

POLITECNICO DI TORINO

Exercises on Data Link Traffic Forwarding

Fulvio Rizzo



May 31, 2016

License

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License.

You are free:

- **to Share:** to copy, distribute and transmit the work
- **to Remix:** to adapt the work

Under the following conditions:

- **Attribution:** you must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).
- **Noncommercial:** you may not use this work for commercial purposes.
- **Share Alike:** if you alter, transform, or build upon this work, you may distribute the resulting work only under the same or similar license to this one.

More information on the Creative Commons website (<http://creativecommons.org>).



Acknowledgments

The author would like to thank all the people that contributed to those exercises.

Contents

I. Intro	6
1. Symbols	7
2. Methodology	8
2.1. Throughput	8
II. Exercises	10
3. Filtering database	11
3.1. Exercise n. 1	11
3.2. Exercise n. 2	12
4. Traffic Analysis	13
4.1. Exercise n. 3	13
4.2. Exercise n. 4	14
4.3. Exercise n. 5	15
4.4. Exercise n. 6	16
4.5. Exercise n. 7	17
4.6. Exercise n. 8	18
4.7. Exercise n. 9	19
4.8. Exercise n. 10	20
4.9. Exercise n. 11	21
4.10. Exercise n. 12	22
4.11. Exercise n. 13	23
4.12. Exercise n. 14	24
4.13. Exercise n. 15	25
5. Performance	26
5.1. Exercise n. 16	26
5.2. Exercise n. 17	27
III. Solutions	28
6. Filtering database	29
6.1. Solution for exercise n. 1	29
6.2. Solution for exercise n. 2	30

7. Traffic Analysis	31
7.1. Solution for exercise n. 3	31
7.1.1. Frames generated on the network	31
7.1.2. ARP cache of the all hosts	31
7.1.3. Ports of the switch involved in receiving/transmitting frames	32
7.1.4. Filtering database of the switch	32
7.2. Solution for exercise n. 4	33
7.2.1. Frames forwarded by the switches	33
7.2.2. ARP cache of the all hosts	33
7.2.3. Ports of the switch involved in receiving/transmitting frames	33
7.2.4. Filtering database of the switch	34
7.3. Solution for exercise n. 5	35
7.3.1. Possibility to continue the PING	35
7.3.2. Filtering database of the switch SW-1	35
7.4. Solution for exercise n. 6	36
7.4.1. Possibility to continue the PING	36
7.4.2. Filtering database of the switch SW-1	36
7.5. Solution for exercise n. 7	37
7.6. Solution for exercise n. 8	38
7.6.1. Possibility to continue the PING	38
7.6.2. Filtering database of the switches	38
7.7. Solution for exercise n. 9	39
7.8. Solution for exercise n. 10	40
7.8.1. Frames generated on the network	40
7.8.2. Filtering Database	40
7.9. Solution for exercise n. 11	42
7.9.1. Frames forwarded by the switch	42
7.9.2. Ports of the switch involved in receiving/transmitting frames	42
7.9.3. Filtering database of the switch	42
7.10. Solution for exercise n. 12	44
7.10.1. Frames forwarded by the switch	44
7.10.2. Ports of the switch involved in receiving/transmitting frames	44
7.10.3. Filtering database of the switch	44
7.11. Solution for exercise n. 13	45
7.11.1. Frames forwarded by the switch	45
7.11.2. Ports of the switch involved in receiving/transmitting frames	45
7.11.3. Filtering database of the switch	45
7.12. Solution for exercise n. 14	46
7.12.1. Frames forwarded by the switch	46
7.12.2. Ports of the switch involved in receiving/transmitting frames	46
7.12.3. Filtering database of the switch	47
7.13. Solution for exercise n. 15	48
7.13.1. Hosts connected to the same switch (scenario 1)	48
7.13.2. Hosts connected to different switches (scenario 2)	48
8. Performance	49
8.1. Solution for exercise n. 16	49

8.2. Solution for exercise n. 17	50
8.2.1. Aggregate Bandwidth	50
8.2.2. Throughput of the switch	50

Part I.

Intro

1. Symbols




Repeater - Hub



Bridge - Switch



Router



Host

2. Methodology

The solution to these exercises can be easily obtained through the following steps:

1. Determine the frames transmitted on the network, considering all the links as shared medium (e.g., shared Ethernet)
2. Starting from the first frame generated on the network:
 - a) use the MAC source address to populate/update the filtering database of the first switch encountered by the frame on its journey toward the destination
 - b) use the MAC destination address as a lookup key for the filtering database; if the MAC address is present, forward the frame only on the interface contained in that entry; otherwise, forward the frame on all the ports of the switch, except the port on which the frame has been received
 - c) move to the next switch (or switches, e.g., in case the frame has been flooded by the previous switch) that received the current frame and repeat the process from step (a), until the current frame disappears from the network
 - d) when the current frame disappeared from the network and it is no longer transmitted by any intermediate device (hub, bridge), move to the next frame in the trace and repeat the process from step (a)

Please remember also that:

1. entries in the filtering database remain till the *max_age* parameter expires, unless refreshed
2. those entries are refreshed only if another frame coming from the MAC source address under consideration is received
3. due to the topology of the network, a unicast frame may not be able to refresh the filtering databases in all the switches present in the network, while a broadcast frame does
4. broadcast frames are always issued when an ARP cache is found empty; however, beware of the different expiration times we may have in the ARP cache and in the filtering database, which may lead an entry to expire on one cache and to be still valid in the other table.

2.1. Throughput

Since some exercise focuses on the throughput of network devices, we give here a couple of definitions that may be useful:

1. **Aggregate Bandwidth:** the aggregated bandwidth of a network device is the total amount of traffic that can be forwarded in the unit of time. For instance, the total amount of traffic of a switch with two FastEthernet ports operating in half-duplex mode is 100 Mbps (i.e., the traffic from one port gets forwarded to the other port). If ports are operating in full-duplex mode, the aggregated bandwidth will become 200 Mbps (100 Mbps can flow from port 1 to port 2, while

other 100 Mbps can flow from port 2 to port 1). In other words, the aggregated bandwidth is the maximum load that can be handled by the network devices and hence it does not depend on the actual traffic sent/received on its network interface.


2. **Throughput:** the throughput of the network device is the total amount of traffic that is actually forwarded in the unity of time. Differently from the aggregated bandwidth, the throughput depends on the workload offered on the network ports of the device; a network switch with two FastEthernet ports operating in full-duplex mode, with traffic flowing from port 1 to port 2 at 50 Mbps will have a throughput of 50 Mbps.

Part II.
Exercises

3. Filtering database

3.1. Exercise n. 1

Assuming the network topology depicted below, a 48-ports switch (SW-1) is connected to 47 hosts, while the last port is used to connect to the other switches of the LAN. Globally, the total number of hosts connected to the LAN is 100. Determine the number of entries we can expect in the filtering database of the switch SW-1, considering that all the hosts are reasonably active all the time and that exchange data with all the other hosts in the network.



3.2. Exercise n. 2


In a LAN, a 48-ports switch (SW-1) is connected to 47 hosts, while the remaining port (uplink) connects to another switch. Globally, the network includes 98 hosts, one router (that all the stations use to connect to the Internet) and a server that exchanges only traffic from and to the Internet and is always busy. The global topology is depicted in the picture below.

Assuming that:

- The ARP cache on all the stations (hosts, server and router) expires after 20 mins
- The aging time of the filtering database of the switches is set to the default value

Calculate the number of entries we can expect in average in the filtering database of the switch SW-1, considering:

1. a first scenario in which all the hosts are reasonably active all the time and exchange data with all the other hosts (i.e., Hosts talk to each other, while do not exchange traffic with S and R) in the network;
2. a second scenario in which the hosts exchange data mostly with the Internet and/or the server S and do not exchange any data between themselves.

