Network Security

BME-TMIT
Information and Network Security
feher.gabor@tmit.bme.hu

ISO/OSI - Internet

- ISO/OSI 1983
  - International Standards Organization Open Systems Interconnection Basic Reference Model

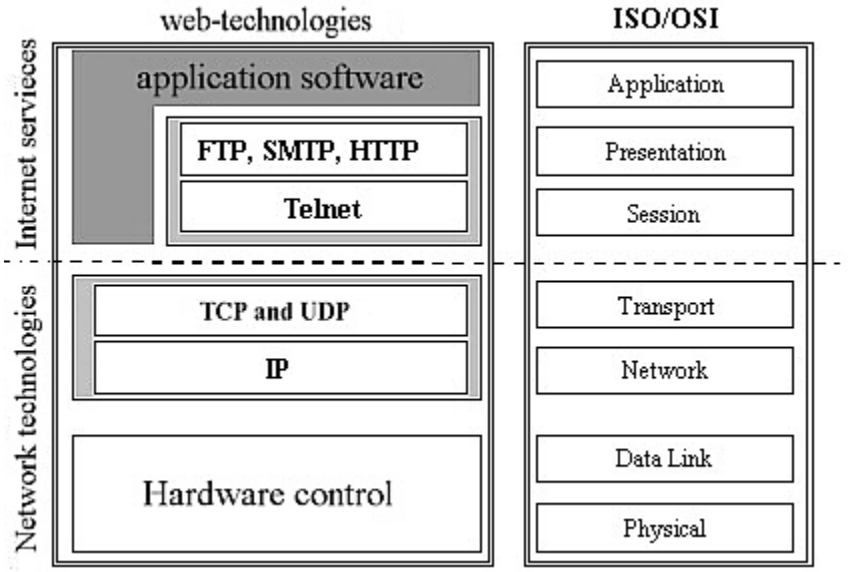

LAYER 1 & 2 SECURITY

Network topologies

BUS

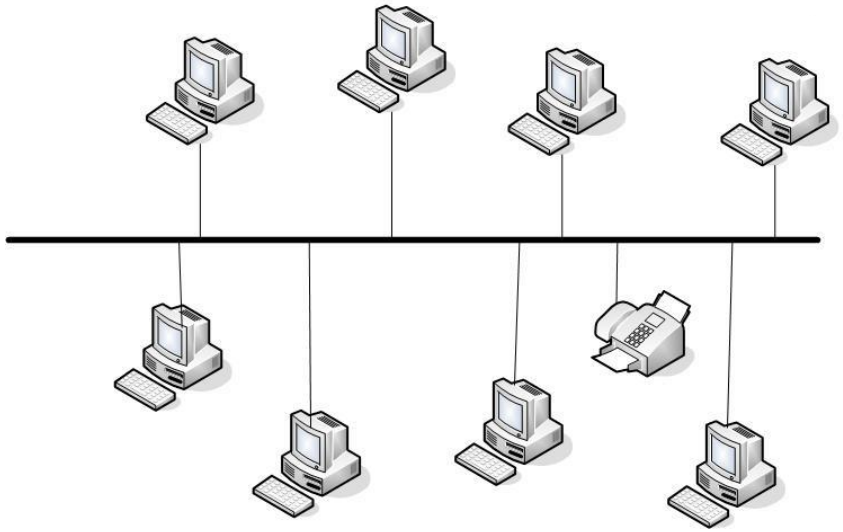

In the case of the bus topology, everybody could see others' traffic

In the case of the star topology, users are separated.

