Layer Attacks — STP Attack

Spanning tree protocol (STP) defines a loop-free tree that spans all the switches in a network. It forces certain data links into a blocked state and keeps other links in a forwarding state.

Link Layer Attacks — STP Attack

Spanning tree protocol (STP) defines a loop-free tree that spans all the switches in a network. It forces certain data links into a blocked state and keeps other links in a forwarding state.

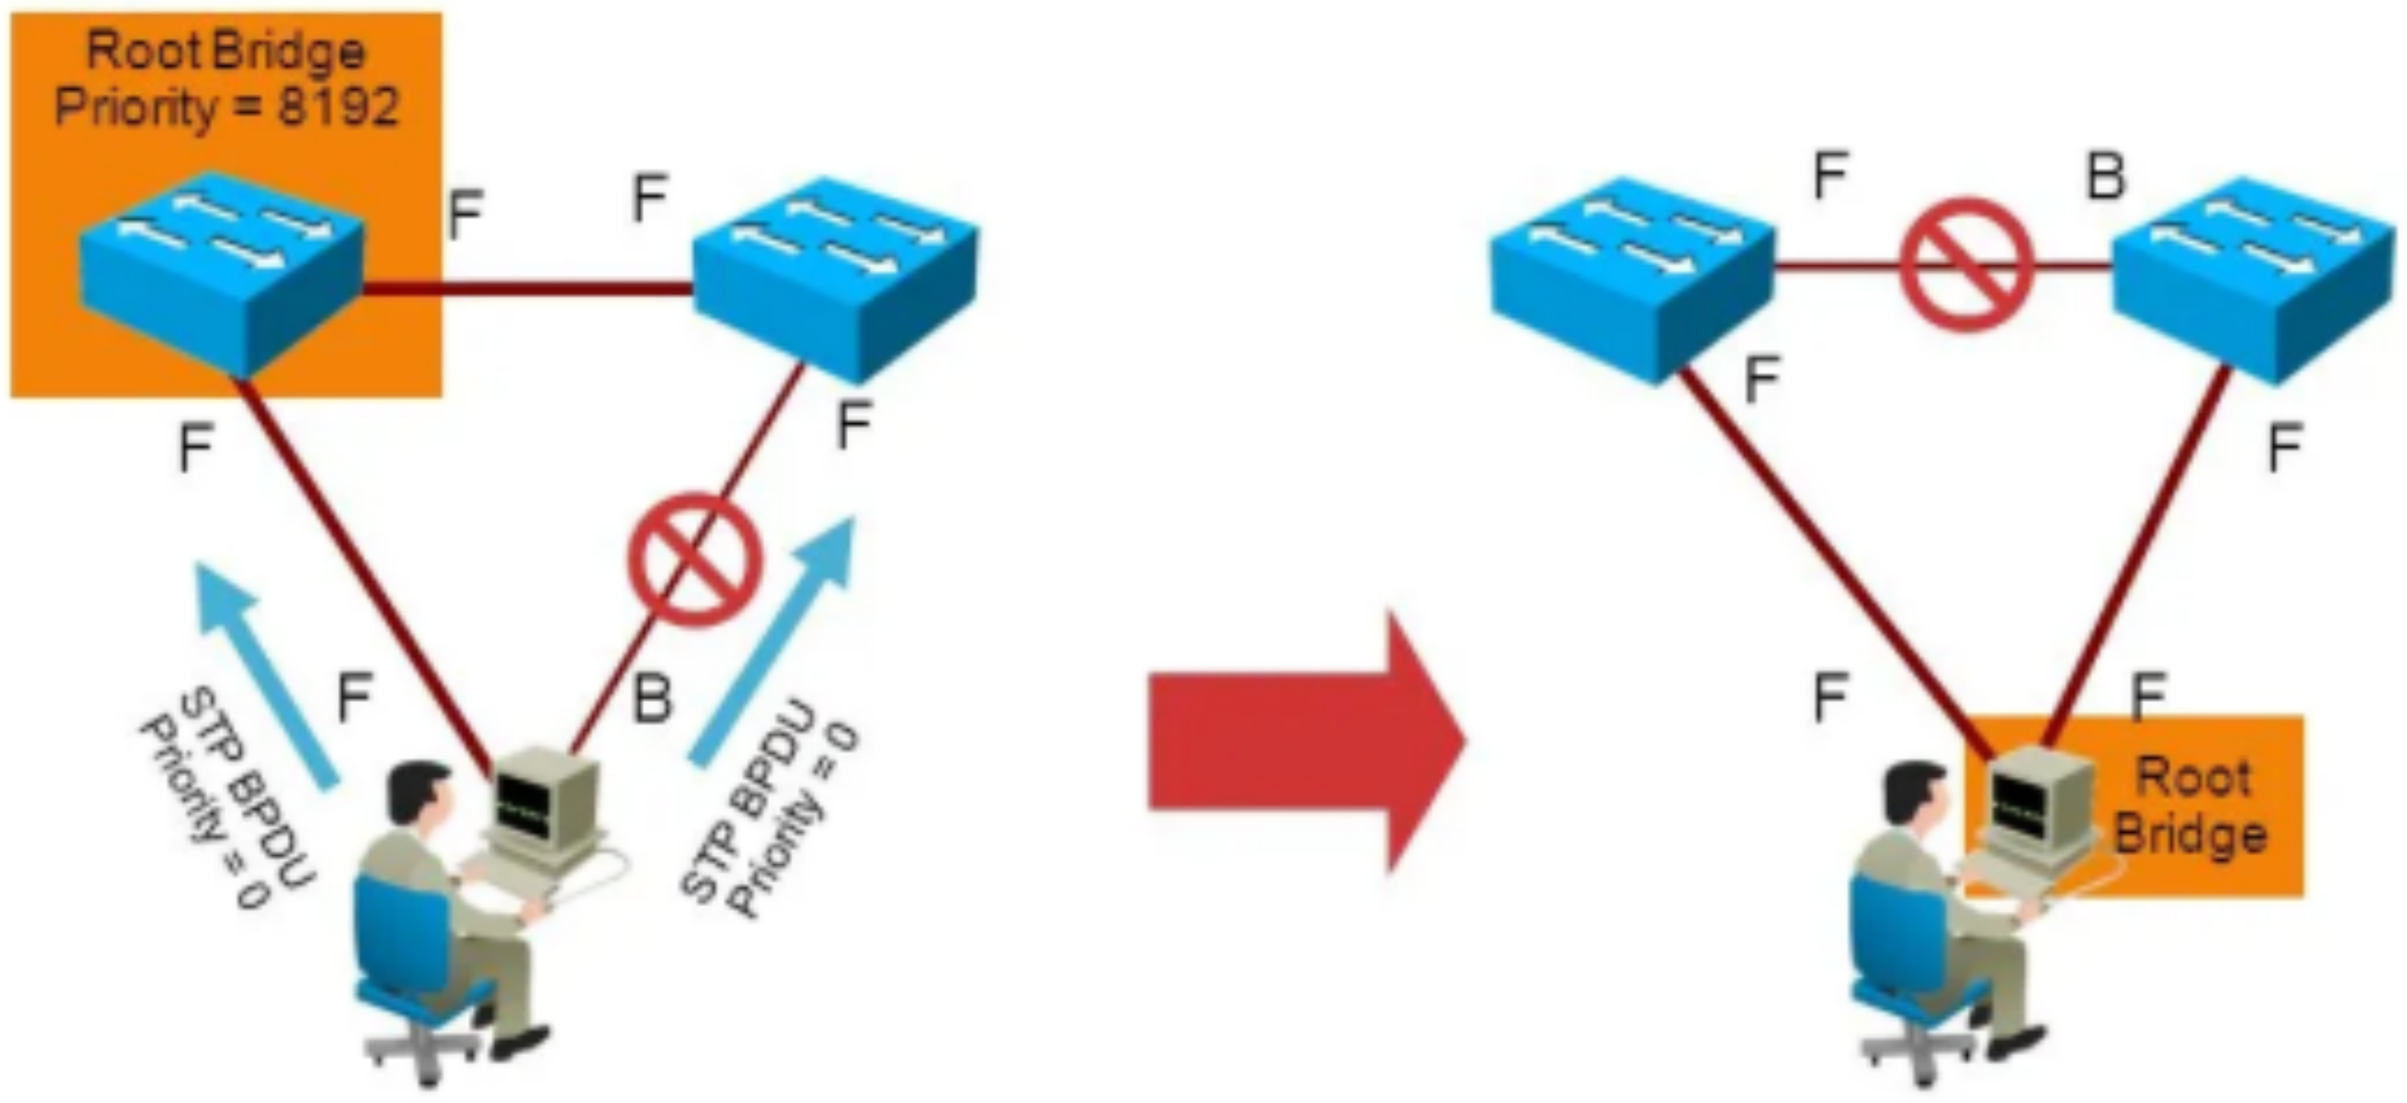

Link Layer Attacks — STP Attack

Spanning tree protocol (STP) defines a loop-free tree that spans all the switches in a network. It forces certain data links into a blocked state and keeps other links in a forwarding state.

