ETH zürich

Networked Systems
ETH Zürich — seit 2015

Communication Networks

Prof. Laurent Vanbever

Solution: Exercise 4 – Ethernet, Switching & Internet Protocol (IP)

Ethernet & Switching

4.1 Spanning-Tree (Exam Style Question)

Consider this network composed of 12 Layer 2 (Ethernet) switches.

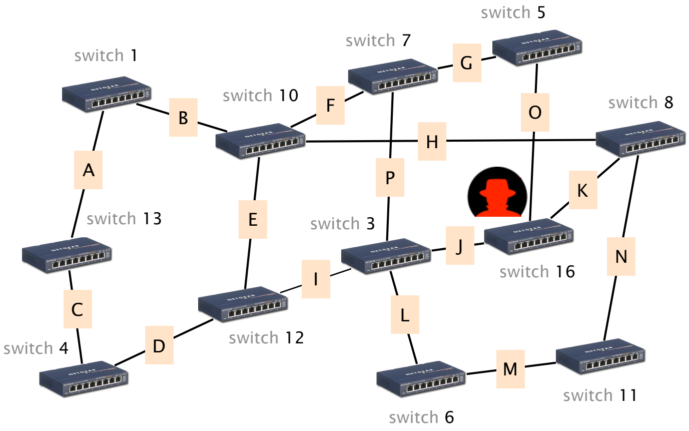

Compute a valid spanning tree, with and without hacker

a) Use the Spanning-Tree Protocol (STP) described in the lecture to compute a spanning tree. The numbers next to each switch indicate the switches identifier (switch 1 has ID "1"). Each link is labeled with a letter. Indicate the set of links (the letters, in alphabetical order) that are not part of the STP after the protocol has converged.

Solution: [D,I,J,M,O] since tie-breaking is done based on the switch ID.

b) As described in the course, STP is not the most secure protocol. Assume now that a hacker managed to take over switch 16 and starts pretending that the switch ID is “1”. Concretely, there are now two switches with ID “1” in the network. Indicate the set of links that will now be part of the attacker’s spanning tree, once the protocol has converged. Is the network still connected?

Solution: [I,J,K,L,N,O,P]. And, *no*, the network is not connected anymore.

4.2 Moving Target (Exam Style Question)

Consider the switched network depicted in the figure below. It is composed of 5 Ethernet switches, two hosts (connected to switch 3 and 4, respectively) and one IP router acting as default gateway for the hosts. For redundancy reasons, the network exhibits cycles and each switch therefore runs the Spanning Tree Protocol (STP). All links have a unary cost. When equal-cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better).

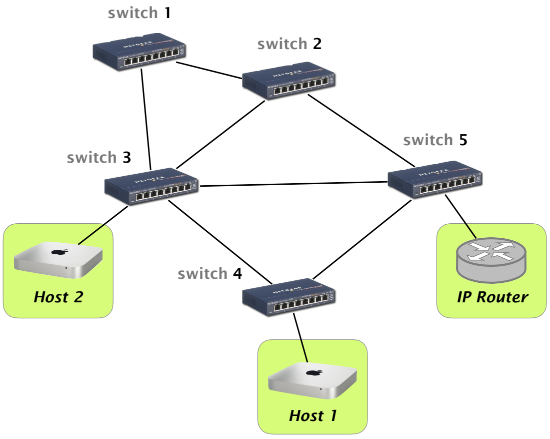

An Ethernet network running the spanning tree protocol.

a) In the figure below, indicate all the links that end up being **deactivated** in the final state, once all the switches have converged towards the final spanning tree.

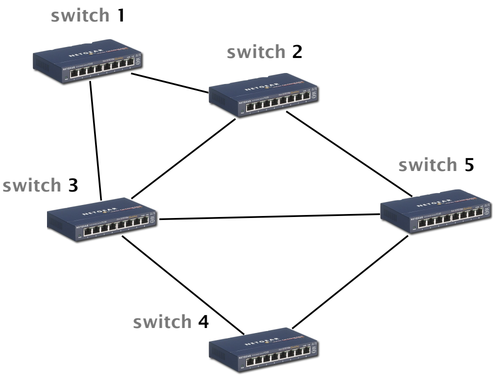

Solution: Links $(4,5)$, $(3,5)$ and $(2,3)$ end up disabled.

b) Perhaps unsurprisingly, a *lot* of traffic is exchanged between Host 1 (resp. Host 2) and Internet destinations. Briefly explain **two distinct reasons** why this configuration is not optimal in terms of network utilization/throughput.

Solution: Any communication between Host 1 (resp. Host 2) and IP router goes over 4 (resp. 3) links. Plus, these links are shared meaning Host 1 and Host 2 will be competing for throughput.

c) Realizing that there is a problem with their configuration, the network operators ask you (a fresh network engineer!) to help them improve their network performance. Briefly explain how you would adapt the configuration of the spanning tree protocol (i.e., the switches identifier and/or the link costs) so as to maximize the throughput between Host 1 (resp. Host 2) and Internet destinations.

