

# Lec 23: DoS and Firewalls

CSED415: Computer Security  
Spring 2024

Seulbae Kim

**POSTECH**  
POHANG UNIVERSITY OF SCIENCE AND TECHNOLOGY

# Recap

- Compromising confidentiality and integrity
  - Many requirements
    - Skilled attacker
    - Vulnerable program or policy
    - Entry points to target system
    - Channel for exfiltration of sensitive data

How about availability? → Today's topic

# Denial of Service (DoS)

# Availability and Denial of Service

---

- Availability
  - Making a service on the network available for legitimate users
- Denial of Service (DoS)
  - An attack that disrupts availability of a service, making it unavailable for legitimate users

# Availability and Denial of Service

---

- Reasons for a DoS attack
  - Competitors DoS each other to benefit their own services
  - Criminals might DoS services and ask for a ransom
  - People might DoS services to make a political statement
  - Entities might DoS each other as part of warware tactics
  - Just for fun or revenge

LILY HAY NEWMAN

SECURITY

MAR 1, 2018 11:01 AM


## GitHub Survived the Biggest DDoS Attack Ever Recorded

On Wednesday, a 1.3Tbps DDoS attack pummeled GitHub for 15–20 minutes. Here's how it stayed online.

ON WEDNESDAY, AT about 12:15 pm EST, 1.35 terabits per second of traffic hit the developer platform GitHub all at once. It was the most powerful distributed denial of service attack recorded to date—and it used an increasingly popular DDoS method, no botnet required.

GitHub briefly struggled with intermittent outages as a digital system assessed the situation. Within 10 minutes it had automatically called for help from its DDoS mitigation service, Akamai Prolexic. Prolexic took over as an intermediary, routing all the traffic coming into and out of GitHub, and sent the data through its scrubbing centers to weed out and block malicious packets. After eight minutes, attackers relented and the assault dropped off.

TRENDING ↗



What is a [  
Them?

## Eight-Hour DDoS Attack Struck AWS Customers

Google Cloud Platform suffered issues around the same time as Amazon Web Services but claims they were not caused by DDoS.



Dark Reading Staff, Dark Reading

October 24, 2019

🕒 1 Min Read

A significant distributed denial-of-service (DDoS) attack lasting approximately eight hours affected Amazon Web Services yesterday, knocking its S3 service and other services offline between 10:30 a.m. and 6:30 p.m. PDT.

The attack struck AWS's Router 53 DNS Web service, which [led to](#) outages for other services that require public DNS resolution: Elastic Load Balancing, Relational Database Service, and Elastic Compute Cloud. AWS alerted customers while the attack was ongoing to inform them of "intermittent errors with resolution of some AWS DNS names." Starting at 5:16 p.m., a small number of specific DNS names experienced a higher error rate. The issues have been resolved.

### Editor's Choice



#### Blinken: Digital Solidarity Is 'North Star' for US Policy

by Karen Spiegelman, Features Editor

MAY 7, 2024

3 MIN READ

CYBER RISK

# DoS in real life

Daryna Antoniuk

January 19th, 2024

News Briefs

Nation-state



Get more insights with the  
Recorded Future  
Intelligence Cloud.

[Learn more.](#)

## Swiss websites hit by DDoS attacks during World Economic Forum in Davos

Swiss websites were hit by a wave of distributed denial-of-service (DDoS) attacks this week, likely orchestrated by pro-Russian hackers.

According to the Swiss National Cybersecurity Centre (NCSC), the attacks temporarily **disrupted** access to several websites run by the Federal Administration — the government's executive branch.

“The cyberattack was promptly detected and the Federal Administration's specialists took the necessary action to restore access to the websites as quickly as possible,” said NCSC’s statement.

DDoS attacks are aimed at making websites unavailable but do not result in any data being lost or compromised.

A Russian politically-motivated hacker group known as NoName **claimed responsibility** for the attacks on its Telegram channel.



# DoS attack strategies

- Exploiting program flaws
  - Software vulnerabilities can cause a service to go offline
  - e.g., exploit a buffer overflow and execute “shutdown” command
- Resource exhaustion
  - Every computing system has limited resources
  - The attacker consumes all the limited resources so legitimate users cannot use them
  - Exhausting system’s bottleneck is enough for DoS

# DoS attack strategies

---

- Bottlenecks
  - Different parts of the system have different resource limits
  - The attacker only needs to exhaust the bottleneck
    - The part of the system with the least resources
    - e.g., the component with the least CPU cycle allocated, the component with the smallest amount of memory allocated, ...

# DoS targets

---

- Application-level DoS
  - Target the high-level application running on the host
- Network-level DoS
  - Target network protocols to affect the host's internet

# Application-level DoS

# Application-level DoS

---

- Target the resources that an application uses
- Exploit features of the application itself
- Some attacks exploit asymmetry:
  - Small amount of input from the attack
  - Large amount of resource consumed to handle the input

# Resource consumption

- Idea: Force the server to consume all its resources
  - Potential payloads

```
int fd = open("/tmp/junk", O_CREAT | O_RDWR, 0664);  
char buf[4096];  
while (1) {write(fd, buf, 4096);}
```

