Hálózat biztonság

BME - TMIT
Médiabiztonság
feher.gabor@tmit.bme.hu

ISO/OSI - Internet

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

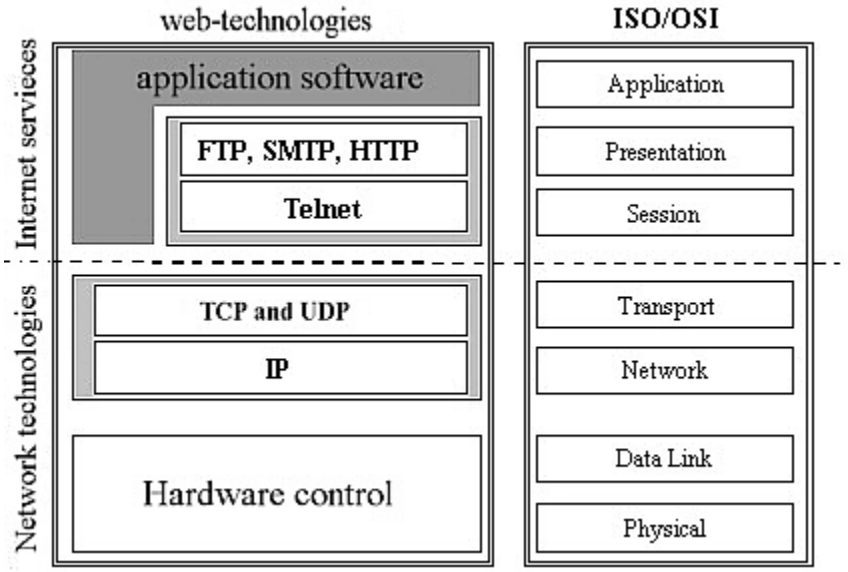

Ethernet topologies

BUS

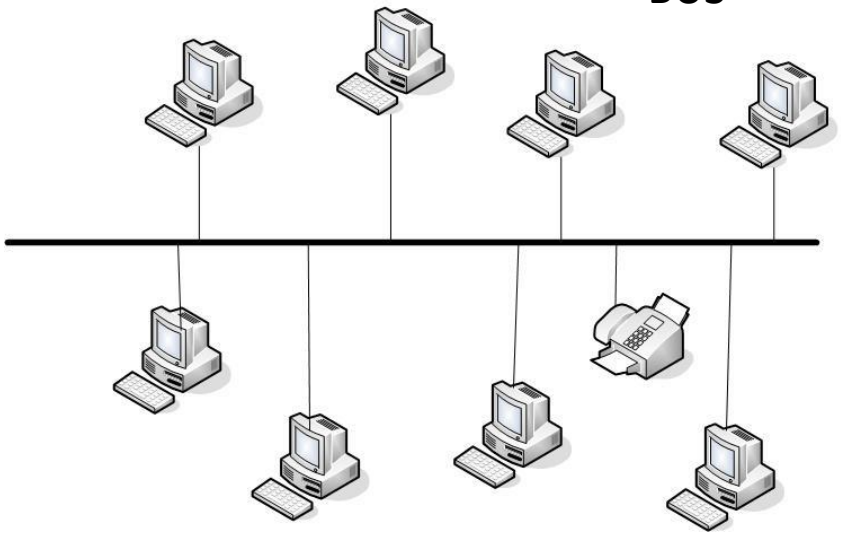

STAR

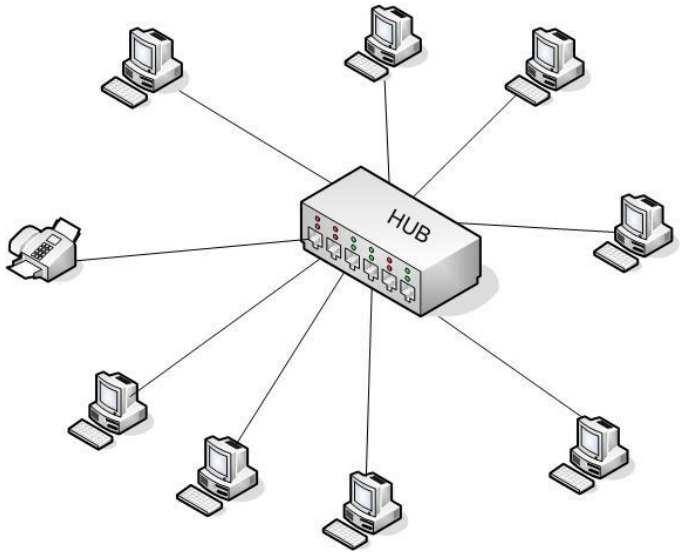

Attacks: PHY / LAN eavesdropping – Identity theft

- PHY
  - Wiretapping – telephone tapping
    - E.g.: Clipper chip – 1993
  - Cable damage
- Identity theft
  - MAC cloning against MAC ACL protection
  - Protection: 802.1X

Ethernet switches

- CAM table (Content Addressable Memory)