STAR

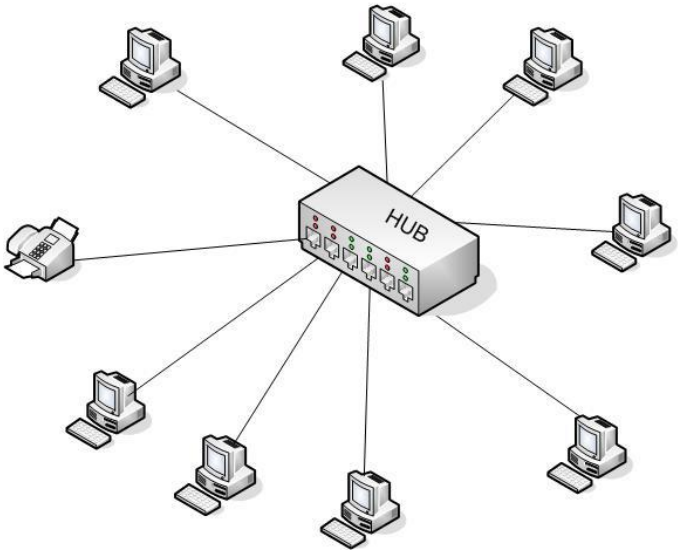

There are also other topologies, like ring or mesh

Today LAN networks

- Local Area Networks

- Ethernet (very cheap, high speed)
  - Star topology based (formerly it was a bus)
  - Network elements: Repeater, bridge, hub, switch
    - Extending, joining networks, forwarding frames
- WiFi (cheap, reasonable speed)
  - Logical star topology (in fact is like a bus)
  - Network elements: access point
    - Mobility between access points
- Others: radio, optical…

Attacks: PHY / LAN eavesdropping – Identity theft

- Physical network (Layer 1)
  - Wiretapping – telephone tapping
    - E.g.: Clipper chip – 1993
  - Cable damage, jamming (service disruption)

- Identity theft
  - MAC address: unique address for identification in local communication
    - Usually fixed by factory, but be changed
  - MAC cloning: Take others' MAC addresses

Ethernet switches

- CAM table (Content Addressable Memory)
  - Speed up frame forwarding

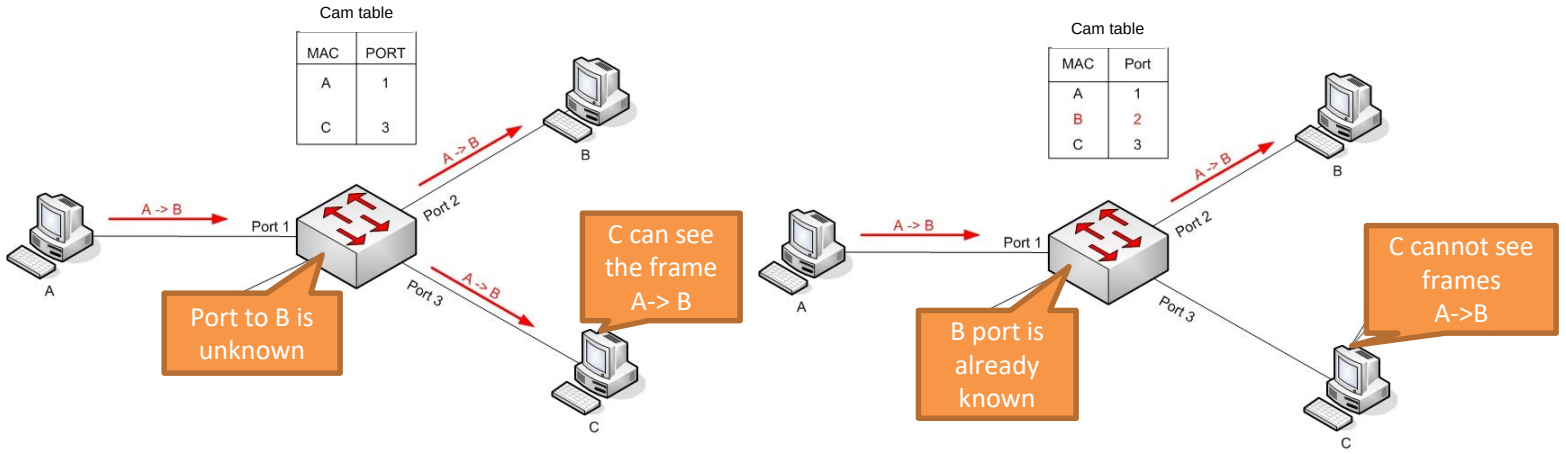

  - CAM has finite size (~100.000)