Exhausts filesystem space

```
while (1) {malloc(100000000000);}
```

Exhausts main memory

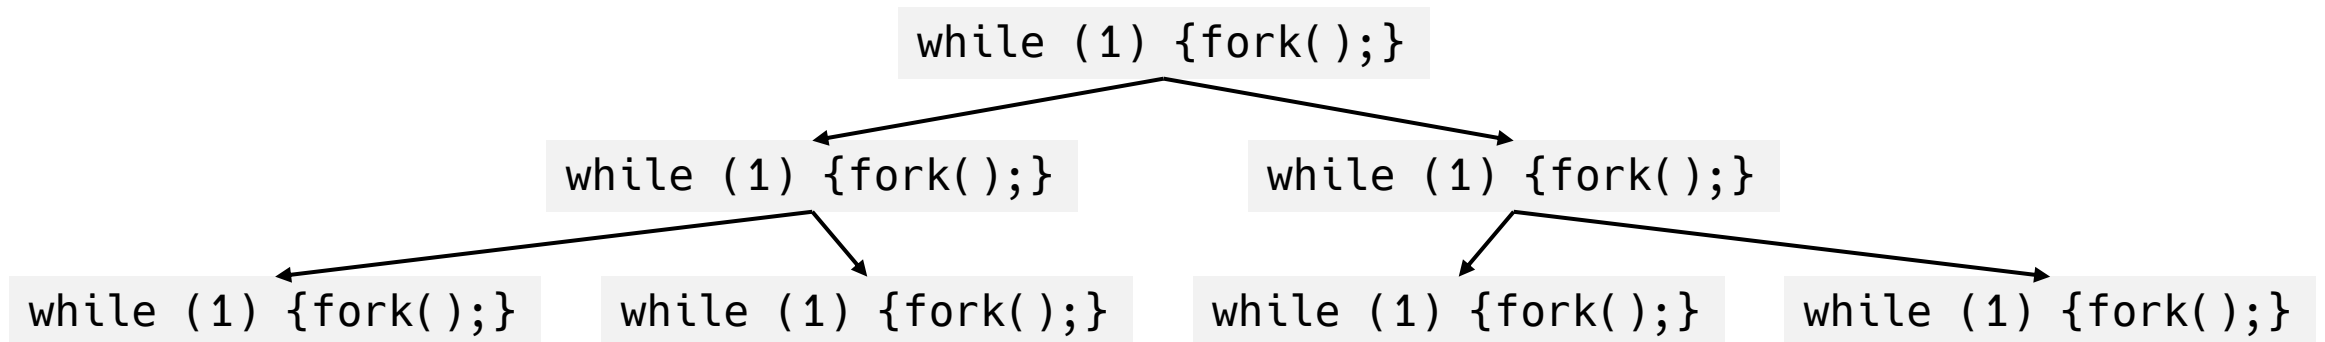
```
while (1) {fork();} Fork bomb!
```

# Resource consumption

- Idea: Force the server to consume all its resources
  - Fork bomb explained

```
$ man fork
```

fork() creates a new process by **duplicating** the calling process. The new process is referred to as the child process. The calling process is referred to as the parent process. The child process and the parent process run in **separate memory spaces**. The child process is an **exact duplicate** of the parent process.



Exhausts OS process table, CPU cycles, and memory

# Algorithmic complexity attacks

- Supplying inputs that trigger worst-case complexity of algorithms and data structures
  - e.g., consider an application that runs quicksort on user data
    - Average time complexity:  $O(n \log(n))$
    - Worst-case:  $O(n^2)$  // when?



# Application-level DoS – Defenses

- Isolation
  - Ensure that one user's actions do not affect other user's experience
  - e.g., modularize services and put in containers (e.g., Docker)
- Quota allocation
  - Ensure that users can only access a certain amount of resources
    - e.g., ulimit and rlimit

# Application-level DoS – Defenses

- Require proof-of-work: Force users to spend some of their resources to issue a request
  - Idea: Make DoS expensive for attackers
  - e.g., add a CAPTCHA
- Overprovisioning: Allocate a huge amount of resources
  - A costly solution
  - Applicable if the importance of availability far outweighs the money
  - e.g., Cloudflare CDN


# Network-level DoS

# Network-Level DoS

- Attack network protocols to affect the victim's internet access
  - Overwhelm the victim's bandwidth