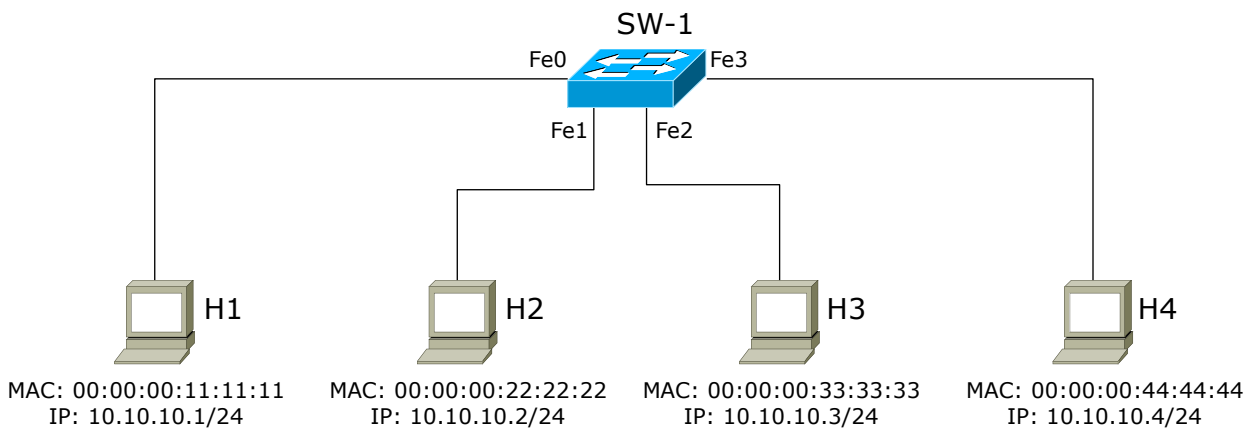


4. Traffic Analysis

4.1. Exercise n. 3

Referring to the network topology depicted below, answer to the following questions:

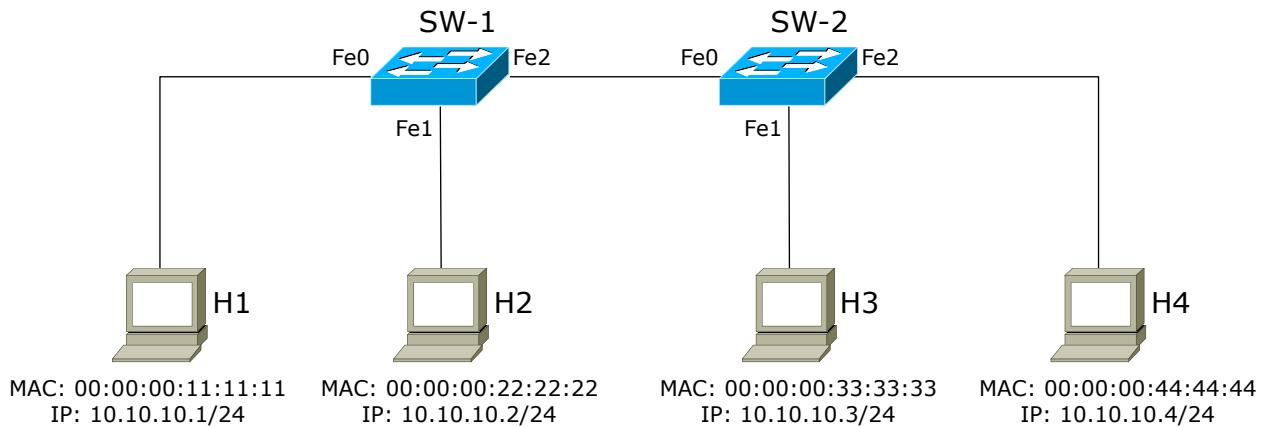
- List all the frames forwarded by the switch on the network when the user on host H1 types “ping 10.10.10.2”, assuming that the ARP caches of all the hosts and the filtering database of the switch are empty.
- List the ARP cache of all the hosts when the ping program terminates.
- Describe, for each frame, which port of the switch will be involved in receiving and/or sending the frame out.
- List the filtering database of the switch when the ping program terminates, ignoring the values associated to the ageing time.



4.2. Exercise n. 4

Referring to the network topology depicted below, answer to the following questions:

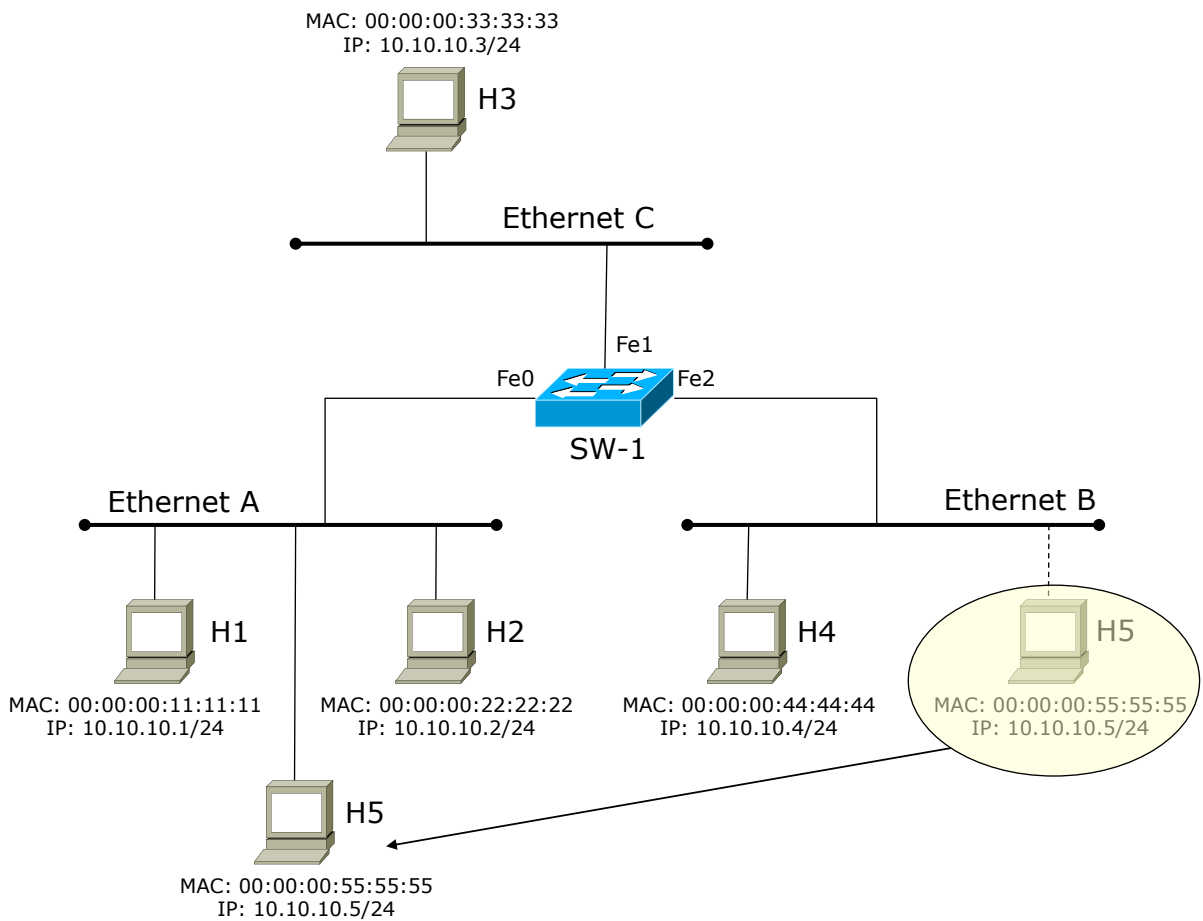
- List all the frames forwarded by the switches when host H1 types “ping 10.10.10.2”, assuming that the ARP cache of all the hosts and the filtering database of the switch are empty.
- List the ARP cache of all the hosts when the ping program terminates.
- Describe, for each frame, which port of the switch will be involved in receiving and/or sending the frame out.
- List the filtering database of the switches when the ping program terminates, ignoring the values associated to the ageing time.



4.3. Exercise n. 5

Referring to the network topology depicted below, let us suppose that while a continuous stream of ICMP packets (generated by executing the command “ping -t 10.10.10.5” on host H1) is in progress, host H5 is moved from Ethernet B to Ethernet A. Assume that the entries in the ARP caches have infinite lifetime.