An attack spoofs the root bridge and broadcasts out an STP configuration/ topology change, causing a denial-of-service attack

IP Layer Attacks — IP Spoofing

All IP packets contain a header that precedes the packet's body and includes the source/destination address. If the packet has been spoofed, the source address will be forged

IP Layer Attacks — IP Spoofing

All IP packets contain a header that precedes the packet's body and includes the source/destination address. If the packet has been spoofed, the source address will be forged

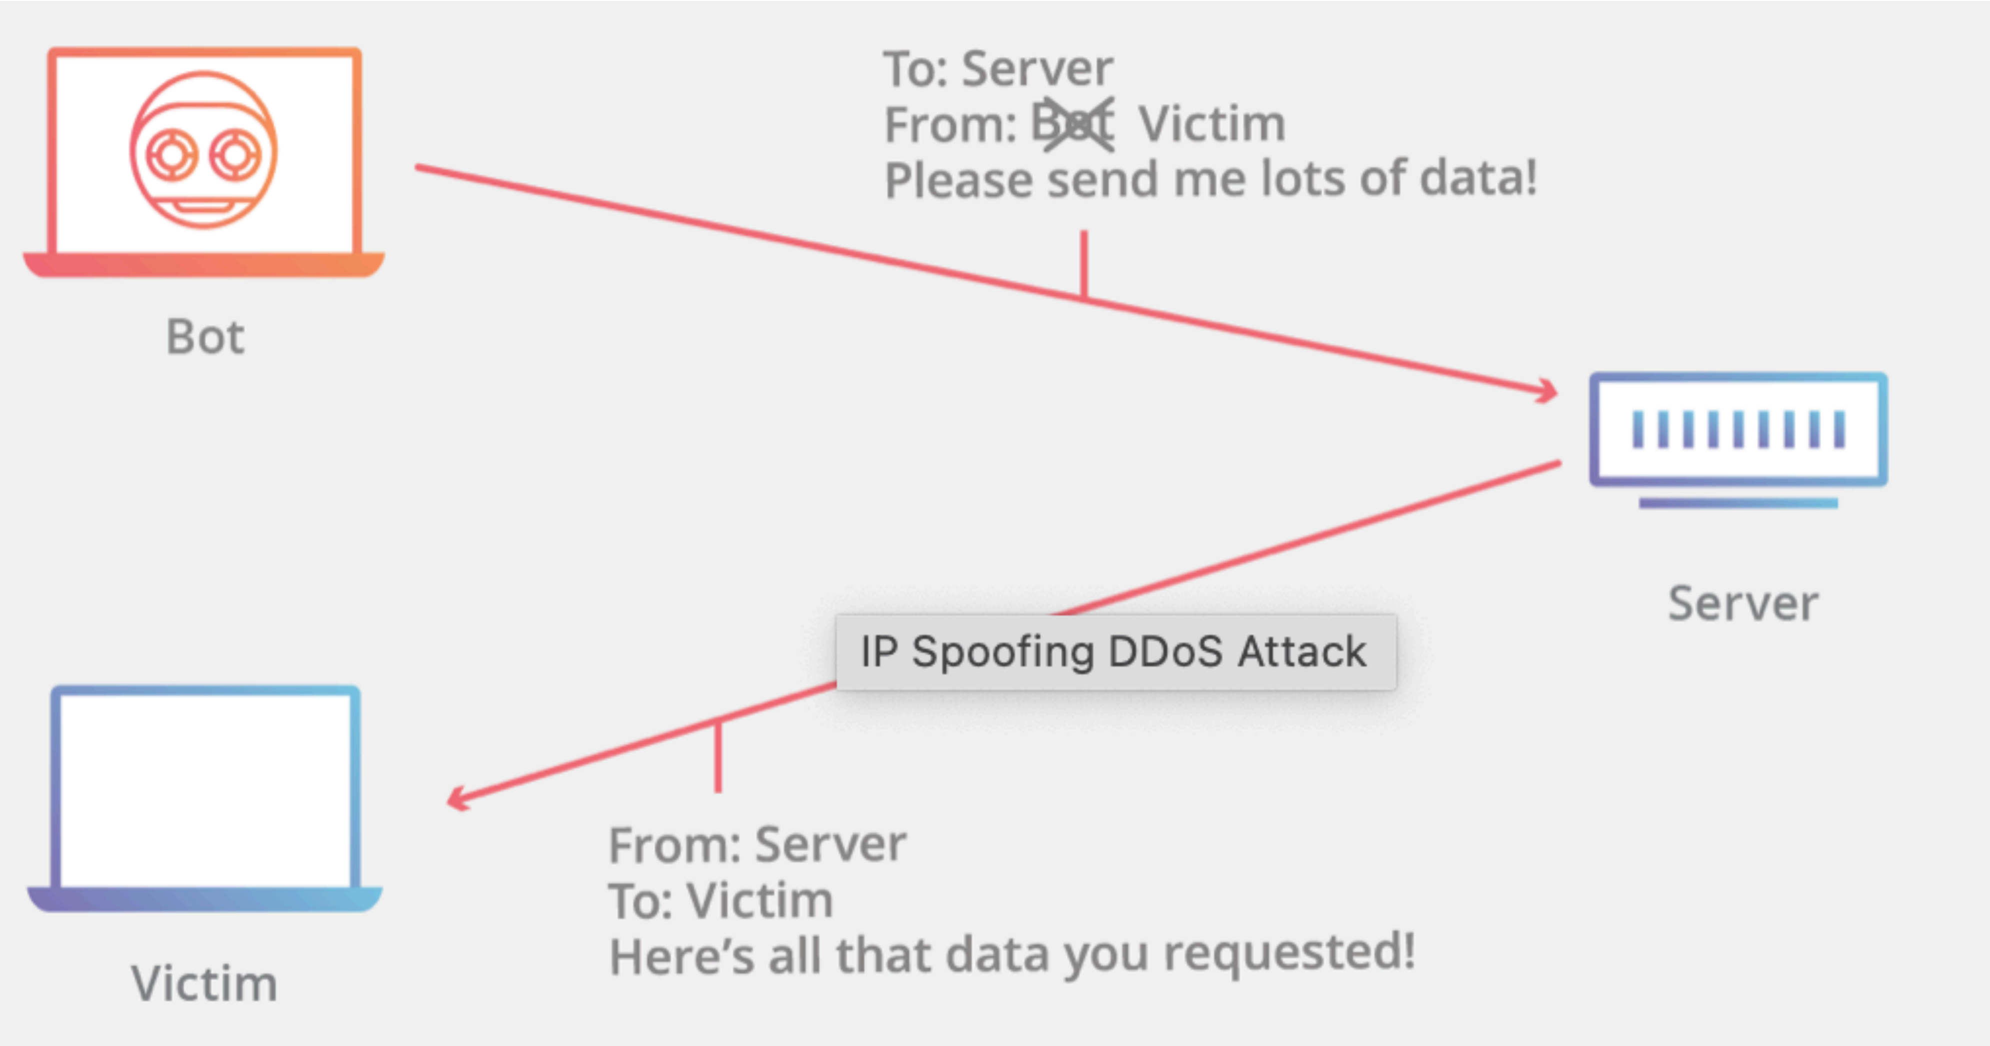

IP Layer Attacks — IP Spoofing

All IP packets contain a header that precedes the packet's body and includes the source/destination address. If the packet has been spoofed, the source address will be forged

An attack can flood the target with an overwhelming volume of traffic, causing a denial-of-service attack

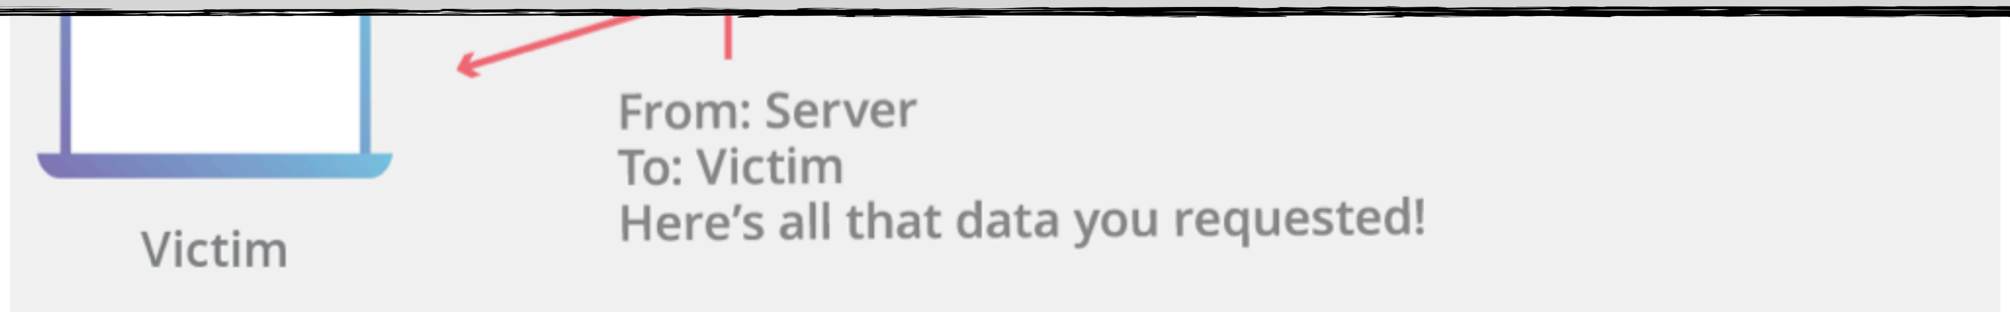

IP Layer Attacks — Route Hijacking

BGP is built on the assumption that interconnected networks are telling the truth about which IP addresses they own.

IP Layer Attacks — Route Hijacking

BGP is built on the assumption that interconnected networks are telling the truth about which IP addresses they own.