    - Bandwidth: The amount of data that can be uploaded and downloaded through a channel in a given time
    - Attacker can send many **maximum-sized** packets
  - Overwhelm the victim's packet processing capability
    - Example: The server can process 50 packets per second. Attacker sends the server 500 packets per second.
    - Attacker can send many **minimum-sized** packets

# Distributed Denial-of-Service (DDoS)

- DDoS: Using multiple systems to overwhelm the target system
  - Every system/channel has a bandwidth
  - Attacker can gain a huge amount of bandwidth by controlling many systems (e.g., botnet!)
  - Sending packets from many sources can fool packet filters to distinguish DDoS traffic from normal traffic
    - Blocking several IPs is not effective




# Amplification DoS


- Amplification DoS: Using an amplifier to overwhelm the target more effectively
  - Some services produce a significantly larger response compared to the size of the request
  - Attackers can send a small request with spoofed source IP address (e.g., disguising the sender as the victim) to exploit such services, causing a large volume of data to be sent to the victim

# Amplification DoS

- Example: Domain Name Server (DNS) amplification
  - DNS lookup



- Attack



# Network-Level DoS – Defenses

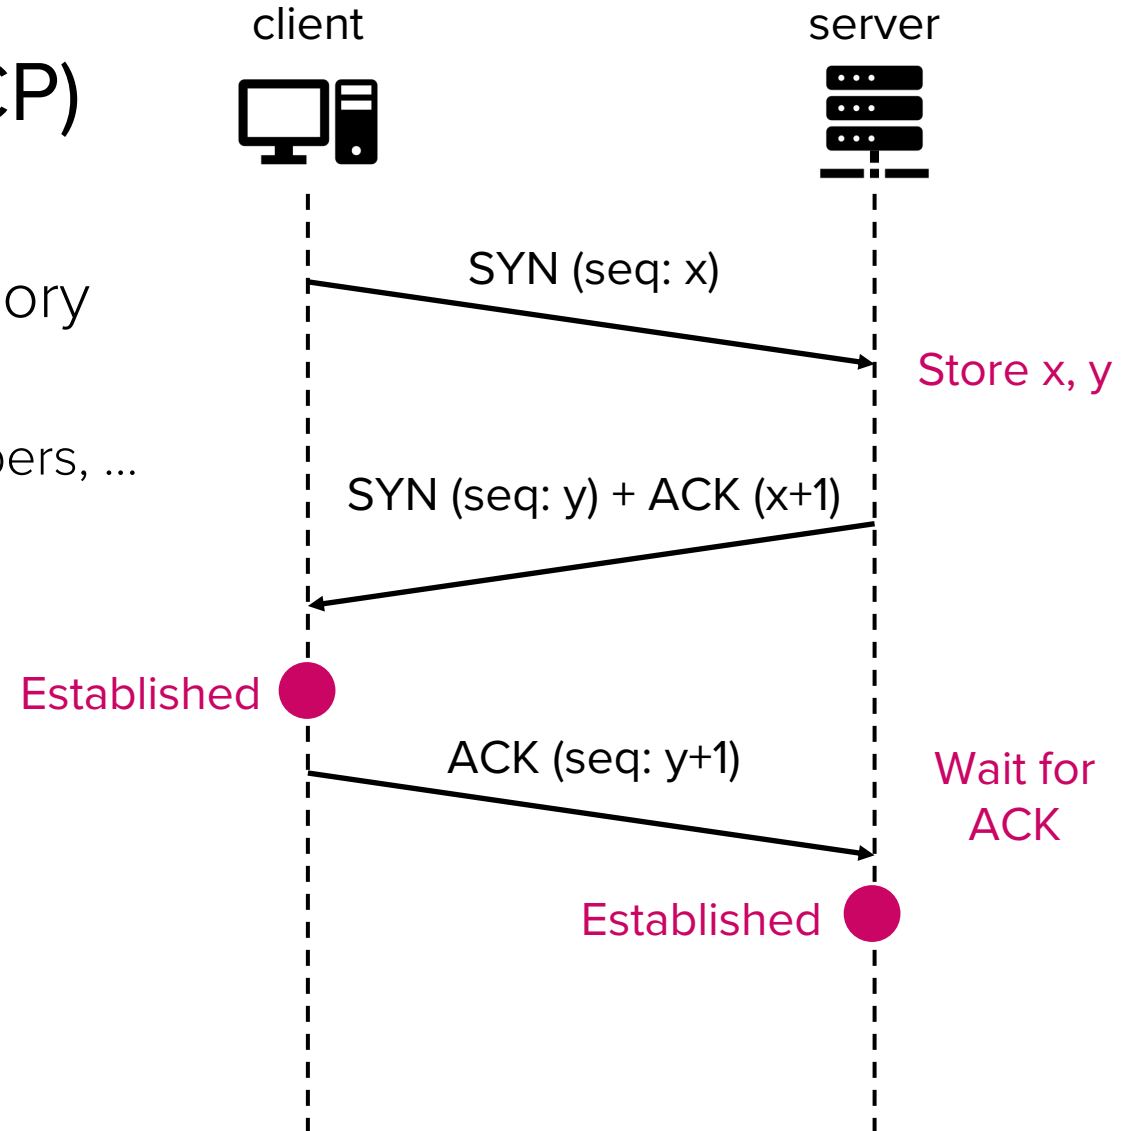
- Packet filter: Drop packets that are part of the DoS attack
  - Discard packet if source IP == attacker's IP
    - Utilize known banlist
    - Add unusual traffic IPs to the banlist
- Subverting packet filters
  - Send packets from multiple IP addresses (DDoS)
  - If attacking from one machine, spoof DoS packets to make them look like they are coming from multiple IP addresses
    - Makes packet filter ineffective