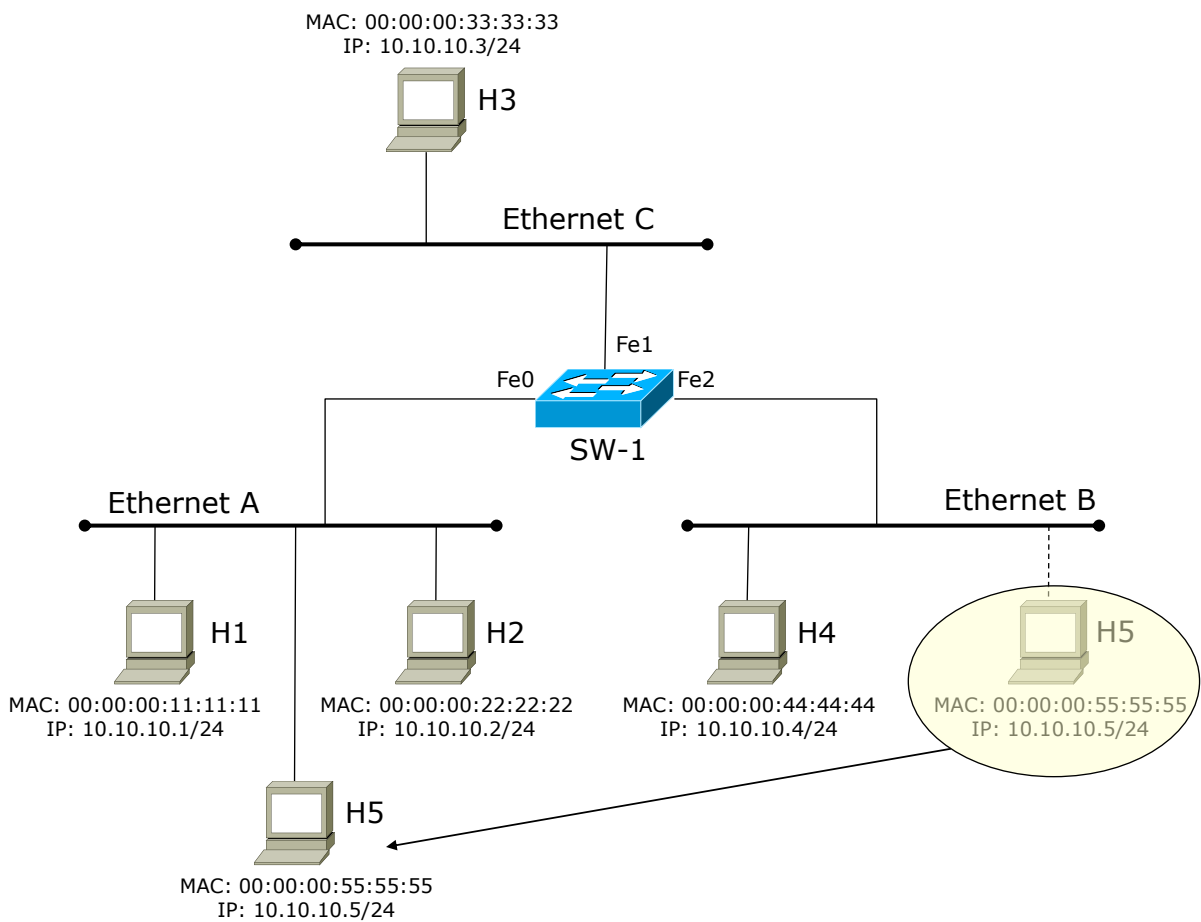
- Determine if host H5 can still receive the ICMP Echo Request from host H1.
- List the entries that are present some seconds after host H5 is moved in the filtering database of the switch SW-1.



4.4. Exercise n. 6

Referring to the network topology depicted below, let us suppose that while a continuous stream of ICMP packets (generated by executing the command “ping -t 10.10.10.5” on host H3) is in progress, host H5 is moved from Ethernet B to Ethernet A. Assume that the entries in the ARP caches have infinite lifetime.

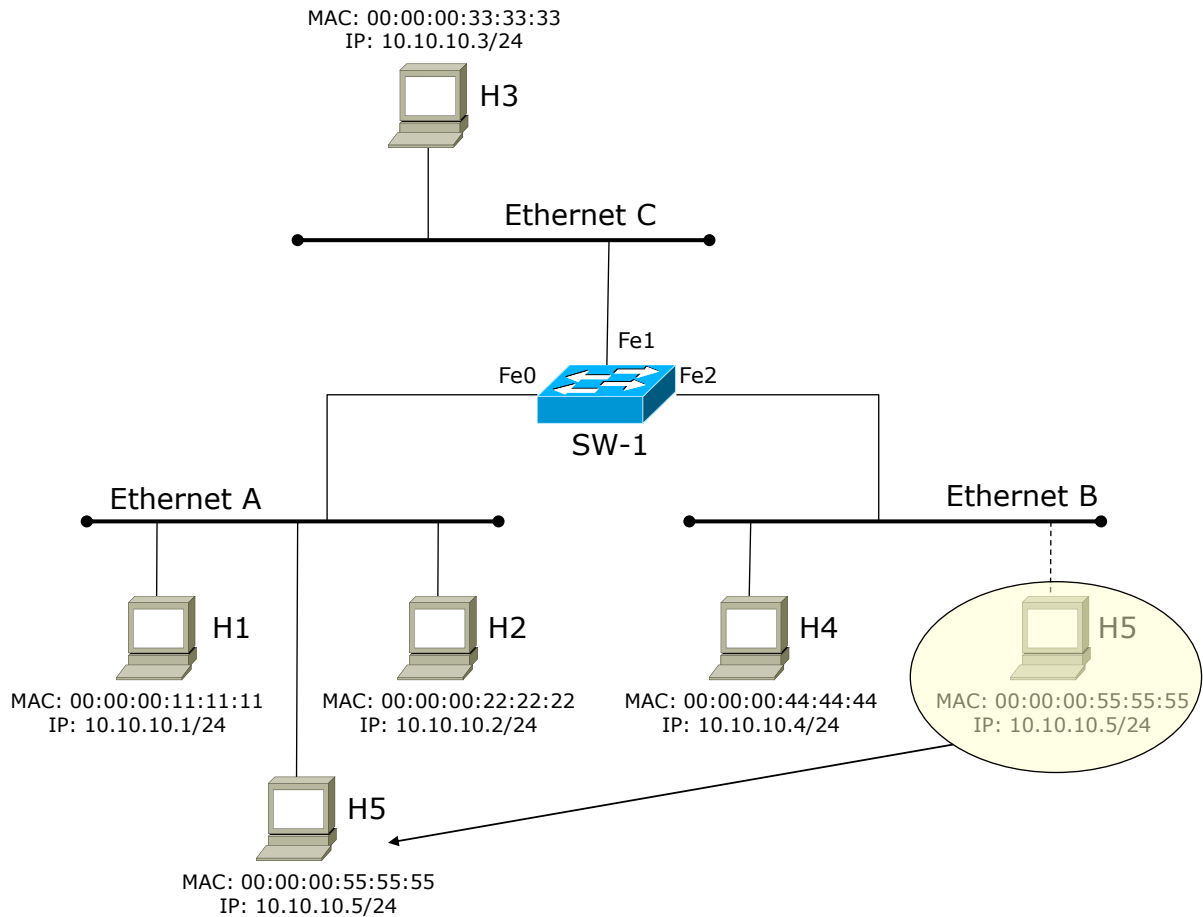
- Determine if host H5 can still receive the ICMP Echo Request from host H3.
- List the entries in the filtering database a couple of minutes after host H5 is moved and after about 10 minutes, including a possible value for the Ageing Time.



4.5. Exercise n. 7

Referring to the network topology depicted below, let us suppose that while a continuous stream of ICMP packets (generated by executing the command “ping -t 10.10.10.5” on host H3) is in progress, host H5 is moved from Ethernet B to Ethernet A. Assume that the entries in the ARP caches have a timeout of 120 seconds.

