CryptoWorks21

Fundamentals of Network Security

4. Offensive and defensive network security

CryptoWorks21 • July 15, 2021

Dr Douglas Stebila

UNIVERSITY OF WATERLOO

https://www.douglas.stebila.ca/teaching/cryptoworks21

Fundamentals of Network Security

- Basics of Information Security
  - Security architecture and infrastructure; security goals (confidentiality, integrity, availability, and authenticity); threats/vulnerabilities/attacks; risk management
- Cryptographic Building Blocks
  - Symmetric crypto: ciphers (stream, block), hash functions, message authentication codes, pseudorandom functions
  - Public key crypto: public key encryption, digital signatures, key agreement
- Network Security Protocols & Standards
  - Overview of networking and PKI
  - Transport Layer Security (TLS) protocol
  - Overview: SSH, IPsec, Wireless (Tool: Wireshark)
- **Offensive and defensive network security**
  - **Offensive: Pen-tester/attack sequence: reconnaissance; gaining access; maintaining access (Tool: nmap)**
    - **Supplemental material: denial of service attacks**
  - **Defensive: Firewalls and intrusion detection**
- Access Control & Authentication; Web Application Security
  - Access control: discretionary/mandatory/role-based; phases
  - Authentication: something you know/have/are/somewhere you are
  - Web security: cookies, SQL injection
  - Supplemental material: Passwords

Assignment 2

2a) Offensive network security

- Use nmap to scan services running on your computer
  - Will be scanning from guest Kali Linux virtual machine to host machine using a simulated network

2b) Defensive network security

- Set up firewall rules in your Kali to prevent certain types of outbound traffic (egress filtering)

Assignment 0
Downloading and installing VirtualBox and Kali Linux

https://www.douglas.stebila.ca/teaching/cryptoworks21/

OFFENSIVE NETWORK SECURITY: PENETRATION TESTING

Types of hackers

- **White hat:** breaks security for non-malicious reasons, for example while working with a client
  - **"Ethical hacker"**
  - **"Penetration testing"**
- **Black hat:** breaks security for malicious reasons or for personal/commercial gain
- **Grey hat:** breaks security for mostly non-malicious reasons, but often without permission

Penetration Testing ("pen testing")

- Authorized attack on a (sometimes simulated) computer system that looks for security weaknesses.

- Could be illegal in some contexts without permission.

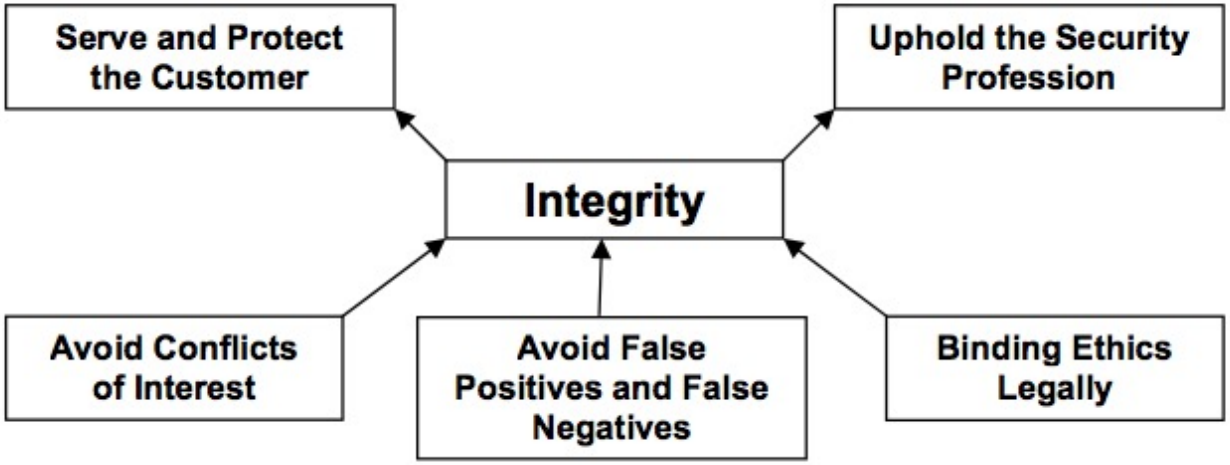

Penetration Testing Professional Ethics: a conceptual model and taxonomy.
http://dl.acs.org.au/index.php/ajis/article/view/52/39

Pen testing

- Only do penetration testing with express authorization (written consent).
- Work on an isolated network to avoid affecting legitimate users (unless working on a real network is part of the testing, in which case obtain permission from the network operator, and take all steps possible to avoid damaging legitimate users).
- Avoid collecting personal information.
- Notify immediately of any severe vulnerabilities that could endanger human life.
- Results of social engineering should be delivered in summarized, statistical form to avoid implicating individuals.
- Maintain confidentiality of the results with your client.

Ethics and Legality

- Don’t try this out on Waterloo systems.
- Don’t try this out on Learn.
- Don’t try this out on Google/Microsoft/Facebook/Apple/….
- Don’t try this out on my computer.

- Use virtualization and isolated networks wherever possible.

"Certified Ethical Hacker"

THE LEARNING TRACK

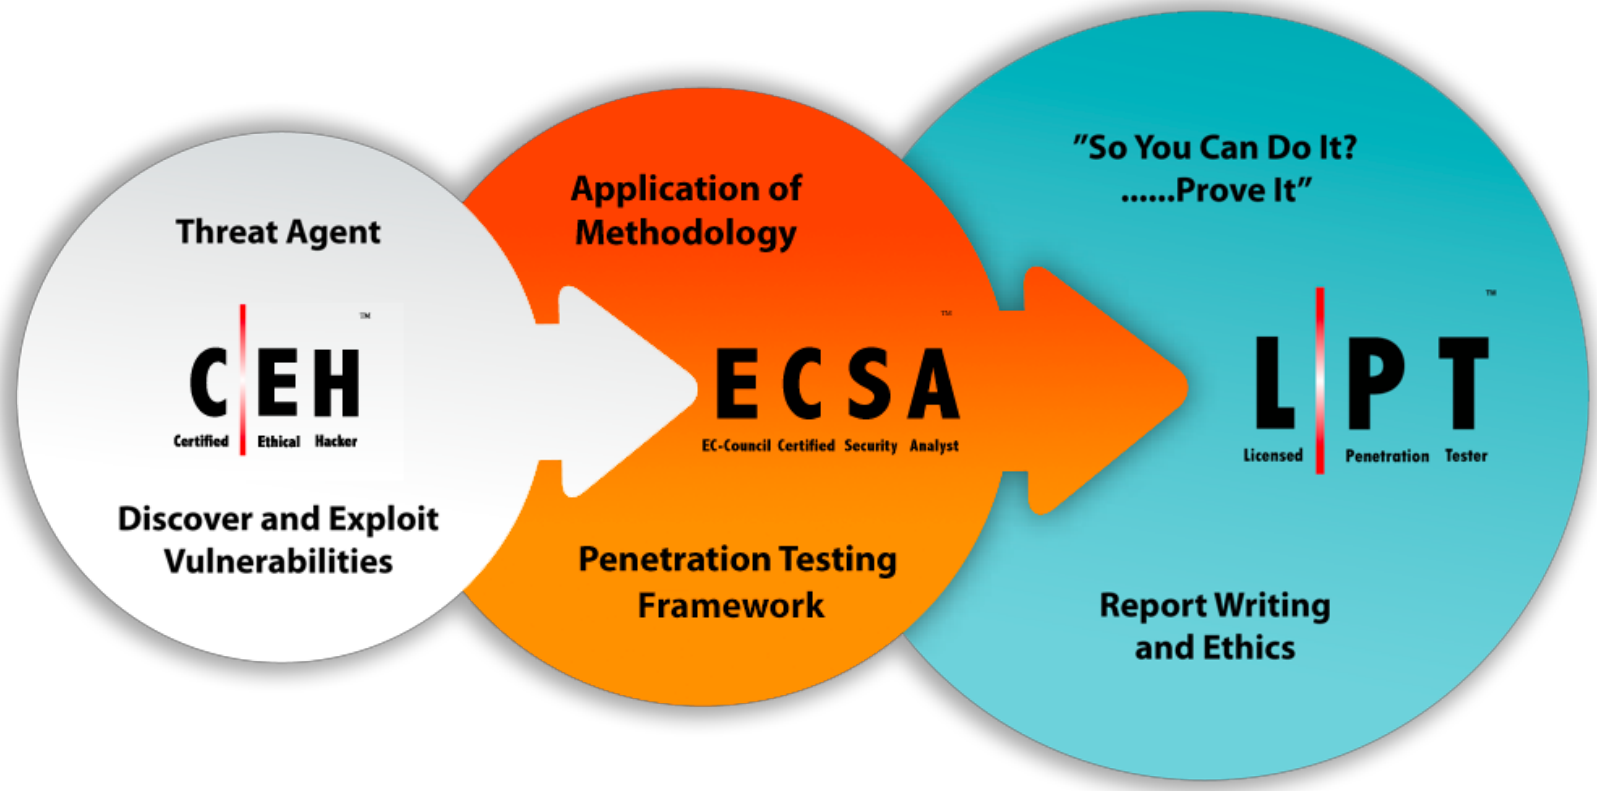

Knowledge Skill Ability

https://www.eccouncil.org/programs/certified-ethical-hacker-ceh/

Pen-tester / attack sequence

- General sequence of operations followed by an attacker to gain access to a system

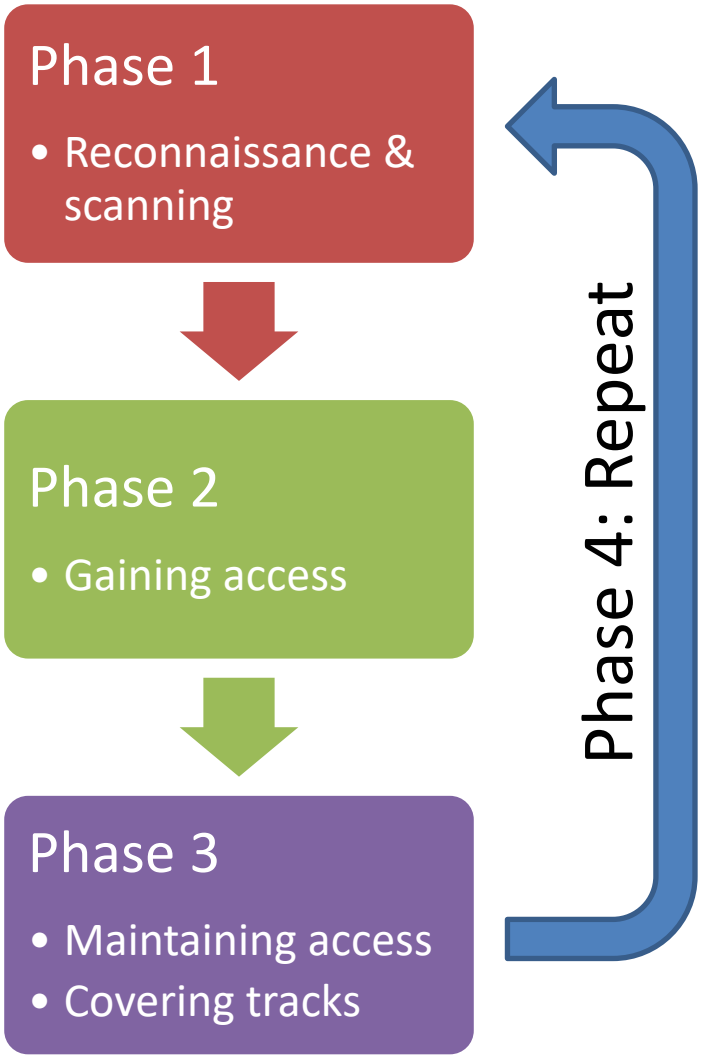

Phase 1: Reconnaissance & scanning

Goal: Gain information about the victim's network and configuration.

Passive reconnaissance

- Gather information about victim's network without victim's knowledge

Active reconnaissance

- Gather information by probing the victim's network
- Possibly detectable

Outcomes from Phase 1

- Network information
  - External IP addresses, domain names
  - Internal IP addresses, private/testing websites
  - Firewalls & intrusion protection configuration
  - VPN gateways
- Operating system information
  - Versions
  - Computer names and purposes
  - Users and groups
- Organization information
  - Locations
  - Key employees
  - Contact information (email, phone numbers)
  - Security policies

Passive reconnaissance

- Eavesdropping wireless network connections
  - Packet sniffing: `wireshark`
- Dumpster diving
- Search engines
- Research network configuration
- Social engineering
- Watching employees from the parking lot

Passive reconnaissance

- Publicly available network information
  - Domain owner: `whois` command, https://whois.net/
  - Servers for a domain name: `nslookup`
  - Network routes: `traceroute`
  - Assigned IP addresses:
    - https://www.iana.org/numbers
    - http://whois.arin.net/rest/ip/130.113.64.65
  - Databases of known services:
    - https://dnsdumpster.com/
    - https://www.shodan.io/
  - Security assessment tools:
    - https://www.ssllabs.com/ssltest/
  - `dmitry` command

Active reconnaissance

- Probing the victim's network

1. Determine which hosts are online
   - IP scanners (nmap, zmap, …)
2. Determine which services are active on which hosts
   - Port scanners (nmap, …)
3. Scan services for vulnerabilities
   - Vulnerability scanners (Nessus, …)

Active reconnaissance

- Probing the victim's network

This is active interaction with the target, and should not be done outside of a test environment without a written agreement with the target!