Attacks: MAC flooding

- Attack against switches
- MAC flooding
  - Fill the CAM table with random MAC addresses
    - -> Switch becomes a hub
  - Protection
    - fix MAC tables (Cisco port security)
    - disable the port on overflow
    - IEEE 802.1X

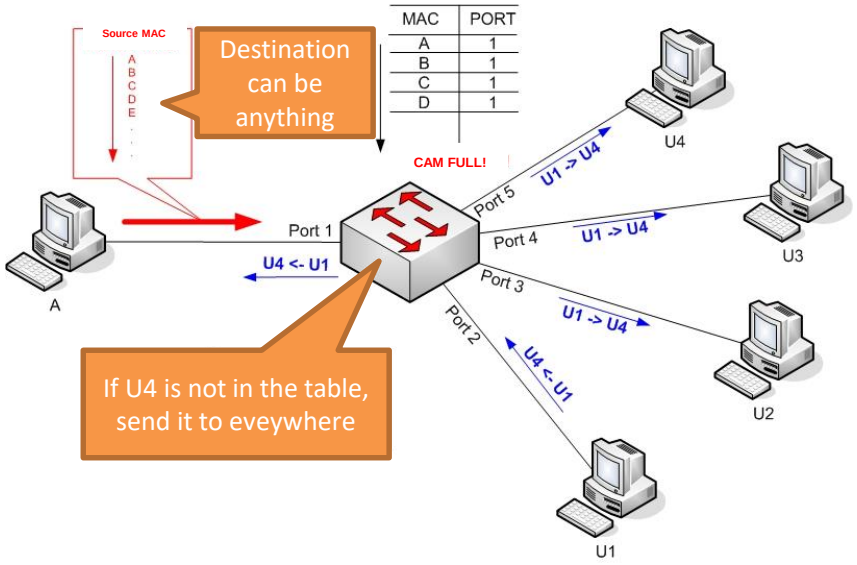

802.1X requires port authentication before the line could be used. Users can be authenticated.

LAYER 2 & 3 SECURITY

Internet

- Advanced Research Projects Agency - ARPA
  - ARPANET: October 29, 1969 UCLA and SRI International
    - University of California, Los Angeles and Stanford Research Institute
  - January 1, 1983 TCP/IP in ARPANET
    - DARPA / Defense Advanced Research Projects Agency
    - TCP/IP: Vinton Cerf and Robert Kahn (Stanford) ~ 1973

IP addresses

- IPv4, IPv6 (v5? – ST2)
- IPv9 RFC 1606 ☺

- Addressing
  - Multicast (Class D)
  - Experimental (Class E)
  - Private
  - Loopback (A) (127. …)
  - Zero addresses (A) (0. …)

Class	Leftmost bits	Start address	Finish address
A	0xxx	0.0.0.0	127.255.255.255
B	10xx	128.0.0.0	191.255.255.255
C	110x	192.0.0.0	223.255.255.255
D	1110	224.0.0.0	239.255.255.255
E	1111	240.0.0.0	255.255.255.255

Name	IP address range	number of addresses	*classful* description	largest CIDR block	defined in
24-bit block	10.0.0.0 – 10.255.255.255	16,777,216	single class A, 256 contiguous class Bs	10.0.0.0/8	RFC 1597 (obsolete), RFC 1918
20-bit block	172.16.0.0 – 172.31.255.255	1,048,576	16 contiguous class Bs	172.16.0.0/12	
16-bit block	192.168.0.0 – 192.168.255.255	65,536	single class B, 256 contiguous class Cs	192.168.0.0/16	