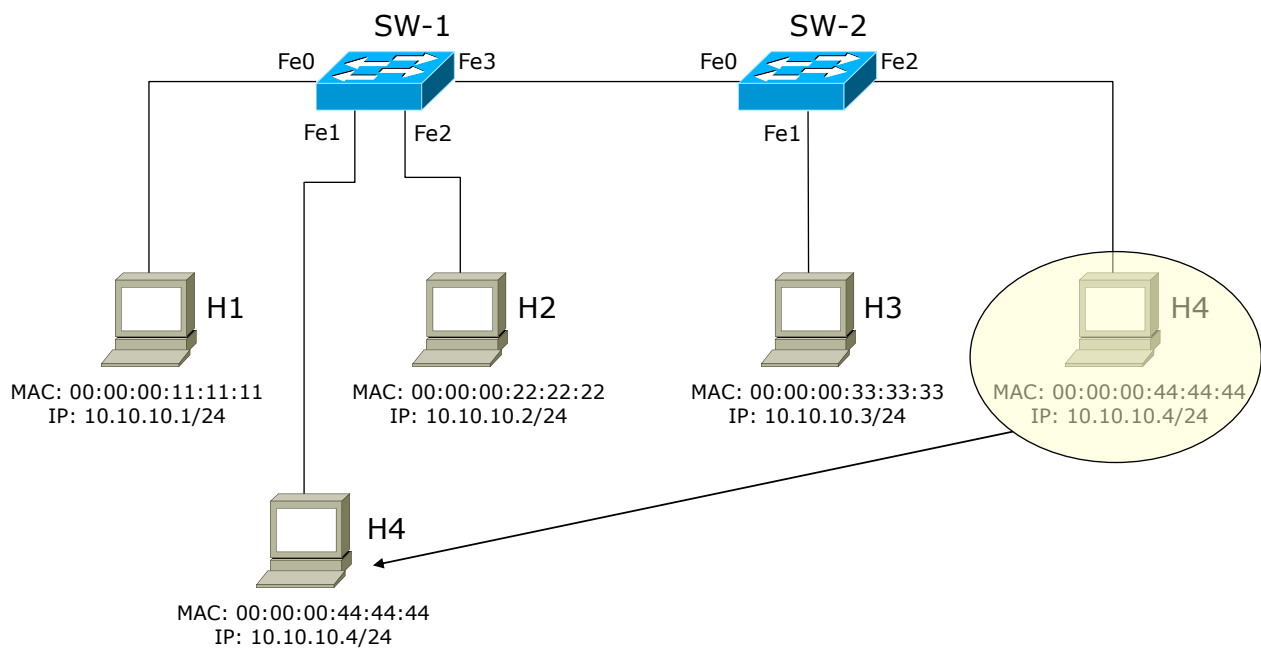
Determine if host H5 can still receive the ICMP Echo Request from host H3.



4.6. Exercise n. 8

Referring to the network topology depicted below, let us suppose that while a continuous stream of ICMP packets (generated by executing the command “ping -t 10.10.10.1” on host H4) is in progress, that host is moved from switch SW-2 to switch SW-1. Assume that the entries in the ARP caches have infinite lifetime.


- Determine if host H4 can still receive the ICMP Echo Reply from host H1.
- List the entries in the filtering databases 2 minutes after host H4 is moved, including a possible value for the Ageing Time.



4.7. Exercise n. 9

Referring to the network topology depicted below, let us suppose that while a continuous stream of ICMP packets (generated by executing the command “ping -t 10.10.10.4” on host H1) is in progress, the cable that connects host H4 to the network breaks just after the ARP Reply.

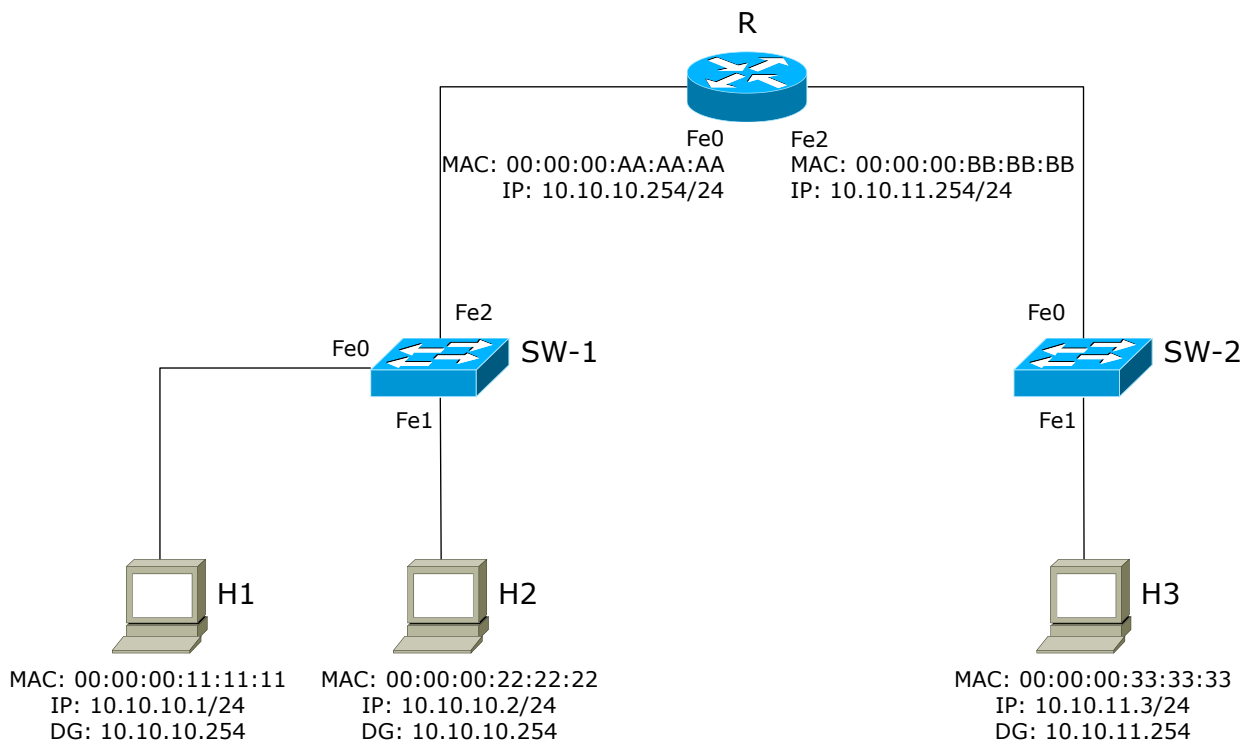
Describe what happens to the ICMP Echo packets sent by host H1 to host H4, assuming that the entries in the ARP caches have infinite lifetime.



4.8. Exercise n. 10

Referring to the network topology depicted below and assuming that the router and the hosts are correctly configured, answer to the following questions:

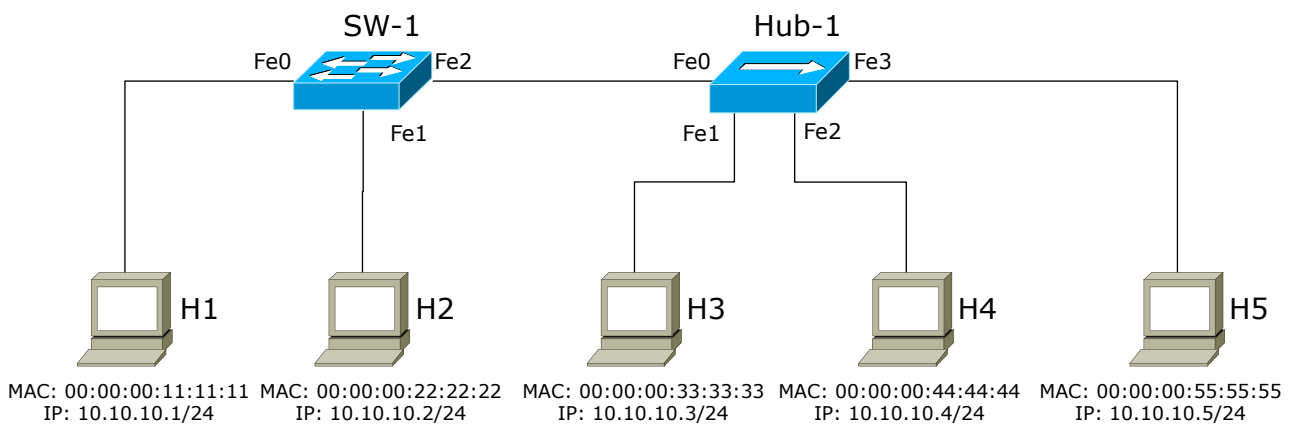
- List all the frames generated on the network when host H1 pings host H3, assuming that all the ARP caches are empty.
- List the filtering database of the switches when the ping program ends, ignoring the values associated to the ageing time.