# Stateful DoS

# SYN flooding


- Transmission Control Protocol (TCP)
  - Stateful protocol (recall: Lec 13)
    - The server need to allocate some memory for each connection established
      - Store states: sequence numbers, ack numbers, ...
    - The server expects an ACK packet



# SYN flooding

- SYN flooding attack

- Attacker establishes many connections with the server
  - Cause the server to consume a lot of memory
- Advantage: The initial SYN packet is enough to get the server's resources wasted
  - Attacker does not waste its own resources



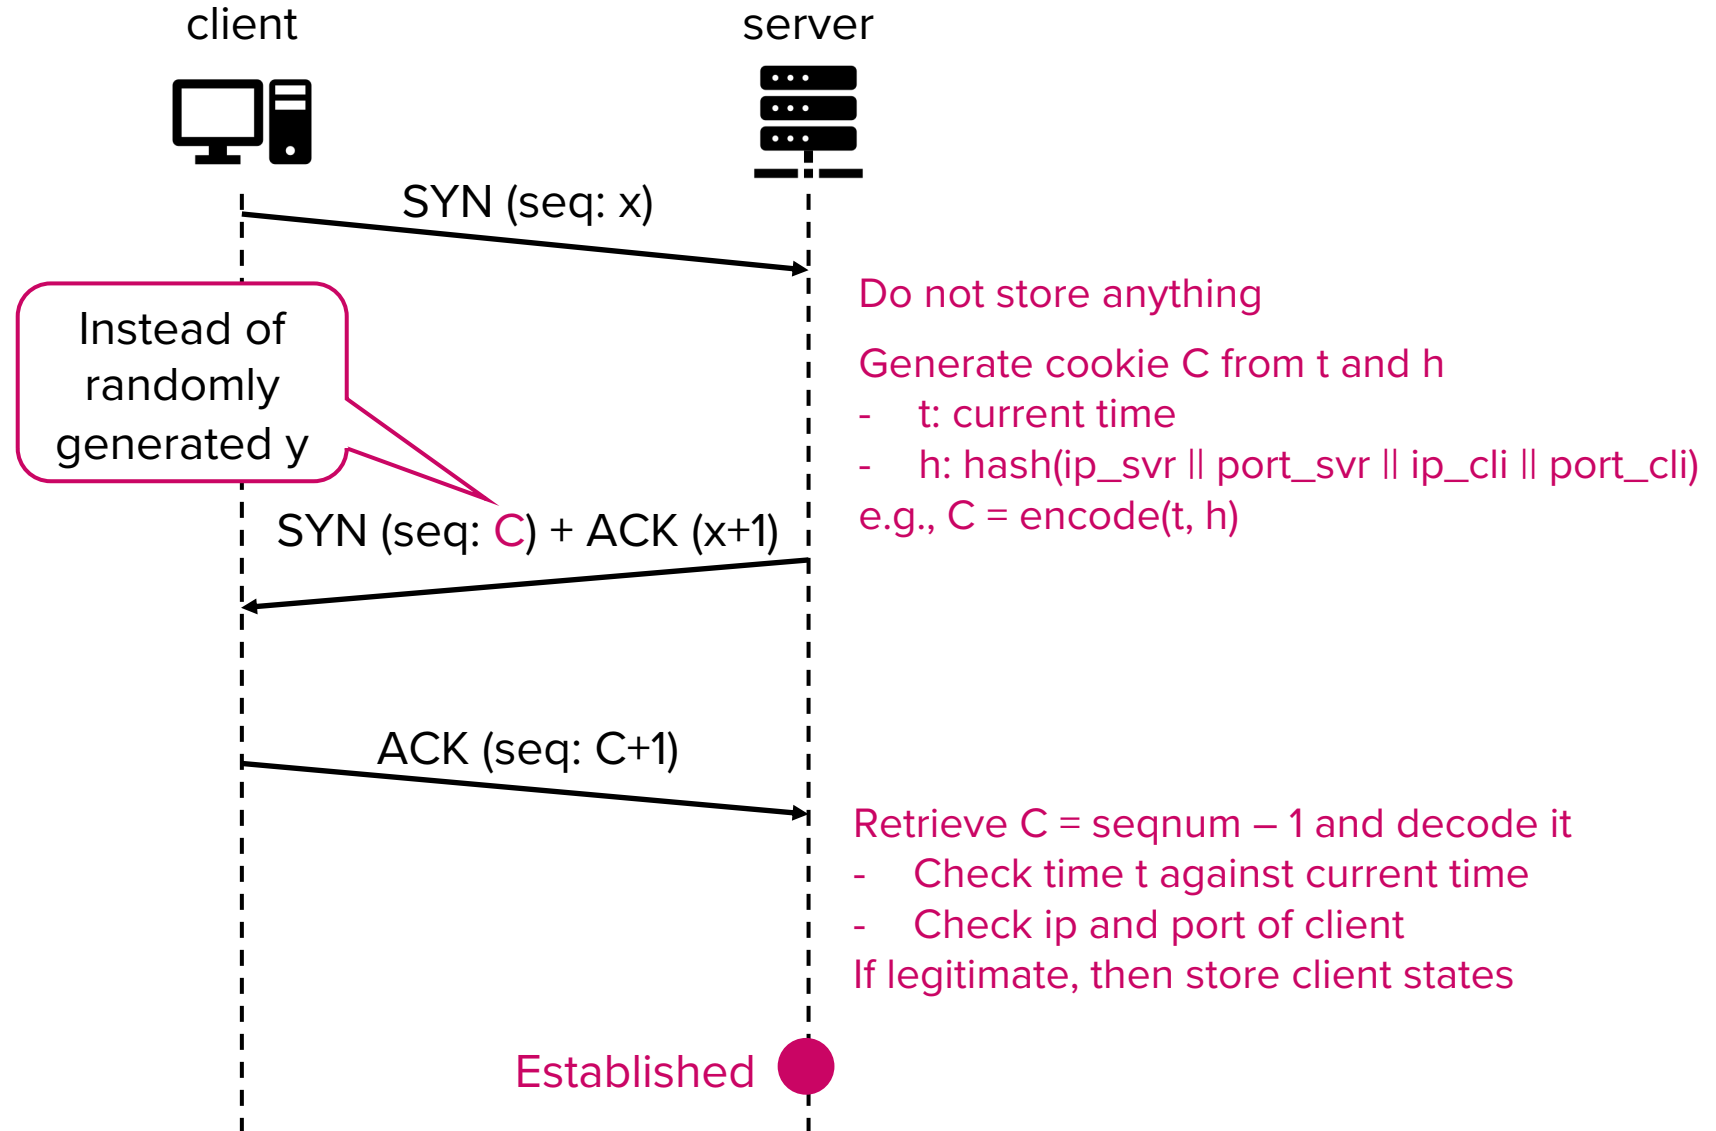
# SYN flooding – defenses

---

- Overprovisioning: Ensure the server has a plenty of memory
  - Can be expensive
- Filtering: Ensure that only legitimate connections will create state
  - Can be effective if packets are not spoofed
  - Same drawbacks as network-level DoS prevention
- SYN cookies: Do not store state in the first place