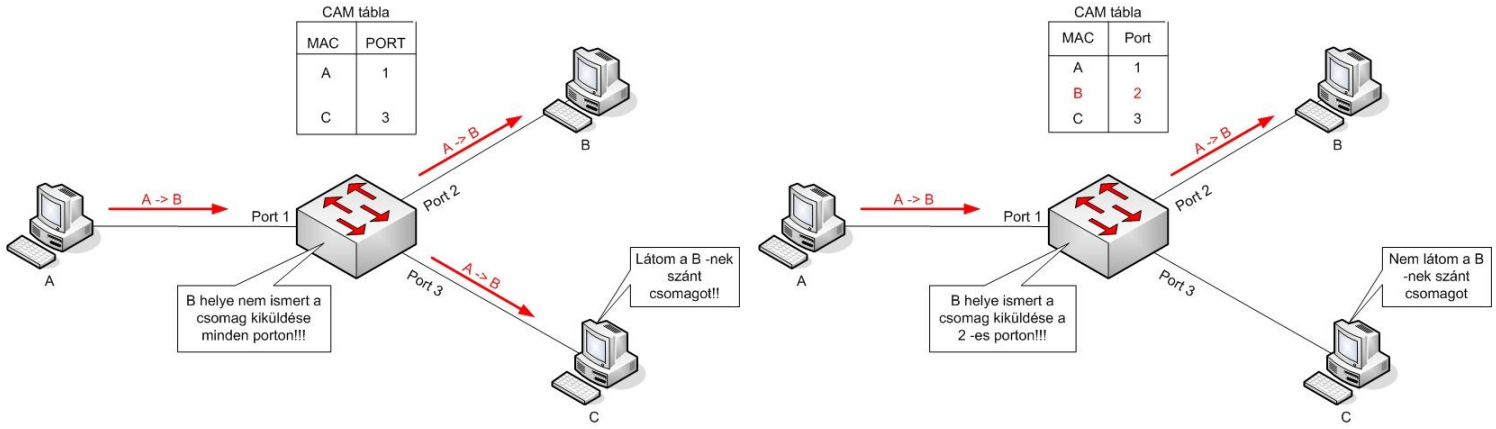

  - CAM has finite size (~100.000)

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

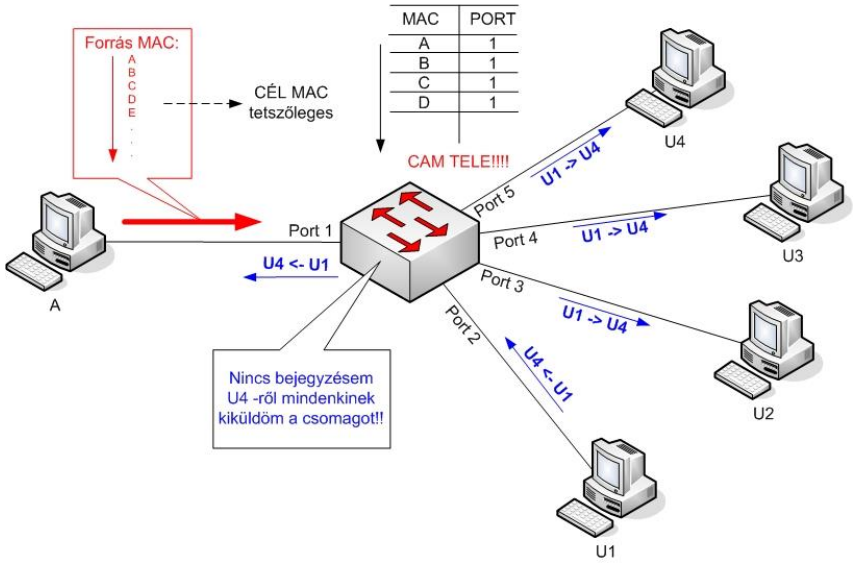

VLAN / Virtual LAN

- VLAN ports and trunks
  - DTP - Dynamic Trunking Protocol: IEEE 802.1Q
    - Security, QoS, Administration (Separate location and LAN)