  - Port scanners (nmap, …)

3. Scan services for vulnerabilities
  - Vulnerability scanners (Nessus, …)

Active reconnaissance
1. Determine which hosts are online

a) Use list of assigned IP addresses from passive reconnaissance
b) Scan IP addresses to determine which hosts respond to network requests
    a) ping, hping3
    b) nmap
    c) zmap

Active reconnaissance
2. Determine which services are active

("Port scanning")

a) Try to connect to network services on each live IP address
    a) nmap
b) Check common (and uncommon) TCP and UDP ports

nmap

Most popular port scanner available

Offers many different scanning techniques:

- Scan for hosts that are up
- TCP ports
- UDP ports
- Other IP Protocols

Can identify software, version, some configuration details

nmap -A -T4 127.0.0.1

```
Starting Nmap 7.70 ( https://nmap.org ) at 2019-08-05 11:25 EDT
Nmap scan report for localhost (127.0.0.1)
Host is up (0.00021s latency).
Not shown: 992 closed ports
PORT     STATE SERVICE       VERSION
22/tcp   open  ssh           OpenSSH 7.9 (protocol 2.0)
| ssh-hostkey:
|   2048 04:fb:61:13:ec:cf:9b:6e:6c:84:6b:7c:e8:9f:97:9e (RSA)
|   256 1d:8b:0c:6b:f2:bf:79:f7:bc:f7:61:b5:e3:17:ca:8c (ECDSA)
|_  256 91:d4:be:be:25:ed:ba:31:e8:68:da:23:64:72:a6:1c (ED25519)
88/tcp   open  kerberos-sec  Heimdal Kerberos (server time: 2019-08-05 15:25:19Z)
445/tcp  open  microsoft-ds?
631/tcp  open  ipp           CUPS 2.2
|_http-title: Home - CUPS 2.2.9
3306/tcp open  mysql?
| mysql-info:
|   Protocol: 10
|   Version: 8.0.16
|   Thread ID: 226
|   Capabilities flags: 65535
|   Some Capabilities: Support41Auth, LongPassword, SupportsCompression, InteractiveClient,
DontAllowDatabaseTableColumn, Speaks41ProtocolOld, SwitchToSSLAfterHandshake, SupportsTransactions,
LongColumnFlag, IgnoreSigpipes, Speaks41ProtocolNew, ODBCClient, IgnoreSpaceBeforeParenthesis,
ConnectWithDatabase, FoundRows, SupportsLoadDataLocal, SupportsMultipleStatments,
SupportsMultipleResults, SupportsAuthPlugins
|   Status: Autocommit
|   Salt: fC}\x7F\x18\x07Ju\\x17#\x12#\x06GArDV\x0C
|_  Auth Plugin Name: 79
3689/tcp open  daap          Apple iTunes DAAP 12.9.5.5
8080/tcp open  http          Apache httpd 2.4.39 ((Unix) PHP/7.3.7)
|_http-open-proxy: Proxy might be redirecting requests
|_http-server-header: Apache/2.4.39 (Unix) PHP/7.3.7
8888/tcp open  http          Apache httpd 2.4.39 ((Unix) PHP/7.3.7)
|_http-server-header: Apache/2.4.39 (Unix) PHP/7.3.7
Service Info: OS: OS X
```

Vulnerability Assessment Tools

Collection of tools for determining possible security holes

Port-scanning + additional checks on ports for:

- Software packages actually running
- Versions of those packages
- Crosscheck vulnerability databases to identify possible vulnerabilities on these versions
- Possibly other components
  - Check for weak passwords
  - Check for general patch levels

Example

- Port scanning may find port 21 listening, ftp
- OS fingerprint – Linux 2.2 kernel
- Service query – identifies ftp as wu-ftpd version 2.4.2
- What specific vulnerabilities does wu-ftpd 2.4.2 have?

OpenVAS

- OpenVAS: Open Vulnerability Assessment System
- In Kali:
  - Need to install:
    - https://www.kali.org/penetration-testing/openvas-vulnerability-scanning/
- Or download separate VirtualBox VM:
    - http://www.openvas.org/vm.html

OpenVAS

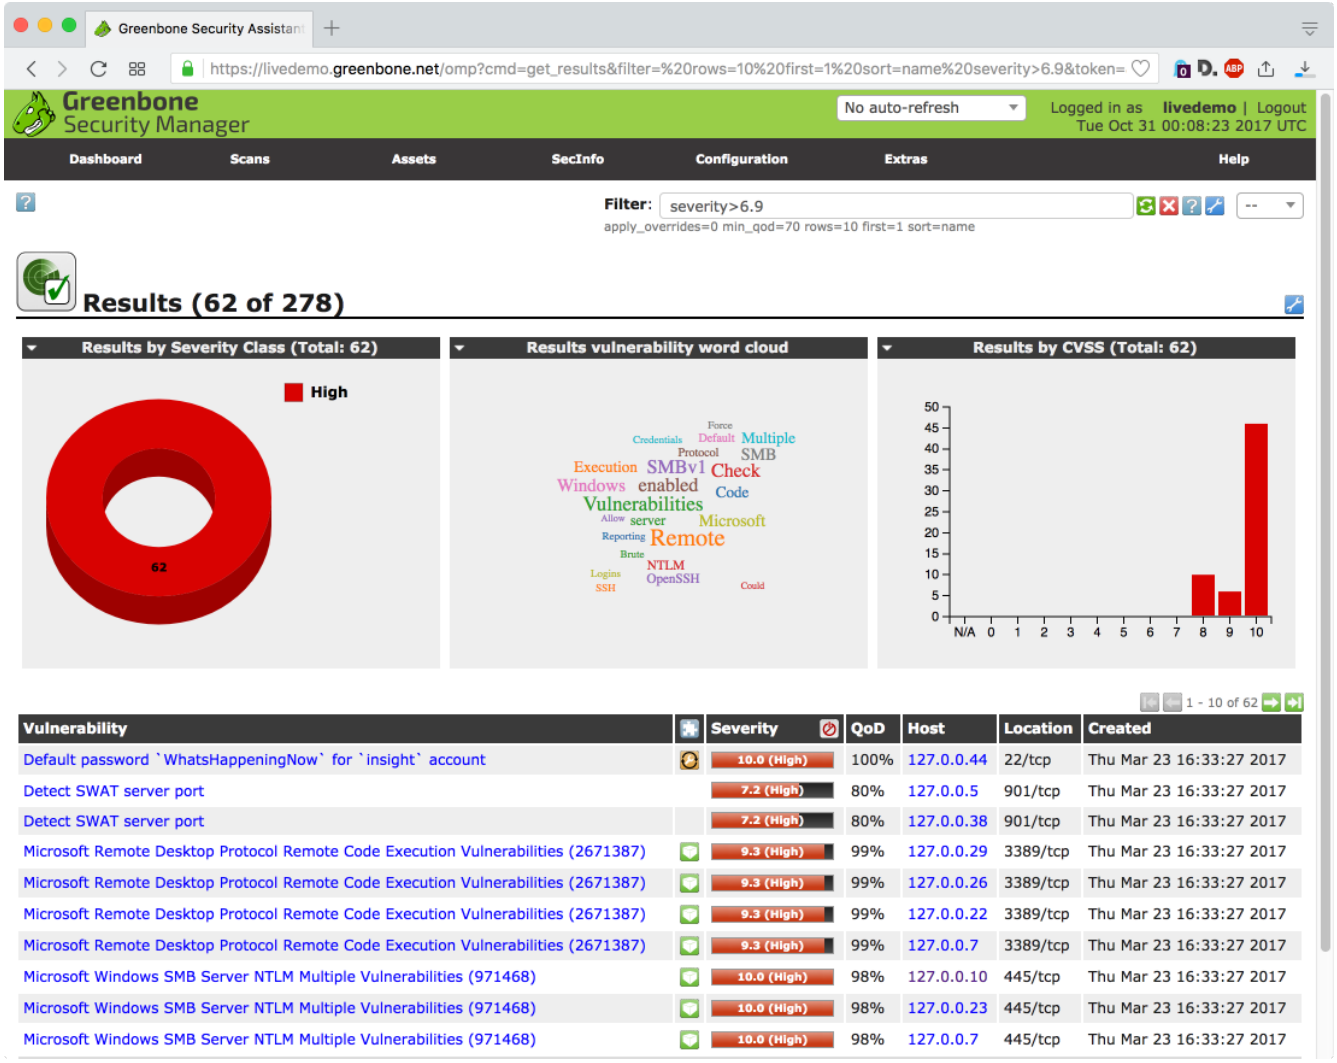

https://livedemo.greenbone.net/

Nessus Vulnerability Report

Commercial product.

Identifies active services and their versions, matches against database of known vulnerabilities, then tests for exploitability using plugins.

192.168.150.100

Scan Information

Start time:	Wed Mar 21 14:40:36 2012
End time:	Wed Mar 21 15:01:02 2012

Host Information

Netbios Name:	WINDOWS2000
IP:	192.168.150.100
MAC Address:	00:0c:29:f7:55:ea
OS:	Microsoft Windows 2000 Service Pack 4

Results Summary

Critical	High	Medium	Low	Info	Total
14	2	5	0	3	24

Results Details

0/tcp

12213 - TCP/IP Sequence Prediction Blind Reset Spoofing DoS

Synopsis

It may be possible to send spoofed RST packets to the remote system.

Description

The remote host might be vulnerable to a sequence number approximation bug, which may allow an attacker to send spoofed RST packets to the remote host and close established connections. This may cause problems for some dedicated services (BGP, a VPN over TCP, etc...).

Solution

See http://www.securityfocus.com/bid/10183/solution/

Risk Factor

Medium

CVSS Base Score

5.0 (CVSS2#AV:N/AC:L/Au:N/C:N/I:N/A:P)

CVSS Temporal Score

4.1 (CVSS2#AV:N/AC:L/Au:N/C:N/I:N/A:P)

References

BID	10183
CVE	CVE-2004-0230
XREF	OSVDB:4030
XREF	IAVA:2004-A-0007

Ports

tcp/0

25/tcp

45517 - MS10-024: Vulnerabilities in Microsoft Exchange and Windows SMTP Service Could Allow Denial of Service (981832) (uncredentialed check)

Synopsis

The remote mail server may be affected by multiple vulnerabilities.

Description

Pen-tester / attack sequence

- General sequence of operations followed by an attacker to gain access to a system

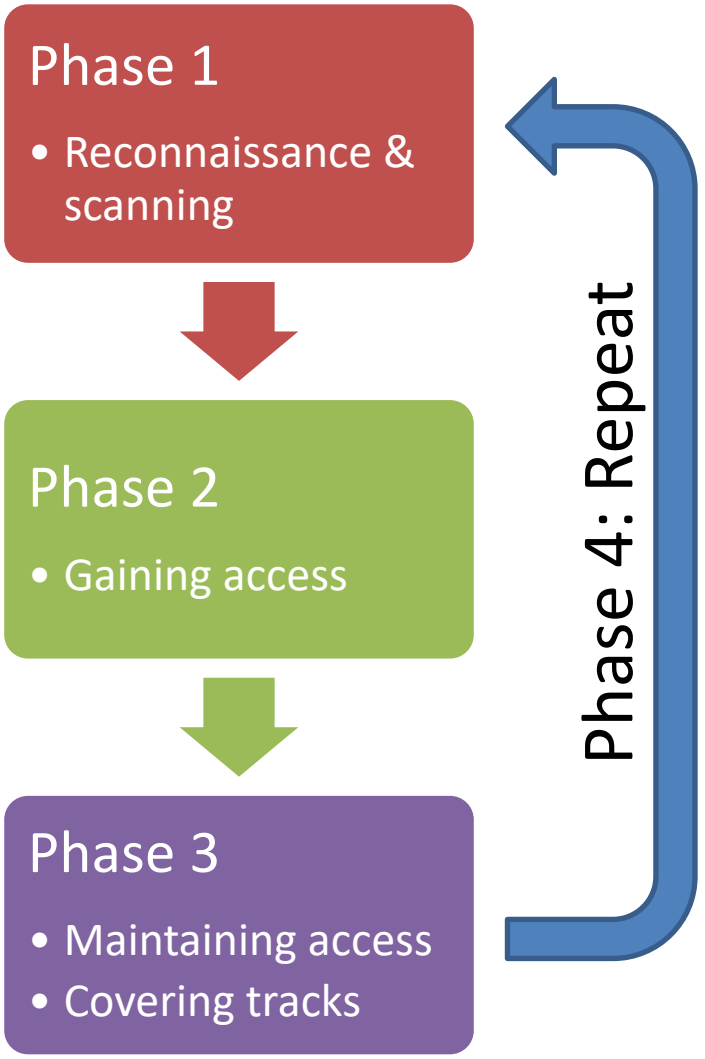

Phase 2: Gaining access

Goal: Be able to login to a system, and ideally obtain root/admin privileges

2.a) Gaining basic access

- Via legitimate user credentials
- Via an exploit

2.b) Privilege escalation

- Obtain root/admin privileges

2.a) Gaining access … via legitimate user credentials

Goal: obtain username & password of legitimate user

- Approaches:
  - Social engineering, phishing
    - Target key users identified in phase 1
  - Password breaches
    - Search password breach data on dark web
  - Guessing
    - Automated tools for remote login: ncrack

Remote password guessing

- **Ncrack**: high speed network authentication cracking tool

- Password guessing against remote servers
  - ftp, ssh, http, email, telnet, Windows file sharing, …
  - Provide possible usernames
  - Provide password dictionary
    - E.g. https://wiki.skullsecurity.org/Passwords
  - Tune rate to avoid triggering server lockout

2.a) Gaining access … via an exploit

Goal: get command-line access via an exploit in an installed program/service

- Make use of results of vulnerability scan from phase 1
- Automated tools
  - Metasploit framework

Metasploit workflow

1. Choose and configure an **exploit** for the victim system
   - **Exploit**: vector for penetrating the system
   - Metasploit contains 1600+ exploits for bugs in Windows, Unix, Linux, macOS

2. Choose and configure the **payload**
   - **Payload**: Code to be executed on the victim system
   - Metasploit contains 450+ payloads
   - Often start up a remote command-line shell / GUI server

Metasploit workflow

3. Choose an **encoding technique** to avoid detection by an intrusion protection system / anti-virus

4. **Execute** the exploit
   - Session: connection obtained from successful exploit