# SYN cookies

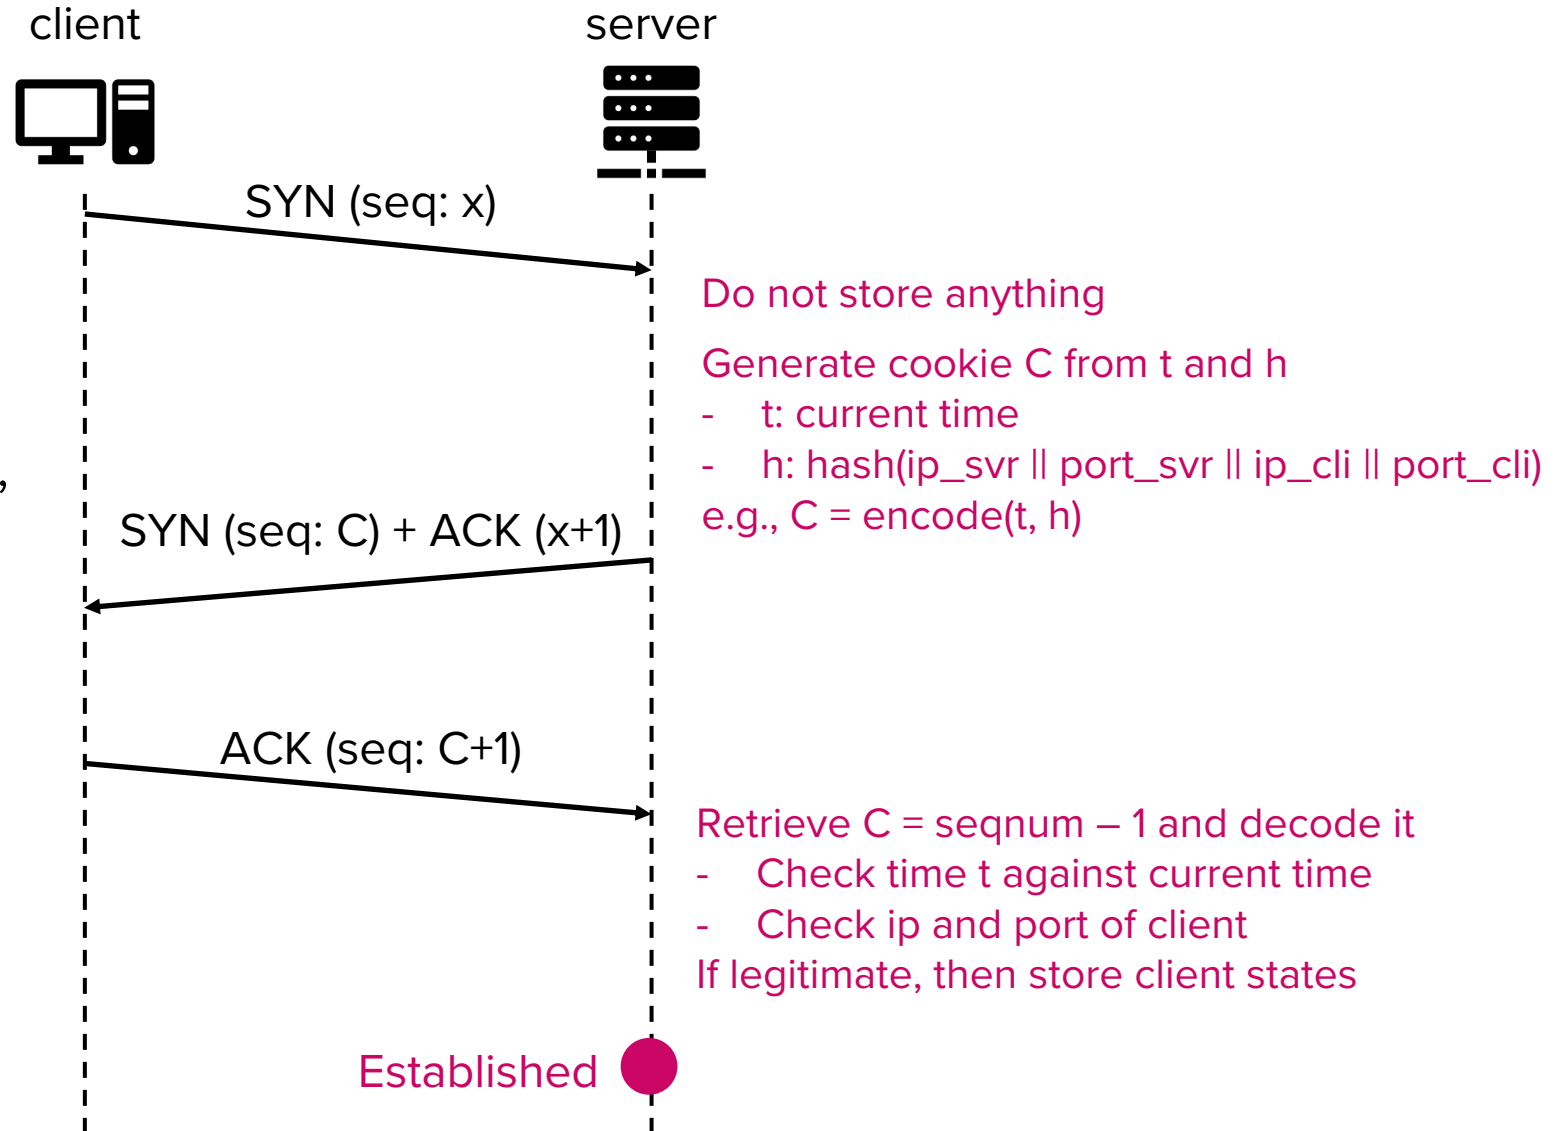
- Workflow



# SYN cookies

- Observation


- Attacker no longer has asymmetric advantage
  - Must send SYN, receive and parse resp, and send ACK



# Case study: CSED415

# CSED415 lab server


- Single-server, single-user






# CSED415 lab server

- Single-server, single-user, **malicious students are behind!**



# CSED415 lab server

- Our mitigation strategy:
  - rlimit, and control groups (cgroups) inside jail



# CSED415 lab server

- Try:
  - `while (1) { fork(); }`
  - `while (1) { malloc(0x100000000); }`
  - `$ fallocate -l 100G file`
  - `$ sha1sum /dev/zero | sha1sum /dev/zero | sha1sum /dev/zero | sha1sum /dev/zero | sha1sum /dev/zero | sha1sum /dev/zero`

# Firewalls


# Motivation: Scalable defenses

- How can we protect a set of systems against external attacks?
  - e.g., a company network with many servers and desktops
- Attack surface is very large
  - More network services == more risk
    - ssh, ftp, http, printer, ...
  - More network enabled machines == more risk
    - Different environments (OS, software, ...)
    - Different user behaviors (phishing, social engineering, ...)

We want to secure the entire network  
rather than securing individual machines


# Firewalls and security policies

- Firewall idea: Interpose a single point of access in and out of the network with a security monitor
  - Any traffic must pass through the firewall
- Security policies control network access
  - What traffic is allowed to enter the network (inbound policy)
  - What traffic is allowed to exit the network (outbound policy)



# Firewalls and security policies

- Security policies for a standard home network:
  - Outbound: **Allow all**
    - Users inside the home network can connect to any service
  - Inbound: **Allow some**
    - Allow inbound traffic in response to an outbound connection
    - Allow inbound traffic to certain trusted services (e.g., SSH)
    - **Deny all other inbound traffic**



# Firewalls and security policies

- Possible default security policies
  - Default-allow: Allow all traffic, but deny blacklisted traffic
    - As problems arise, add them to the blacklist
  - Default-deny: Deny all traffic, but allow whitelisted traffic
    - When need be, e.g., when users complain, add them to the whitelist
- Which policy is better?
  - Default-allow is more flexible, but flaws are catastrophic
  - Default-deny is more conservative, but less prone to flaws
  - Considering the “fail-safe defaults” principle, default-deny is generally deemed more secure



# Two types of firewalls

---

- Stateless firewall
- Stateful firewall

# Stateless firewall

- Implemented as stateless packet filters
- Inspect each incoming and outgoing network packet
  - All decision is made using only the information in the packet itself
    - Src IP, dst IP, src port, dst port, and protocol
  - Do not consider history
- Example
  - No inbound connection to low ports except
    - 20, 21 for FTP / 23 for Telnet / 25 for SMTP / 80 for HTTP

# Stateless firewall

- Stateless packet filtering examples
  - Allow inbound from some ports
    - Allow port 20, 21 for FTP (if running file server)
    - Allow port 22 for SSH (if running secure shell server)
    - Allow 23 for Telnet
    - Allow 25 for SMTP (if running mail server)
    - Allow 80 and 443 for HTTP, HTTPS (if running web server)
  - Allow inbound from src IP 8.8.8.8
    - Google's DNS
  - Deny all other inbound IP and ports