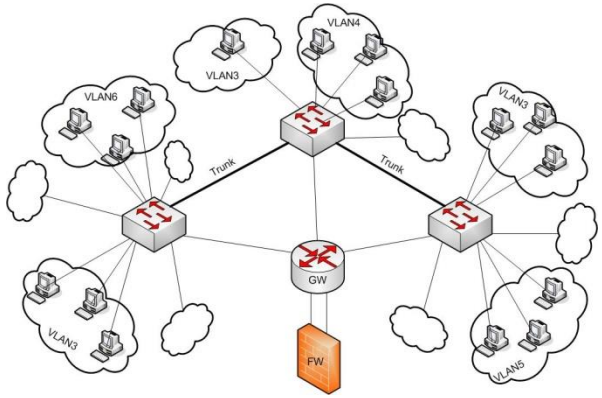

VLAN Attacks

- VLAN hopping
  - Sending fake DTP message
  - Protection: disable auto mode

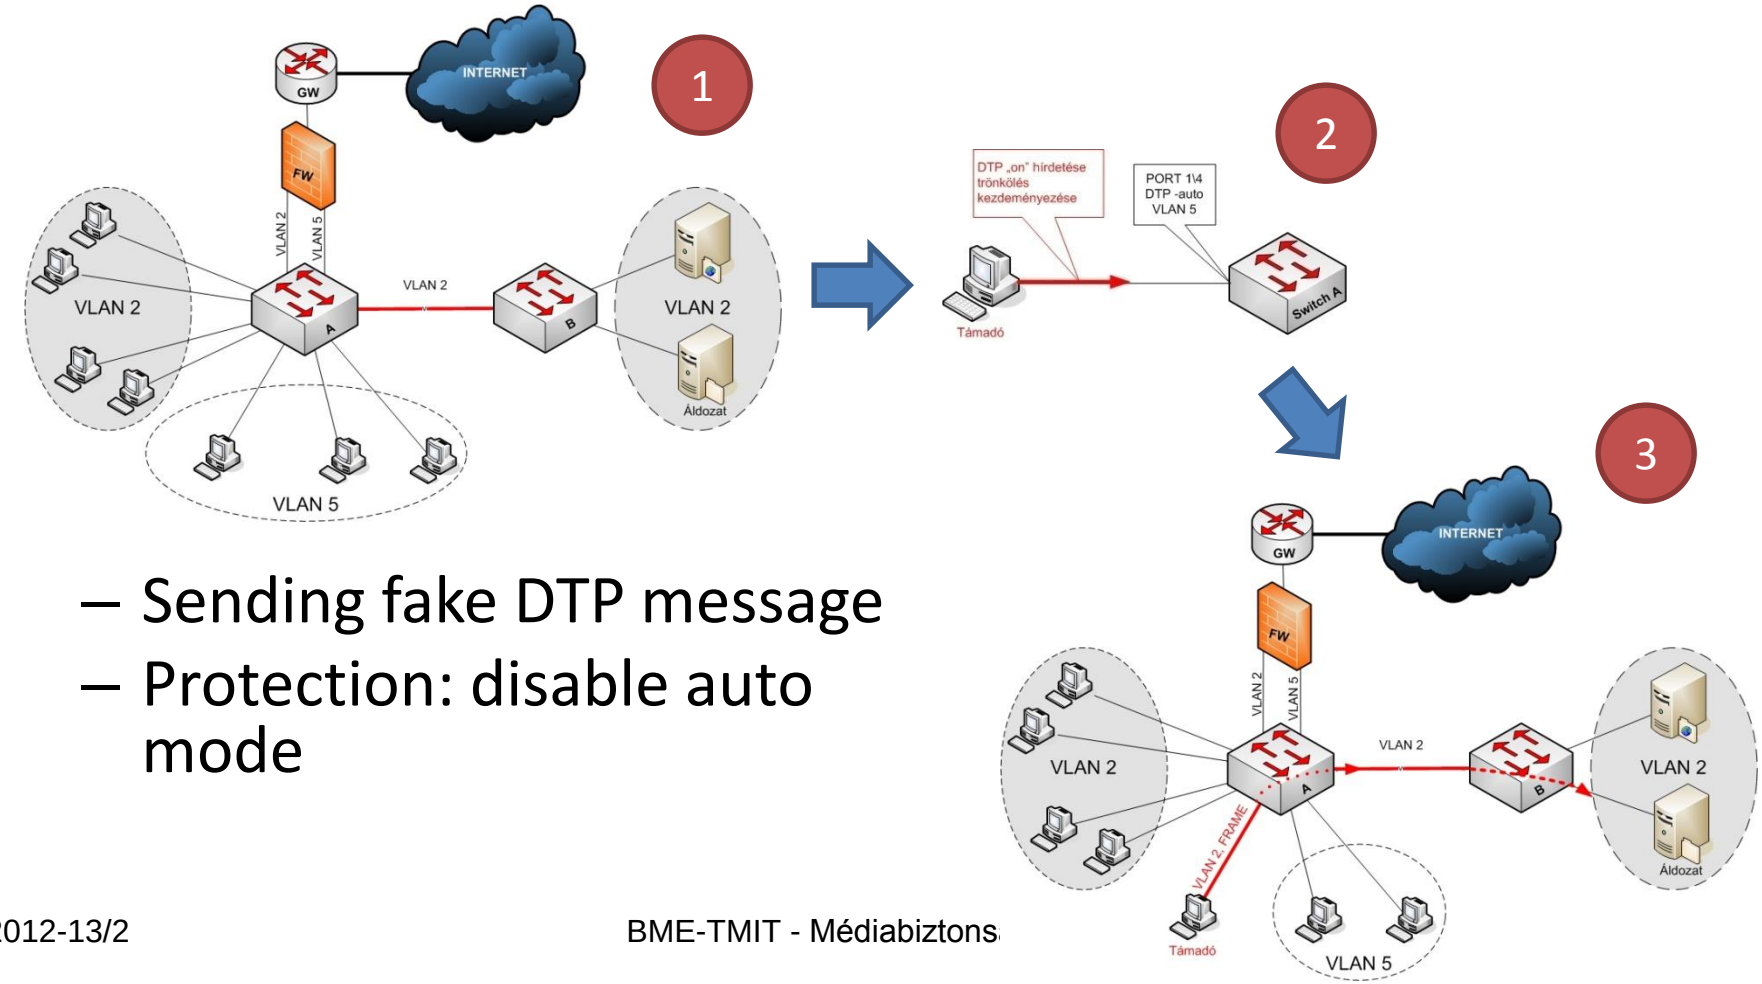

Attack: Double-Encapsulated 802.1Q

- VLAN hopping 2.
- Native VLAN support for compatibility reason (802.3) -> Native VLAN = VLAN without tagging
  - VLAN tag of the attacker and the native VLAN tag should be the same
  - Attacker to victim only