IPv6 addresses

- Addresses
  - Anycast
  - ::/96 IPv4 compatibility
  - ::/128 unspec
  - ::1/128 loopback

```
2001:0db8:3c4d:0015:0000:0000:abcd:ef12
_____________|____|__________________
global prefix subnet  Interface ID
```

```
IPv6 Prefix    Allocation
-------------------------------------
0000::/8       Reserved by IETF
2000::/3       Global Unicast
FC00::/7       Unique Local Unicast
FE80::/10      Link Local Unicast
FF00::/8       Multicast

These blocks are reserved for examples
and documentation
-------------------------------------
3fff:ffff::/32
2001:0DB8::/32   EXAMPLENET-WF
```

Attack: IP spoofing

- Sender sets fake IP address
  - Non-blind spoofing: attacker sees the return packets. Assuming same subnet
  - Blind spoofing: attacker does not see the return packets. Response should be predicted!

- Protection: filtering at the router
  - Ingress
  - Egress

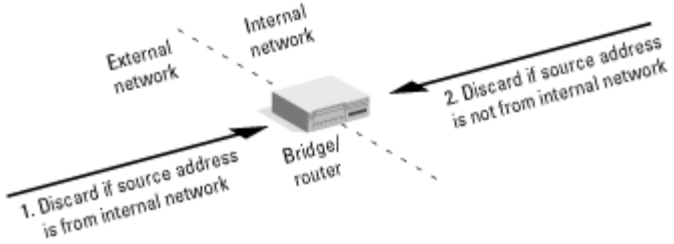

IP & LAN addresses

- ARP - Address Resolution Protocol
  - IP -> LAN address (HW address / Protocol address)
  - Gratuitous ARP (resolve own IP address)
    - Find conflicts (Should be no answer)
    - Update old entries at other machines
  - Cache addresses
  - Always 28 bytes for Ethernet like
    - | HW MAC | proto type (Eth=1) | Proto addr. Type (IP=0x800) | MAC size | IP size | operation (1=request, 2=reply) | SRC MAC | SRC IP | DST MAC | DST IP |

- Inverse ARP
  - ATM
- Reverse ARP
  - BOOTP, DHCP

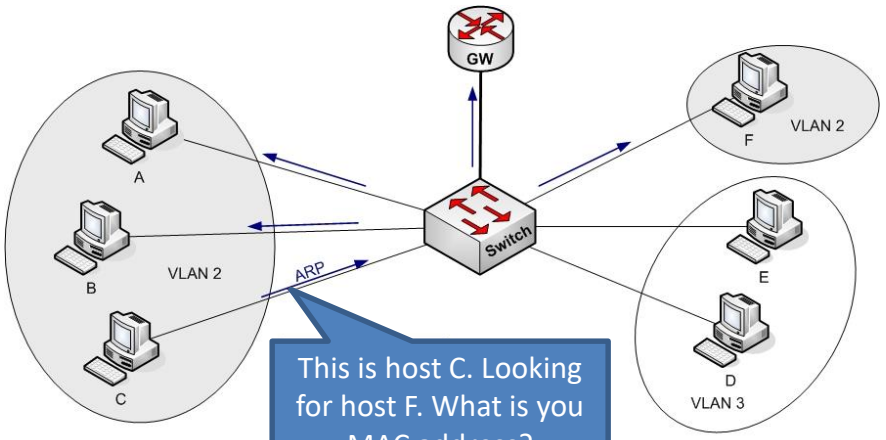

Attack: ARP poisoning

- Perform Man-in-the-Middle (MiM) attack
- Send fake ARP packets
- Route traffic to destination
- Protection:
  - Inspect ARP traffic
  - Lock changes

