### YΣ13 - Computer Security

### **Network Security**

Κώστας Χατζηκοκολάκης

- Computers connected in a network
  - but also: smartphones, fridges, IoT devices, ...
- Each device has an IP address
- Packets are routed via intermediate nodes
- Still using IPv4 (almost 40 year old!)
  - very very hard to replace

#### Attacker model

- Intercept packets
- Modify packets
- Inject packets
- Control some routers
- Participate in any protocol
- Useful to consider combinations of the above

Four layers (7 in the OSI model)

- Link
  - Physical addresses
  - Physical aspects of communication
- Internet
  - Addressing (source/dest IP)
  - Routing
  - Time to live

Four layers (7 in the OSI model)

- Transport
  - Source/dest ports
  - Ordering of packets (Sequence numbers)
  - ACKs, checksums

#### Application

- The "real data", application-dependent

#### Protocols

- Link
  - Ethernet
  - WiFi
  - DSL
  - ...
- Internet
  - IP
  - ICMP
  - ...

#### Protocols

- Transport
  - TCP
  - UDP
  - ...
- Application
  - HTTP/HTTPS
  - SSH
  - SMTP
  - ...

#### Packet example

Physical Layer: eth the 2 MAC addresses + IP indication			<u>Network Layer</u> : IP IP addresses, TTL, checksum, fragmentation		
0000	00 0f db 4d	1 77 95 00 0d 93	b0 a3 24 08 00 45 00	▶ Mw\$E.	
0010	01 75 c8 de	40 00 40 06 44	dd c0 a8 01 2e 40 ec	.u@.@.D@.	
0020	29 05 e6 10	00 50 15 86 10	4d 25 b6 67 ed 80 18	)PM%.q	
0030	ff ff 2d 2f	00 00 01 01 08	0a 2f 41 cf 64 62 38	//A.db8	<u>Transport Laver</u> : TCP
0040	81 9a 47 45	54 20 2f 63 6e	6e 2f 32 30 30 36 2f	GET /cnn/2006/	Danta Can Aali avaa baaa
0050	55 53 2f 30	32 2f 32 37 2f	6b 61 74 72 69 6e 61	US/02/27/katrina	Ports, Seq Ack numbers,
0060	2e 70 6f 6c	: 6c 2f 74 31 2e	32 31 33 35 2e 6d 6f	.poll/t1.2135.mo	checksum timestamns
0070	6e 2e 62 65	61 64 73 2e 61	70 2e 6a 70 67 20 48	n.beads.ap.jpg H	checkaum, umeatampa
0080	54 54 50 2f	31 2e 31 0d 0a	41 63 63 65 70 74 3a	TTP/1.1Accept:	
0090	20 2a 2f 2a	d 0a 41 63 63	65 70 74 2d 4c 61 6e	*/*Accept-Lan	
00a0	67 75 61 67	65 3a 20 65 6e	0d 0a 41 63 63 65 70	guage: enAccep	
00b0	74 2d 45 6e	e 63 6f 64 69 6e	67 3a 20 67 7a 69 70	t-Encoding: gzip	
00c0	2c 20 64 65	66 6c 61 74 65	0d 0a 52 65 66 65 72	, deflateRefer	Application Laver: HTTP
0000	65 72 3a 20	68 74 74 70 3a	2f 2f 77 77 77 2e 63	er: http://www.c	
00e0	be be 2e 63	6f 6d 2f 0d 0a	55 73 65 72 2d 41 67	nn.com/User-Ad	- Request: GET
OUTU	65 be /4 3a	20 4d 6E 7a 69	6C 6C 61 2E 35 2E 30	ent: Mozilia/5.0	Deguest LIDI
0100	20 28 40 61	. 63 69 68 74 6I	73 68 3D 20 55 3D 20	(Macintosh; U;	nequest Oni
0110	20 20 43 20	40 61 63 20 41	53 20 58 3D 20 65 6e	PPC Mac US X; en	Referrer
0120	29 20 41 70	1 70 60 65 57 65	62 4D 69 74 21 34 31	) Apprevebric/41	TIGIGIT GI
0130	37 2e 39 20	28 40 48 54 40	40 20 20 60 69 60 65	7.9 (KHTML, 11Ke	User-agent info
0140	20 47 00 00	00 01 29 20 53	60 65 63 74 69 6F 60	17 9 Connection	
0150	3a 20 6b 65	65 70 2d 61 6c	69 76 65 0d 0a 48 6f	: keen=alive Ho	Connection info
0170	73 74 3a 20	69 20 61 20 63	6e 6e 2e 6e 65 74 0d	st. i a cnn net	
0180	Oa Od Oa	00 20 01 20 00			