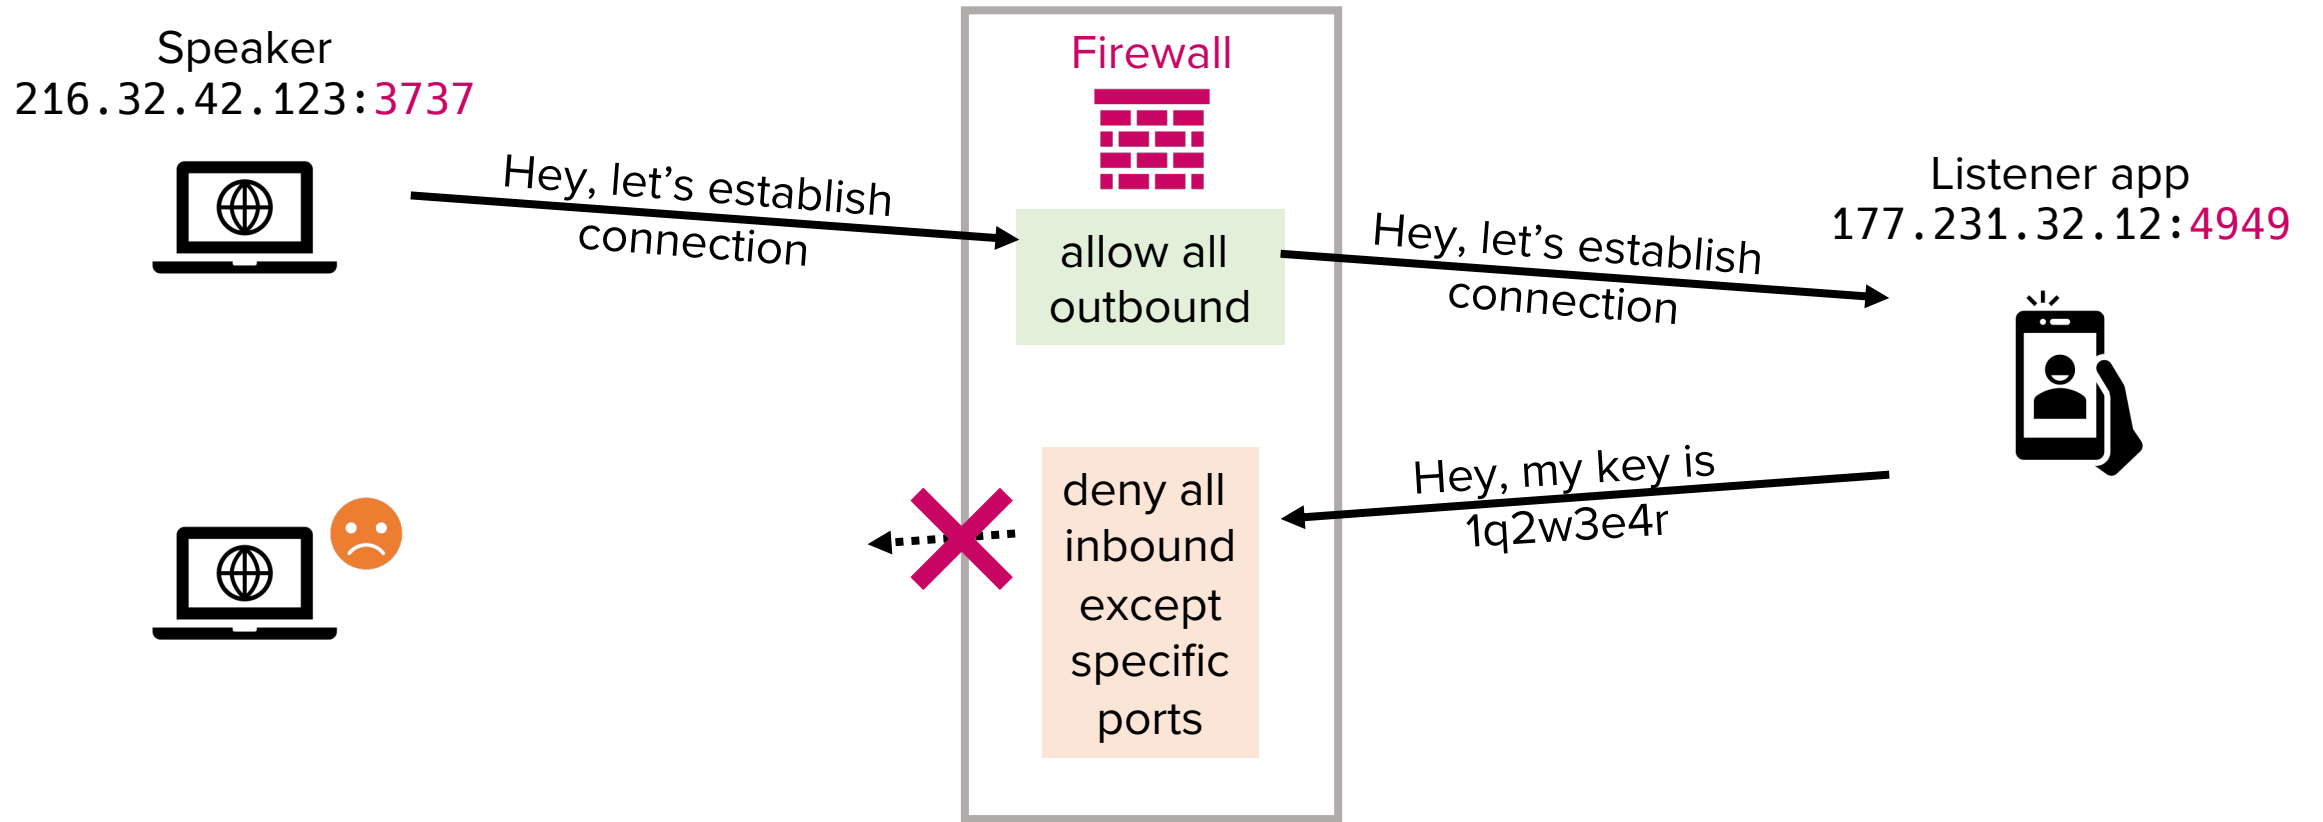
# Stateless firewall

- Applying simple packet filter to the home network example
  - Allow all outbound  Allow all
  - Allow inbound traffic to certain trusted services (e.g., SSH)  Using portnum
  - Allow inbound traffic in response to an outbound connection
  - Deny all other inbound traffic  Deny all

Issue: How do we know if inbound traffic is in response to an outbound connection without keeping track of states?