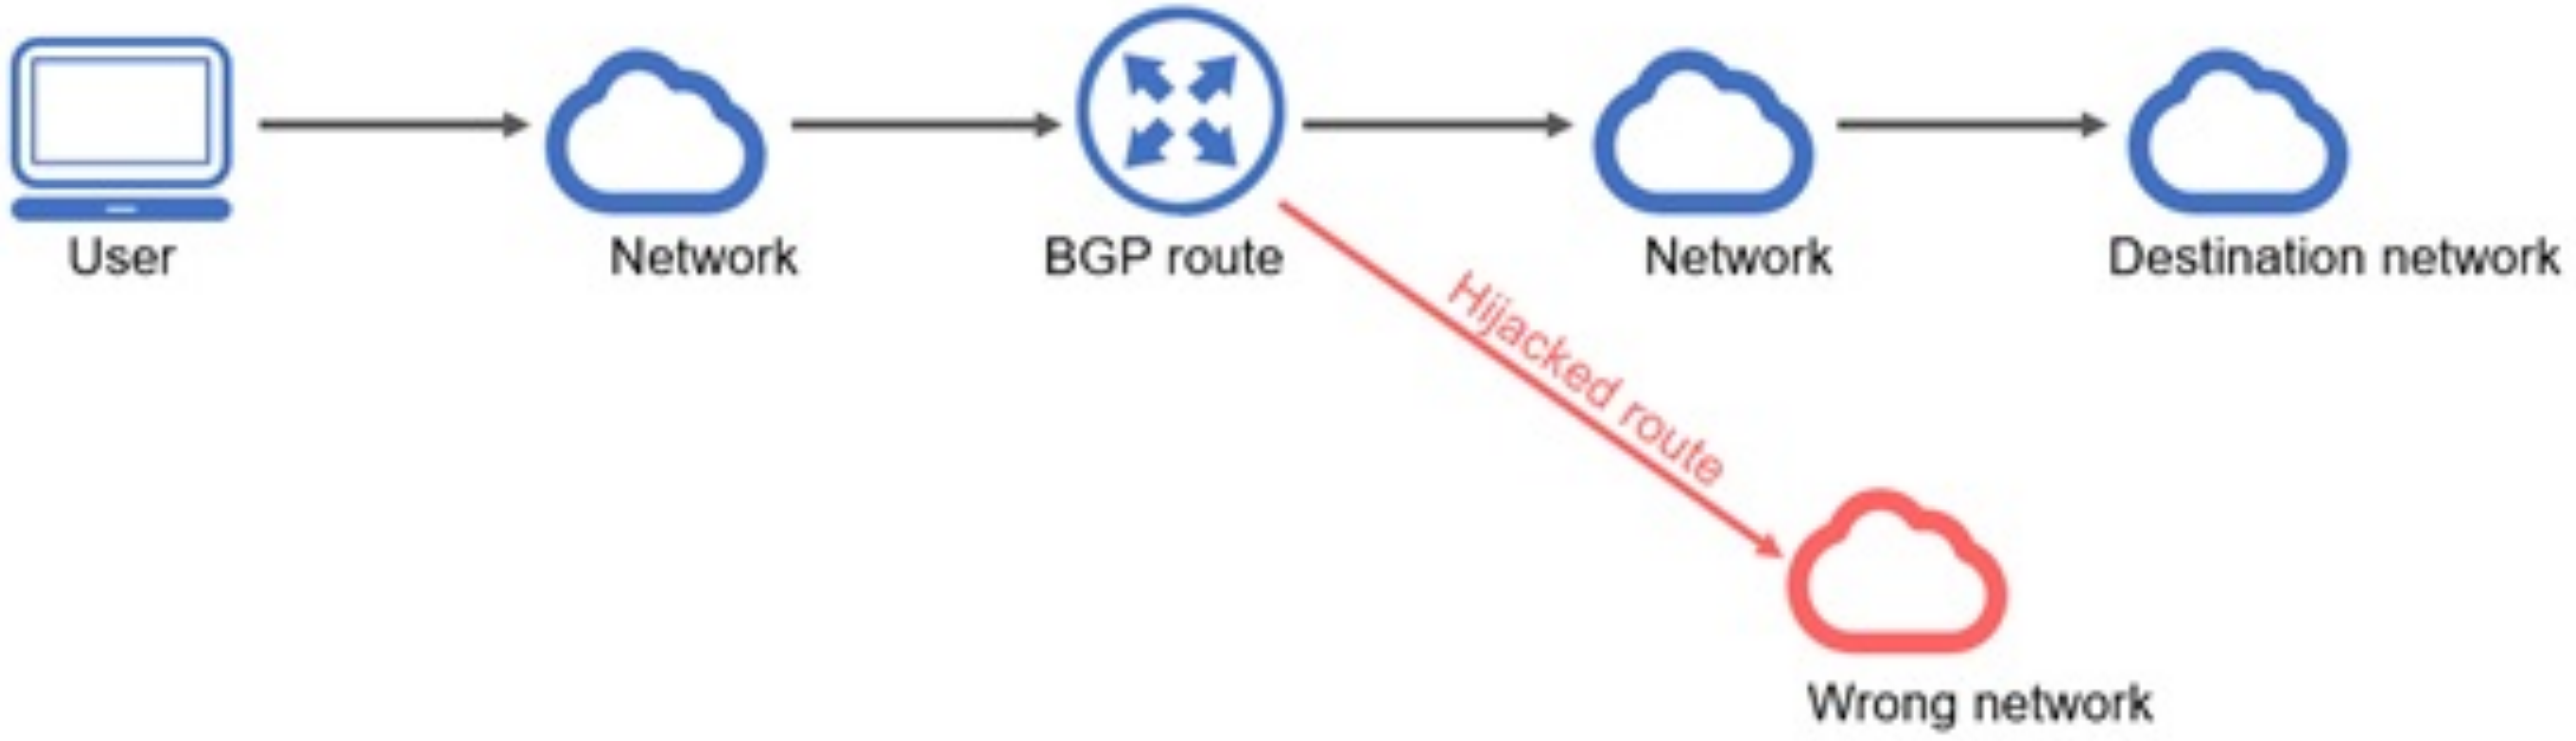

IP Layer Attacks — Route Hijacking

BGP is built on the assumption that interconnected networks are telling the truth about which IP addresses they own.

An attacker sends BGP announcements for prefixes you do now own or cannot reach. As a result, traffic would just take an unnecessarily long route or be redirect to fake websites, etc.

Transport Layer Attacks — SYN Flood

The TCP connection establishment phase starts with a standardized three-way handshake. The client sends an SYN packet. The server responds with a SYN-ACK

Transport Layer Attacks — SYN Flood

The TCP connection establishment phase starts with a standardized three-way handshake. The client sends an SYN packet. The server responds with a SYN-ACK

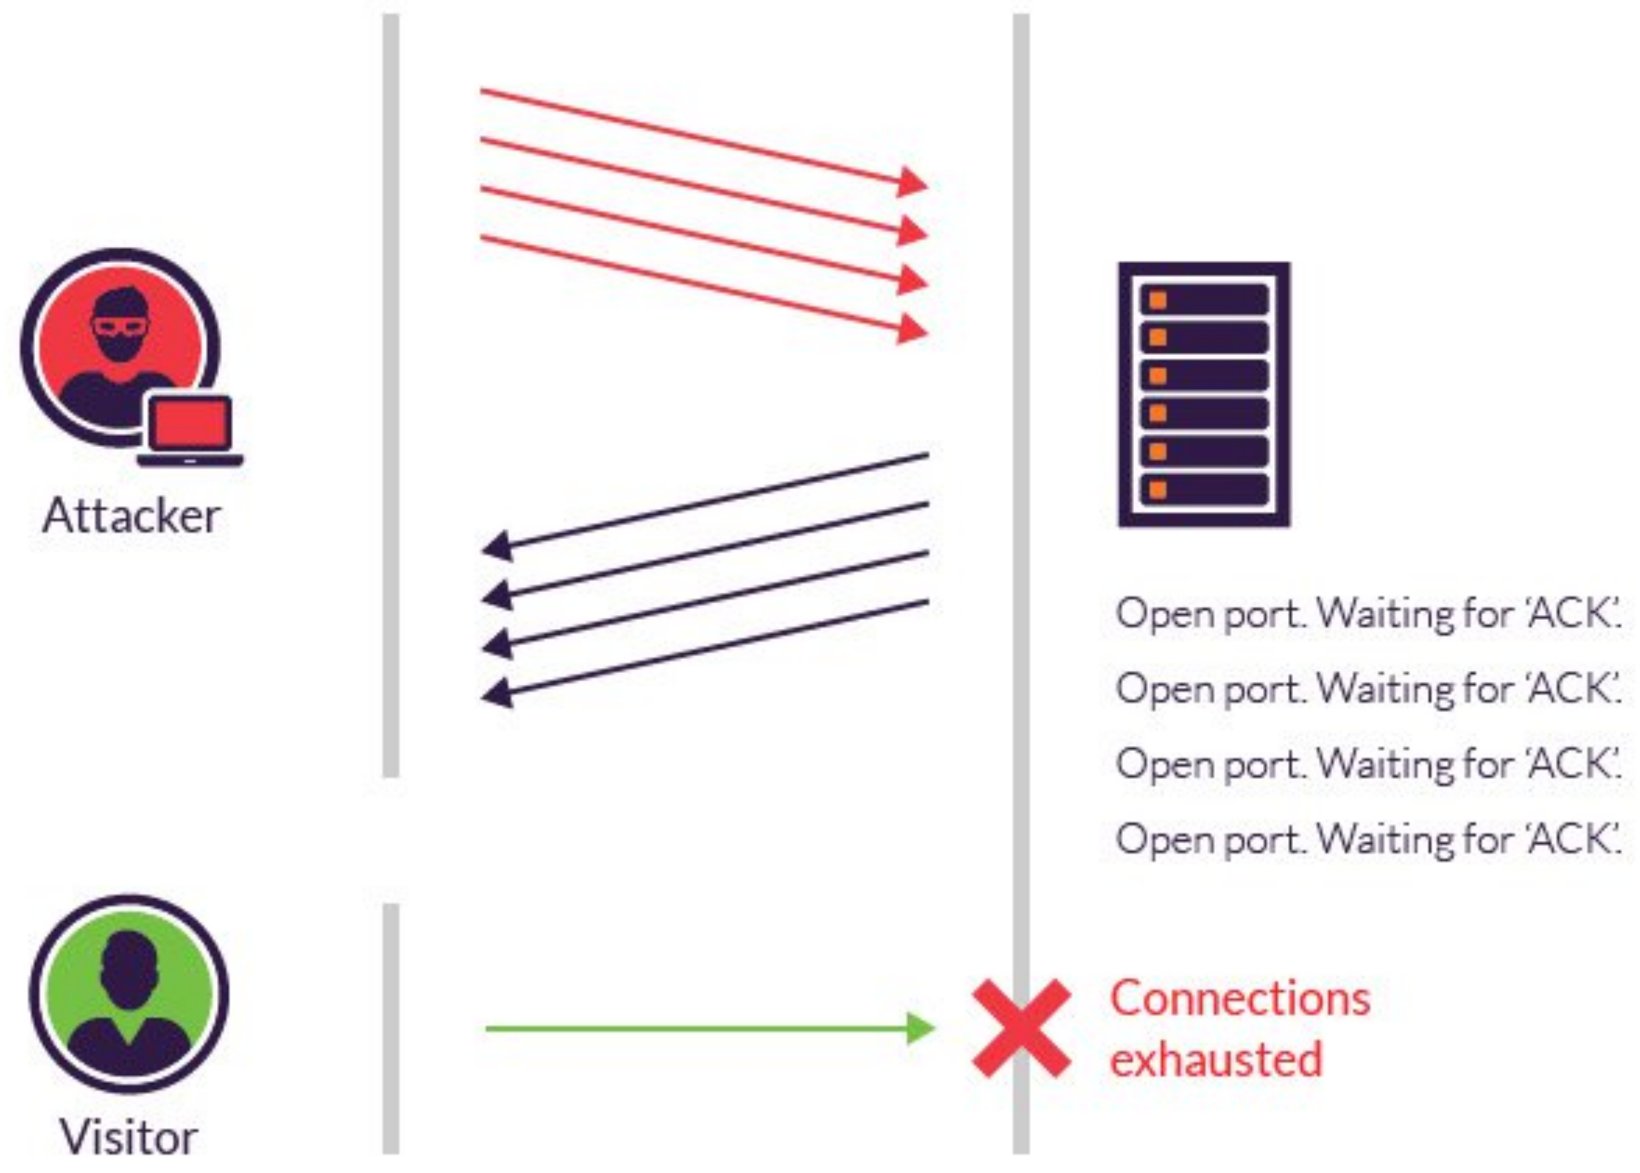

Transport Layer Attacks — SYN Flood

The TCP connection establishment phase starts with a standardized three-way handshake. The client sends an SYN packet. The server responds with a SYN-ACK