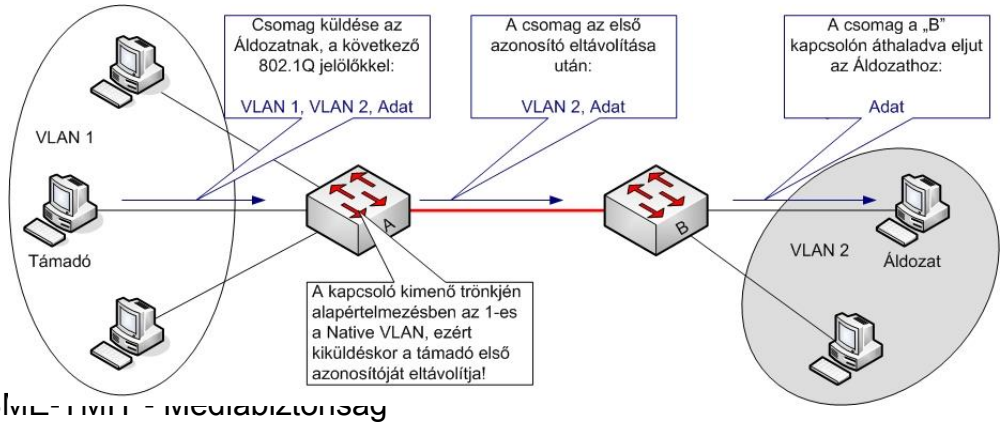

Spanning tree

- IEEE 802.1D
- Spanning tree creation
  1. Elect root
  2. Select root port on switches (sw -> root on minimal cost)
  3. Select designated port on switches (lan -> root on minimal cost)
  4. Block other ports

- Using BPDU – Bridge Protocol Data Unit
  - Root ID, Sender ID
  - ID: priority (0-32767) and MAC
  - Multicast address: 01-80-c2-00-00-00

- CST (Common ST) and PVST (Per VLAN ST)
  - 1 common for all vs. 1 VLAN 1 ST

Spanning tree creation

- Elect root
  - Send messages with root = self on every port
  - Stop sending if there is a higher priority (or MAC if equal)
- Select root port
  - Based on cost
- Select designated port
  - If there are hubs as well…
  - On every segment
- Block other ports