# Stateless firewall

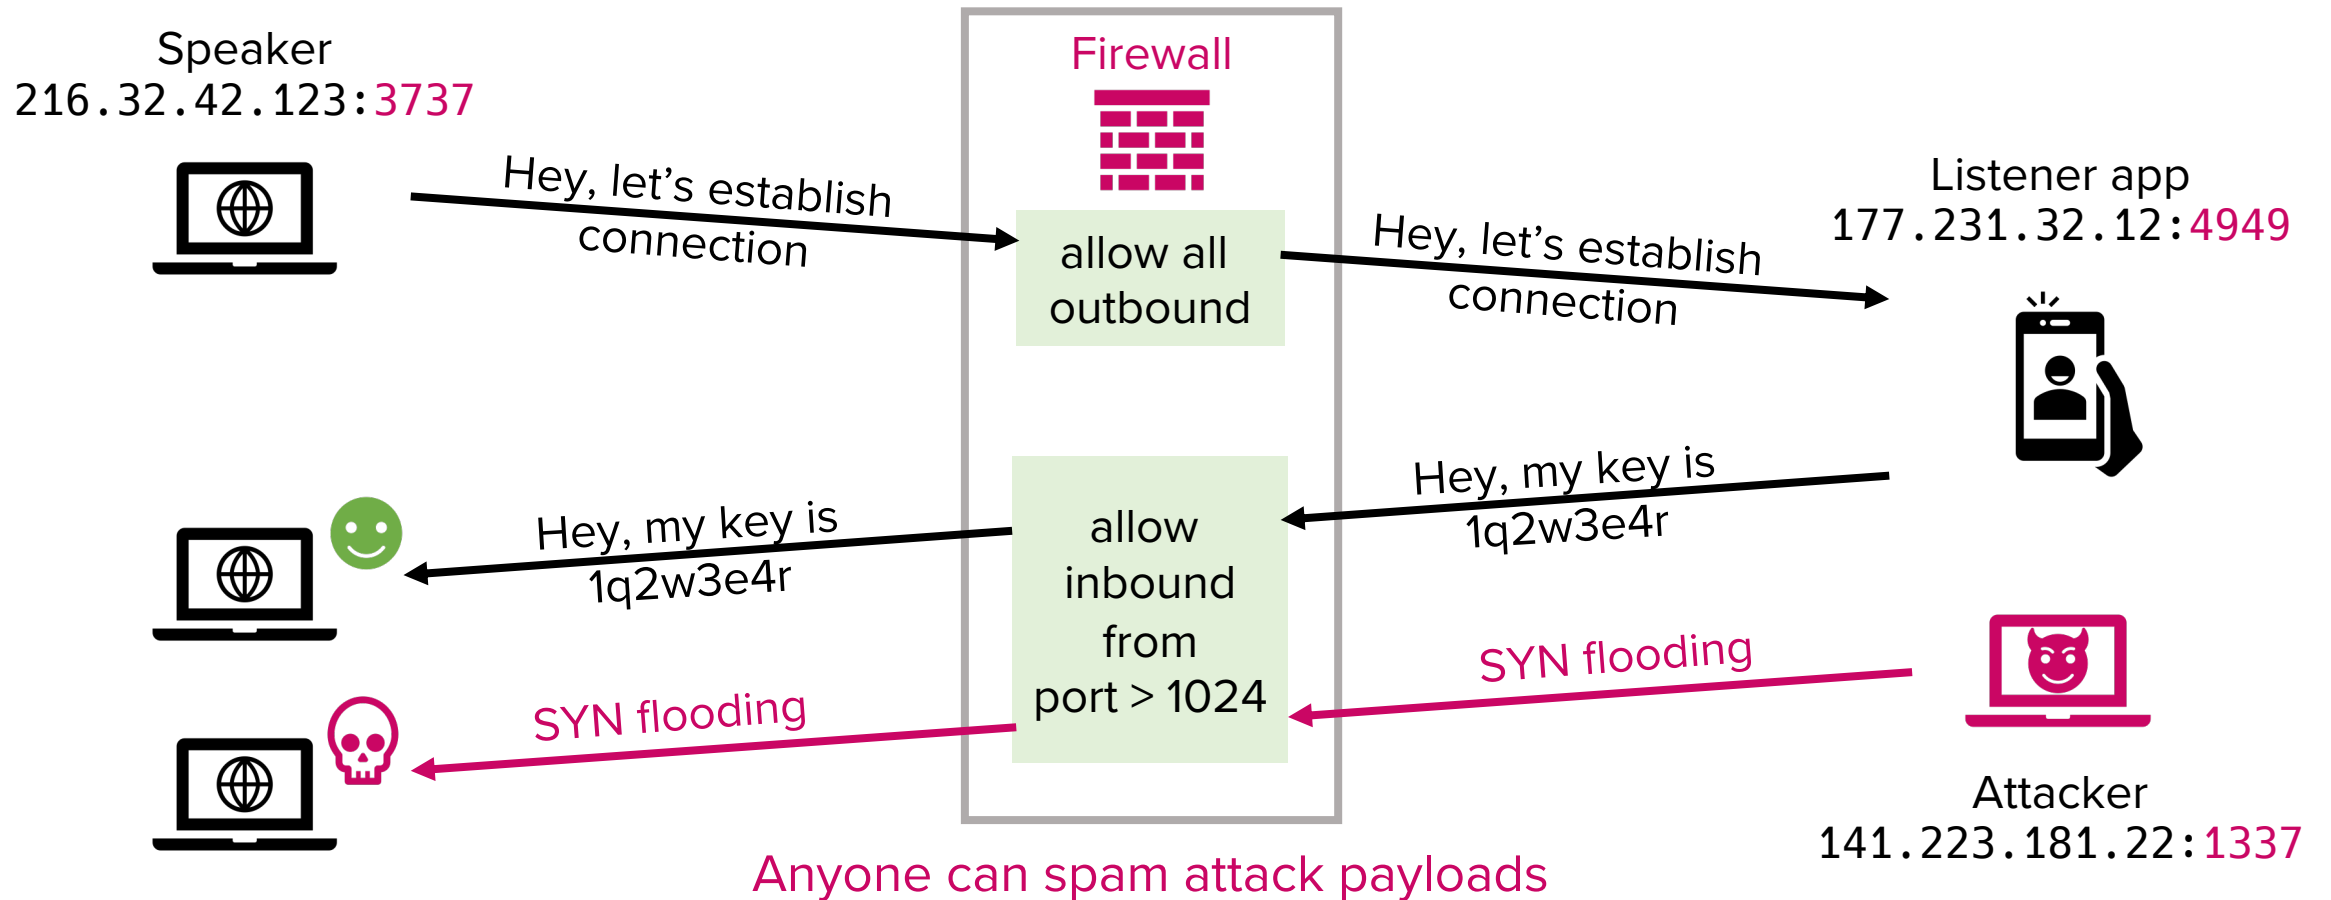
- Example: Temporary P2P connection with **strict** policy



Need to add rules, e.g., allow inbound from ports > 1024


# Stateless firewall

- Example: Temporary P2P connection with **relieved** policy



# Stateful firewall

- A better idea: Store and use full context of connections
  - Create a directory of outbound connections
  - Allow packets from the destination of the recorded outbound connections
  - Use connection information along with strict stateless rules



# Stateless vs Stateful firewalls

---

- Stateless
  - Fast and scalable to large volumes of traffic
  - Cheaper than stateful firewalls
  - Less secure
- Stateful
  - More secure and versatile
    - Don't need to allow a range of ports
    - Dynamically configured
  - Slower and more expensive than stateless firewalls



# Firewall pros and cons

---

- **Pros**

- Centralized management of security policies (single point of control)
- Transparent operation to end users
- Mitigates security vulnerabilities