4.9. Exercise n. 11

Referring to the network topology depicted below, answer to the following questions:


- List all the frames forwarded by the switch when the user on host H1 types “ping 10.10.10.4”, assuming that the ARP caches of all the hosts and the filtering database of the switch are empty.
- Describe, for each frame, which port of the switch and of the hub will be involved in receiving and/or sending the frame out.
- List the filtering database of the switch/hub when the ping program terminates, ignoring the values associated to the ageing time.



4.10. Exercise n. 12

Referring to the network topology depicted below, answer to the following questions:


- List all the frames forwarded by the switch when the user on host H1 types “ping 10.10.10.2”, assuming that the ARP caches of all the hosts and the filtering database of the switch are empty.
- Describe, for each frame, which port of the switch and of the hub will be involved in receiving and/or sending the frame out.
- List the filtering database of the switch when the ping program terminates, ignoring the values associated to the ageing time.



4.11. Exercise n. 13

Referring to the network topology depicted below, answer to the following questions:


- List all the frames forwarded by the switch when the user on host H3 types “ping 10.10.10.2”, assuming that the ARP caches of all the hosts and the filtering database of the switch are empty.
- Describe, for each frame, which port of the switch and of the hub will be involved in receiving and/or sending the frame out.
- List the filtering database of the switch when the ping program terminates, ignoring the values associated to the ageing time.



4.12. Exercise n. 14

Referring to the network topology depicted below, answer to the following questions:

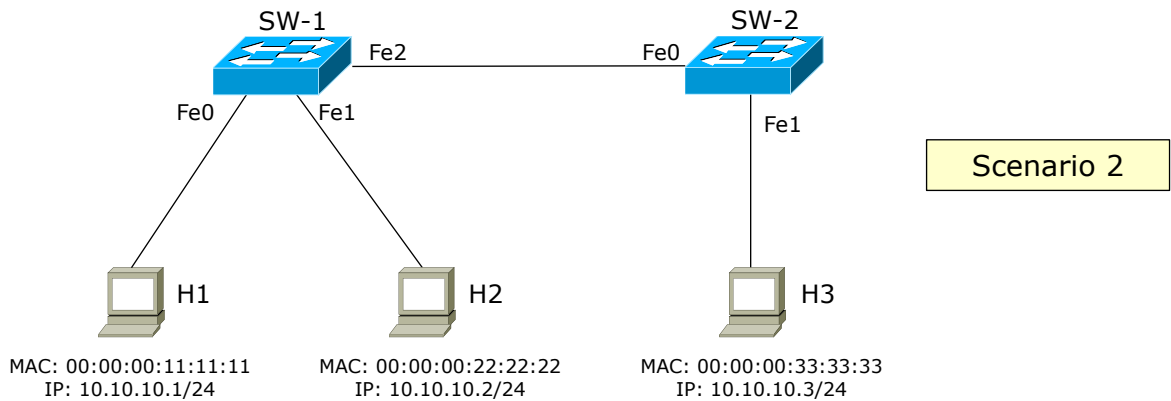

- List all the frames forwarded by the switch when the user on host H3 types “ping 10.10.10.5”, assuming that the ARP caches of all the hosts and the filtering database of the switch are empty.
- Describe, for each frame, which port of the switch and of the hub will be involved in receiving and/or sending the frame out.
- List the filtering database of the switch when the ping program terminates, ignoring the values associated to the ageing time.



4.13. Exercise n. 15

Assuming the network topology depicted below (scenario 1), let us suppose that host H3 wants to capture all the traffic exchanged between hosts H1 and H2. Is it possible?

If the network topology changes in such a way that hosts H2 and H3 are connected to different switches (scenario 2), does the behavior of the network change?




5. Performance

5.1. Exercise n. 16

In a network where clients exchange traffic mainly with the departmental server S, a 48-ports hub operating at 10Mbps is replaced by an equivalent switch, operating at the same speed.

Determine if the network clients experience a better service after the upgrade.




5.2. Exercise n. 17

A small network is made up of a 48-switch whose ports operate at 100Mbps. This network includes 47 clients and one server; clients generate a continuous stream of UDP traffic to the server, and the server re-sends each received packets back to the sender using the same protocol.

Determine:

- The aggregate bandwidth of the switch, both in case ports are operating in Half-Duplex and Full-Duplex mode
- The throughput of the switch in the given network scenario, both in case ports are operating in Half-Duplex and Full-Duplex mode



Part III.

Solutions

6. Filtering database

6.1. Solution for exercise n. 1

The filtering database keeps track of all the MAC addresses present in the LAN, independently from the position of the hosts, provided that the hosts are reasonably active (i.e., that they are able to refresh the entries in the database, which is one of the assumptions in the text).

Therefore we expect the filtering database of Switch SW-1 to contain 100 entries; 47 entries point to stations connected to its ports, while the remaining 53 entries will be associated to port **Fe47**, i.e., the uplink toward SW-2.

6.2. Solution for exercise n. 2

Scenario 1: hosts reasonably active

The filtering database keeps track of all the MAC address present in the LAN, independently from the position of the hosts. However, entries must be periodically refreshed in order to keep them valid.

In our case, Server S will send/receive a continuous traffic to/from the router, which means that the switch SW-3 will always have MAC(S) in its filtering database, but this may not be valid for switches SW-1 and SW-2.

Given that the validity of the ARP cache is 20 mins, at some point the ARP caches of S and R will still be valid, while switches SW-1 and SW-2 will purge the MAC addresses associated to S and R from their filtering databases. During this phase, switches SW-1 and SW-2 will have 98 entries in their filtering databases, while SW-3 will have 100 entries (this switch will receive also the traffic of S and R).

At some point, the ARP cache mapping will expire in either S or R. In this case, we may have two cases:

- The ARP cache of the server S expires first. In this case, S will broadcast an ARP request toward R in order to know MAC(R), which will update the filtering database of switches SW-1 and SW-2 bringing their total number of entries to 100 for a limited amount of time.
- The ARP cache of the router R expires first. In this case, R will broadcast an ARP request toward S in order to know MAC(S), which will trigger a direct (unicast) answer from S to R. This answer will not be propagated to the rest of the network, bringing the total number of entries in the filtering database of switches SW-1 and SW-2 to 99.

In general, the filtering database of the switch SW-1 will contain a number of entries that oscillates between 98 and 99, while being 98 for the most part of the time (the filtering database expires after 5 minutes, leaving at least 15 minutes of vacancy, even if the server S will issue the ARP Request first).

Scenario 2: hosts communicate mostly with S and/or R

The behavior of S and R does not change in this second scenario. However, the behavior of the hosts does, since in this case they do not exchange traffic with all the other hosts, but only with S and R which are both located on switch SW-3. Therefore, only hosts that are attached to switch SW-1 will need to traverse the entire network in order to reach S and R (and therefore their traffic will update the filtering database of all the switches), while other hosts will not be able to update all the filtering databases unless they send a broadcast packet.

Given that the broadcast ARP Request is relatively rare (due to the high expiration time of the ARP cache), we can conclude that the number of entries in the filtering databases will be:

- **Switch SW-1:** 49 entries (47 hosts, plus S and R) most of the time; more entries when some additional MAC addresses are seen;
- **Switch SW-2:** 96 entries (47 hosts from SW-1, 47 hosts from SW-2, plus S and R) most of the time; more entries when some additional MAC addresses are seen
- **Switch SW-3:** 100 entries most of the time, as in the previous scenario.