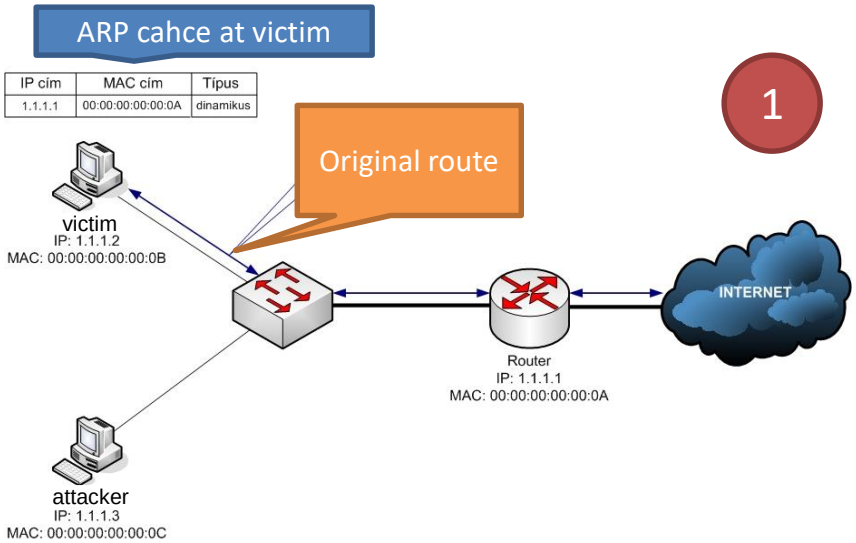

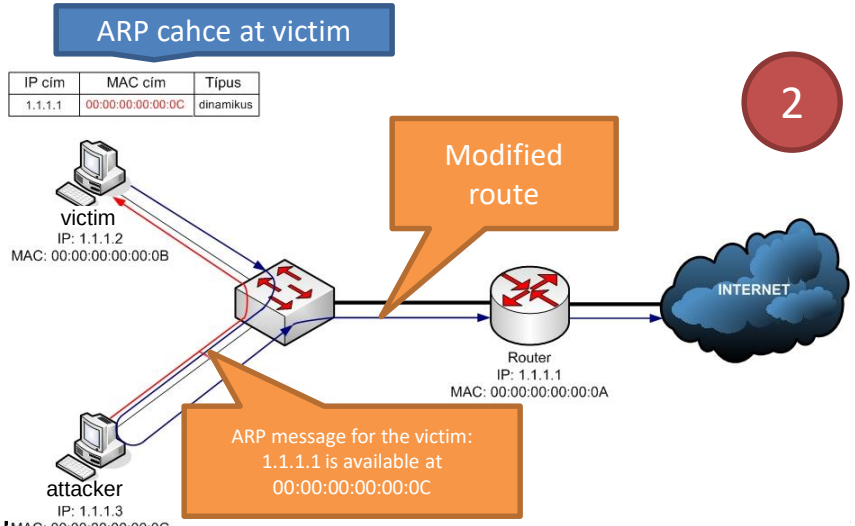

BOOTP & DHCP

- BOOTP (Bootstrap protocol)
  - Static MAC and IP pairs
  - Request (MAC address) and reply (IP address)
    - IP, Server IP, GW IP, Server host name, boot filename
- DHCP (Dynamic Host Configuration Protocol)
  - Pool (or static) IP for a possibly unknown MAC
  - BOOTP+options
  - DHCPDISCOVER -> DHCPOFFER -> DHCPREQUEST …

- Request to broadcast MAC
- Wait for the first reply

Attack: DHCP starvation + rogue DHCP server

- Attacker asks IPs for spoofed MACs
  - IP address pool will be exhausted
  - Legitimate users can’t get IP addresses
- DoS (Denial of Service) or rogue DHCP setup
  - Fake GW and DNS
- Protection:
  - limit MAC addresses per port
  - IEEE 802.1X
  - Certificates (configuration files)

Routing

- Routing IGP and EGP / Interior and Exterior Gateway Protocol
  - AS - Autonomous System (2 byte -> 4 byte)
    - IANA Internet Assigned Numbers Authority and the Regional Internet Registries (RIR)
  - IGP: OSPF (link state), RIP (distance vector)
  - EGP: BGP