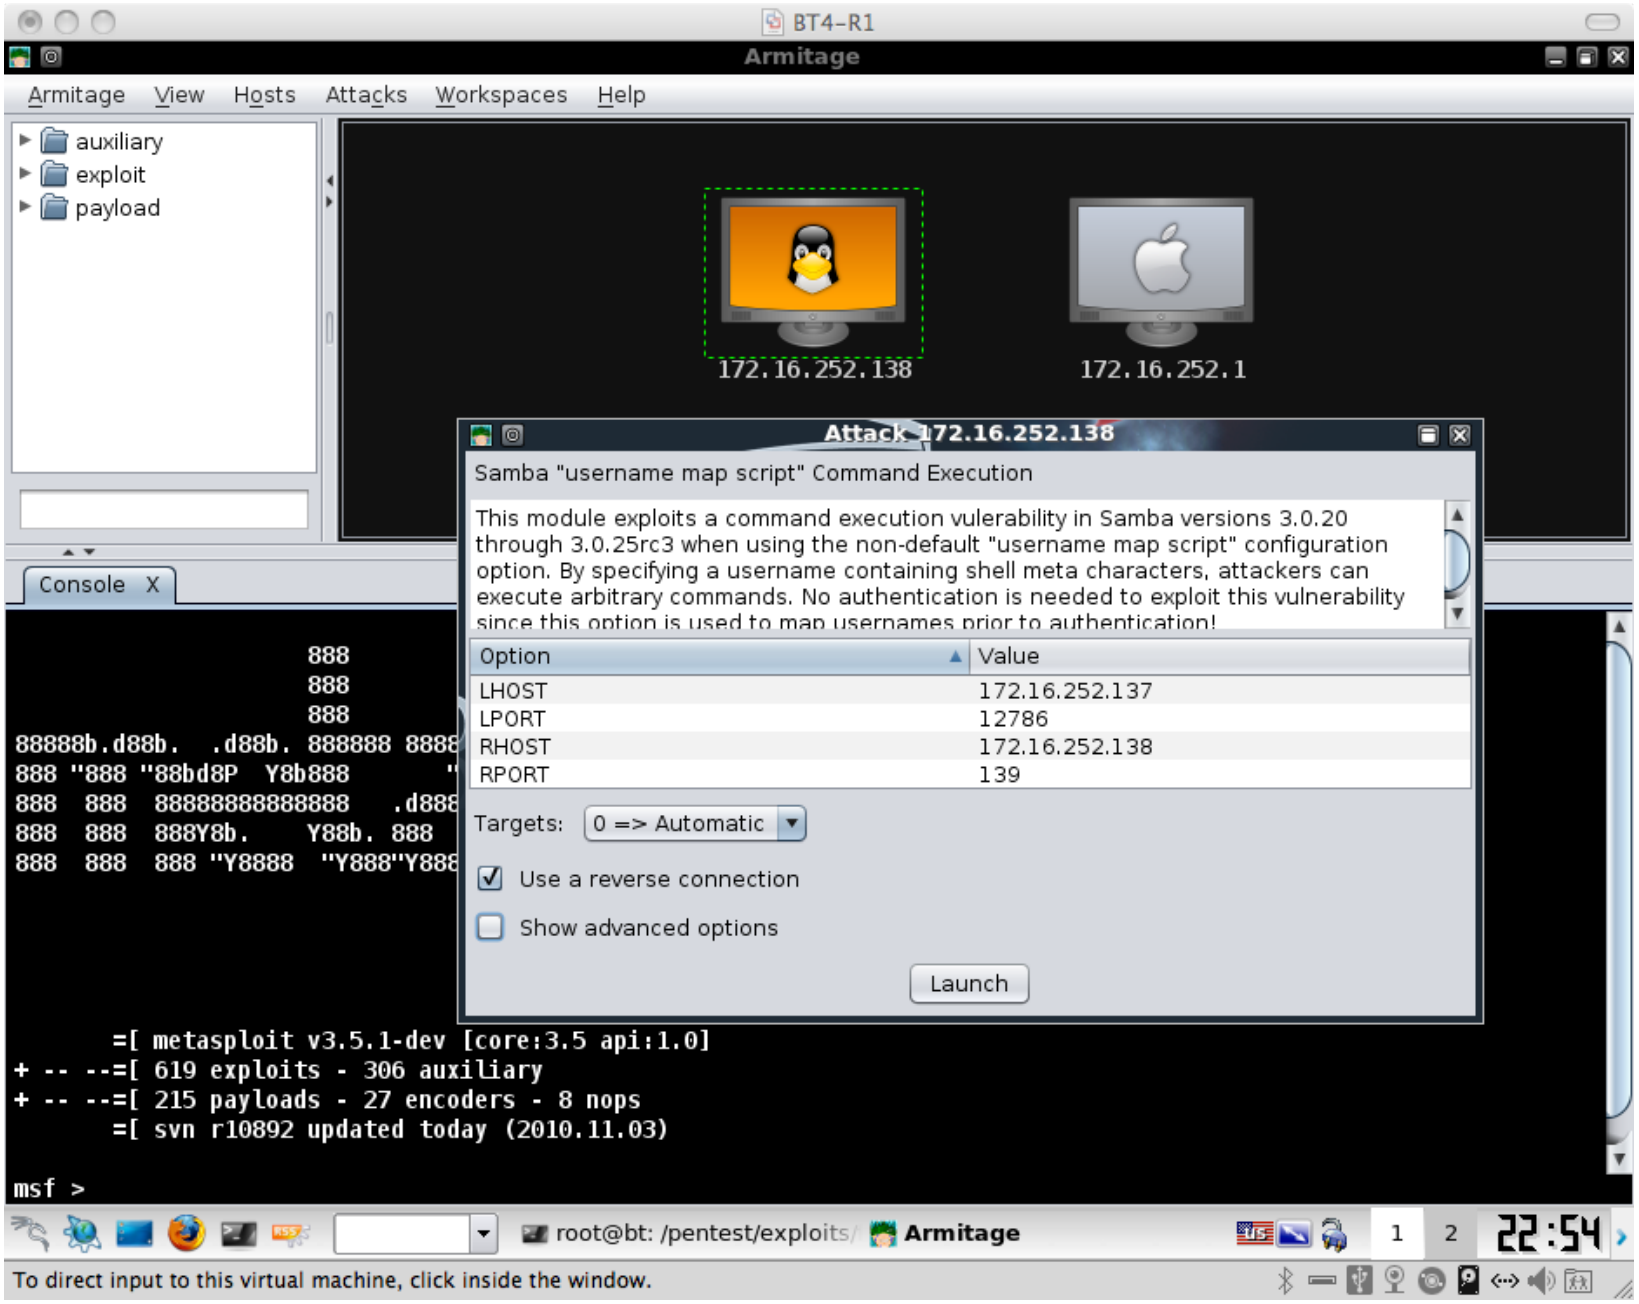

http://www.fastandeasyhacking.com/images/screenshots/launch.png

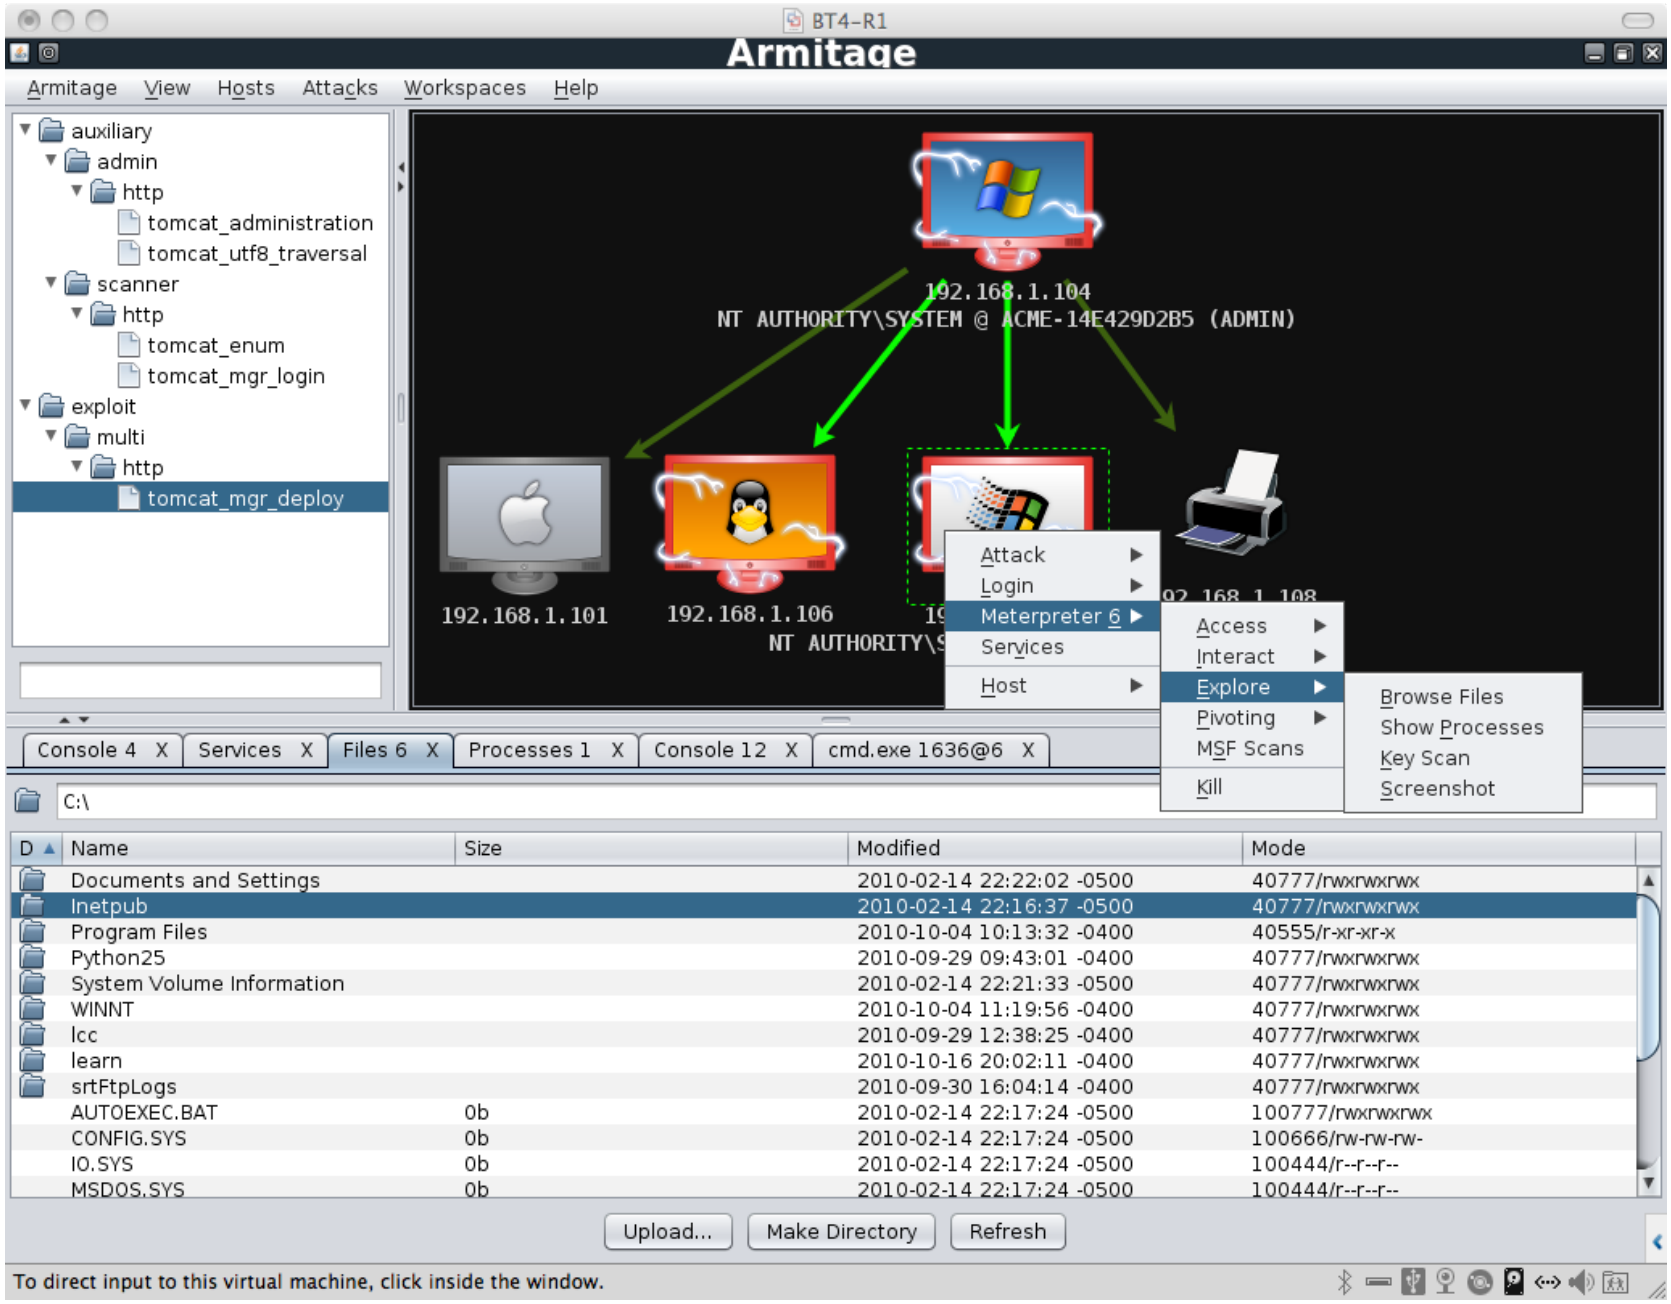

http://www.fastandeasyhacking.com/images/screenshots/armitage4.png

2.b) Privilege escalation

- Often gain access to a basic user account
- Want to obtain root/admin privileges

- Apply same basic approach for originally gaining access to get root/admin, but now with extra abilities because you can run code locally, not just rely on network services:
  - Via legitimate user credentials
    - Extra abilities:
      - Get local password hashes and cracking
  - Via an exploit
    - Extra abilities:
      - Exploit OS vulnerabilities
      - Exploit filesystem permission mistakes

2.b) Privilege escalation

Horizontal privilege escalation

- Attempt to gain access to account of another user who has roughly the same privileges as the current account

Vertical privilege escalation

- Attempt to elevate privileges of the account
- Or attempt to gain access to a higher-privileged account

Pen-tester / attack sequence

- General sequence of operations followed by an attacker to gain access to a system

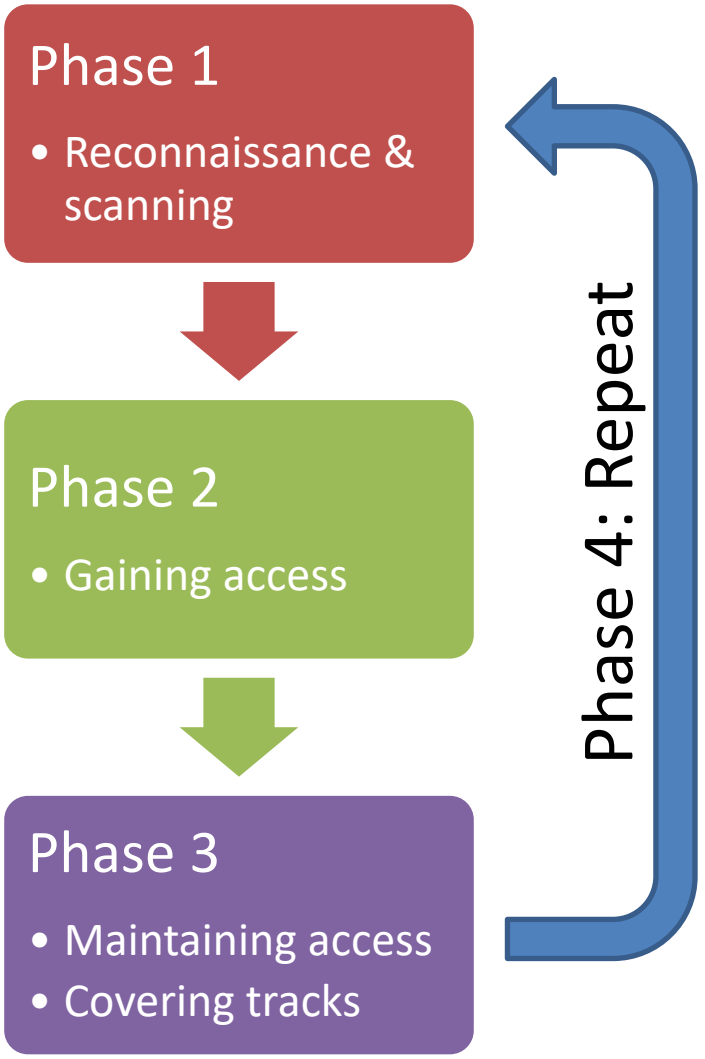

3) Maintaining access & covering your tracks

Maintaining access

- Don't want to have to go through complicated exploitation again
  - Especially in case vulnerabilities get patched
- Plant a backdoor to be able to obtain access again later
  - Enable direct remote access
  - Or install software that periodically queries a "command-and-control" server

Covering your tracks

- Disable auditing / logging
- Hide files
- Cloaking: Replace standard monitoring programs to hide presence of backdoor
  - Replace Windows task manager / process monitor / Unix `ps` command with a lookalike that doesn't show your backdoor process

4) Repeat

- Use current access to try to gain new access
- Example:
  - First attack gets you command-line access on the front-line web server
  - Now that you're on the network, try to attack
    - The database server
    - The corporate network
      - The sysadmin's computer
      - The CEO's computer
      - The credit card processing system
      - The nuclear plant control system

Firewalls • Intrusion detection systems

DEFENSIVE NETWORK SECURITY

Firewalls

- Placed at the interface between two networks with differing security requirements
- Aims to control network traffic flowing between a protected network and other networks

- Frequently used to prevent unauthorized Internet users from accessing private networks (Intranet).
  - All messages entering or leaving the intranet pass through the firewall
  - Each message is examined, and those that do not meet the specified security criteria are blocked.