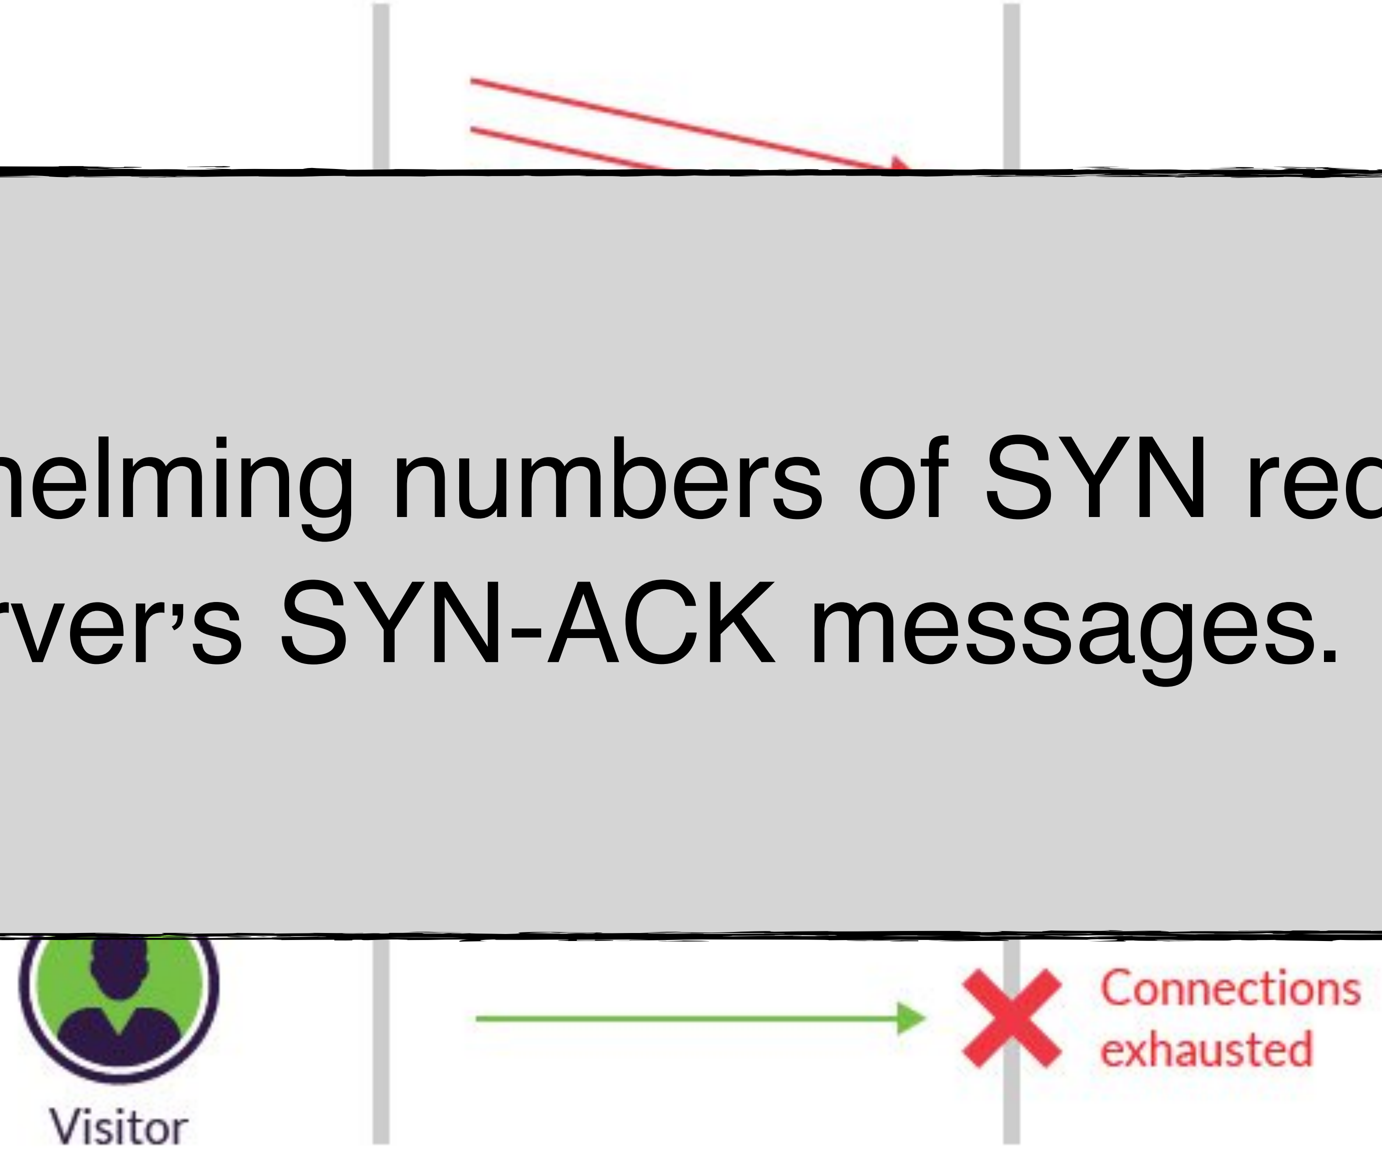

An attacker sends overwhelming numbers of SYN requests and intentionally never responds to the server's SYN-ACK messages.

Transport Layer Attacks — Session Hijacking

TCP is a connection-oriented protocol. Upon receiving a packet, if (1) <source IP, destination IP, source port, destination port> is matched and (2) sequence number is appropriate, the packet is accepted

Transport Layer Attacks — Session Hijacking

TCP is a connection-oriented protocol. Upon receiving a packet, if (1) <source IP, destination IP, source port, destination port> is matched and (2) sequence number is appropriate, the packet is accepted