- Routing
  - Next hop
  - Metric

- Classless
  - CIDR: Classless Inter Domain Routing

Routing Information Protocol
RIP v1 & v2

- Next hop: distance
  - Distance vectors
  - Broadcast (v1) or multicast (v2) to everybody
    - Update in every 30 sec
  - Routers calculate new routing tables
    - Shorter, newer, aging

- RIP v1 (1988) problems
  - Max 15 hop
  - Shortest is not the fastest
  - Huge signaling traffic
  - Slow convergence
  - Counting to infinity

- Security in RIP v1
  - not at all

RIP v2 (1994)

- Split Horizon
  - No backward advertisement
- Triggered Updates , Poison Reverse, Hold-down timer
  - Broken links are signaled immediately, set distance to infinity, no refresh for a certain time amount
- Introducing subnet masks
- RIP v2 Security
  - Open password
  - Message authentication with extensions

Open Shortest Path First: OSPF (1997)

- LSA
  - Link State Advertisement
  - LSA flooding
  - Describe a metric (speed or anything)
- Routers calculate shortest path based on all LSA using Dykstra

- IR - Internal Router, ABR - Area Border Router, BR - Backbone Router, ASBR - Autonomous System Boundary Router, DR - Designated Router, BDR - Backup Designated Router

OSPF Security

- Messages are authenticated
  - Null (just integrity protection)
  - Password
  - Crypto
    - Password + monotone increasing counter + hash
    - Protection against fraud, modification, reply

BGP (1995: BGP-4)

- De facto EGP routing
- Policy based!
  - Often replaced by static routing tables
- Huge resources
  - Many entries, big tables
- BGP security
  - Message authentication
  - BUT
    - No address space authorization check
    - No AS authentication
    - No BGP router authentication
- S-BGP
  - Communication protection using IPSec
  - PKI for router, AS and address space authentication
    - Attestation: Authorization over an address space
    - Signed digitally
    - Protection against bad refresh
  - Implementation: PC plus router
  - Not deployed

YouTube and routing

- On February 2008:
  - Pakistan was about to ban YouTube
    - Pakistan Telecom created a black hole routing for YouTube
  - YouTube became unreachable form most of the Internet
    - ISP in Hong Kong get this routing info and propagated to the world
    - All YouTube traffic went to Pakistan (nice self-DoS)
    - Pakistan was shut down during the fix
- How could happen?
  - BGP routers trust each other without testing!
  - Pakistan Telecom set up a small subnet in the BGP table. Using longest match it became the winner against the still existing YouTube block.

TOR

- Provides anonymity
  - EFF sponsor - Electronic Frontier Foundation
    - defending free speech, privacy, innovation, and consumer rights

- Currently ~1000 routers (volunteers)