Firewalls

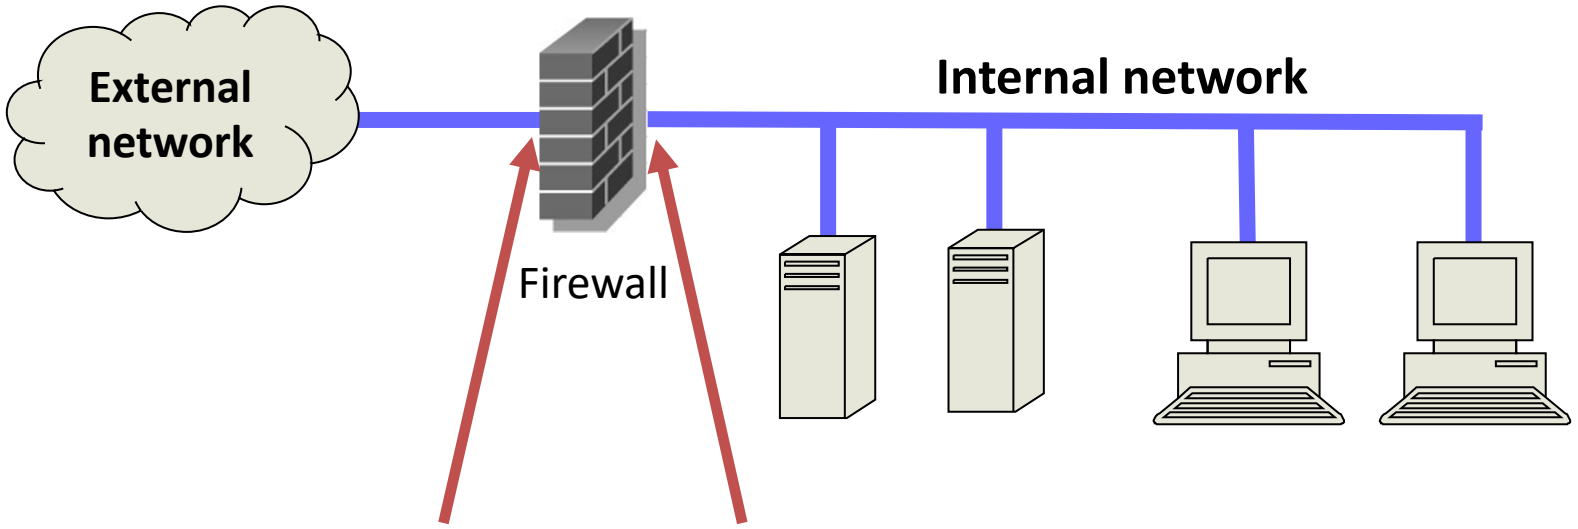

Firewall has two network interfaces:
One for external traffic, one for internal traffic

Firewall policies

- Enforce a security policy established by an administrator on all network traffic passing the boundary

- Two policy approaches:
  - **Default permit**: allow all traffic except that which is expressly prohibited (blacklist)
  - **Default deny**: block all traffic except that which is expressly permitted (whitelist)

Firewall policies

- Can apply different policies to traffic in different directions
  - **Ingress filtering**: applied to traffic coming from less trusted (external) network
  - **Egress filtering**: applied to traffic coming from more trusted (internal network)

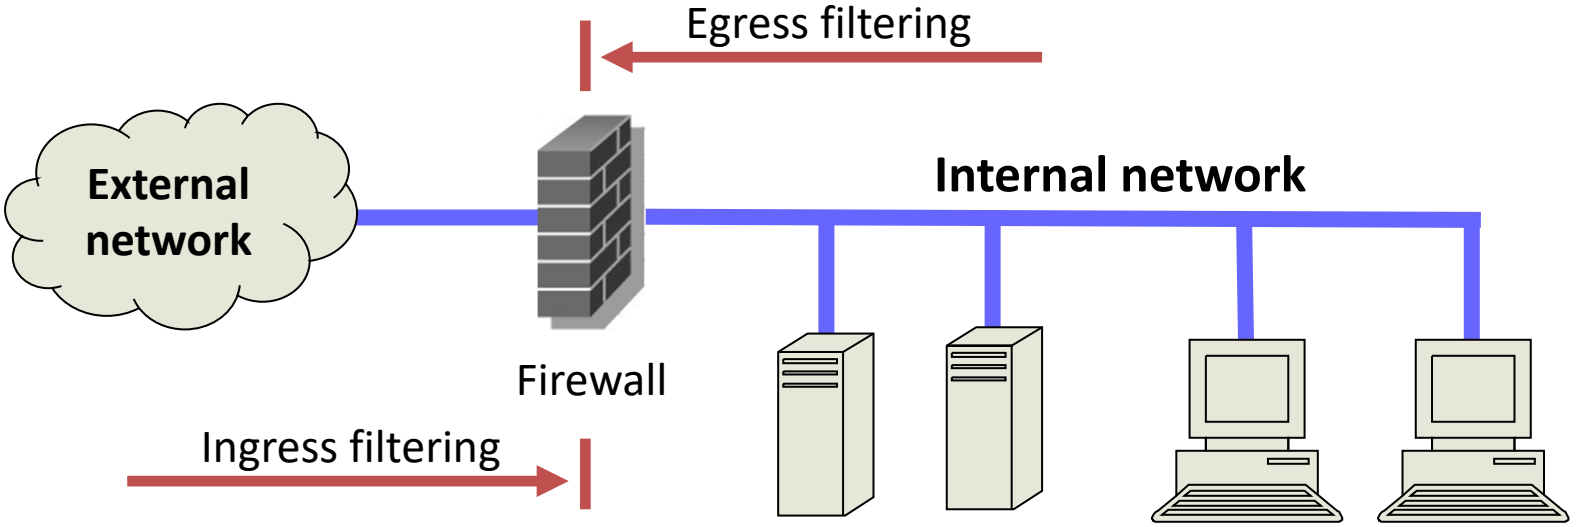

Firewall types

Packet filters

- Two types: **stateless** and **stateful**
- Primitive, high performance firewalls
- Examines data associated with lower levels of network stack (network & transport layers)
  - e.g. IP source/destination address, TCP port number
- Does not understand the upper (application) layer

Application proxies

- Performs deep packet inspection on application data
  - e.g. prevent any virtual private network connections
- Slower performance
- Must be customized for each application protocol

Packet filters

- Operate at the network or transport layer
- Makes decisions based on information in packet headers, such as
  - **IP headers**: source or destination IP address
  - **Protocol**: TCP, UDP, or ICMP
  - **TCP headers**: source or destination port numbers
  - **Direction** of travel (into/out of the internal network)

Packet filters

- A rule table specifies how to filter network traffic:
  - Each rule consists of conditions and an action
  - For each packet, the **first matching rule** is found
  - Two possible actions: allow or block

- Example rule table: inbound traffic to email (SMTP) server 10.0.2.6

Prot.	Src IP	Src port	Dest IP	Dest port	Action	Comment
TCP	4.5.6.7	*	10.0.2.6	25	Block	Block specific spammer
TCP	*	*	10.0.2.6	25	Allow	Inbound SMTP mail
TCP	10.0.2.6	25	*	*	Allow	Outgoing SMTP responses
*	*	*	*	*	Block	Default deny

Stateless packet filters

- **Stateless:** Examine each packet independently of other packets
  - Even if they are part of the same connection

- High speed
- Low memory

Stateful packet filters

- **Stateful** packet filters operate in the same way as stateless packet filters:
  - examining headers and comparing to ruleset to see if the packet transmission is allowed under the firewall rules
- But stateful packet filters also keep a **state table** noting the state of each connection:
  - Is the connection being established, in use, or terminated?
- Stateful packet filters examine the state in the context of the of the conversation
  - If header values contradict the expected state, the packet will be dropped

Packet filters

Strengths

- Low overhead
- High throughput
- Operates at lower layers, so supports almost any application

Weaknesses

- Do not examine application layer data/commands
  - May allow insecure operations to occur
  - Cannot perform content filtering or user authentication
- Allow direct connections between hosts inside & outside firewall
- Stateless packet filters only:
  - less secure (can be susceptible to IP spoofing)
  - more difficult to write complex rules

Application proxy

a.k.a. **application proxy gateway**, a.k.a. **bastion host**

- Operate at the application layer
- Makes decisions based on information in packet body, i.e., application data
  - "Deep packet inspection"

- Examples:
  - Censorship of web browsing
  - Filtering adult content at schools
  - Anti-virus scanning of email attachments

Application proxy gateway

- Usually configured to support only specific applications or specific features of an application:
  - Each application (email, web browser) must have its own proxy (specific gateway) in the firewall
  - If proxies are designed specifically for that protocol, they understand whether the traffic flowing is following the protocol and allowed by the policy rules
- Application layer firewalls have proxies for the most commonly used protocols

Application proxy gateways

Strengths

- Provides potential for best security through control of application layer data/commands
- Better logging and audit of traffic
- Allows content filtering and user authentication

Weaknesses

- Slower than packet filters – requires time to examine packet data in details, so may be unsuitable for real-time applications
- Limited support for new applications – additional time requirement for vendor to write new gateways for new applications
- Requires one additional connection (including processing resources) for each new connection

Comparing firewall types

Stateless packet filter	Stateful packet filter	Application proxy
Inspects single packets	Tracks state across many packets	Tracks state across many packets
Examines IP and TCP headers	Examines IP and TCP headers	Examines application data
High speed	Medium speed	Low speed
Simplest rules	Simple rules	Complex rules
Little/no auditing/logging	Auditing/logging possible	Auditing/logging likely