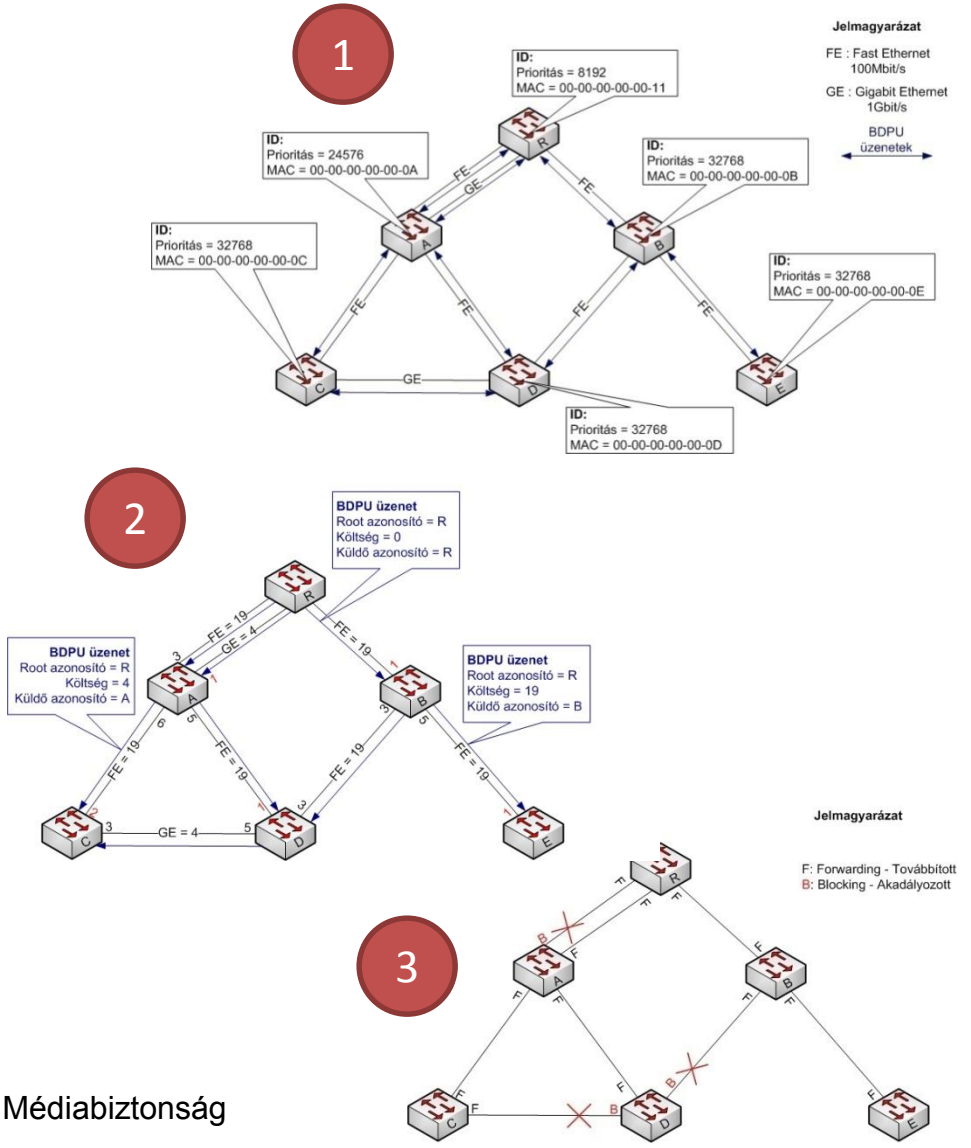

Attack: STP dual-homed root

- Fake BPDU messages to elect new root
  - Reroute traffic
- Protection:
  - BPDU guard
    - No BPDU messages on a given port
  - Root guard
    - Impossible to become root on a given port

1

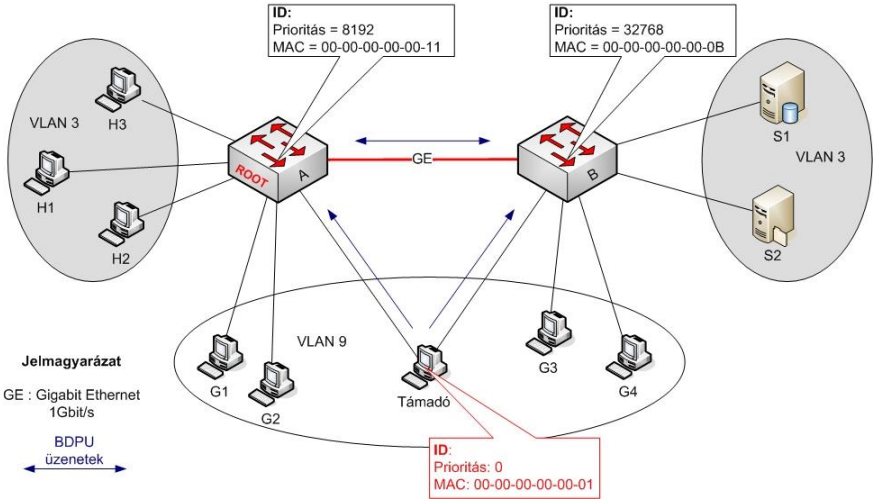

2

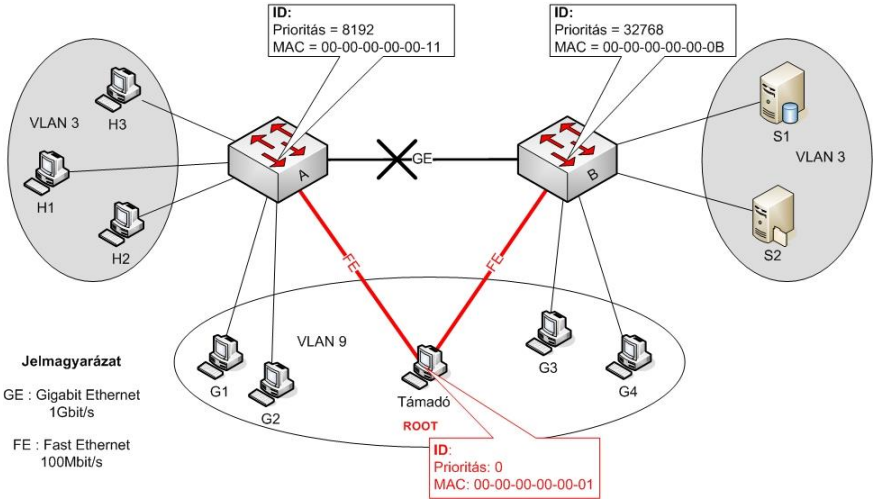

Attack: STP single-homed root

- Fake BPDU messages to elect new root
  - Reroute traffic
- Protection:
  - BPDU guard
    - No BPDU messages on a given port
  - Root guard
    - Impossible to become root on a given port