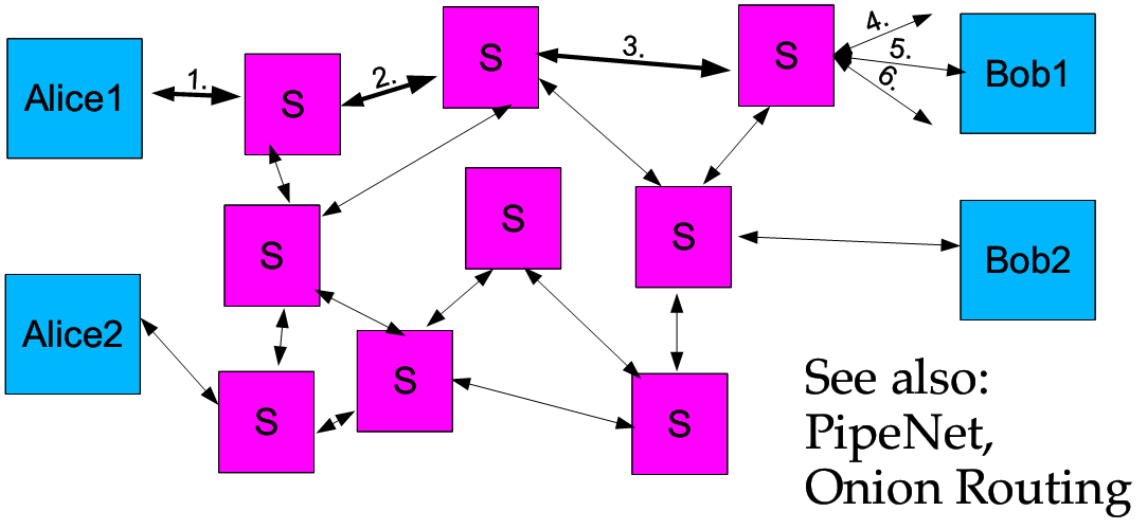

Fragmentation & Security

- Overlapping fragments
  - Different fragments -> what packet does the destination get?
  - Implementation problems -> teardrop (Win95/NT, Linux 2.0)
- Buffer Full exploit
- Fragment overrun
  - Write over the max. size of the packet (usually 65535) -> Implementation problems
- Incomplete datagram
  - Too many incomplete fragments

- Bypass IDS, firewalls
  - Problem if security devices mishandle fragments

ICMP

- Internet Control Message Protocol
  - Echo
  - Destination Unreachable
  - Router Advertisement
  - Traceroute
  - …
- SMURF
  - DoS attack with spoofed ping to broadcast IP
  - Fixes
    - do not respond to ping (broadcast ping)
    - Routers do not forward packets to broadcast addresses
  - SMURF amplifier – network that can generate large number of echo responses to spoofed source
- Fake ICMP messages
  - “Time exceeded” or ”destination unreachable”
- Ping of death
  - Send an oversized ping packet (crash due to implementation problem)
  - Fixed around 1998

DNS

- Resolve names to IP addresses
  - Using UDP (port 53)
- Security problems
  - Attacker can send anything…
    - Eavesdropping on the LAN
    - Guessing the DNS query (16 bit nonce & source port)
  - Cache poisoning
    - False info in the response
- DNSSEC
  - Signature on DNS actions
    - Origin authentication of DNS data
    - Data integrity
    - Authenticated denial of existence
  - Deployment? Not really…
- International domain names

TCP/IP

- Transport Control Protocol (RFC 793)
  - Reliability
  - Congestion avoidance
    - additive-increase-multiplicative-decrease (AIMD)
    - Slowstart
    - TCP Tahoe and Reno, Vegas, New Reno, CUBIC, SACK, fast TCP, …
  - Window size
- Ports
  - Server ports
  - Client ports
- IP ID
  - Counter or random