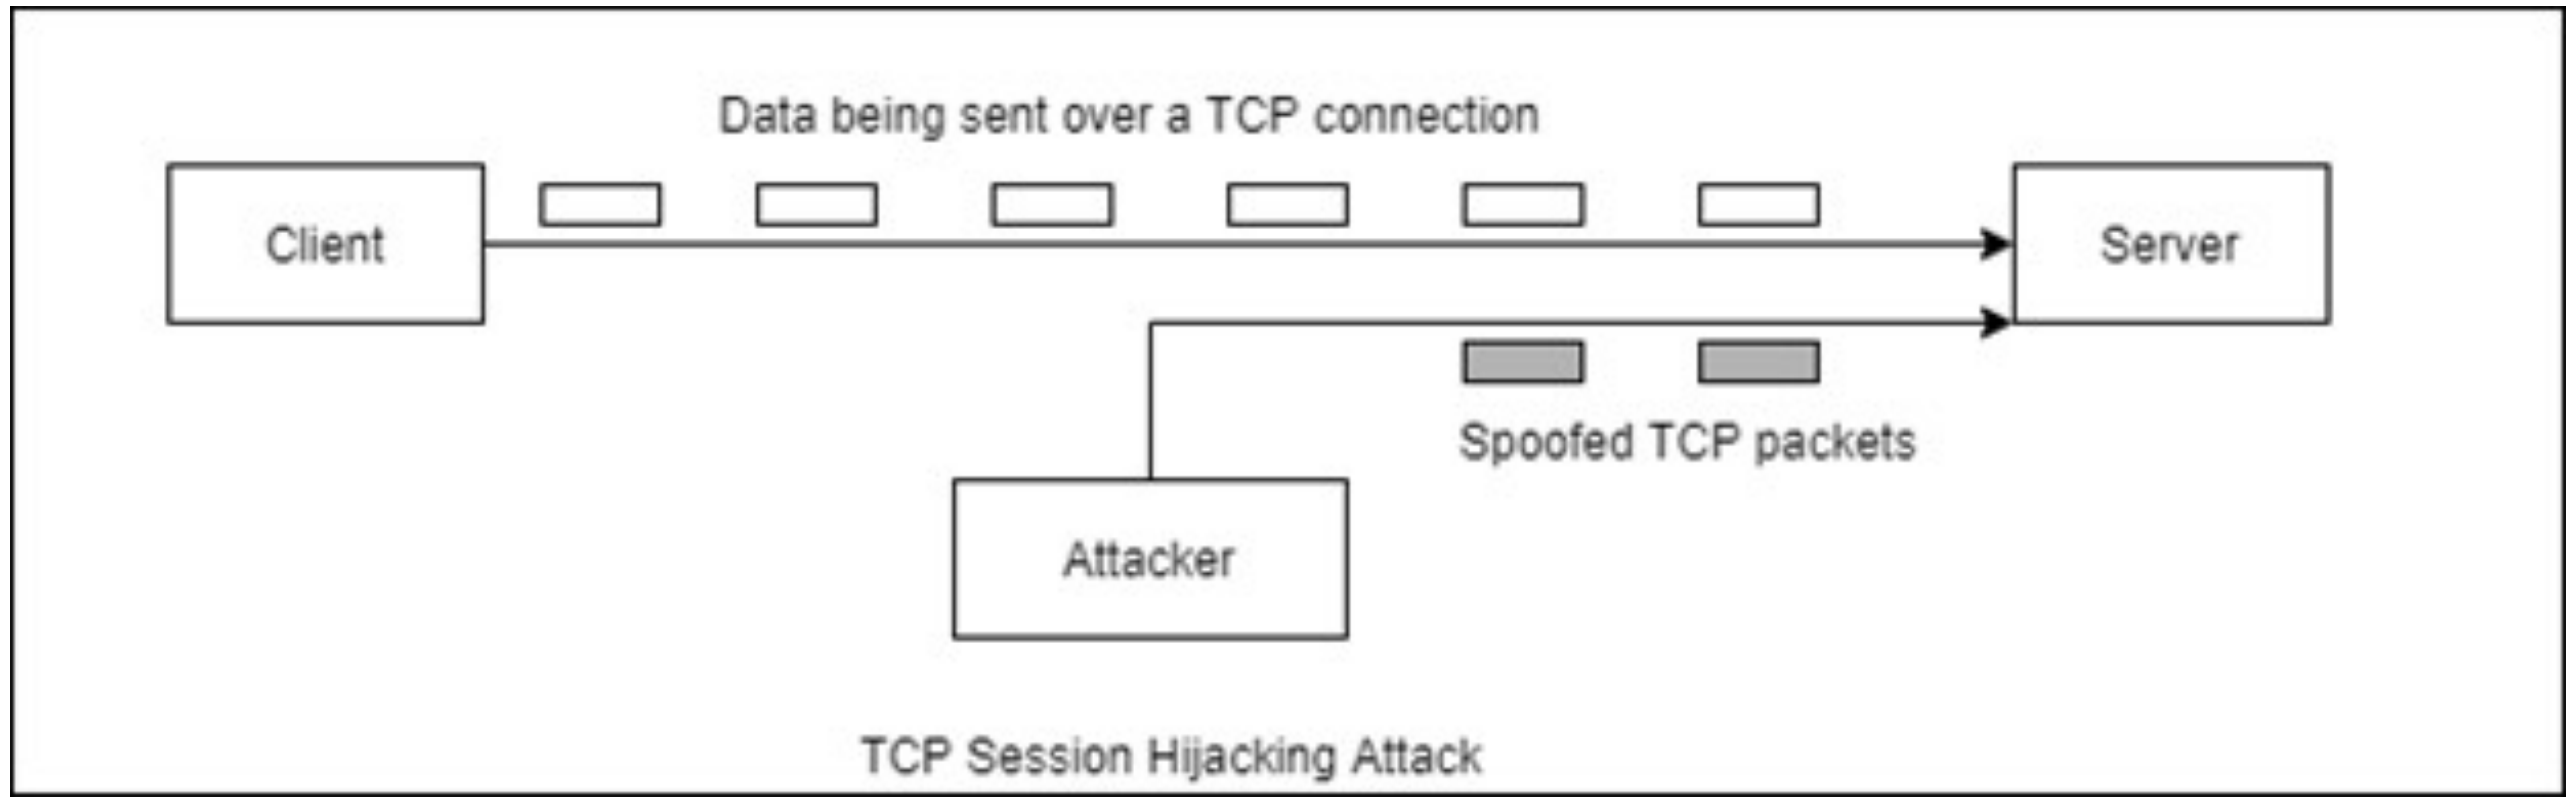

Transport Layer Attacks — Session Hijacking

TCP is a connection-oriented protocol. Upon receiving a packet, if (1) <source IP, destination IP, source port, destination port> is matched and (2) sequence number is appropriate, the packet is accepted

An attacker is able to gain the control of the session between the sender and receiver and do anything over the existing connection.

TCP Session Hijacking Attack

Application Layer Attacks — DNS Hijacking

DNS translates human-friendly URLs into machine-friendly IP addresses.

Application Layer Attacks — DNS Hijacking

DNS translates human-friendly URLs into machine-friendly IP addresses.

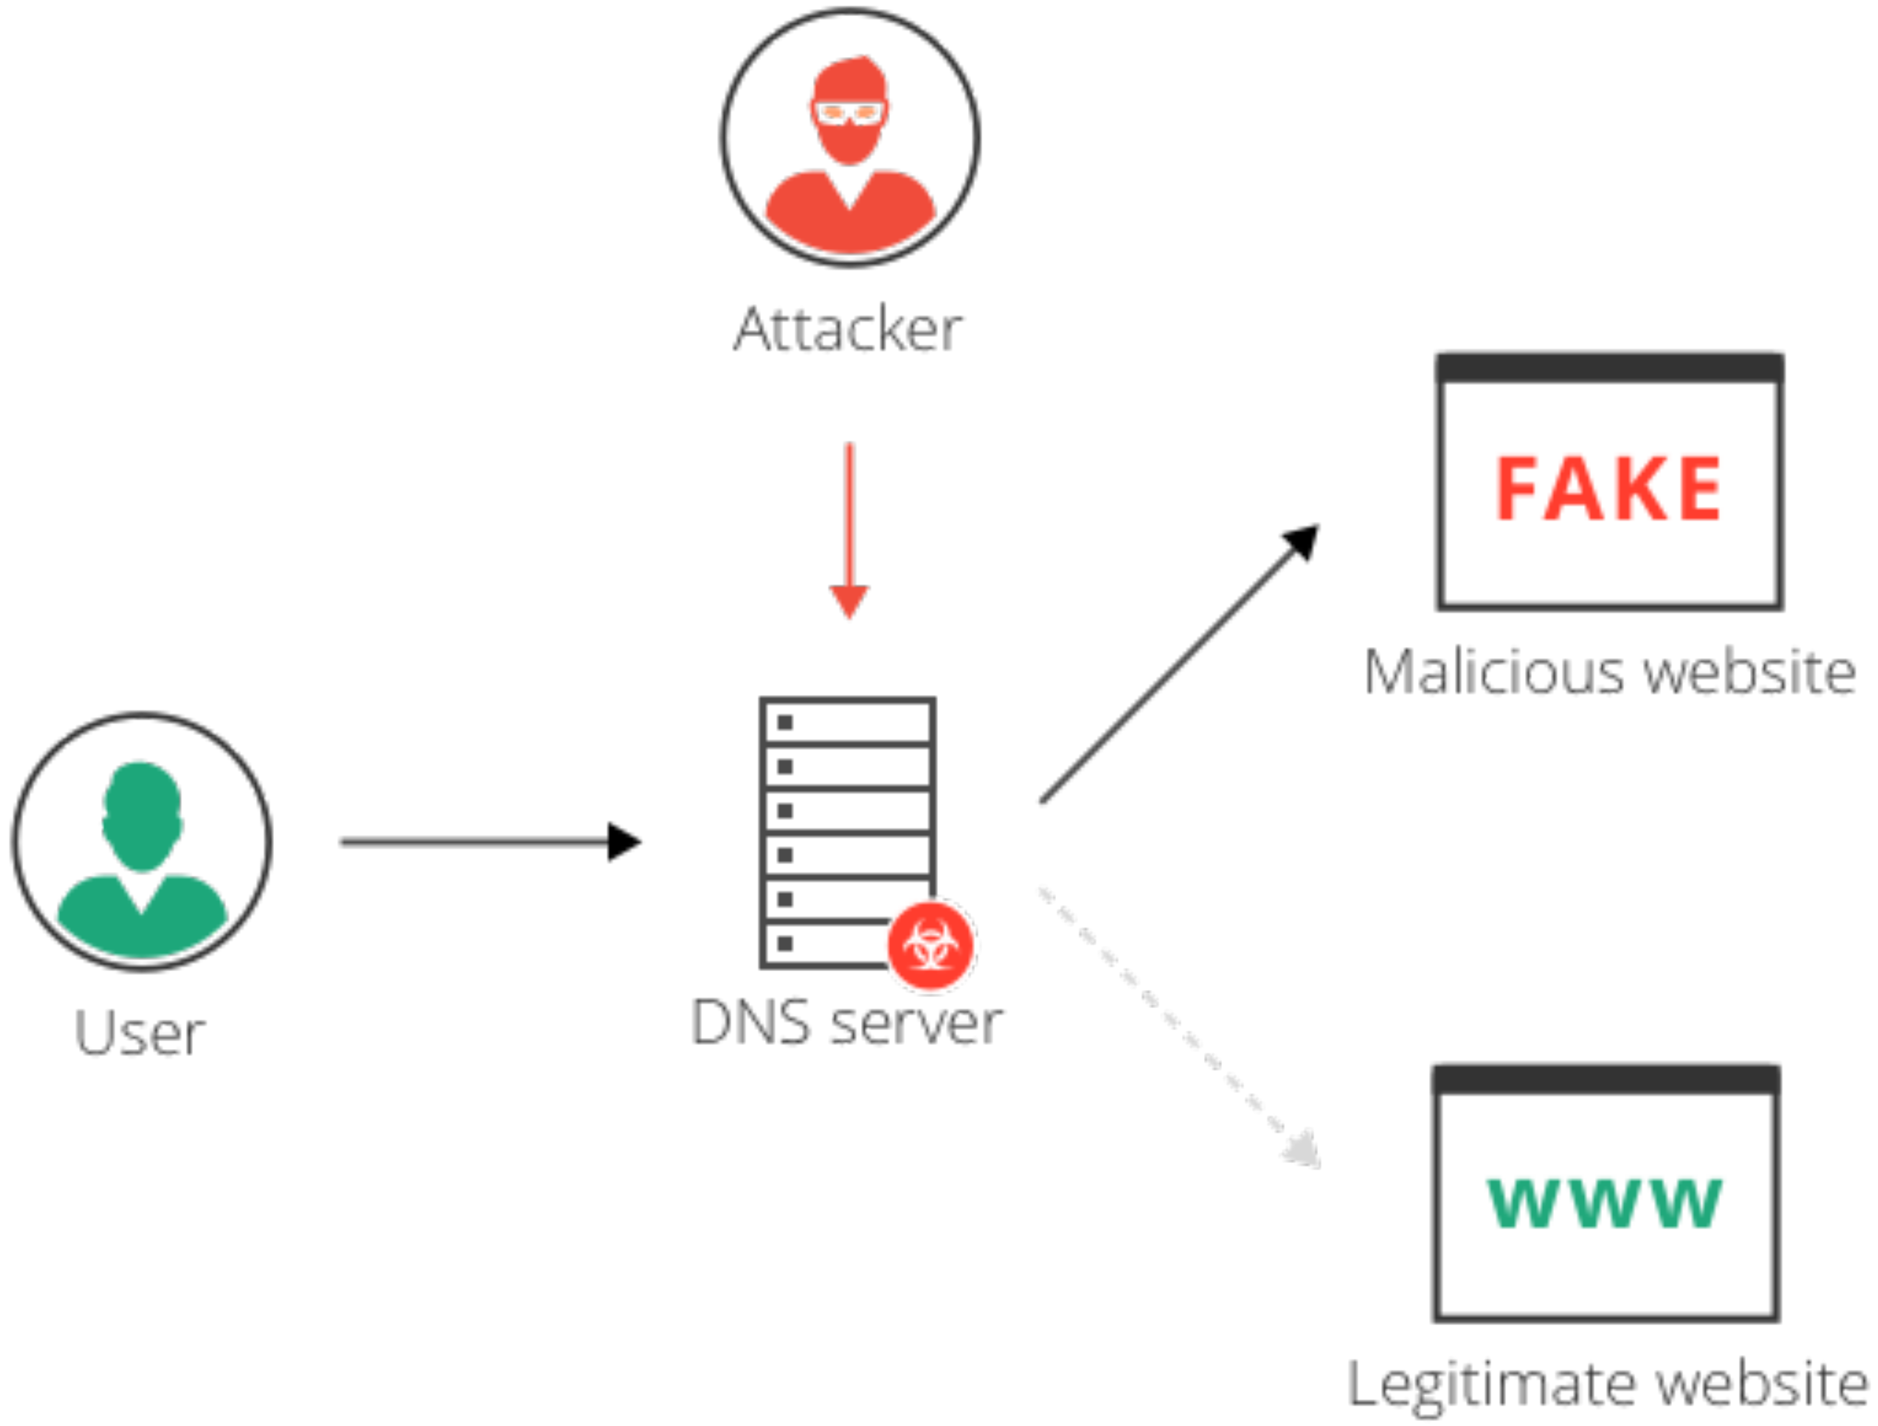

Application Layer Attacks — DNS Hijacking

DNS translates human-friendly URLs into machine-friendly IP addresses.