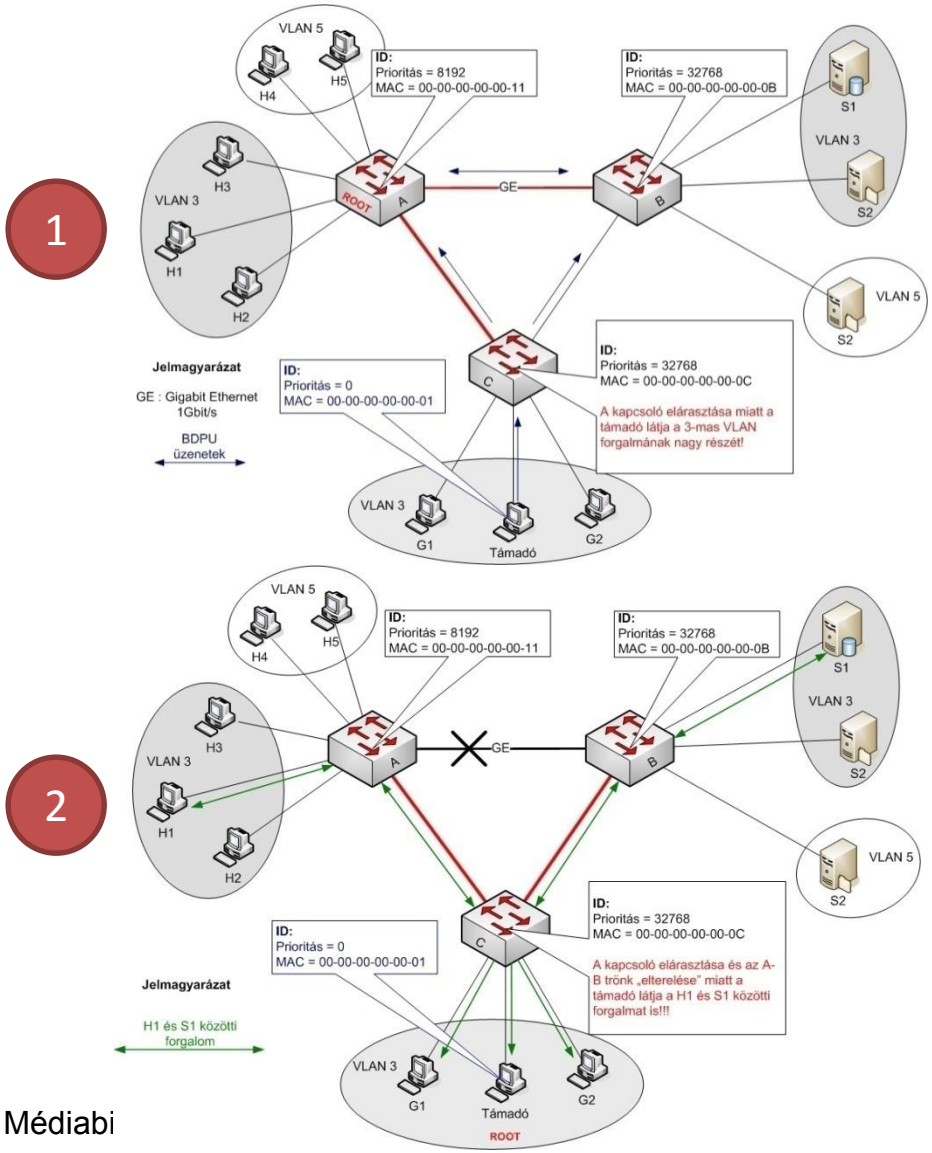

Internet

- 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

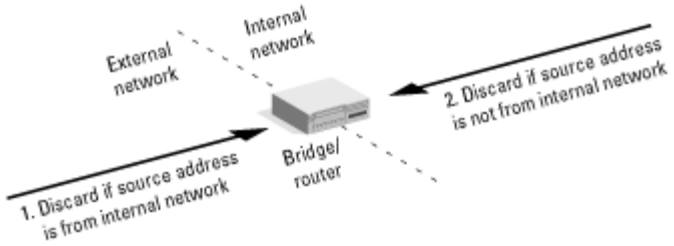

IP & LAN addresses

- 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

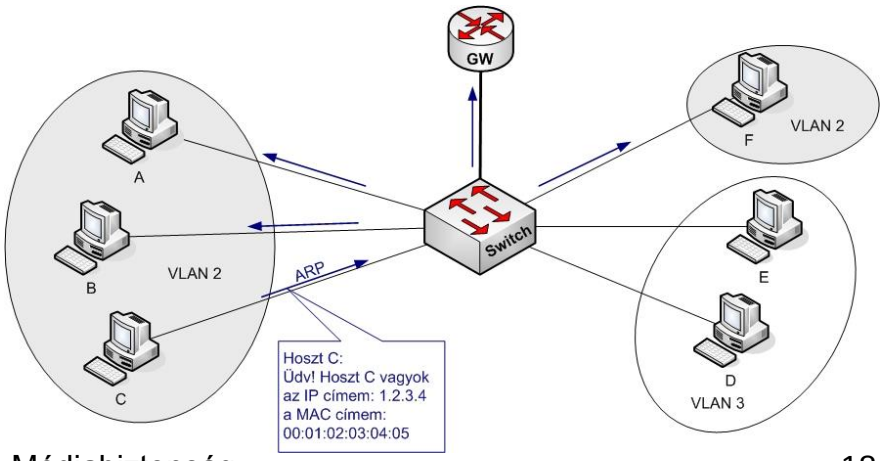

Attack: ARP poisoning

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