7. Traffic Analysis

7.1. Solution for exercise n. 3

7.1.1. Frames generated on the network

The frames generated on the network are the following:

N.	L2	L3	Appl-layer protocol	Description
1	00:00:00:11:11:11 → FF:FF:FF:FF:FF:FF	—	ARP Request	Who has IP=10.10.10.2 please reply with its MAC address
2	00:00:00:22:22:22 → 00:00:00:11:11:11	—	ARP Reply	Host 10.10.10.2 has MAC = 00:00:00:22:22:22
3	00:00:00:11:11:11 → 00:00:00:22:22:22	10.10.10.1 → 10.10.10.2	ICMP	ICMP Echo Request
4	00:00:00:22:22:22 → 00:00:00:11:11:11	10.10.10.2 → 10.10.10.1	ICMP	ICMP Echo Reply
5-10	Packets 3 and 4 are replicated 3 times			

7.1.2. ARP cache of the all hosts

Assuming that the operating system adds an entry in its ARP cache only if it takes an active part in the transaction (i.e., the host that requires the ARP resolution and the requested host), the ARP caches of the hosts will be the following:

Host H1

IP address	MAC address
10.10.10.2	00:00:00:22:22:22

Host H2

IP address	MAC address
10.10.10.1	00:00:00:11:11:11

Host H3, Host H4

IP address	MAC address
-	-

7.1.3. Ports of the switch involved in receiving/transmitting frames

The ports of the switch involved in receiving/transmitting frames are the following:

Frame#	Description	Fe0	Fe1	Fe2	Fe3
1	ARP Request	IN	OUT	OUT	OUT
2	ARP Reply	OUT	IN	-	-
3	ICMP Echo Request	IN	OUT	-	-
4	ICMP Echo Reply	OUT	IN	-	-

7.1.4. Filtering database of the switch

The filtering database of the switch is the following:

MAC address	Interface
00:00:00:11:11:11	Fe0
00:00:00:22:22:22	Fe1

7.2. Solution for exercise n. 4

7.2.1. Frames forwarded by the switches

The frames forwarded by the switch SW-1 correspond to the frames generated on the network:

N.	L2	L3	Appl-layer protocol	Description
1	00:00:00:11:11:11 → FF:FF:FF:FF:FF:FF	—	ARP Request	Who has IP=10.10.10.2 please reply with its MAC address
2	00:00:00:22:22:22 → 00:00:00:11:11:11	—	ARP Reply	Host 10.10.10.2 has MAC = 00:00:00:22:22:22
3	00:00:00:11:11:11 → 00:00:00:22:22:22	10.10.10.1 → 10.10.10.2	ICMP	ICMP Echo Request
4	00:00:00:22:22:22 → 00:00:00:11:11:11	10.10.10.2 → 10.10.10.1	ICMP	ICMP Echo Reply
5-10	Packets 3 and 4 are replicated 3 times			

Vice versa, the frames forwarded by the switch SW-2 are only those that are sent in flooding on the network:

N.	L2	L3	Appl-layer protocol	Description
1	00:00:00:11:11:11 → FF:FF:FF:FF:FF:FF	—	ARP Request	Who has IP=10.10.10.2 please reply with its MAC address

In fact, the ARP Reply coming from host H2 will already find a proper entry in the filtering database of switch SW-1, and therefore is forwarded directly to the destination (and not in flooding).

7.2.2. ARP cache of the all hosts

The ARP caches of the hosts are the same of the previous exercise, since the behavior of end-systems at the IP layer does not change when the topology changes at layer 1 (physical) and/or 2 (data-link).

7.2.3. Ports of the switch involved in receiving/transmitting frames

The ports of the switch SW-1 involved in receiving/transmitting frames are the following:

PKT#	Description	Fe0	Fe1	Fe2
1	ARP Request	IN	OUT	OUT
2	ARP Reply	OUT	IN	-
3	ICMP Echo Request	IN	OUT	-
4	ICMP Echo Reply	OUT	IN	-

The ports of the switch SW-2 involved in receiving/transmitting frames are the following:

PKT#	Description	Fe0	Fe1	Fe2
1	ARP Request	IN	OUT	OUT
2	ARP Reply	-	-	-
3	ICMP Echo Request	-	-	-
4	ICMP Echo Reply	-	-	-

7.2.4. Filtering database of the switch

The filtering database of the switch SW-1 is the following:

MAC address	Interface
00:00:00:11:11:11	Fe0
00:00:00:22:22:22	Fe1

The filtering database of the switch SW-2 is the following:

MAC address	Interface
00:00:00:11:11:11	Fe0

7.3. Solution for exercise n. 5

7.3.1. Possibility to continue the PING

Host H5 can still receive the ICMP Echo Request sent by host H1 because, after its relocation, they are both connected to the same shared network segment. In fact, each ICMP Echo Request sent by host H1 is received by all the network interfaces connected to the Ethernet A (including host H5), and the same happens for each ICMP Echo Reply sent by host H5.

It is worthy noticing how the filtering database of the switch SW-1, immediately after the move, has a wrong entry with respect to host H5, which still appears on the port Fe2. This entry will be updated as soon as host H5 generates a frame on the network.

7.3.2. Filtering database of the switch SW-1

The content of the filtering database at the end of the process will be the following:

MAC address	Interface
00:00:00:11:11:11	Fe0
00:00:00:55:55:55	Fe0

7.4. Solution for exercise n. 6

7.4.1. Possibility to continue the PING

Host H5 can no longer receive the ICMP Echo Request sent by host H3 because, after its relocation, the corresponding entry in the filtering database of the switch SW-1 is wrong (in fact, it associates the MAC 00:00:00:55:55:55 to port Fe2).

Therefore, any frame directed to host H5 will be forwarded on the port Fe2 by the switch and will never reach host H5, which then will not reply to the ICMP Echo Request.

This misbehavior will be recovered when the aging time of the above entry in the filtering database will expire, i.e., after 300 seconds. At that time, the switch SW-1 will no longer know the exact position of the host H5 and will forward the ICMP Echo Request frame on all its ports, and hence the frame will reach host H5 as well.

Therefore, after a blackout of approximately 300 seconds in which host H5 appears unreachable, ICMP Echo Reply frames will be delivered again to host H3, which will then see the answers to its ping requests.

7.4.2. Filtering database of the switch SW-1

The content of the filtering database 2 min after the relocation can be the following:

MAC address	Interface	Ageing time
00:00:00:33:33:33	Fe1	0
00:00:00:55:55:55	Fe2	120

Vice versa, the filtering database after 10 minutes can be the following:

MAC address	Interface	Ageing time
00:00:00:33:33:33	Fe1	0
00:00:00:55:55:55	Fe0	0

7.5. Solution for exercise n. 7

As in the previous exercise, host H5 can no longer receive the ICMP Echo Request sent by host H3 because, after its relocation, the corresponding entry in the filtering database of the switch SW-1 is wrong (in fact, it associates the MAC 00:00:00:55:55:55 to port Fe2).

However, host H3 will continue the transmission of its ICMP Echo Request packets until the ARP cache expires, i.e., for 2 minutes (120 seconds). When this timeout occurs, host H3 will refresh the ARP entry by sending a new ARP request toward H5 in order to discover its MAC address. This broadcast message will be flooded by the switch SW-1 and therefore will be received by host H5 as well, despite the new location. The ARP Reply coming from host H5 will refresh the ARP cache of host H3, but it will also update the filtering database of SW-1 with the new position of host H5.

Therefore, host H3 will not receive answers to its ping requests for 2 minutes, then the problem is recovered and host H5 will begin to answer again.

Please note that in this case the ageing time of the filtering database will not have any effect, since it lasts longer than the ARP cache.

7.6. Solution for exercise n. 8

7.6.1. Possibility to continue the PING

In this case the *ping* will continue because the host H4 is the originator of the ICMP Echo Request. Therefore, it will update the filtering database of the switch SW-1 immediately after its relocation, when it sends its first new ICMP Echo Request.

Since that frame is directed to host H1, which is known by the switch SW-1, the frame will arrive correctly to host H1 and consequently the ICMP Echo Reply will be sent back immediately.

It is worthy noticing how the filtering database of the switch SW-2 will not be updated by the frames generated by host H4, and hence it will contain an entry with the previous location of host H4 until it expires due to Max Age.

7.6.2. Filtering database of the switches

The content of the filtering database at the end of the process will be the following:

Switch SW-1

MAC address	Interface	Ageing time
00:00:00:11:11:11	Fe0	0
00:00:00:44:44:44	Fe1	0

Switch SW-2

MAC address	Interface	Ageing time
00:00:00:11:11:11	Fe0	120
00:00:00:44:44:44	Fe2	120

In this case, the last two frames seen by switch SW-2 are the two Echo Request/Reply immediately before the relocation, which (as per the exercise text) have been seen 2 minutes ago.