Simple firewall architecture

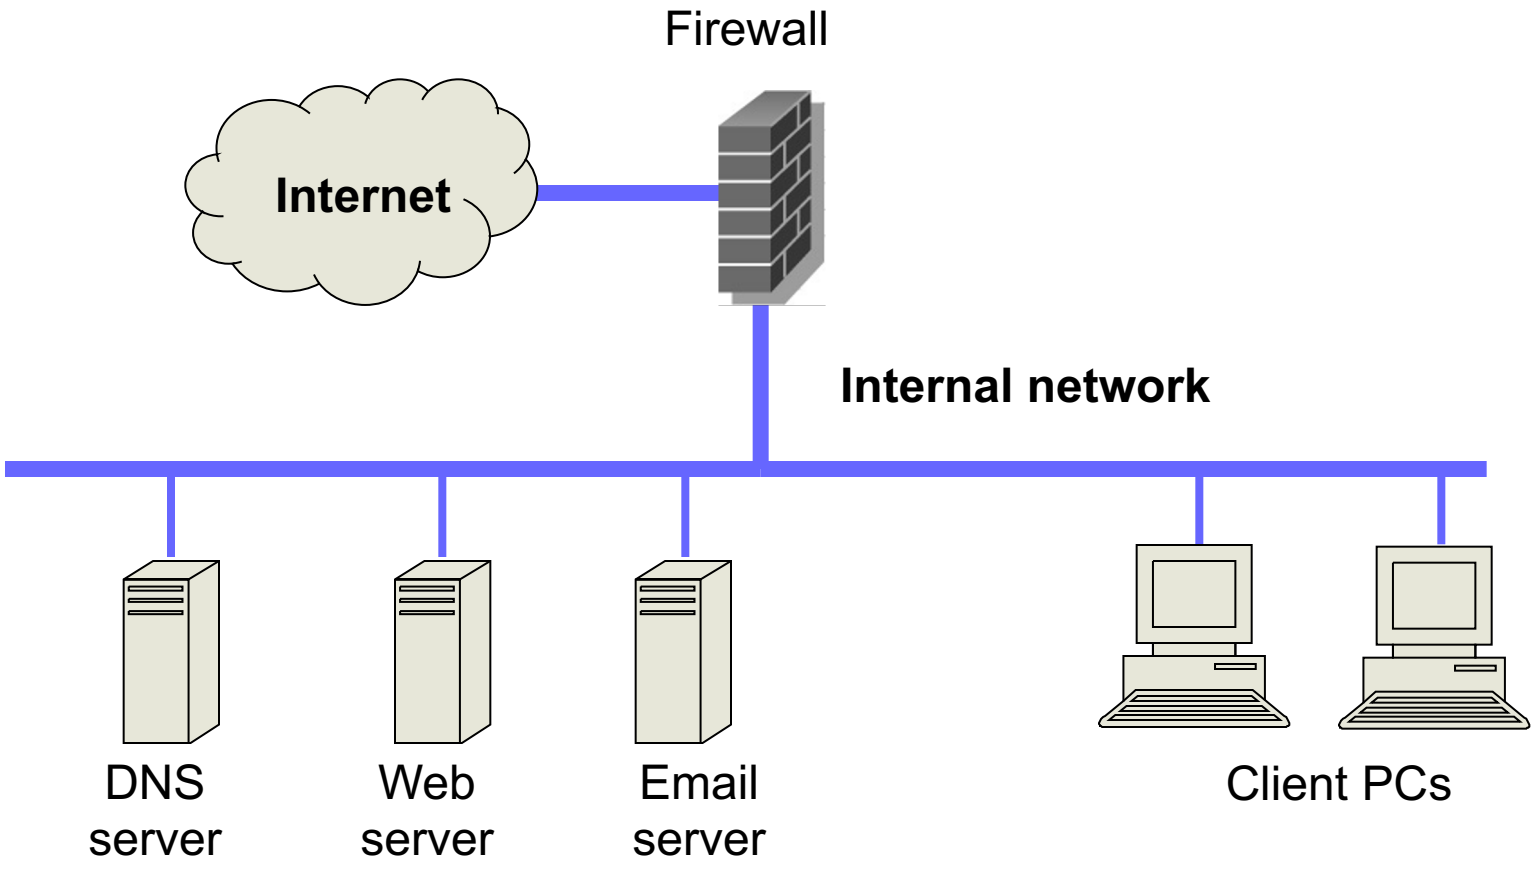

DMZ firewall architecture

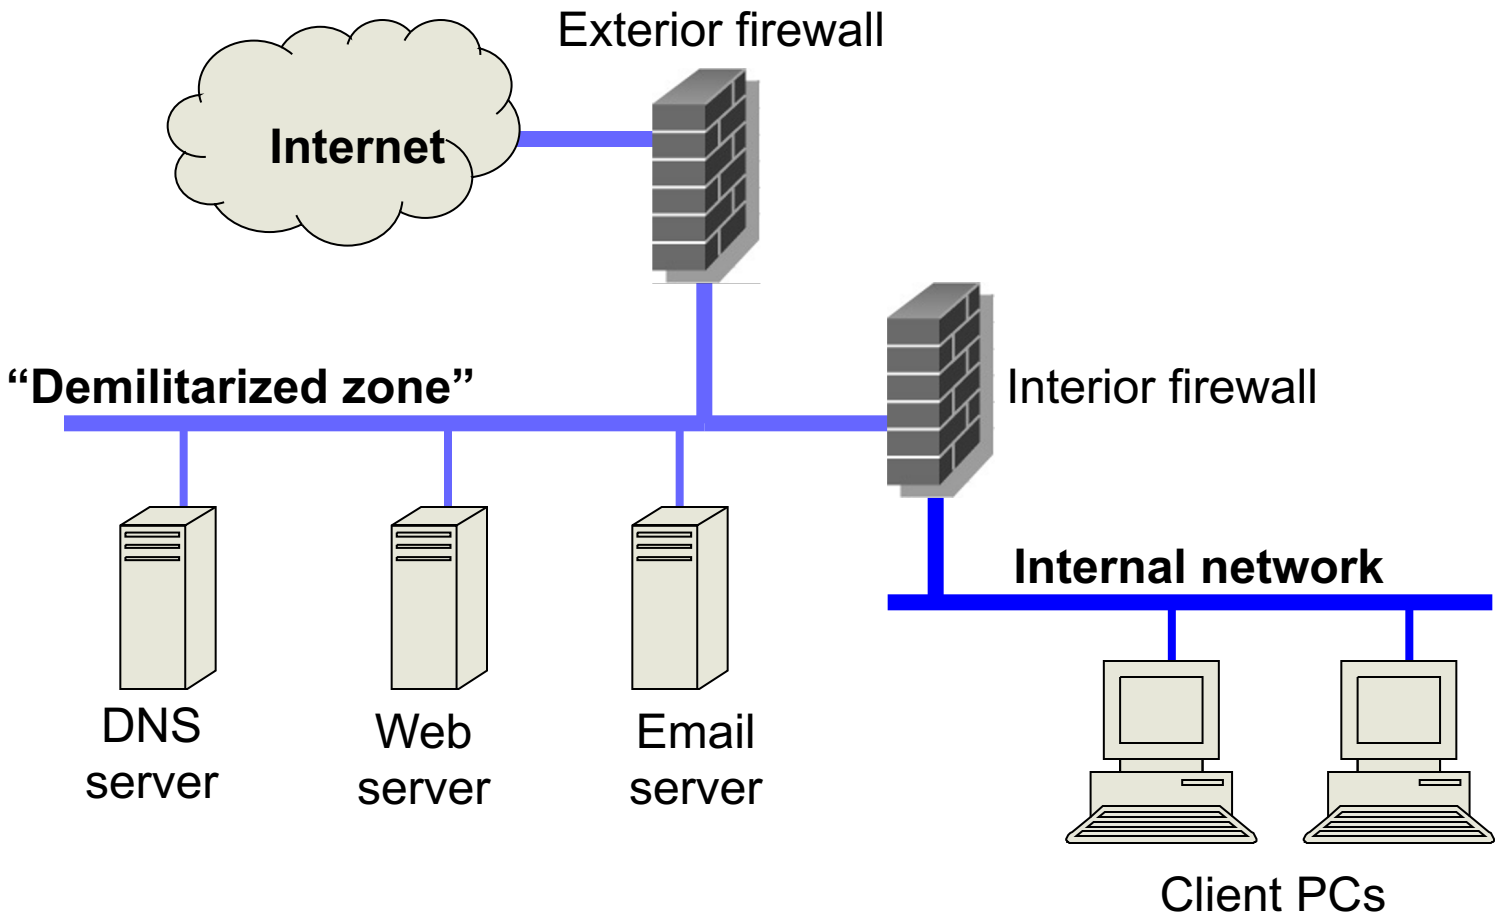

Personal firewalls

- A personal firewall is a software program that is designed to protect the computer **on which it is installed**.
  - Frequently used by home users to provide protection against unwanted Internet traffic.
- Usually these are stateful packet filters.

- Examples:
  - Windows, Ubuntu, and macOS all include a personal firewall
  - Commercial personal firewalls: ZoneAlarm, Symantec, Little Snitch, …
    - Some include anti-virus software as well

Firewalls in Linux

- **Netfilter**: framework in Linux kernel for registering Kernel modules for manipulating networking functions
- **iptables**: kernel module and user-space program for defining packet flow rules
  - iptables rules can be used to construct a firewall, router, …
- **nftables**: next-generation version of iptables
- **ufw**: "uncomplicated firewall", a wrapper around iptables, originally designed for Ubuntu

Challenges with firewalls

Technical

- Trade-off:
  - Simple packet filters have high performance
  - Application level gateways offer more comprehensive filtering
- Hard to configure; policy errors are common
- Need to be kept up to date
- Often ways to bypass

Non-technical

- Rely on well-formulated security policy
- Firewall != Security
  - Perimeter security is often bypassed
- Training human operators

IDS and IPS

- **Intrusion detection systems (IDS)** aim to detect attempts to break in to networks
- **Intrusion prevention systems (IPS)** aim to stop attempts to break in to networks

- Monitors logs and sniffs packets in real time to detect
  - traffic that matches known attack signatures
  - anomalies compared to normal behaviour
  - stateful analysis of protocol and program behaviour
- E.g., Snort

IDS and IPS

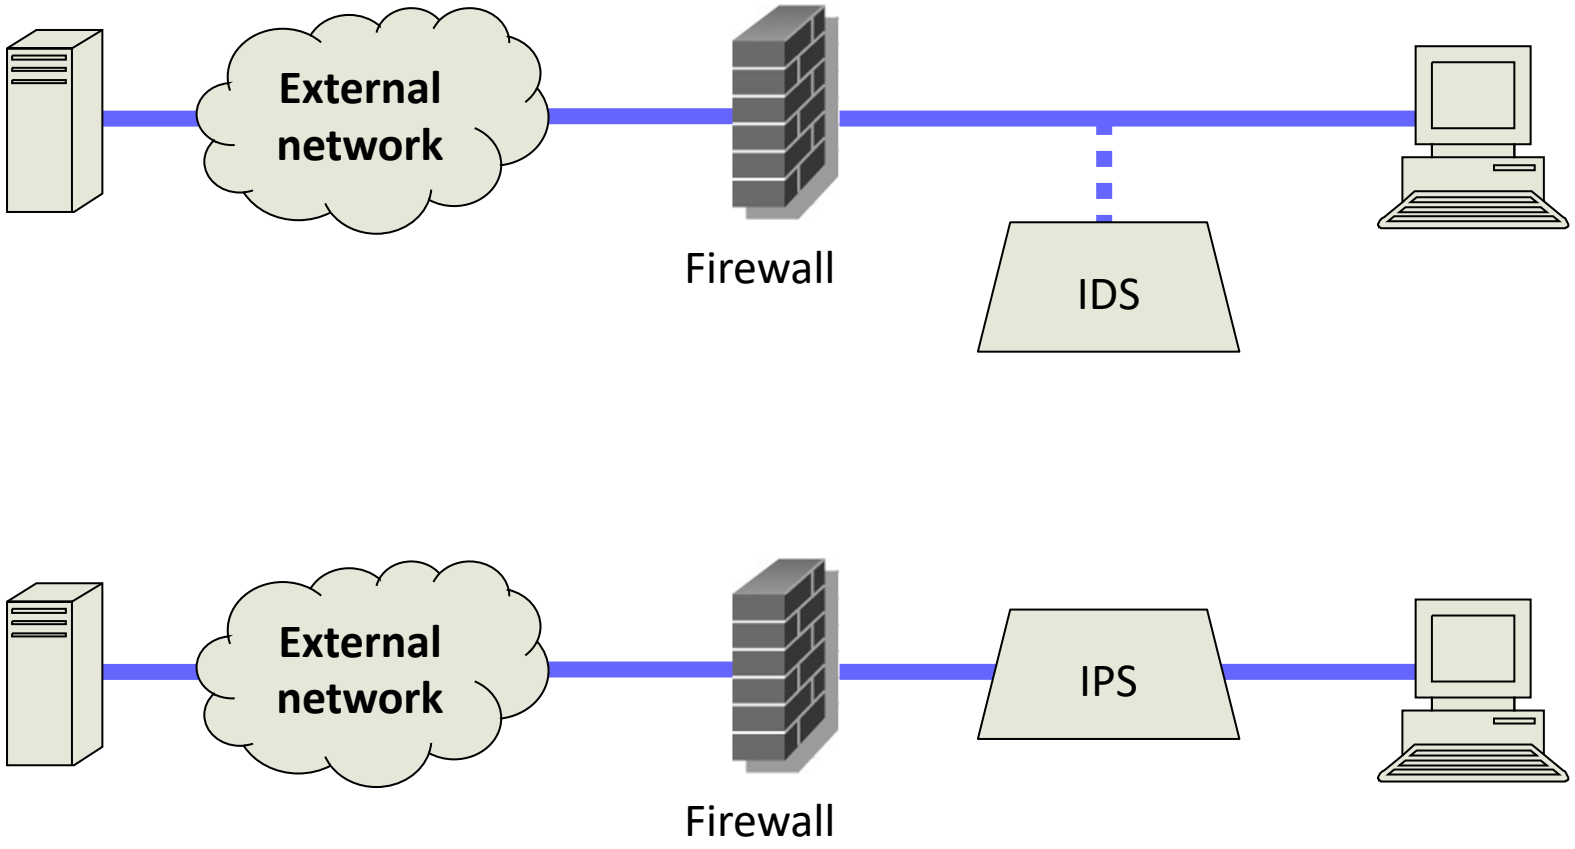
External network
Firewall
IDS
External network
Firewall
IPS

Model of IDS / IPS

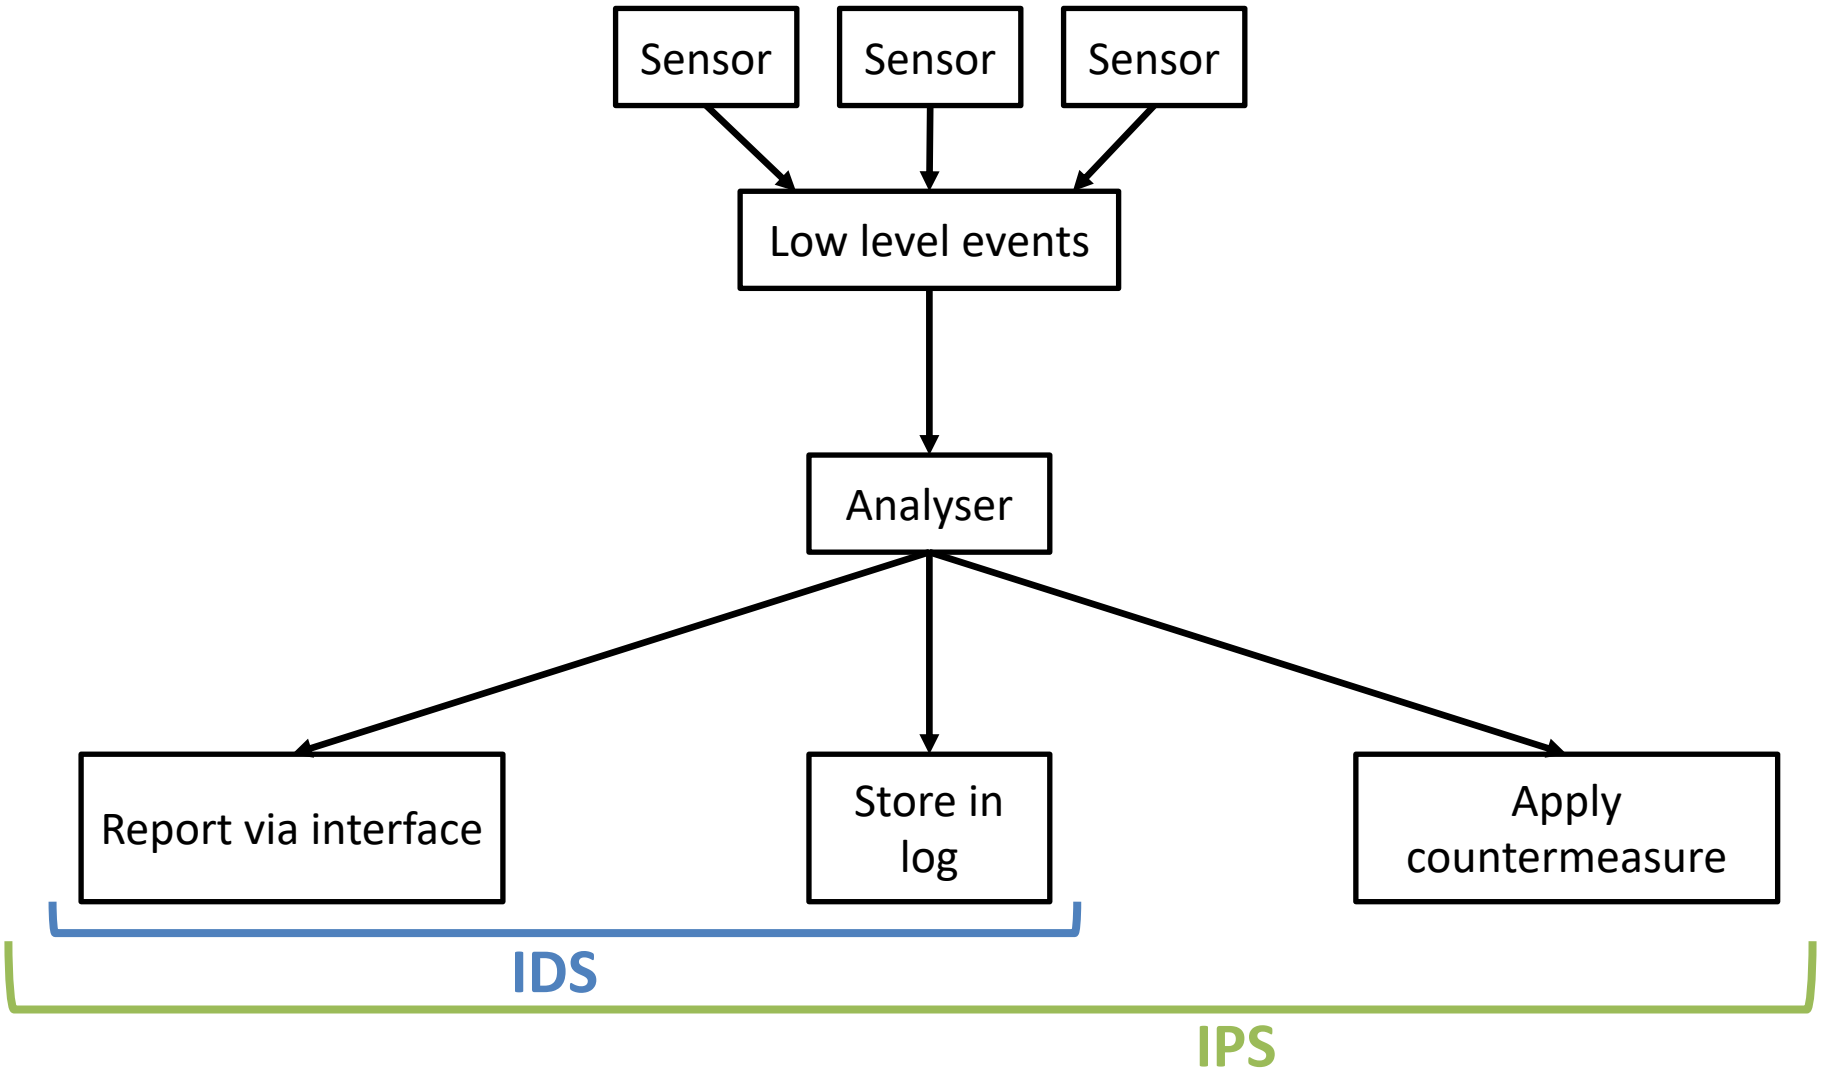

Types of IDS / IPS classified by input sources

Host-based IDS/IPS

- Runs on a single computer
- Input sources:
  - Behaviour of applications on that host
  - System characteristics of that host

Network-based IDS/IPS

- Input sources:
  - Network traffic from various points in the network

Infrastructure IDS/IPS

- Combines both host-based and network-based
- Input sources:
  - Application behaviour & system characteristics from many hosts
  - Network traffic

Types of analyses

Signature- or misuse-based detection

- detects pattern or signature matching known misuse or threat

Anomaly- or heuristic-based detection

- detects deviation from normal
  - Network Behaviour Analysis
  - Stateful Protocol Analysis

Limitations of analysis types

Signature- or misuse-based detection

- Ineffective against novel (zero-day) attacks where misuse pattern is unknown
- Ineffective against polymorphic attack code

Anomaly-based detection

- Requires training or learning “normal” profile
- High false-positive rate

Firewalls vs. IDS vs. IPS

Packet filter	Application proxy	IDS	IPS
Preventive	Preventive	Detective	Preventive
Examines packet headers		Examines packet headers	Examines packet headers
	Examines application data	Examines application data	Examines application data
Drops packets not matching policy	Drops packets not matching policy		Drops packets not matching policy
		Logs / raises alerts for data matching criteria	Logs / raises alerts for data matching criteria
			Applies countermeasures

Simple, fast

Complex, slow

Assignment 2

2a) Offensive network security

- Use nmap to scan services running on your computer
  - Will be scanning from guest Kali Linux virtual machine to host machine using a simulated network

2b) Defensive network security

- Set up firewall rules in your Kali to prevent certain types of outbound traffic (egress filtering)

Assignment 0
Downloading and installing VirtualBox and Kali Linux

https://www.douglas.stebila.ca/teaching/cryptoworks21/

SUPPLEMENTAL MATERIAL: DENIAL OF SERVICE ATTACKS

Denial of Service (DoS) attacks

- **Goal: Defeat availability.**
  - Deny users access to authorized services or data.
  - Extort random from service providers by threatening/denying availability of their service.