Solution: Flipping the switch IDs so that the now-switch 5 becomes the root (e.g. making it switch 1 and the now-switch 1, switch 5).

d) The network operators are happy with your changes. But they now realize that Host 1 and Host 2, in addition to exchanging a lot of Internet traffic, also exchange a lot of traffic between themselves. The network operators ask for your help again! They ask you to find a spanning tree configuration such that: *(i)* the number of hops between any of these three hosts (Host 1 and 2, and the router) is equivalent; and at the same time *(ii)* the number of hops is minimal considering the given topology.

Briefly explain how you would configure the spanning tree protocol to achieve these requirements, or why these requirements are impossible to achieve.

Solution: Requirements are impossible to get: Either the hosts are using their direct link with each other, or with the router. But they cannot all use the direct link between themselves as otherwise that would cause a loop which would be prevented by the spanning tree protocol anyway.

4.3 VLAN

The network below consists of 9 switches and hosts in two different VLANs (blue and red).

Solution:

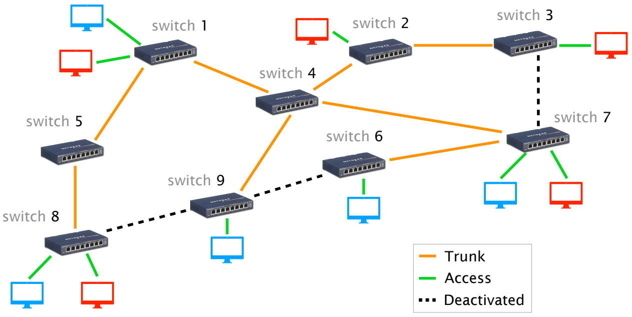

L2-network with hosts in two different VLANs (blue and red).

a) Compute a spanning tree in the network using switch 1 as root. When equal-cost paths to the root are encountered, switches break the tie based on the sender ID (lower is better). Clearly indicate the type of each link (trunk, access or deactivated).

Solution: See figure above.

b) Using the previously computed spanning tree, which path will the red host connected to switch 7 use to communicate with the red host connected to switch 1?

Solution: 7-4-1

c) Using the previously computed spanning tree, which path will the red host connected to switch 7 use to communicate with the blue host connected to switch 8?

Solution: Not possible. A host in the blue VLAN cannot directly communicate with a host in the red VLAN. Traffic would have to go over a layer 3 router to reach the other destination.

d) Compute now two per-VLAN spanning-trees (one for each VLAN) such that each link is active in at least one spanning-tree. Hint: you can adapt the switch IDs for each VLAN.

Solution: Multiple solutions are possible, for example:

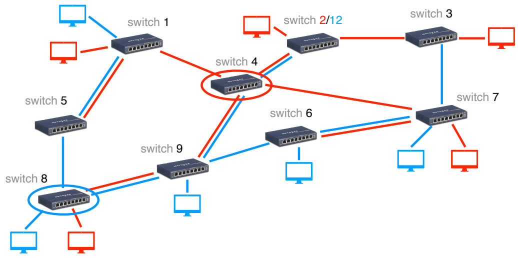

Internet Protocol (IP)

4.4 IPv4 Calculations

Each row in the following table describes an IP network. Fill in the missing values.

Solution:

Slash-notation	Netmask-notation	First usable address	Last usable address	Broadcast address
10.0.0.0/24	10.0.0.0/255.255.255.0	10.0.0.1	10.0.0.254	10.0.0.255
126.127.128.0/17	126.127.128.0/255.255.128.0	126.127.128.1	126.127.255.254	126.127.255.255
12.34.32.0/19	12.34.32.0/255.255.224.0	12.34.32.1	12.34.63.254	12.34.63.255
222.208.0.0/12	222.208.0.0/255.240.0.0	222.208.0.1	222.223.255.254	222.223.255.255
123.45.67.224/27	123.45.67.224/255.255.255.224	123.45.67.225	123.45.67.254	123.45.67.255

4.5 IPv4 vs. IPv6

a) In the lecture you heard about IPv4 and IPv6. Why was IPv6 introduced? What is the main difference?

Solution: The main motivation for IPv6 is the IPv4 address exhaustion. Even though Network Address Translation (NAT) could temporarily solve the problem, there are no longer enough IPv4 addresses / subnets for all the devices connected to the Internet. The main difference is the higher number of bits for each IP address (128 instead of 32). Furthermore, IPv6 also handles e.g. fragmentation or header options in a different way.

b) How many IPv4 and IPv6 addresses exist? Is it possible to use all the addresses for hosts in the Internet?

Solution: IPv4: $2^{32} \approx 4.3 * 10^9$

IPv6: $2^{128} \approx 3.4 * 10^{38}$

No, it is not possible to use all the addresses. Some address spaces are reserved e.g. for private addresses. Other addresses are used to identify the network/router or as broadcast addresses.

c) Assuming there are 7.8 billion people in the world, how many IPv4/IPv6 addresses are theoretically available per person? According to Wikipedia[a] an average human body contains around 7×10^{27} atoms. How many IPv6 addresses do we have per atom belonging to all 7.8 billion people?