- Connectionless communication
  - using only source/dest IP addresses
- Routing
  - communication across network boundaries
  - routing tables kept by routers
  - no authentication
- Fragmentation & reassembly
  - No reliability

- Connection-based communication
  - identified by source/dest IP + port (multiplexing)
- Server process "listens" to a port
  - Often determined by the application protocol (HTTP, SMTP, etc)
- Client process connects to dest IP+port
  - Source port selection usually random
- Connection established by handshake
- Reliability

### UDP

- Connectionless communication over IP
- Fast alternative to TCP
  - Only 8 bytes overhead, no handshakes
  - Stateless
- Some higher-level features
  - addressing based on IP+port (multiplexing)
  - checksums
- But many missing
  - No ACKs (unreliable)
  - No ordering
- Often used for "streaming"-like applications

traceroute to google.com (216.58.215.46), 30 hops max, 60 byte packets 1 gateway (195.134.67.1) 0.715 ms 0.789 ms 0.884 ms uoa-ilisia-1-gw.kolettir.access-link.grnet.gr (62.217.96.172) 0.763 ms 0.796 ms 0 2 grnet-ias-geant-gw.mx1.ath2.gr.geant.net (83.97.88.65) 1.574 ms 1.630 ms 1.620 ms 3 ae0.mx2.ath.gr.geant.net (62.40.98.140) 31.556 ms 31.650 ms 31.547 ms 4 ae2.mx1.mil2.it.geant.net (62.40.98.150) 25.654 ms 27.861 ms 27.793 ms 5 6 72.14.203.32 (72.14.203.32) 25.593 ms 25.766 ms 25.500 ms 7 108.170.245.73 (108.170.245.73) 64.548 ms 108.170.245.89 (108.170.245.89) 73.238 m 209.85.142.221 (209.85.142.221) 72.001 ms 72.14.238.21 (72.14.238.21) 71.999 ms 6 8 9 216.239.35.201 (216.239.35.201) 78.302 ms 78.299 ms 78.277 ms 10 209.85.251.217 (209.85.251.217) 54.466 ms 72.14.238.54 (72.14.238.54) 54.472 ms 1 11 108.170.245.1 (108.170.245.1) 52.509 ms 52.443 ms 50.669 ms 108.170.235.15 (108.170.235.15) 54.116 ms 51.975 ms 51.967 ms 12 13 par21s17-in-f14.1e100.net (216.58.215.46) 51.943 ms 54.241 ms 54.202 ms

- Time to live (TTL)
  - IP header
  - Decreased at every hop
  - If 0 the router discards and notifies the originator (ICMP time exceeded)
- Traceroute: repeatedly send packets (ICMP echo request)
  - with TTL = 1, 2, ...
  - 3 packets for every value
  - Until we reach the host (or a threshold)
  - Routers might not respond

# TCP 3-way handshake

- Connection identified by source/dest address/port
- Sequence numbers (SN) in every message
- Handshake
  - SYN(SNc)
  - SYN(SNs)-ACK(SNc)
  - ACK(SNs)
  - Data-exchange (bidirect.)



# TCP 3-way handshake

- Connection identified by source/dest address/port
- Sequence numbers (SN) in every message
- Handshake
  - SYN(SNc)
  - SYN(SNs)-ACK(SNc)
  - ACK(SNs)
  - Data-exchange (bidirect.)
- What can go wrong here?