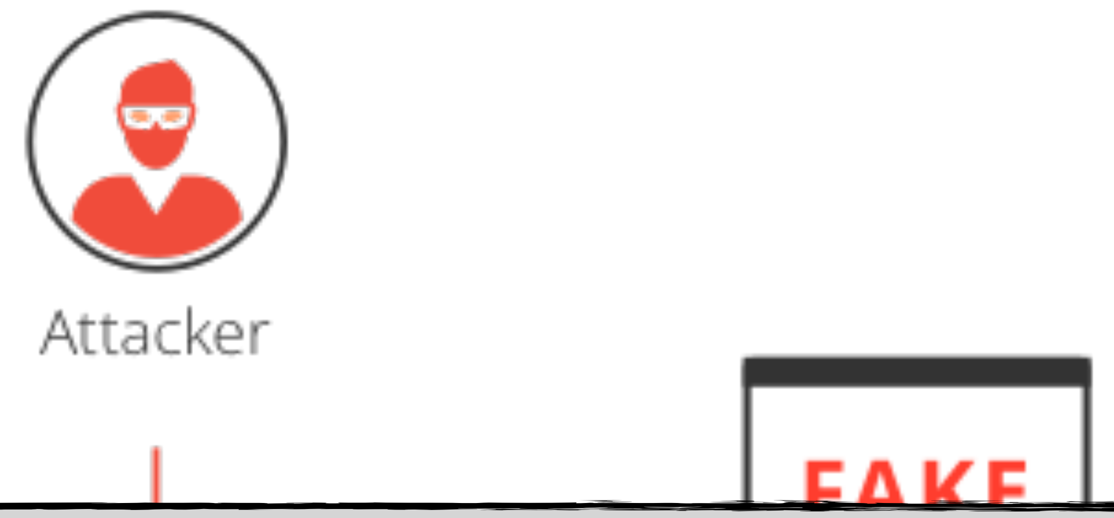

An attack can resolve domain names to IP address for servers with malicious code or phishing sites

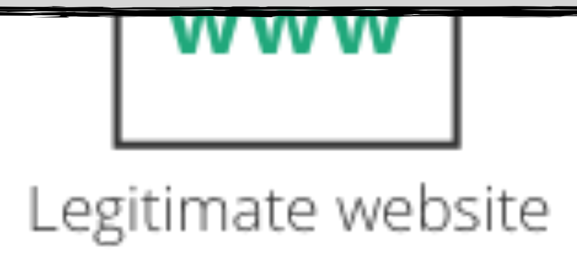

Application Layer Attacks — HTTP Flood

The web server replies based on the HTTP requests. A HTTP GET returns the fetched data. A HTTP POST pushes data into a persistent layer

Application Layer Attacks — HTTP Flood

The web server replies based on the HTTP requests. A HTTP GET returns the fetched data. A HTTP POST pushes data into a persistent layer

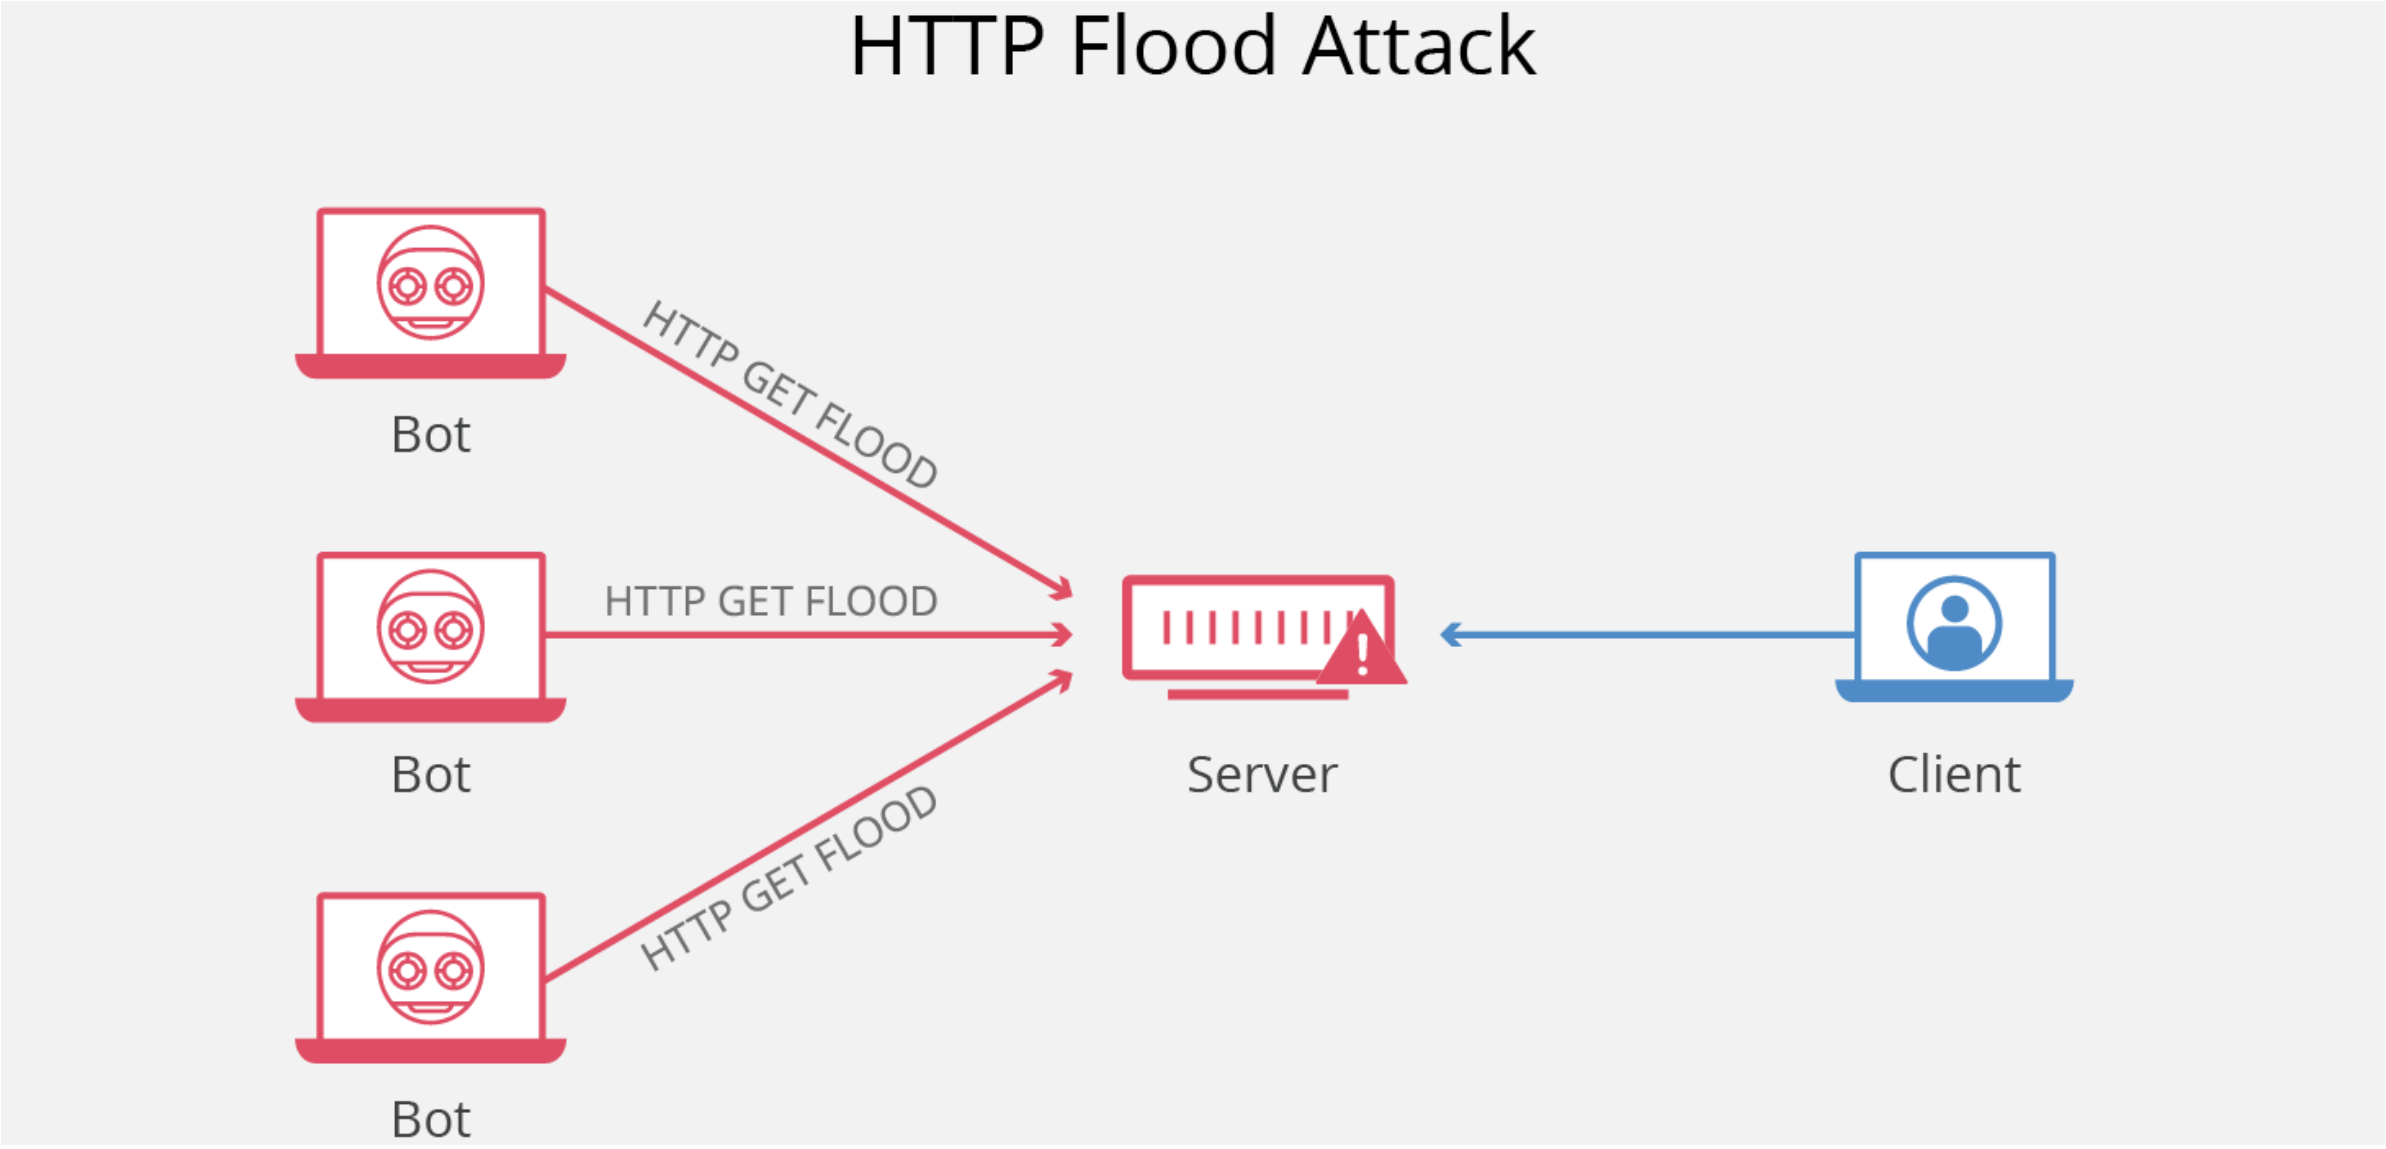

Application Layer Attacks — HTTP Flood

The web server replies based on the HTTP requests. A HTTP GET returns the fetched data. A HTTP POST pushes data into a persistent layer