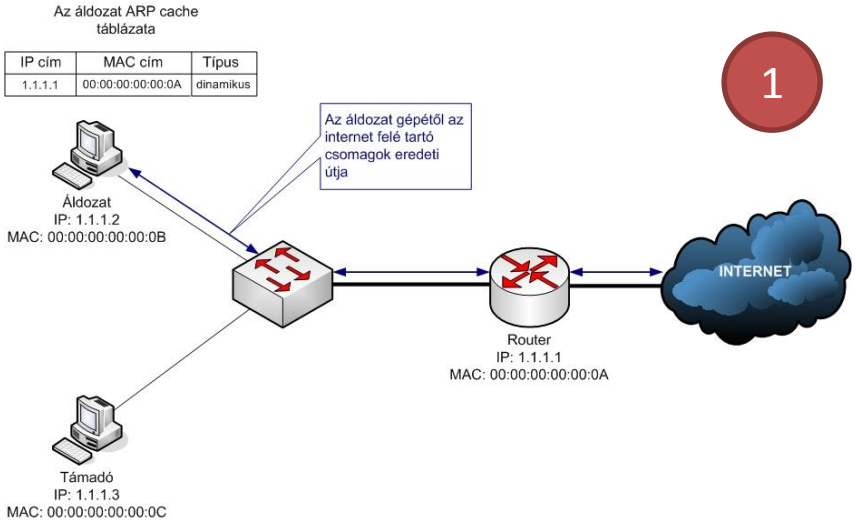

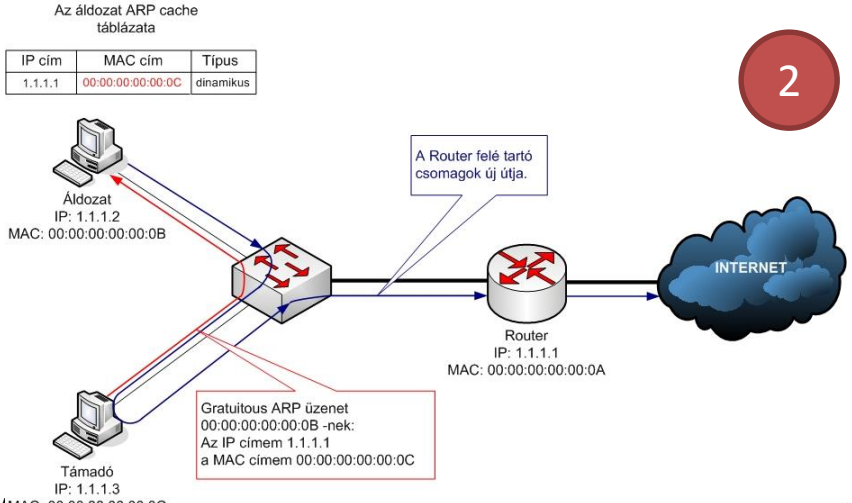

BOOTP & DHCP

- 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 + fake 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 rouge 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

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
- 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