# SYN flood

- Flood the server with SYNs
- But no ACK!
- Connections stay "half-open" on the server until they timeout
  - Keeping state consumes resources
  - Can lead to Denial of Service (DoS)



**TCP-SYN Flood Attack** 

# SYN flood

- Flood the server with SYNs
- But no ACK!
- Connections stay "half-open" on the server until they timeout
  - Keeping state consumes resources
  - Can lead to Denial of Service (DoS)
- Can the server limit the number of SYNs from the same host?
  - No! the attacker can easily "spoof" the sender IP



TCP-SYN Flood Attack

• Can we impersonate a client?



- Can we impersonate a client?
  - Trivial if we control an intermediate router!
  - If we don't?



- Can we impersonate a client?
  - Trivial if we control an intermediate router!
  - If we don't?
- We can still send packets with a spoofed IP, without access to the replies



- Can we impersonate a client?
  - Trivial if we control an intermediate router!
  - If we don't?
- We can still send packets with a spoofed IP, without access to the replies
  - It's sufficient to guess SNs for the ACK!
  - $A(C) \rightarrow S : SYN(SNa)$
  - S  $\rightarrow$  C : SYN(SNs)-ACK(SNa)
  - A(C)  $\rightarrow$  S : ACK(SNs)



- Can we guess the server's SN?
- Initial Sequence Number
  - Counter incremented over time and for every new connection
  - Predictable!

- Bypass IP-based authorization
  - Still widely-used today
  - SMTP, web-services, firewall IP white/black-listing, etc

- Bypass IP-based authorization
  - Still widely-used today
  - SMTP, web-services, firewall IP white/black-listing, etc
- Inject data to existing connection
  - DNS response (UDP, no SN at all!)

- Bypass IP-based authorization
  - Still widely-used today
  - SMTP, web-services, firewall IP white/black-listing, etc
- Inject data to existing connection
  - DNS response (UDP, no SN at all!)
- Reset existing connections (RST)
  - SNc is needed, but only approximately
  - Denial of service, or exploit to break some other protocol

Solution?

- Routers expect ISN to be increasing
  - Protocol bugs are hard to fix (compared to implementation bugs)

Solution?

- Routers expect ISN to be increasing
  - Protocol bugs are hard to fix (compared to implementation bugs)
- Different ISN for each client! How?

Solution?

- Routers expect ISN to be increasing
  - Protocol bugs are hard to fix (compared to implementation bugs)
- Different ISN for each client! How?
- RFC 6528:

ISN = Timer + H(localip, localport, remoteip, remoteport, secretkey)

- Why we include secretkey?

Solution?

- Routers expect ISN to be increasing
  - Protocol bugs are hard to fix (compared to implementation bugs)
- Different ISN for each client! How?
- RFC 6528:

ISN = Timer + H(localip, localport, remoteip, remoteport, secretkey)

- Why we include secretkey?
- But still not perfect
  - 32bit space possible to guess
  - Still trivial if we control routers
  - Conclusion : For serious security we need to build on top of TCP

- Remotely consume a resource of the server
  - Bandwidth,
  - CPU
  - Memory
  - ...
- Until the resource is depleted
  - no more clients can connect

- Typically involves some sort of flooding
- SYN flooding

- Typically involves some sort of flooding
- SYN flooding
- Ping flooding
  - ICMP echo request

~\$ ping google.com
PING google.com (216.58.215.46) 56(84) bytes of data.
64 bytes from par21s17-in-f14.1e100.net (216.58.215.46): icmp\_seq=1 ttl=47 time=52

- Low-level protocol (no use of TCP), can send packets fast
- The server sends a reply back

- Typically involves some sort of flooding
- SYN flooding
- Ping flooding
  - ICMP echo request

~\$ ping google.com
PING google.com (216.58.215.46) 56(84) bytes of data.
64 bytes from par21s17-in-f14.1e100.net (216.58.215.46): icmp\_seq=1 ttl=47 time=52

- Low-level protocol (no use of TCP), can send packets fast
- The server sends a reply back
- We need more resources than the server
  - Use many senders at once

## Smurf attack

- Send an Echo to a broadcast address
- Spoof the sender IP with the sender's
- All machines flood the victim



### Distributed Denial of Service (DDoS)

- Compromise hosts via virus, worm, etc
- Coordinate the attack
- Hard to distinguish from legitimate users



### Fork Bomb

- Another kind of DoS
- Fork, and keep forking in the children
  - exponential growth
- Consumes OS resources for process management
- Try this in your own machine!