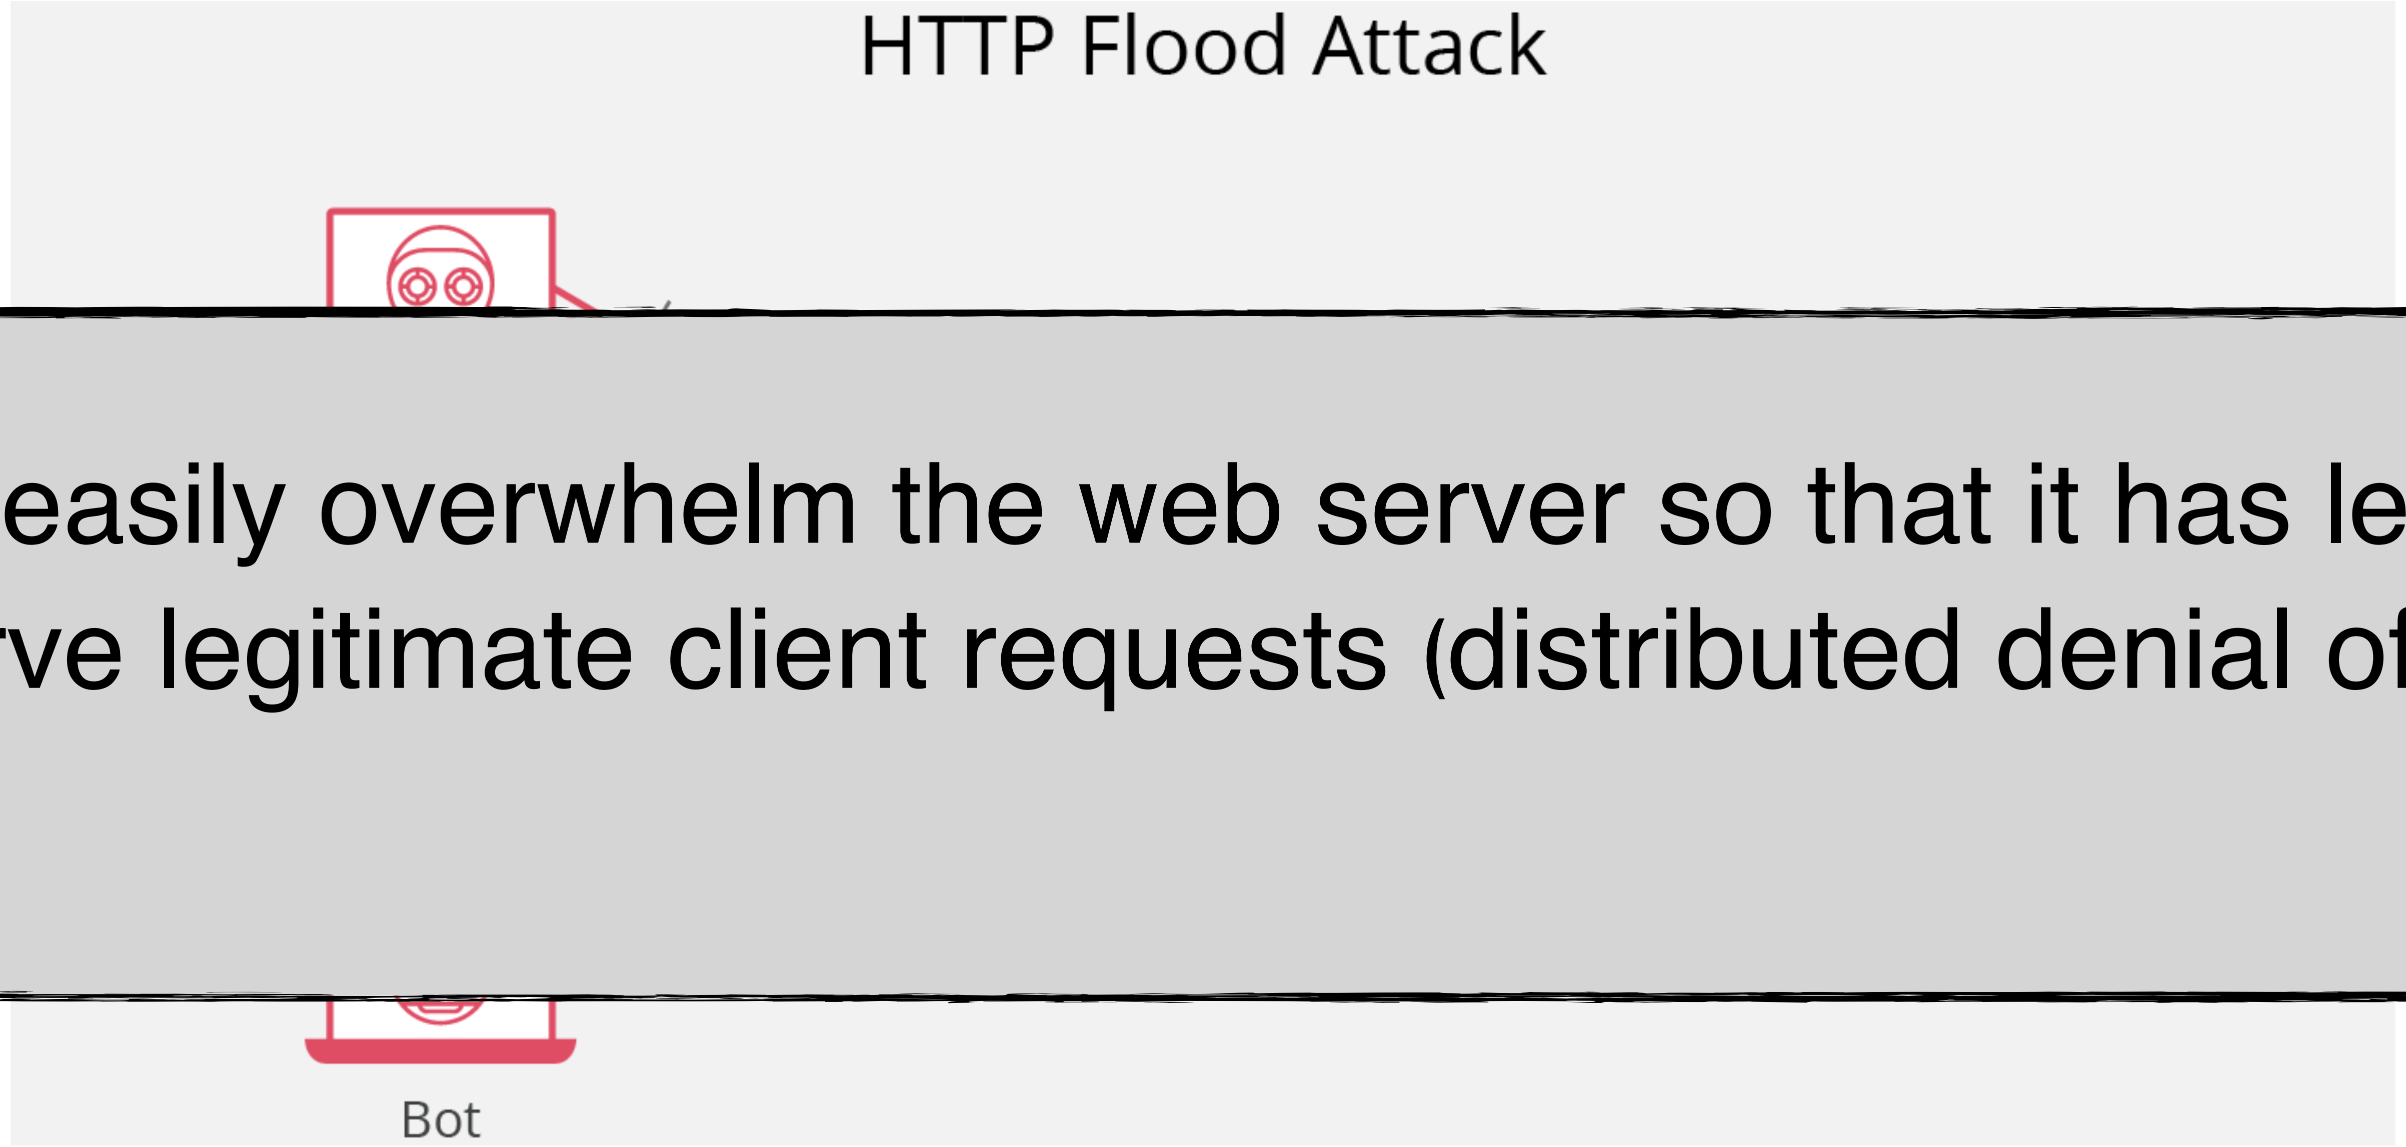

An attacker can easily overwhelm the web server so that it has less computing headroom to serve legitimate client requests (distributed denial of service or DDoS)

Q: How do networking attacks happen?

A: Everywhere across the stack.

Q: How does the network defense work?