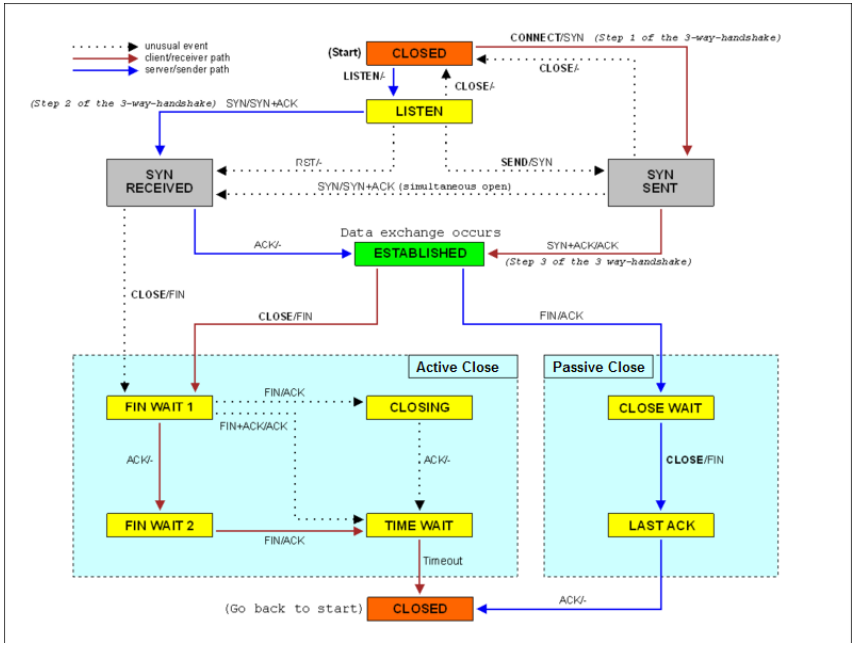

TCP connection

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

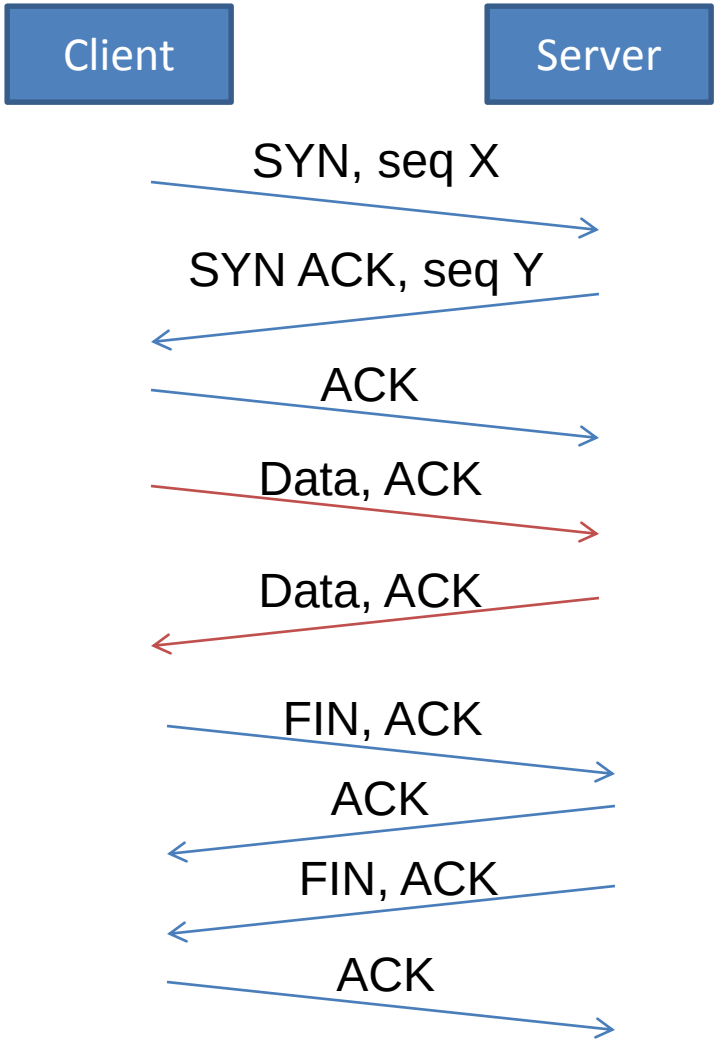

TCP session hijacking

- 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

TCP SYN támadás

- SYN elárasztásos DoS (SYS floofing)
  - SYN kérelmek küldése hamis (spoofed) IP címről
    - A szerver helyet foglal a kapcsolatnak (backlog)
    - A kapcsolat félig nyitott állapotba kerül (half open state)
  - A támadó sohasem nyugtázza a választ
    - Hamis IP esetén nem is tudja megtenni
    - A szervernél lévő memória tár véges (128, 1024, …), A bejegyzések ideiglenesek, de percekre maradnak (ismétlések a feltételezett hiba miatt)
    - Amikor a memóriaterület betelik, a szerver nem tud több kapcsolatot fogadni
  - Sikeres támadás esetén nem senki sem tud TCP kapcsolatot kezdeményezni a szerverhez
  - A hamis TCP SYN folyamnak nem is kell túl gyorsnak lennie
- Megoldás
  - Túlméretezés / TCP SYN proxy
  - SYN Cookie

SYN Cookies

Normál esetben

- A --- SYN ---> B

  Állapot tárolás, kapcsolat feljegyzése, várakozás az ACK csomagra

- A <- SYN/ACK – B

Kapcsolat felépült

- A --- ACK ---> B

  Kapcsolat felépült

SYN Cookies használata

- A --- SYN ---> B

  Kapcsolat információ a Cookie-ban

- A <- SYN/ACK – B

  + Cookie

Kapcsolat felépült

- A --- ACK ---> B

  + Cookie

  Kapcsolat információ a Cookie-ban

TCP SYN cookies

- TCP SYN cookie készítése (D. J. Bernstein)
  - Cél, hogy a SYN ACK állapotra emlékezzen a szerver, de ne tároljon semmit maga
  - A cookie-t visszaküldi a kliensnek, aki a kapcsolat nyugtázásánál újraküldi
  - TCP kompatibilis
    - Nincs szükség új TCP üzenetre, nincs szükség módosításra a kliensben
      - A kliensnek vissza kell küldenie az üzenetet akkor is, ha nem ismeri a cookie-t
- Cookies tárolása a sorszám helyén (SEQ)
  - 5 bit: Timestamp - t mod 32, t az idő számláló, 64 másodpercenként nő
  - 3 bit: MSS index - A leggyakoribb 8 Maximum Segment Size
  - 24 bit: Hitelesítés - MD5 és egy időfüggő kulcs
    - Bemenetek: SRC IP addr, port, DST IP addr, port, t + kulcs
- SYN cookie működés
  - Van hátrány is
    - Kliens által kezdeményezett TCP opciókat nem tud tárolni, így néhány TCP funkció nem működik (pl. large windows)
    - Nincs SYN-ACK újraküldés
  - Csak szükség esetén kell használni!

Kliens rejtvény (puzzle)

- Megoldás DoS problémákra
  - A kliensnek dolgoznia kell a kapcsolatért
    - Matematika rejtvény megoldása
  - Garantálja, hogy az adott kliens nem tud túl sűrűn fordulni a szerverhez
  - A rejtvénynek skálázhatónak kell lennie, az eredményt gyorsan lehessen ellenőrizni

- Rejtvény példa
  - Hash algoritmus használata
    - A kliensnek ki kell találnia, hogy a szerver mely számnak a hash értékét küldte el számára
    - Skálázás a hash érték hosszával

- Már megtalálható mai protokollokban: pl. HIP