Solution: IPv4 per human: $2^{32}/(7.8 * 10^9) \approx 0.55$

IPv6 per human: $2^{128}/(7.8 * 10^9) \approx 4.36 * 10^{28}$

IPv6 per atom in all human bodies:
$2^{128}/(7.8 * 10^9 * 7 * 10^{27}) \approx 6$

[a] https://en.wikipedia.org/wiki/Composition_of_the_human_body#Elemental_composition_list

d) Even though a first IPv6 version has been standardized more than 20 years ago, it still has very limited coverage. What are the reasons why the deployment of IPv6 is so slow?

Solution: Every network device, which has to interact with the network layer, needs to be able to understand the new IPv6 addresses and the corresponding header. It is therefore not possible to switch from IPv4 to IPv6 on one specific day. Upgrading the hardware is costly and especially for end-users there is no real motivation. At the moment, everything seems to work well with IPv4 addresses.

4.6 IPv6 Computations

Answer the following questions to IPv6.

a) Currently, all global unicast IPv6 addresses are inside 2000::/3. Assume that every network in the Internet gets an entire /64 prefix. How many different /64 prefixes can you distribute? How many hosts can be inside one of these /64 prefixes? Compare these numbers with the total amount of IPv4 addresses.

Solution: If every network gets a /64 prefix within 2000::/3. Then, there are $64 - 3 = 61$ bits to use. This allows for $2^{61} = 2.30 * 10^{18}$ /64 prefixes. In each prefix, you can allocate a total of $2^{64} = 1.84 * 10^{19}$ hosts. There are only $2^{32} = 4.29 * 10^{9}$ IPv4 addresses.

b) Simplify the notation of the following IPv6 addresses:

Solution:

Full IPv6 address	**Simplified IPv6 address**
000a:1234:abda:0000:0000:0140:0000:0001	a:1234:abda::140:0:1
0000:0000:0000:0000:0000:0003:0000:0010	::3:0:10
000a:0031:003f:0000:0000:0000:0000:0000	a:31:3f::
0000:0000:0000:0000:0000:0000:0000:000b	::b

c) For the following pairs of IPv6 addresses, find the longest prefix that contains both addresses.

Solution:

Address 1	**Address 2**	**Prefix**
2000::a35a	2000::ac3f	2000::a000/116
2005::2e90:12fa:1	2005::2eb0:0:1	2005::2e80:0:0/90
200a::789:3	200a:5c::	200a::/25

4.7 Putting everything together

Consider the network on the left composed of three Ethernet segments separated by two intermediate routers (A and B). In this network, the server's interface along with the routers' interfaces are configured with static IP addresses. While clients connected to the 3.0.0.0/24 Ethernet segment obtain an IP address via DHCP.

Assuming that the client has just started, with a perfectly empty state, precisely describe all packets that are generated when the command "`ping 1.0.0.2`" is issued (until the server answers back). Among others, your answer *must* include the content of the Layer 2 and Layer 3 headers.

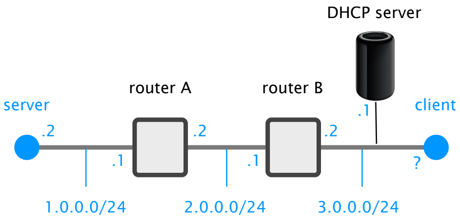

Describe everything that happens to the packets sent between the client and the server

Solution:

- src_mac: *client_mac*, dst_mac: *broadcast*
  src_ip: *0.0.0.0*, dst_ip: *255.255.255.255*
  payload: *DHCP discovery*
- src_mac: *dhcp_server_mac*, dst_mac: *broadcast*
  src_ip: *3.0.0.1*, dst_ip: *255.255.255.255*
  payload: *DHCP offer for 3.0.0.3*

  Note: In addition to the IP address of the client (arbitrarily picked by the DHCP server), the DHCP offer also contains the gateway (3.0.0.2) to use. The DHCP offer can be sent either in broadcast or unicast (both answers accepted). Broadcast is often used as some network stacks will not accept "unicasted" frames unless an IP address is configured first.
- src_mac: *client_mac*, dst_mac: *broadcast*
  payload: *ARP request: Who has 3.0.0.2? Tell 3.0.0.3.*
- src_mac: *routerB_right_mac*, dst_mac: *client_mac*
  payload: *ARP reply: 3.0.0.2 is at routerB_right_mac*
- src_mac: *client_mac*, dst_mac: *routerB_right_mac*
  src_ip: *3.0.0.3*, dst_ip: *1.0.0.2*
  payload: *ICMP echo request*

... (rest of the packets are omitted)

- src_mac: *server_mac*, dst_mac: *routerA_left_mac*
  src_ip: *1.0.0.2*, dst_ip: *3.0.0.3*
  payload: *ICMP echo reply*