- Main approaches:
  - **Flooding attacks**: Overwhelm the victim system.
    - Distributed DoS (DDoS) is a special case.
  - **Crashing attacks**: Exploit some bug to disable the victim system.
  - **Disable communication**: Physically or logically disable or reroute communication.

DoS attacks

Different DoS attacks target different layers of the networking stack

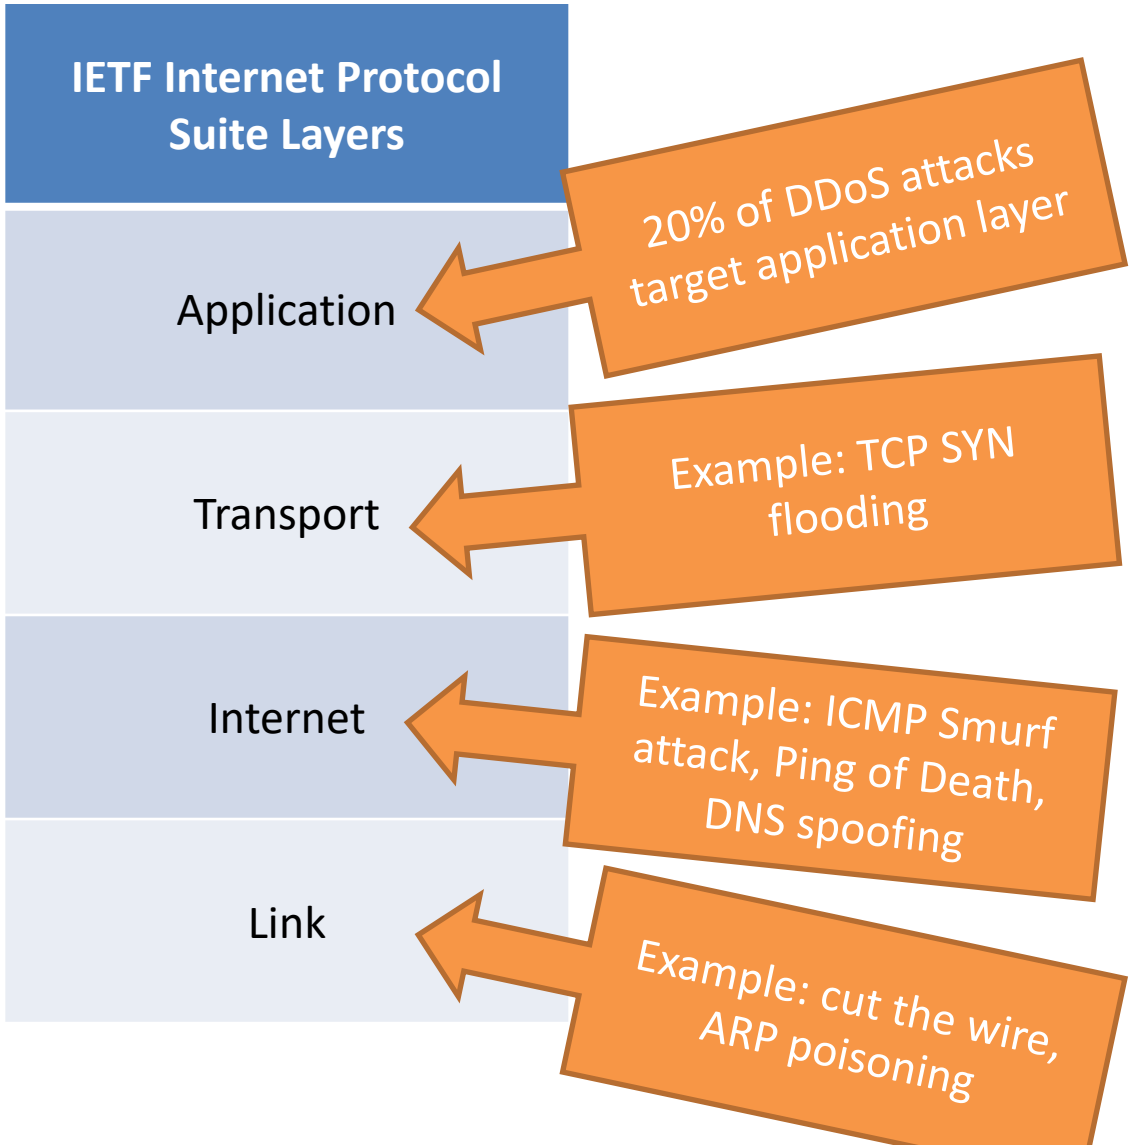

Flooding attacks

- Flooding attacks aim to overwhelm the victim system.
  - Overwhelm resources:
    - Victim system may have some resource restrictions (disk space, number of open sockets, database connections)
  - Overwhelm network capacity:
    - Victim system has limited bandwidth and receives/sends more data than bandwidth available
    - Best attacks are **asymmetric attacks** where attacker doesn't need more bandwidth than victim

Resource exhaustion attacks

Example: TCP SYN flooding

Recall TCP three-way handshake for establishing a reliable, ongoing connection

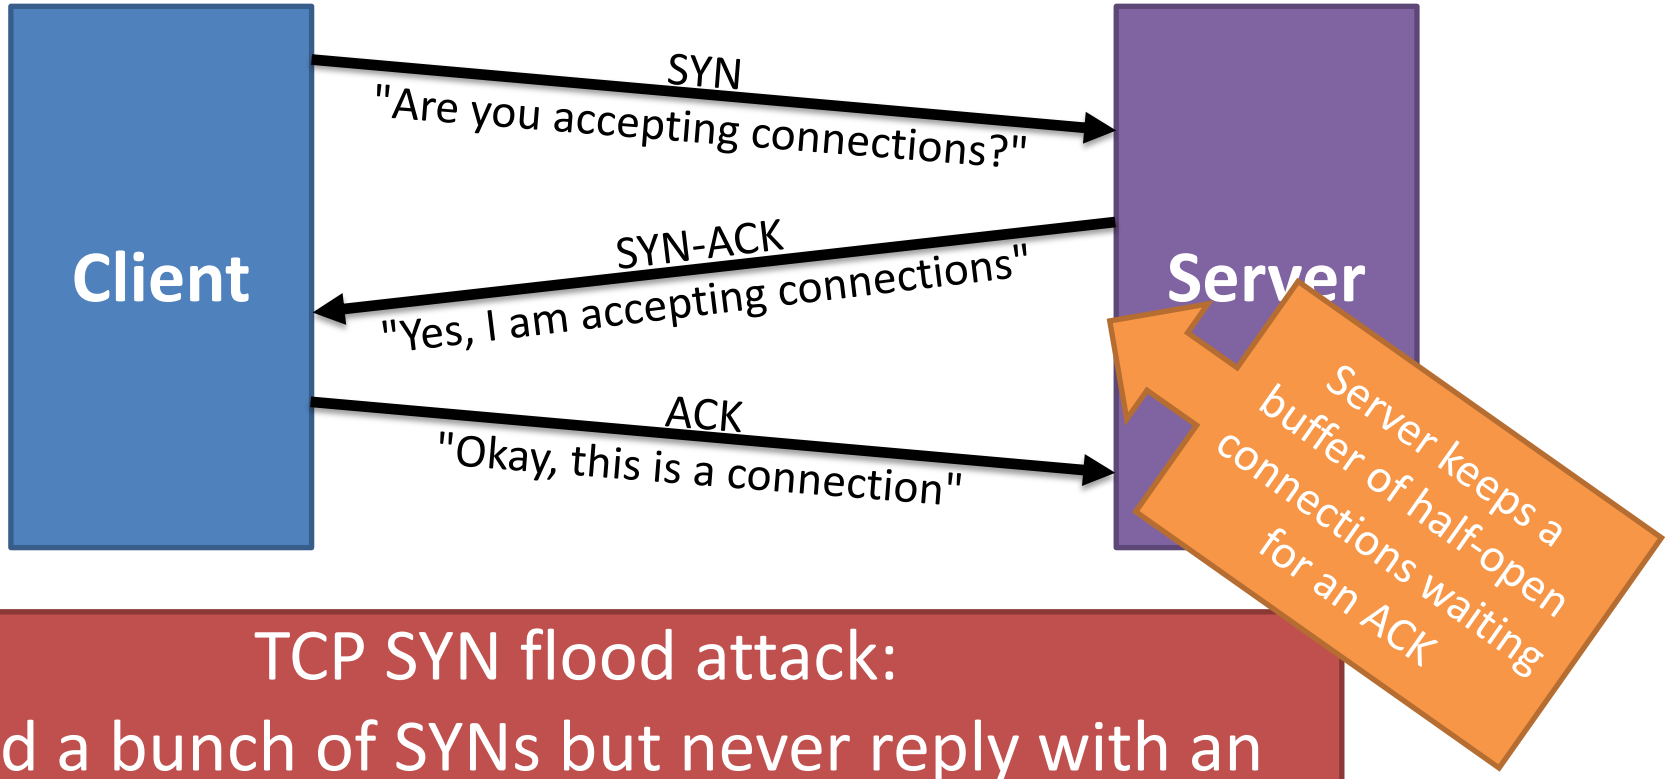

TCP SYN flood attack:
Send a bunch of SYNs but never reply with an ACK. Server's buffer eventually exhausted.

Network flooding attacks: Reflected / spoofed attacks

- Attacker tricks intermediate servers or clients into sending replies to the victim, not the attacker

Network flooding spoofed attacks: Example: Smurf attack

1. Attacker sends Internet Control Message Protocol (ICMP) "echo" (ping) packets with victim's spoofed source IP address to broadcast address
2. Router delivers packets to all recipients of broadcast address
3. Recipients reply to victim's IP address

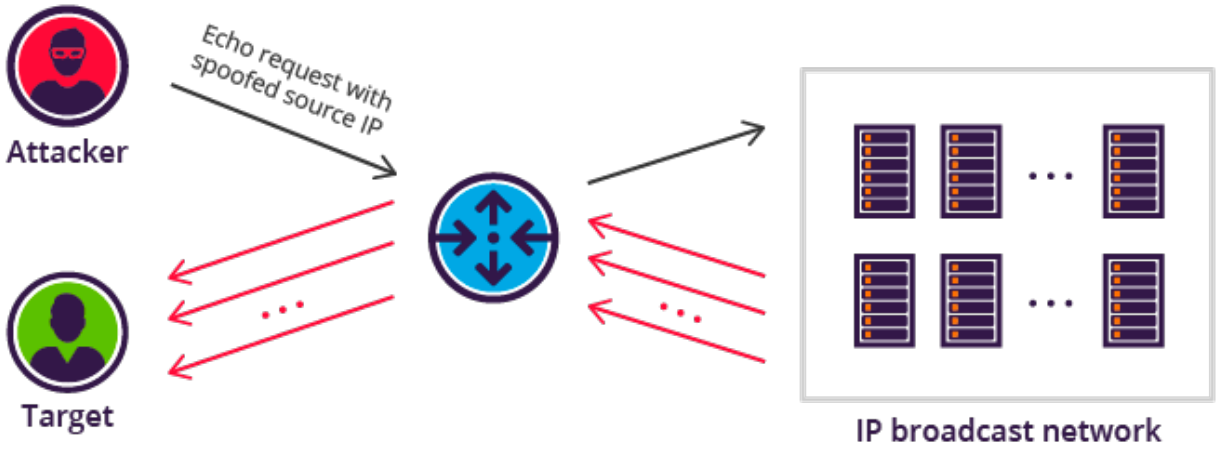

https://www.incapsula.com/ddos/attack-glossary/smurf-attack-ddos.html

Network flooding attacks: Amplification

- Makes use of intermediate services with responses much bigger than the request
- Attacker makes many (small) requests with victim's spoofed IP address
- Service responds with (large) responses to victim's IP address

Protocol	Bandwidth amplification factor
NTP	556.9x
DNS	Up to 179x
Quake Protocol	63.9x
BitTorrent	Up to 54x
SNMPv2	6.3x

Hard to defend against UDP-based attacks since there is no 3-way handshake to verify source address like in TCP.