7.7. Solution for exercise n. 9

After the ARP Request and the ARP Reply went on the network, the filtering databases of the switches is the following:

Switch SW-1

MAC address	Interface	Ageing time
00:00:00:11:11:11	Fe0	0
00:00:00:44:44:44	Fe2	0

Switch SW-2

MAC address	Interface	Ageing time
00:00:00:11:11:11	Fe0	0
00:00:00:44:44:44	Fe2	0

As soon as the cable between SW-2 and host H4 is disconnected, the switch SW-2 detects the fault on Fe2 (*missing carrier on the interface*) and deletes the related entry in its filtering database. Upon receiving an ICMP Echo Request from host H1 to host H4, SW-1 forwards it through Fe2 to SW-2, which, without any filtering database entry related to the MAC address of host H4, forwards it on all interfaces but Fe0 (i.e., the one on which the frame was received).

Eventually, after Aging Time also the entry related to host H4 in the filtering database of SW-1 will expire, and also SW-1 will send the ICMP Echo Requests in flooding mode. In this way, the ICMP Echo Request frames will be forwarded over the entire network.

The process will terminate when the entry related to host H4 in host H1's ARP cache expires. At that time, even host H1 will not have the possibility to send the ICMP Echo Request packet because it does no longer have host H4's MAC address. Host H1 will begin a new phase of address resolution (i.e., ARP Request, which are then sent in flooding on the entire network), until its operating system times out and aborts the `ping` command.

7.8. Solution for exercise n. 10

7.8.1. Frames generated on the network

The frames generated on the network are the following:

N.	L2	L3	Appl-layer protocol	Description
1	00:00:00:11:11:11 → FF:FF:FF:FF:FF:FF	—	ARP Request	Who has IP=10.10.10.254 please reply with its MAC address
2	00:00:00:AA:AA:AA → 00:00:00:11:11:11	—	ARP Reply	Host 10.10.10.254 has MAC = 00:00:00:AA:AA:AA
3	00:00:00:11:11:11 → 00:00:00:AA:AA:AA	10.10.10.1 → 10.10.11.3	ICMP	ICMP Echo Request
4	00:00:00:BB:BB:BB → FF:FF:FF:FF:FF:FF	—	ARP Request	Who has IP=10.10.11.3 please reply with its MAC address
5	00:00:00:33:33:33 → 00:00:00:BB:BB:BB	—	ARP Reply	Host 10.10.11.3 has MAC = 00:00:00:33:33:33
6	00:00:00:BB:BB:BB → 00:00:00:33:33:33	10.10.10.1 → 10.10.11.3	ICMP	ICMP Echo Request
7	00:00:00:33:33:33 → 00:00:00:BB:BB:BB	10.10.11.3 → 10.10.10.1	ICMP	ICMP Echo Reply
8	00:00:00:AA:AA:AA → 00:00:00:11:11:11	10.10.11.3 → 10.10.10.1	ICMP	ICMP Echo Reply
9-16	Packets 3 and 6-8 are replicated 3 times			

Please note that frames 1-3 and 8 are transmitted on the LAN on the left, while frames 4-7 are transmitted on the LAN on the right.

7.8.2. Filtering Database

The filtering database on the two switches at the end of the process will be the following:

Switch SW-1

MAC address	Interface
00:00:00:11:11:11	Fe0
00:00:00:AA:AA:AA	Fe2

Switch SW-2

MAC address	Interface
-------------	-----------

00:00:00:BB:BB:BB	Fe0
00:00:00:33:33:33	Fe1

7.9. Solution for exercise n. 11

7.9.1. Frames forwarded by the switch

The frames forwarded by the switch SW-1 correspond to the frames generated on the network:

N.	L2	L3	Appl-layer protocol	Description
1	00:00:00:11:11:11 → FF:FF:FF:FF:FF:FF	—	ARP Request	Who has IP=10.10.10.4 please reply with its MAC address
2	00:00:00:44:44:44 → 00:00:00:11:11:11	—	ARP Reply	Host 10.10.10.4 has MAC = 00:00:00:44:44:44
3	00:00:00:11:11:11 → 00:00:00:44:44:44	10.10.10.1 → 10.10.10.4	ICMP	ICMP Echo Request
4	00:00:00:44:44:44 → 00:00:00:11:11:11	10.10.10.4 → 10.10.10.1	ICMP	ICMP Echo Reply
5-10	Packets 3 and 4 are replicated 3 times			

7.9.2. Ports of the switch involved in receiving/transmitting frames

The ports of the switch involved in receiving/transmitting frames are the following:

PKT#	Description	Fe0	Fe1	Fe2
1	ARP Request	IN	OUT	OUT
2	ARP Reply	OUT	-	IN
3	ICMP Echo Request	IN	-	OUT
4	ICMP Echo Reply	OUT	-	IN

For the hub (which does not have any frame filtering capabilities), the ports involved in receiving/transmitting frames are the following:

Pkt#	Description	Fe0	Fe1	Fe2	Fe3
1	ARP Request	IN	OUT	OUT	OUT
2	ARP Reply	OUT	OUT	IN	OUT
3	ICMP Echo Request	IN	OUT	OUT	OUT
4	ICMP Echo Reply	OUT	OUT	IN	OUT

7.9.3. Filtering database of the switch

The filtering database of the switch is the following:

MAC address	Interface
00:00:00:11:11:11	Fe0
00:00:00:44:44:44	Fe2

Vice versa, the hub does not have any filtering database.

7.10. Solution for exercise n. 12

7.10.1. Frames forwarded by the switch

The frames forwarded by the switch SW-1 correspond to the frames generated on the network:

N.	L2	L3	Appl-layer protocol	Description
1	00:00:00:11:11:11 → FF:FF:FF:FF:FF:FF	—	ARP Request	Who has IP=10.10.10.2 please reply with its MAC address
2	00:00:00:22:22:22 → 00:00:00:11:11:11	—	ARP Reply	Host 10.10.10.2 has MAC = 00:00:00:22:22:22
3	00:00:00:11:11:11 → 00:00:00:22:22:22	10.10.10.1 → 10.10.10.2	ICMP	ICMP Echo Request
4	00:00:00:22:22:22 → 00:00:00:11:11:11	10.10.10.2 → 10.10.10.1	ICMP	ICMP Echo Reply
5-10	Packets 3 and 4 are replicated 3 times			

7.10.2. Ports of the switch involved in receiving/transmitting frames

The ports of the switch involved in receiving/transmitting frames are the following:

PKT#	Description	Fe0	Fe1	Fe2
1	ARP Request	IN	OUT	OUT
2	ARP Reply	OUT	IN	-
3	ICMP Echo Request	IN	OUT	-
4	ICMP Echo Reply	OUT	IN	-

For the hub, the ports involved are the following:

Pkt#	Description	Fe0	Fe1	Fe2	Fe3
1	ARP Request	IN	OUT	OUT	OUT
2	ARP Reply	-	-	-	-
3	ICMP Echo Request	-	-	-	-
4	ICMP Echo Reply	-	-	-	-

7.10.3. Filtering database of the switch

The filtering database of the switch is the following:

MAC address	Interface
00:00:00:11:11:11	Fe0
00:00:00:22:22:22	Fe1

7.11. Solution for exercise n. 13

7.11.1. Frames forwarded by the switch

The frames forwarded by the switch SW-1 correspond to the frames generated on the network:

N.	L2	L3	Appl-layer protocol	Description
1	00:00:00:33:33:33 → FF:FF:FF:FF:FF:FF	—	ARP Request	Who has IP=10.10.10.2 please reply with its MAC address
2	00:00:00:22:22:22 → 00:00:00:33:33:33	—	ARP Reply	Host 10.10.10.2 has MAC = 00:00:00:22:22:22
3	00:00:00:33:33:33 → 00:00:00:22:22:22	10.10.10.3 → 10.10.10.2	ICMP	ICMP Echo Request
4	00:00:00:22:22:22 → 00:00:00:33:33:33	10.10.10.2 → 10.10.10.3	ICMP	ICMP Echo Reply
5-10	Packets 3 and 4 are replicated 3 times			