- **Cons**

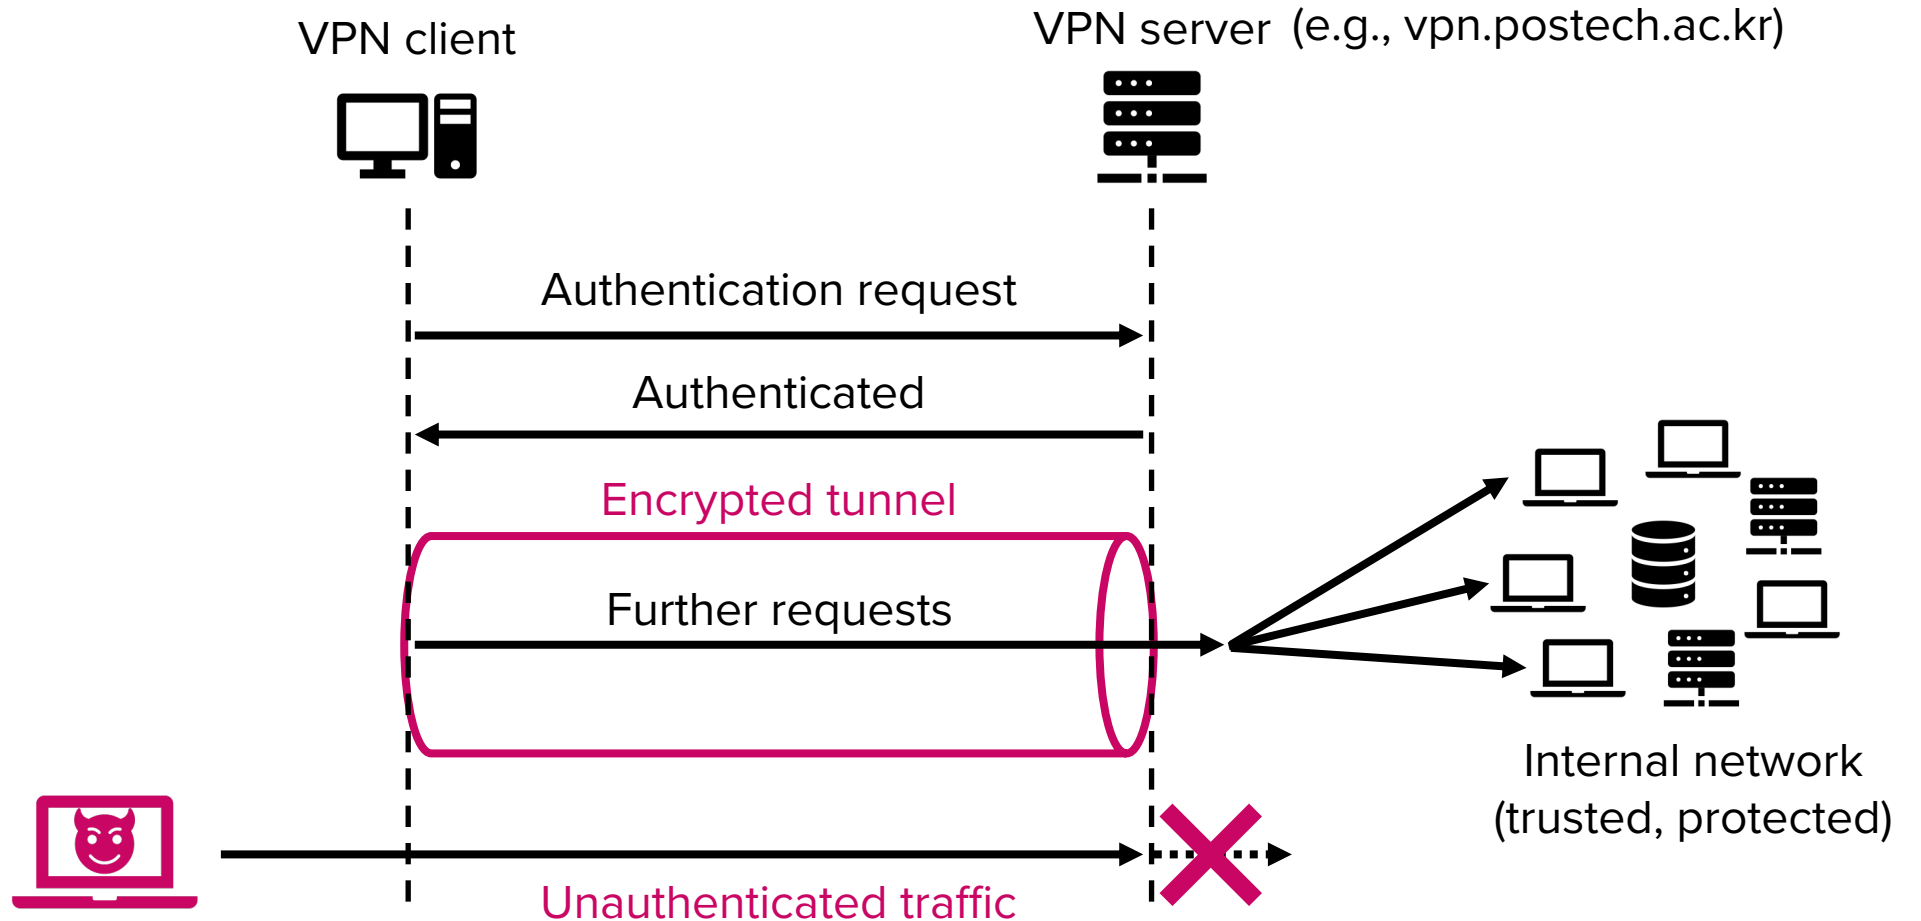
- Reduced network connectivity
- Vulnerable to insider attacks
  - Employees can be bribed or threatened
  - Untrusted devices are brought into the network (e.g., employee laptops)

# Alternative for firewalls

- Virtual private network (VPN)
  - A set of protocols that allows direct access to an internal network from external entities
  - VPN server authenticates user (VPN client)
  - Provide encrypted VPN tunnel
    - Only authenticated clients can send traffic via the tunnel
    - VPN server routes the traffic to individual services

# VPN for perimeter security

- VPN workflow



# Summary

- Availability: Making sure users are able to use a service
- DoS: Attack availability of services
  - Application-level DoS
  - Network-level DoS
  - Stateful DoS
  - Overprovisioning can mitigate DoS but it is very expensive
- Firewalls: Perimeter security to defend a network
  - Utilize stateless and/or stateful packet filters

# Questions?