https://en.wikipedia.org/wiki/Denial-of-service_attack#Amplification

Low Orbit Ion Cannon

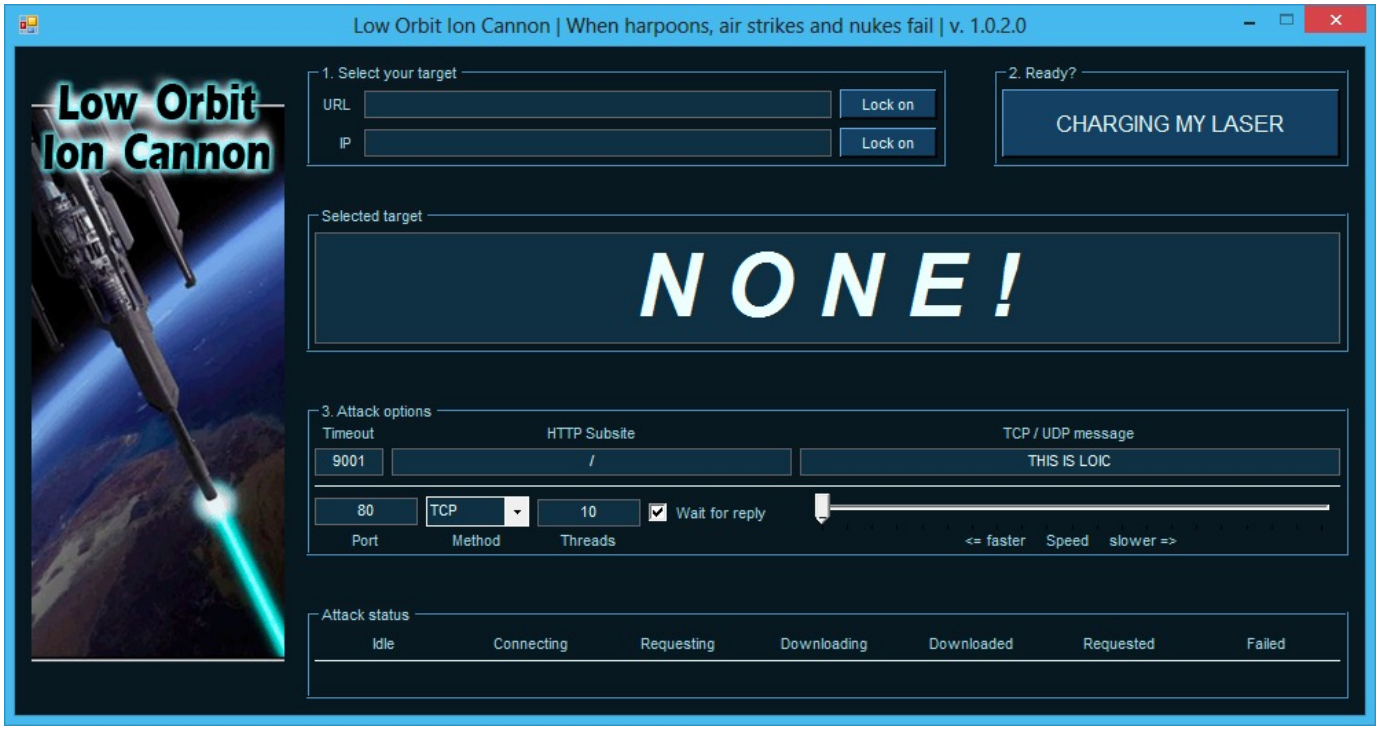

- Open source tool
  - Legitimate use: stress testing your own systems
  - Illegitimate use: denial of service attacks

https://sourceforge.net/projects/loic0/

Crashing attacks
Example: Ping of death

- Maximum size of an IPv4 packet is 65,535 bytes
- Ping is a special type of IPv4 packet using the Internet Control Message Protocol (ICMP)
- Prescribed size for a ping packet is 56 bytes

- **Ping of death**: send a ~60,000 byte ping packet
- Caused crashes in early TCP/IP networking implementations
- 2013 vulnerability in IPv6 ping in Windows

Disabling communication

- **Goal**: disable the communication between parties and their intended peer.

Physically disabling communication:

- Cut the wire
- Jam the wireless signal
- Turn off power

Logically disabling communication:

- Change addressing
- Change routing

Logically disabling communication by changing addressing: **DNS spoofing**

- **Goal**: Spoof responses to DNS queries to redirect queries for a particular domain name to attacker-controlled IP address

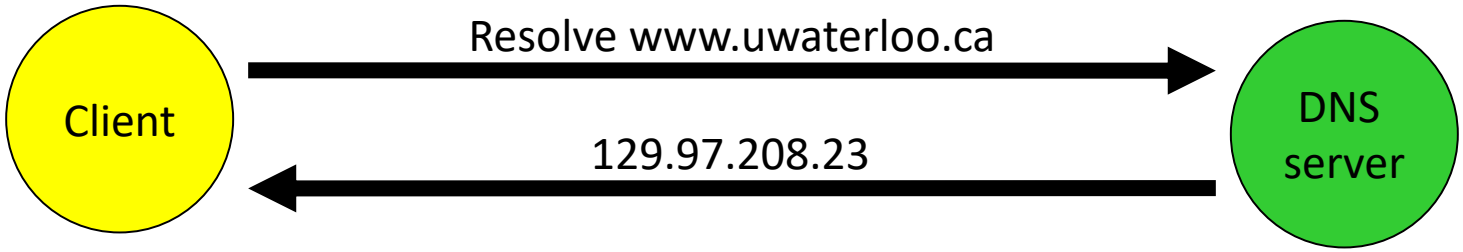

Logically disabling communication by changing addressing: **DNS spoofing**

- Goal: Spoof responses to DNS queries to redirect queries for a particular domain name to attacker-controlled IP address

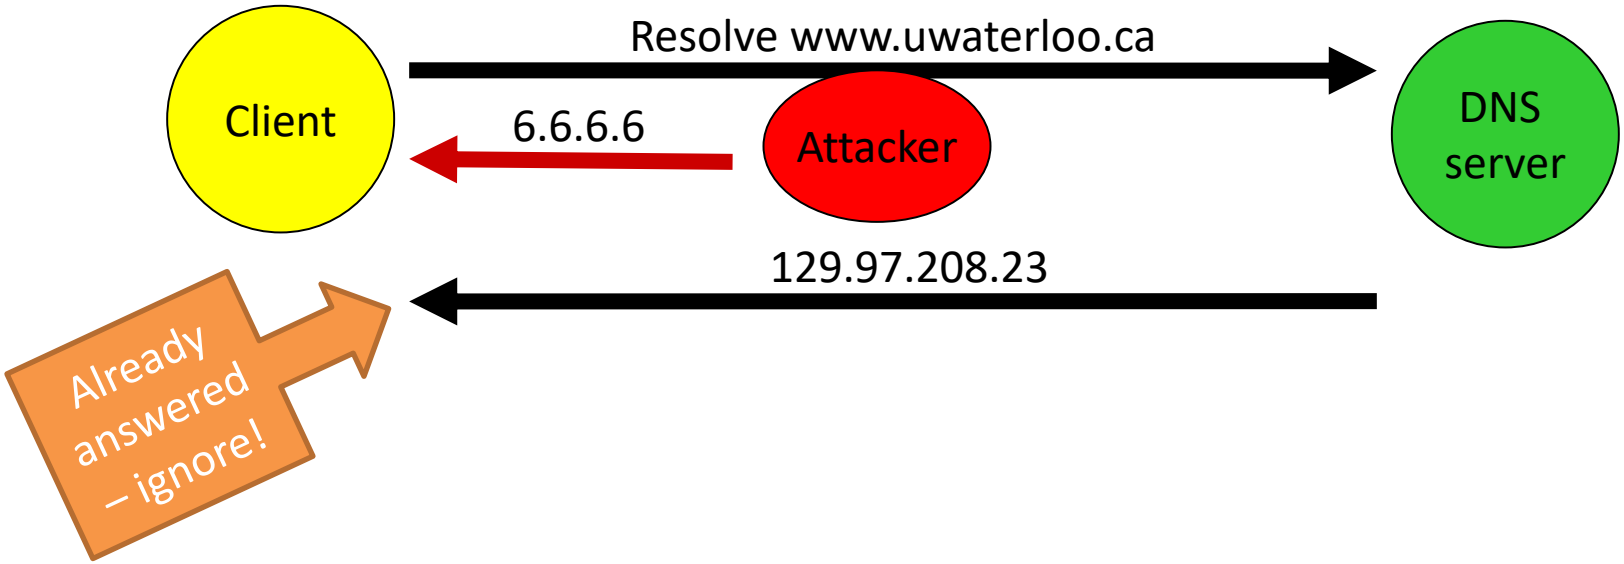

Logically disabling communication by changing addressing: **DNS spoofing**

- **DNS spoofing** targets a single client
  - Have to match a nonce value and get the response in before legitimate server, but possible
  - Works because plain DNS has no cryprographic authentication/integrity mechanism
  - DNSSEC adds cryptographic protection, but not widely deployed
- **DNS cache poisoning** targets intermediate DNS servers that cache responses
  - Successful DNS cache poisoning affects all clients relying on that DNS server

Logically disabling communication by changing routing: **ARP poisoning**

- **Address Routing Protocol (ARP)** works at link layer to map IP addresses to Ethernet MAC addresses

- **ARP poisoning** like DNS spoofing, but goal is to redirect frames for a particular IP address to attacker's MAC address

DoS attacks

Different DoS attacks target different layers of the networking stack

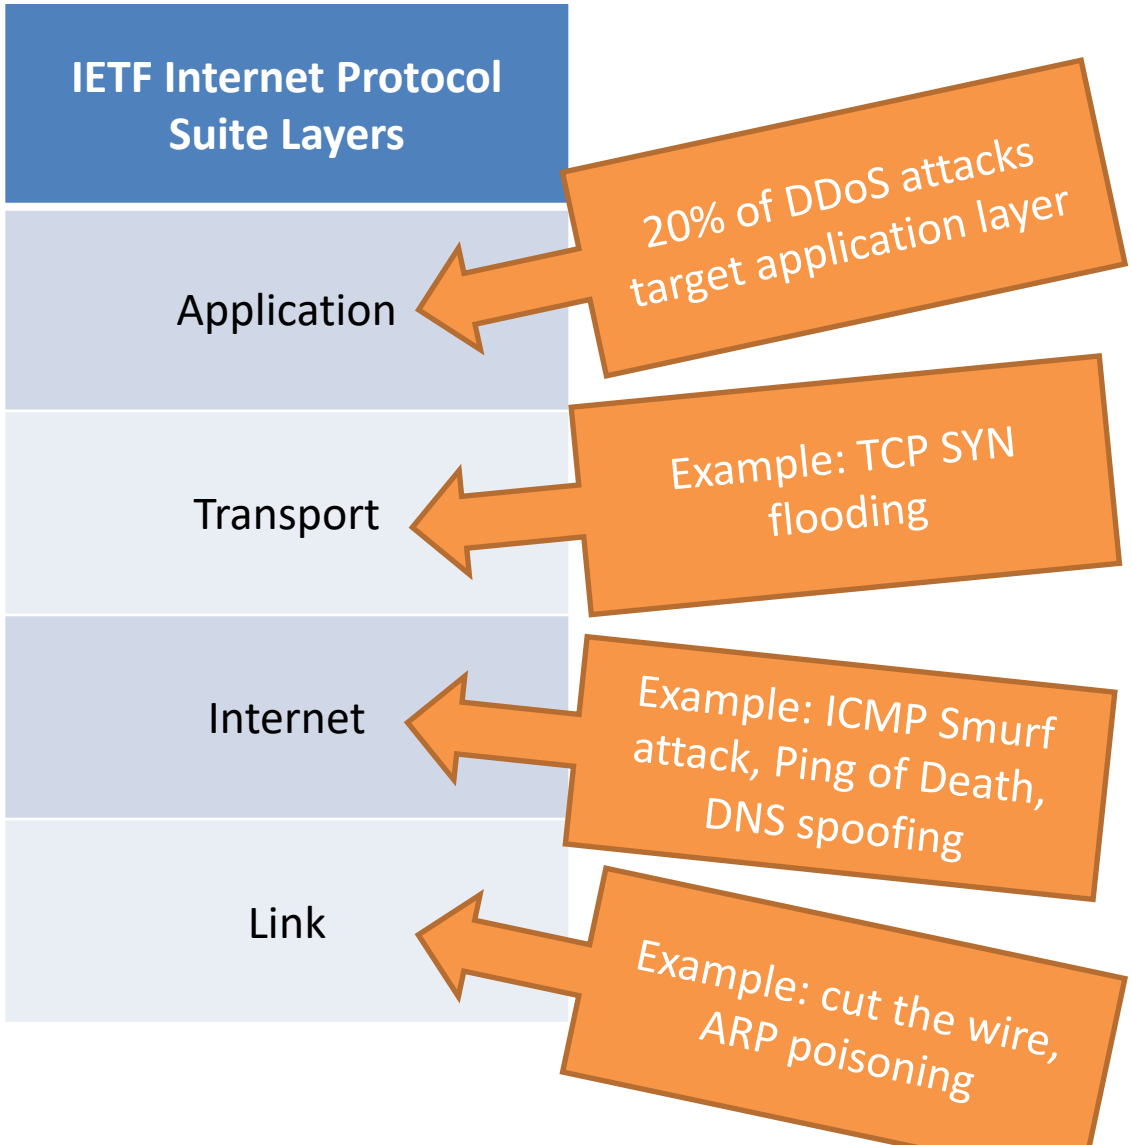

SUPPLEMENTAL MATERIAL : DISTRIBUTED DENIAL OF SERVICE ATTACKS

Distributed denial of service attacks

- To overwhelm the victim via flooding attacker, attacker needs either:
  - More resources/bandwidth than victim
    - Expensive to obtain
  - Asymmetric attack
    - Need to be clever
- In DDoS attack, attacker gets resources & bandwidth by forming a botnet of compromised computers around the world