~\$ :(){ :|:& };:

# some other terminal
~\$ ls
bash: fork: retry: Resource temporarily unavailable

Crucial properties

- Packets need to arrive from multiple source IPs
  - otherwise trivial to filter
- The adversary spoofs the sender IP
  - but does not get replies!
- The server needs to keep state for all fake clients



Three-Way Handshake

Crucial properties

- Packets need to arrive from multiple source IPs
  - otherwise trivial to filter
- The adversary spoofs the sender IP
  - but does not get replies!
- The server needs to keep state for all fake clients

#### Idea

- Make the client store the state!
- Only store state in the server for clients that have proven to get our replies



Three-Way Handshake



Three-Way Handshake

- Encode the state in the SNs send to client
  - Then forget about the connection (no state!)
- Check the SNs contained in the client's ACK
  - Store state only if ok
- Spoofing the source is useless
  - The adversary needs to control the users or the network



- Encode the state in the SNs send to client
  - Then forget about the connection (no state!)
- Check the SNs contained in the client's ACK
  - Store state only if ok
- Spoofing the source is useless
  - The adversary needs to control the users or the network
- One approach
  - SNs = H(ports, ips, key, time) || time



Three-Way Handshake

- One approach
  - SNs = H(ports, ips, key, time) || time
- Protocol compliant
- Problem 1: when ACK is lost
  - The server is supposed to resend SYN-ACK
  - But it won't, it does not store connections!
- Problem 2: options sent in SYN are lost
- Solution: only use under attack



Three-Way Handshake

### Preventing DoS

- Client needs to solve a puzzle to connect
  - Eg: brute-force a hash (within a controlled range of values)
- Generic solution, also used to prevent spam
- But requires changes to both client and server

### Achieving secure communication

- TCP is an inherently insecure protocol
  - no security against an adversary who controls the network
  - limited security against an adversary who simply participates
- Solution
  - Use crypto to build a secure connection over an insecure network
- Most widely used: TLS
  - Also: IPSec, SSH, ...
- We can also tunnel the traffic of an entire network
  - Secure VPN

- Widely used in web-browsers
- Crucial use of crypto:
  - Assymetric-crypt: exchange keys
  - Symmetric crypto: encrypt the main traffic
  - Digital signatures: authentication

### TLS handshake



### TLS handshake

Client		Serve
versio	n random1 TI S BSA WITH	AES 128 CBC SHA
Clie	ntHello)	>
version, r	andom2, session id, cipher,	PKRSA, Sign(SKCA, PK)
	ServerHello	Certificate
	E(PKRSA, premaster	_key)
	ready	
<	ready	
L Extract m	aster_key from random1 + ra	ndom2 + premaster_key

- Make sure to stay up to date
  - SSL 2.0, 1995, Deprecated in 2011
  - SSL 3.0, 1996, Deprecated in 2015
  - TLS 1.0, 1999, Deprecated in 2020
  - TLS 1.1, 2006, Deprecated in 2020
  - TLS 1.2, 2008
  - TLS 1.3, 2018
- But it's hard to do while maintaining compatibility

- POODLE
  - Man in the middle
  - Block the connection until the client tries SSLv3
  - The server will hapilly accept it
  - TLS\_FALLBACK\_SCSV: tell the server we are downgrading

- Ross Anderson, Security Engineering, Chapter 21
- A look back at "security problems in the TCP/IP protocol suite
- SYN cookies
- Bypassing domain control verification with DNS response spoofing