7.11.2. Ports of the switch involved in receiving/transmitting frames

The ports of the switch involved in receiving/transmitting frames are the following:

PKT#	Description	Fe0	Fe1	Fe2
1	ARP Request	OUT	OUT	IN
2	ARP Reply	-	IN	OUT
3	ICMP Echo Request	-	OUT	IN
4	ICMP Echo Reply	-	IN	OUT

For the hub, the ports involved are the following:

Pkt#	Description	Fe0	Fe1	Fe2	Fe3
1	ARP Request	OUT	IN	OUT	OUT
2	ARP Reply	IN	OUT	OUT	OUT
3	ICMP Echo Request	OUT	IN	OUT	OUT
4	ICMP Echo Reply	IN	OUT	OUT	OUT

7.11.3. Filtering database of the switch

The filtering database of the switch is the following:

MAC address	Interface
00:00:00:22:22:22	Fe1
00:00:00:33:33:33	Fe2

7.12. Solution for exercise n. 14

7.12.1. Frames forwarded by the switch

The frames *received* by the switch SW-1 correspond to the frames generated on the network:

N.	L2	L3	Appl-layer protocol	Description
1	00:00:00:33:33:33 → FF:FF:FF:FF:FF:FF	—	ARP Request	Who has IP=10.10.10.5 please reply with its MAC address
2	00:00:00:55:55:55 → 00:00:00:33:33:33	—	ARP Reply	Host 10.10.10.5 has MAC = 00:00:00:55:55:55
3	00:00:00:33:33:33 → 00:00:00:55:55:55	10.10.10.3 → 10.10.10.5	ICMP	ICMP Echo Request
4	00:00:00:55:55:55 → 00:00:00:33:33:33	10.10.10.5 → 10.10.10.3	ICMP	ICMP Echo Reply
5-10	Packets 3 and 4 are replicated 3 times			

However, the exercise asks for the frames *forwarded*, i.e., frames that are received on one port and are repeated on one or more ports. In this case, only one frame (the first) is forwarded by the switch SW-1.

For instance, being (1) a broadcast frame, it is forwarded on all the ports, but this frame allows also the switch SW-1 to learn the position of host H3. Hence, the switch does not forward frame (2) anymore since it knows that that host is reachable through its left port. For the same reason, frame (3) is no longer forwarded to any other port, since the switch SW-1 just learned the position of host H5 from previous frame, which associates it to its left port.

At this point, all the following frames are not forwarded to any port for the same reason.

7.12.2. Ports of the switch involved in receiving/transmitting frames

The ports of the switch involved in receiving/transmitting frames are the following:

PKT#	Description	Fe0	Fe1	Fe2
1	ARP Request	OUT	OUT	IN
2	ARP Reply	-	-	IN
3	ICMP Echo Request	-	-	IN
4	ICMP Echo Reply	-	-	IN

For the hub, the ports involved are the following:

Pkt#	Description	Fe0	Fe1	Fe2	Fe3
1	ARP Request	OUT	IN	OUT	OUT
2	ARP Reply	OUT	OUT	OUT	IN
3	ICMP Echo Request	OUT	IN	OUT	OUT

4	ICMP Echo Reply	OUT	OUT	OUT	IN
---	-----------------	-----	-----	-----	----

7.12.3. Filtering database of the switch

The filtering database of the switch is the following:

MAC address	Interface
00:00:00:33:33:33	Fe2
00:00:00:55:55:55	Fe2

7.13. Solution for exercise n. 15

7.13.1. Hosts connected to the same switch (scenario 1)

Host H3 can launch a MAC flooding attack on the network, saturating the entries of the filtering database of the switch SW-1.

In these conditions, when host H1 sends a frame to H2, the switch does no longer know the MAC address of H2 and then it forwards that frame on all its ports. Please note that at this time the switch learns the MAC address of host H1 and inserts it in the database. However, if the MAC flooding attack continues, it is very likely that the above MAC address will disappear from the filtering database very soon, before the possible answer sent from host H2 to H1, because of the new (fake) entries learned from the attacker H3. Therefore, also the answer above will be broadcasted by the switch on all its ports, since it does no longer know the MAC address of host H1.

7.13.2. Hosts connected to different switches (scenario 2)

The behavior is exactly the same even if the involved hosts are connected to different switches, since the MAC flooding attack is carried out through packets that have source MAC addresses unknown in the network, and hence are flooded throughout the entire network.

The result is that all the filtering databases on the broadcast network will be affected by this problem, which causes host H3 to be able to intercept all the traffic between any hosts on the network.

A final warning: during a MAC flooding attack, the risk of overloading the network is very high. The attack itself is usually not able to saturate the network (although it should be carried out at a pace that causes the filtering database to contain only “fake” entries), but the problem is that all the traffic gets flooded and hence the aggregated bandwidth is definitely smaller than the one available in normal operating conditions. Therefore, the risk that some packets will not be received because of some droppings occurring in intermediate switches is rather high. For this reason, the attacker H3 will probably be able to intercept all the traffic exchanged by H1 and H2 (which are all attached to the same switch), but it may have some trouble in intercepting the traffic exchanged by other remote hosts (although theoretically possible).

8. Performance

8.1. Solution for exercise n. 16

Due to the client-server model of the traffic exchanged in the network (hosts H1-H47 toward S and vice versa), we expect that the bottleneck will be the bandwidth of the server S, which will remain unchanged after the upgrade.

The network will improve its behavior with respect to collisions, which are no longer present. However this may be a secondary problem in this network where traffic is mainly driven by the server.

Therefore we expect that the service perceived by the users will remain almost equivalent after the upgrade.

8.2. Solution for exercise n. 17

8.2.1. Aggregate Bandwidth

The aggregate bandwidth of the switch is calculated by taking into account the maximum traffic that the switch can handle, independently from the actual traffic pattern.

Considering ports operating in Full-Duplex mode, each port can send and receive at its maximum speed (e.g., port 1 sends data to port 2, while port 2 sends data to port 1, the same applies for ports 3 and 4, etc.).

Therefore, the maximum **incoming** load of the switch when ports operate in Full-Duplex is:

$$100 \text{ Mbps} \times 48 \text{ ports} = 4.8 \text{ Gbps}$$

It is worthy noticing that usually data sheets report the aggregated bandwidth as the sum of the *incoming* and *outgoing* traffic, which lead to 9.6 Gbps, although this number looks a little bit exaggerated from the engineering perspective.

In case of ports operating in Half-Duplex mode, each port can either transmit or receive at a time (e.g., port 1 sends data to port 2, the same applies for ports 3 and 4, etc.), leading to a throughput that is one half of the previous one:

$$100 \text{ Mbps} \times 24 \text{ ports} = 2.4 \text{ Gbps}$$

8.2.2. Throughput of the switch

From the given usage scenario, it is evident that the link that connects the server to the network, which operates at 100 Mbps, represents the bottleneck.

In Full-Duplex mode, all the clients taken together are allowed to send globally 100 Mbps to the server, which are then sent back to the clients. Therefore, the switch will have a throughput of 200 Mbps (i.e., it will forward 200 Mbps of traffic).

In Half-Duplex mode things are a little bit more complicated. Being used in shared mode, the bottleneck link does not allow more than 100 Mbps of traffic. On this link, two sources of traffic are present: the switch and the server. Since Ethernet networks do not give precedence to any network interface (i.e., each interface has the same probability to take the ownership of the channel), we can assume that the traffic generated on that link will come half from the switch and half from the server. This means that the switch will forward 50 Mbps of traffic coming from all the clients taken together and directed to the server, while the remaining 50 Mbps will be consumed by the server to send the same data back to the clients. Therefore, the throughput of the switch will be 100 Mbps (i.e., the switch will forward 100 Mbps of traffic).

It is worthy noticing that the offered load on the switch is more than 100 Mbps. In fact, since the traffic is UDP, clients will generate traffic at their maximum rate, which is close to 100 Mbps¹. This traffic reaches the switch but is not *forwarded* in its entirety because of the bottleneck link. Most of the traffic coming from hosts is then lost and hence is not valid for calculating the throughput of the switch.

¹Not exactly 100 Mbps because of some traffic coming back from the server, which consumes some bandwidth on the client links, which operate in half-duplex mode.