Botnets

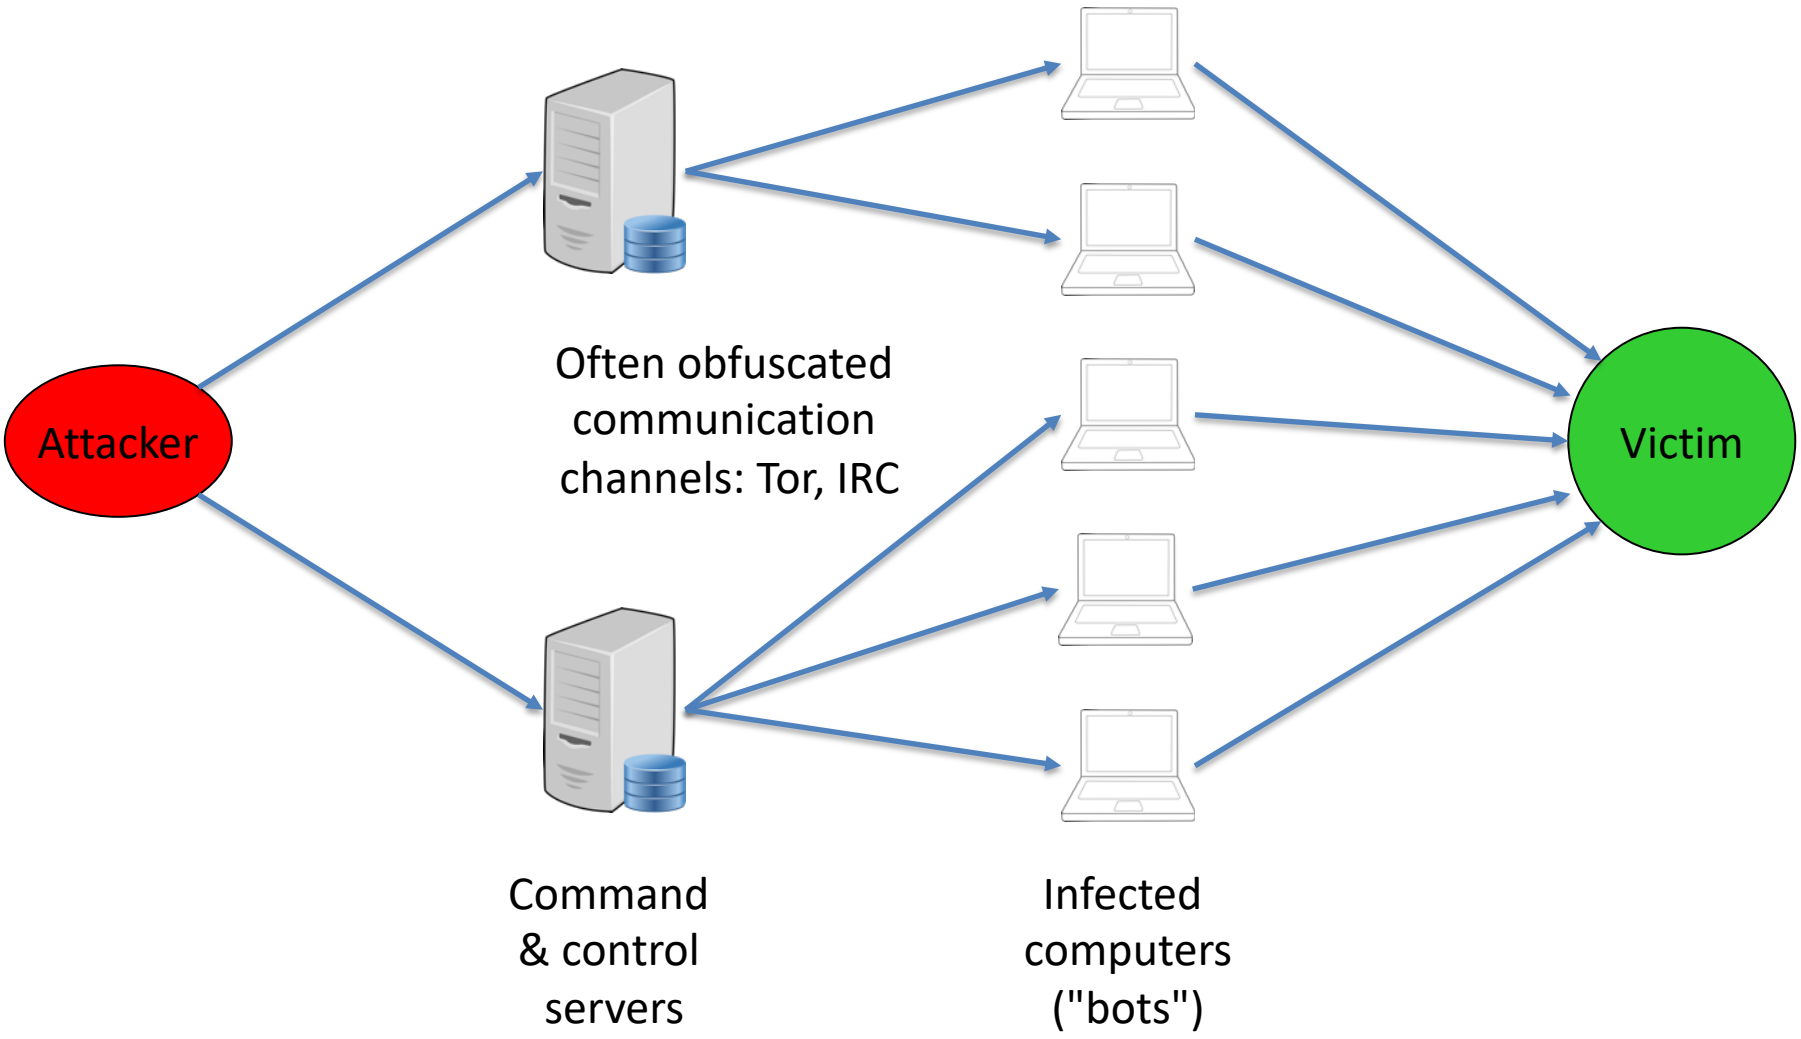
Attacker
Often obfuscated communication channels: Tor, IRC
Victim
Command & control servers
Infected computers ("bots")

DDoS-as-a-service

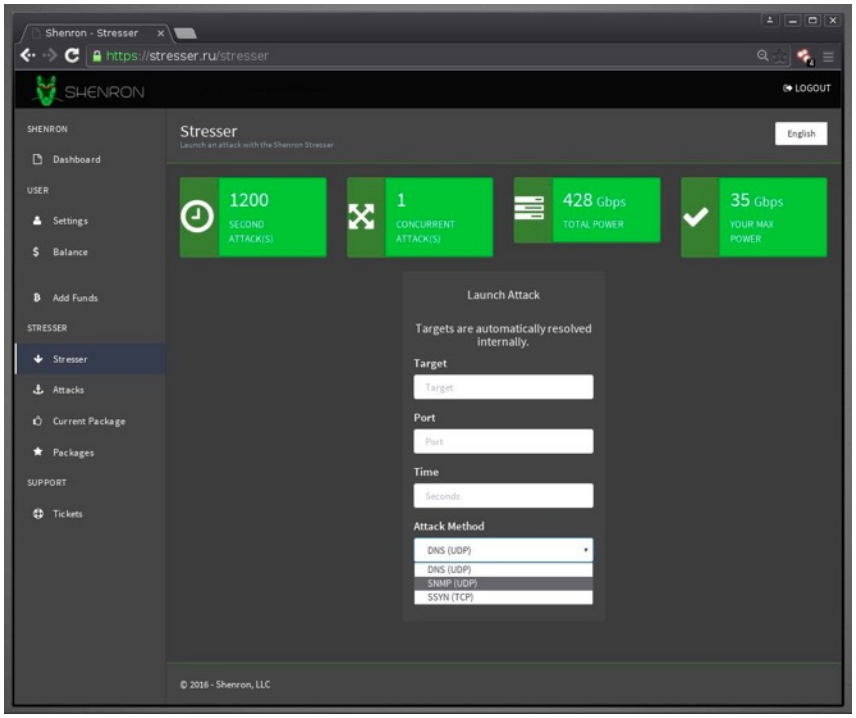

- **LizardStresser –** up to 500 Gbps, prices range between $20 and $1000
- **Bang Stresser** – costs $12 to $100 for up to 1.5 hours' attack duration
- **uStress –** can generate a 20-minute 300 Gbps attack. Prices vary between $15 and $150
- **NetStresser** – Prices range from $10 to $150 for up to 1.5 hours' attack duration
- **vDoS –** over 200 Gbps of multi-vector attacks. Prices vary between $20 and $150

https://security.radware.com/ddos-threats-attacks/threat-advisories-attack-reports/ddos-as-a-service/

Scale of DDoS attacks

- 1999 – first DDoS attack
- 2000 – Yahoo, eBay, Amazon DDoS'ed for hours
- Early 2000's: peak speed 4 gigabit/sec
- 2015:
  - Average speed: 10–60 Gb/sec • peak speed: ≥ 400 Gb/sec
    - One company reported receiving 250x normal bandwidth
  - Average duration: 17 hours
- 2016:
  - Mirai botnet attack on Krebs on Security blog
    - Peak speed ≥ 600 Gb/sec
    - UWaterloo: 6+5Gb/sec commercial, 10Gb/sec universities

http://www.darkreading.com/cloud/inside-a-vicious-ddos-attack/a/d-id/1321286
http://gcn.com/articles/2015/07/27/ddos-attack-mitigation.aspx
https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/
http://dailynews.mcmaster.ca/worth-mentioning/mcmasters-internet-and-research-networks-get-speed-boost/

Selected DDoS attack statistics – 2017
Type of attack

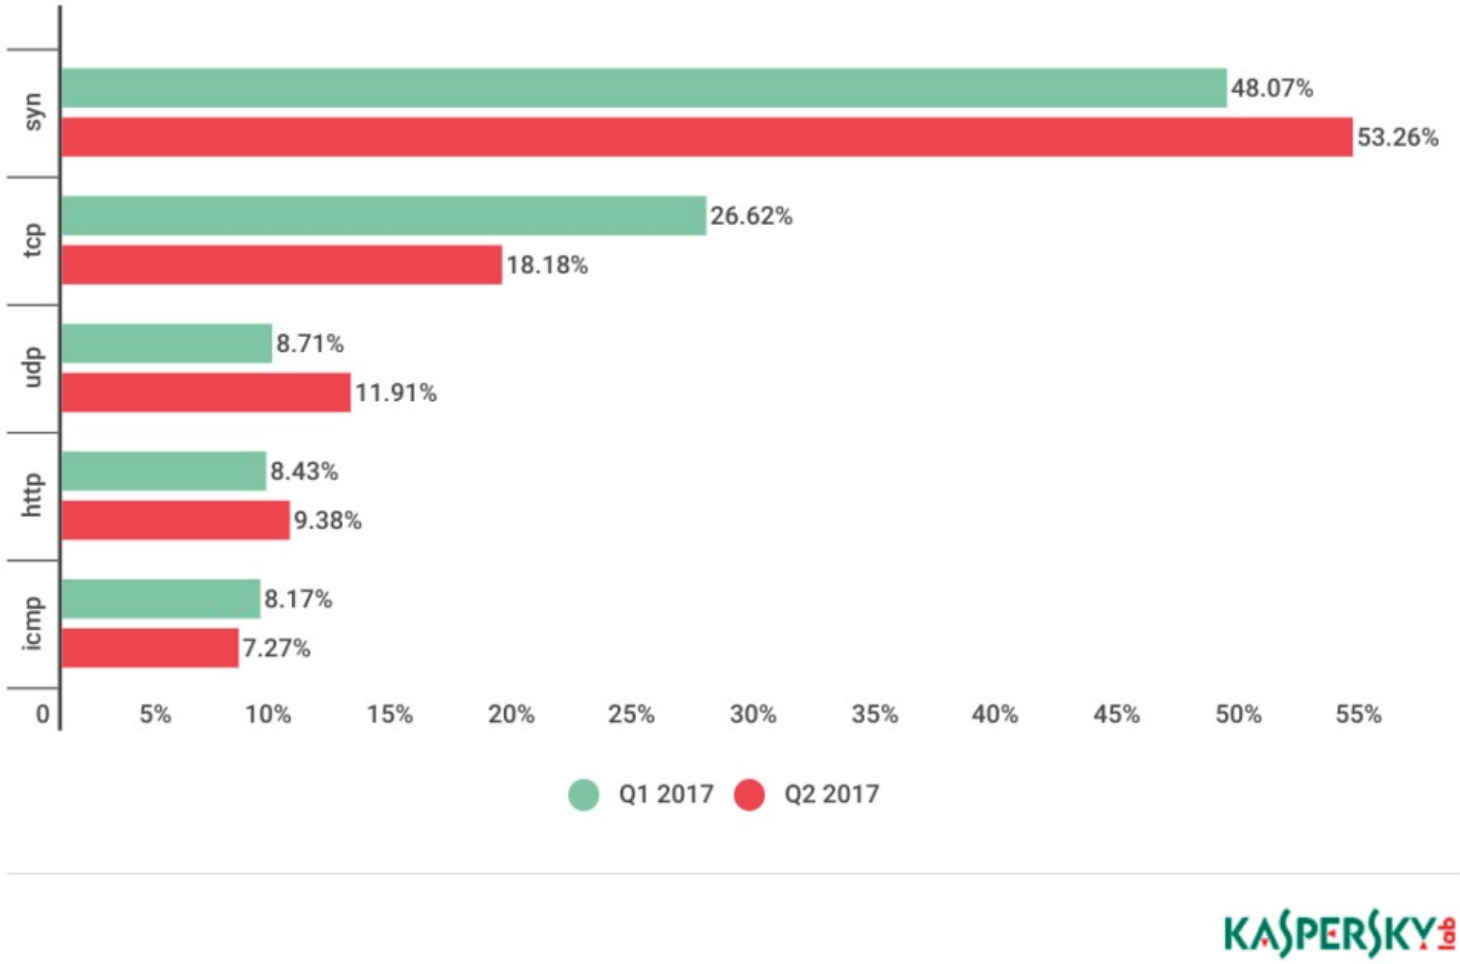

https://securelist.com/ddos-attacks-in-q2-2017/79241/

Selected DDoS attack statistics – 2017
Duration of attack in hours

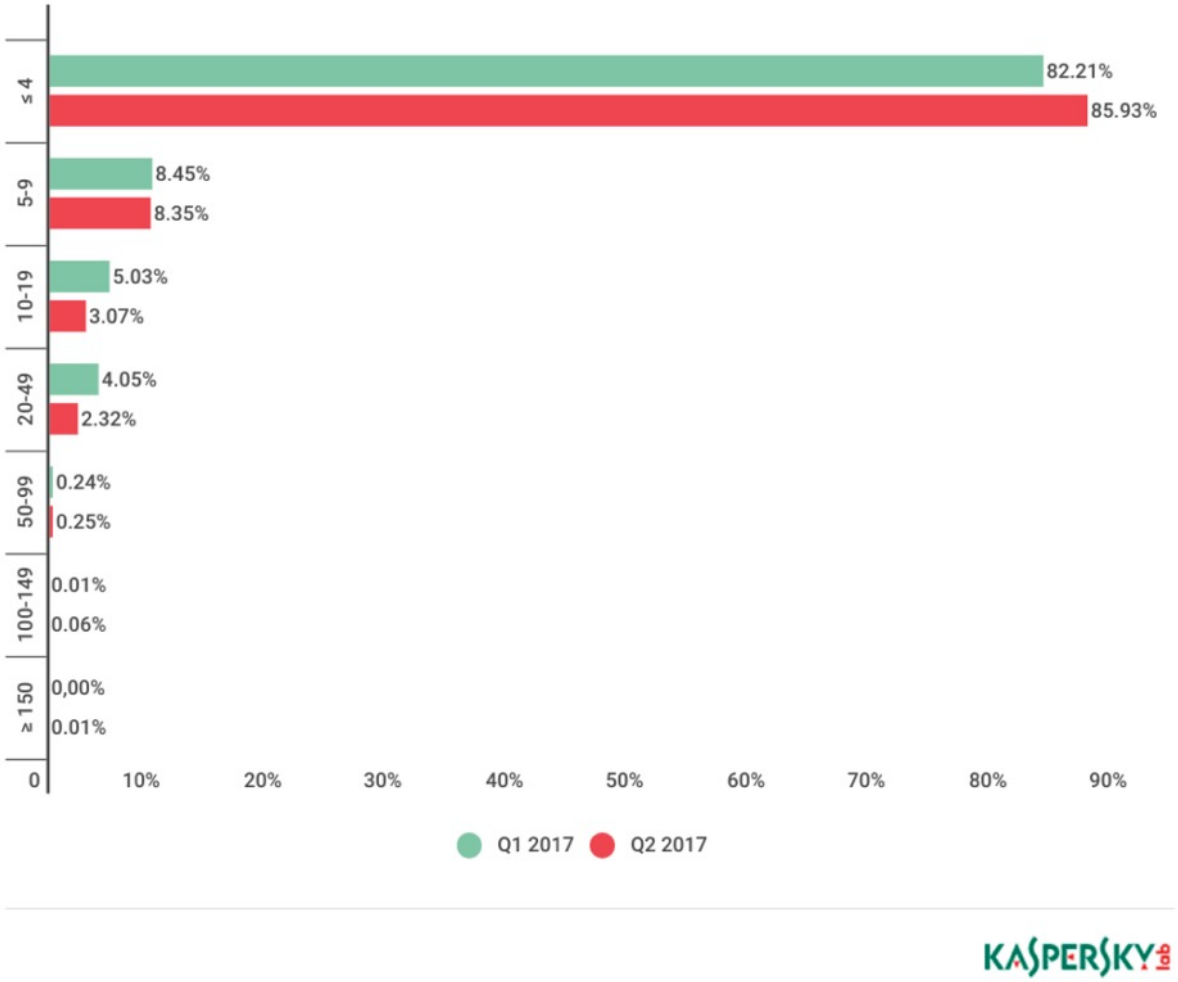

https://securelist.com/ddos-attacks-in-q2-2017/79241/