A: Five approaches:

- #1: Encryption — make sure data remains confidential
- #2: Authentication — identify and assure the origin of information
- #3: Integrity checks — identify if the data is modified
- #4: Access control — selectively restrict users' hight
- #5: Middlebox — firewall, intrusion prevention

#1: Encryption

Encrypt/decrypt algorithm should be

- #1: Public inventing algorithms is hard, so we don't want to have to develop a new one if something is leaked
- #2: Easy to compute with the key: efficient in software and hardware, and on mobile devices which have fewer resources
- #3: Hard to compute without the key -> computers keep getting more powerful

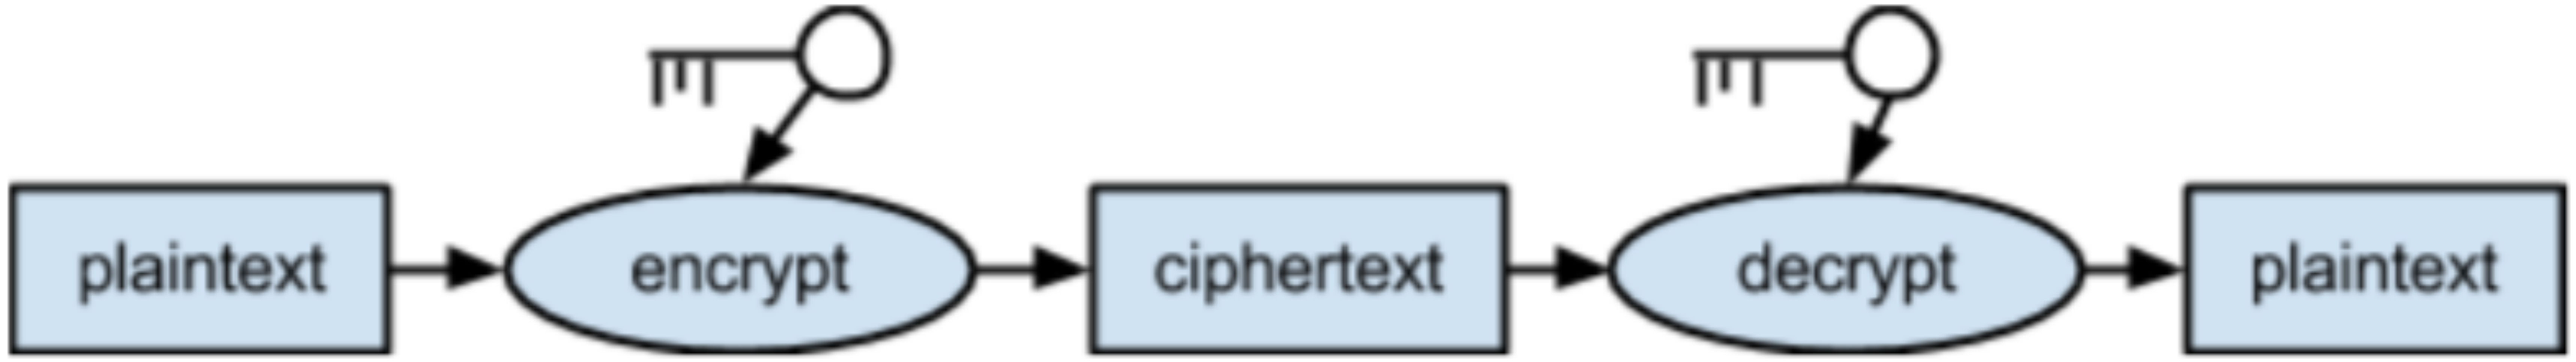

Basic Cryptography

Key should be

- Secret
- Long: length of key often determines “level” of security

Types of functions

- Cryptographic hash: no keys
- Symmetric/secret key: one shared key
- Asymmetric/public key: pair of keys, one public & one private

Cryptographic Hash

Also known as “cryptographic checksum” used to detect if a message has been tampered

Take message m of any length, and produce a smaller message h(m)

Cryptographic Hash

Properties

- #1: preimage resistance: hard to find m given h(m); “one-way” function
- #2: second preimage resistance: given a message m, it is hard to find a message m’ that hashes to the same h(m)
- #3: collision resistance: hard to find any two messages m and m’ such that h(m) = h(m’)
- E.g., SHA-2, MD5

Cryptographic Hash

Standards

- MD5 uses 128 bits; weakness known for a while
- SHA1 uses 160 bits; also not recommended for use
- SHA2 includes a collection of six different hash functions; use 224, 256, 384, or 512 bits
- SHA3 released in 2016

Example: self-certifying names

- File-sharing software (e.g., BitTorrent) names files with h(file data)
- Verify h(downloaded data) = name of file

Symmetric/Secret Key

Sender and receiver share a common key

None of the original structure of the plaintext should exist in the ciphertext

- Otherwise, attackers could look for patterns, e.g., commonly used letters in the English language; HTTP request starts with method (GET, POST, etc.)

Symmetric/Secret Key

Variants

- Data Encryption Standard (DES):
  - 64bit keys (8bits are parity)
  - Easy to recover a key given today's processing power
- Triple DES (3DES):
  - 168bit keys;
  - Slow to implement in software
- Advanced Encryption Standard (AES):
  - 128/192/256bit keys
  - Fast implementations in hardware or software and low memory footprint

Symmetric/Secret Key Challenge

Key distribution

- Physically delivery key:
  - Not practical
- Using an existing key to deliver a new key:
  - Need unique key pair for each pair of endpoints
  - N * (N-1)/2 total keys for n endpoints
- Use key distribution center (KDC):
  - KDC generates session keys and distributes them to pairs of endpoints
  - Need n master keys if you have n endpoints

#5: Middleboxes

Systems in the "middle" of the network that examines and block packets and flows

- Middle = on the path between pairs of communication hosts
- Packets forced to pass through the middlebox based on physical network topology or using SDN

Middleboxes — Firewalls

Basic firewalls

- Apply simple rules to decide whether to forward or block packets
- Rules are based on fields in packet header, e.g., source/destination IPs, transport layer
- Default rule to either forward or block if no other rules match

Middleboxes — Firewalls

Advanced firewalls

- Maintain some states about active connections
  - The current state of TCP connection: SYN sent, SYN+ACK/FIN sent, established, etc.
  - Based on both packet headers and the current state
- Application Awareness
  - Extra protection for services that should not be blocked (e.g., HTTP)
  - E.g., check if the HTTP POST method is allowed
  - E.g., check if clients are requesting a web page from a domain that is on a blacklist because it is known to host malware
  - Often acts as a proxy – terminates TCP connection from the client and establishes separate TCP connection to server

Middleboxes — IDS/IPS

Intrusion detection/prevention system

- IDS: raise alerts
- IPS: raise alerts and block traffic

Perform deep packet inspection

- Look at the payload of the packet, not just the header, to decide if it should be blocked
- Know the format of the packet for many different transport and application protocols – HTTP/SSH/SSL/TLS/FTP/NFS/FTP/NTP/etc.
- Cannot perform deep packet inspection on traffic that is encrypted

Middleboxes — IDS/IPS

Maintain the state about active connections

- Connection info such as src/dst IP, src/dst port, and TCP connection state
- Reassembled payloads – e.g., HTTP reply may be split among multiple packets, so the packets are reassembled into a single memory region that contains the entire reply

IDS/IPS Design and Implementation

Use a set of signatures to detect malicious traffic

- Specific sequences of packet – e.g., TCP SYN+ACK after TCP FIN
- Keywords in payloads – e.g., "root"
- MD5 sum of payload – compare against a database of MD5 sums for known malware
- Large numbers of packets to one host in a short time

The speed of signature matching can significantly impact latency and throughput

- Want efficient pattern-matching algorithms
- Sometimes use custom hardware
- Can be easily parallelized

Terminology

1. Host
2. NIC
3. Multi-port I/O bridge
4. Protocol
5. RTT
6. Packet
7. Header
8. Payload
9. BDP
10. Baud rate
11. Frame/Framing
12. Parity bit
13. Checksum
14. Ethernet
15. MAC
16. (L2) Switch
17. Broadcast
18. Acknowledgement
19. Timeout
20. Datagram
21. TTL
22. MTU
23. Best effort
24. (L3) Router
25. Subnet mask
26. CIDR
27. Converge
28. Count-to-infinity
29. Line card
30. Network processor
31. Gateway
32. Private network
33. IPv6
34. Multicast
35. IGMP
36. SDN
37. (Transport) port
38. Pseudo header
39. SYN/ACK
40. Incarnation
41. Flow
42. SYN flood
43. TCP Segment
44. Window
45. Advertised Window
46. Effective Window
47. TCP Reno
48. Duplicated ACK
49. Congestion Window
50. Congestion Threshold
51. Selective Acknowledgment
52. Active Queue Management (AQM)
53. URL
54. HTML
55. Peer-to-peer (P2)
56. Swarm
57. CDN
58. ARP/IP Spoofing
59. MAC/SYN/HTTP Flooding
60. Route/Session/DNS Hijacking
61. Presage resistance
62. Collision resistance
63. Middlebox
64. Firewall

Principle

1. Layering
2. Minimal States
3. Hierarchy
4. Mechanism/policy separation

Technique

1. NRZ Encoding
2. NRZI Encoding
3. Manchester Encoding
4. 4B/5B Encoding
5. Byte Stuffing
6. Byte Counting
7. Bit Stuffing
8. 2-D Parity
9. CRC
10. MAC Learning
11. Store-and-Forward
12. Cut-through
13. Spanning Tree
14. CSMA/CD
15. Stop-and-Wait
16. Sliding Window
17. Fragmentation and Reassembly
18. Path MTU discovery
19. DHCP
20. Subnetting
21. Supernetting
22. Longest prefix match
23. Distance vector routing (RIP)
24. Link state routing (OSPF)
25. Boarder gateway protocol (BGP)
26. Network address translation (NAT)
27. User Datagram Protocol (UDP)
28. Transmission Control Protocol (TCP)
29. Three-way Handshake
30. TCP state transition
31. EWMA
32. Sliding window

Technique

33. Flow control
34. AIMD
35. Slow start
36. Fast retransmit
37. Fast recovery
38. Nagle’s algorithm
39. Karn/Partridge algorithm
40. TCP Vegas
41. Bit-by-bit Round Robin
42. Fair Queueing (FQ)
43. Random Early Detection (RED)
44. Explicit Congestion Notification (ECN)
45. Domain Name System (DNS)
46. Simple Network Management Protocol (SNMP)
47. HyperText Transfer Protocol (HTTP)
48. Persistent Connection
49. BitTorrent
50. Cryptographic Hash
51. DES/3DES/AES
52. Intrusion detection/prevention system (IDS/IPS)

Summary

Today's takeaways

#1: Network security is a concern to any entities in the network (including client/server)

#2: Network attacks can happen across the stack

#3: There are five common ways to do networking defense

Next lecture

- CS640 Recap