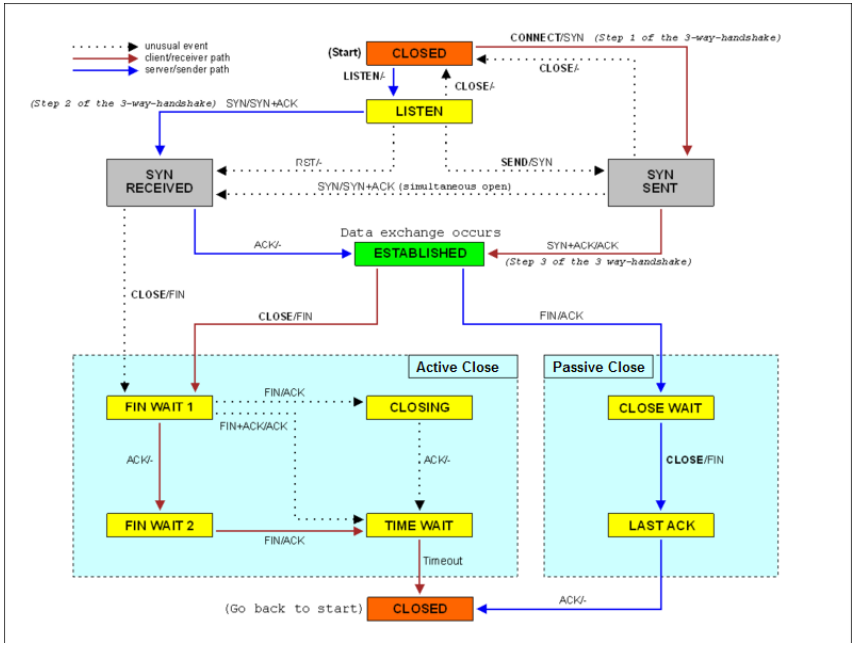

TCP connection

- SYN for synchronizing
  - Set initial sequence number
- ACK for acknowledgements
  - Acks sequence number
- FIN for indicating no more data
  - Client, server or both initiated

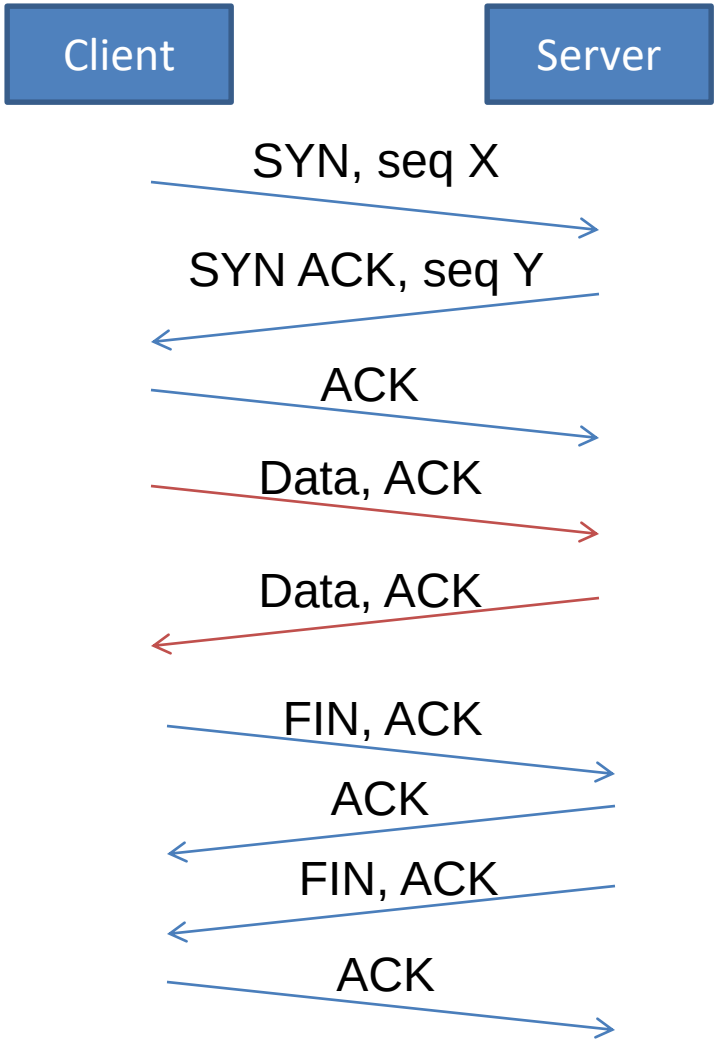

TCP session hijacking

- Hijack a connection server <-> client
  - Using MiM
  - Blind
    - Guess TCP sequence numbers
      - Predictable sequence numbers
        - » 1. Silence client using SYN flood
        - » 2. Get info for SEQ prediction
        - » 3. Use blind packets to perform the attack
    - How to guess?
      - Old TCP stacks has weak sequence number randomness
        - » Win98: SEQ num is the actual time
        - » Usually small increments to the previous SEQ num
        - » Using IP ID