SCIPER: ____________________

First name: ________________________ Family name: ________________________

EXAM
TCP/IP NETWORKING
Duration: 3 hours

Jean-Yves Le Boudec

January 2023

INSTRUCTIONS

1. Verify that you have 4 problems + one figure sheet.
2. Write your solution into this document and return it to us (you do not need to return the figure sheet). You may use additional sheets if needed. Do not forget to write your name on **each of the four problem sheets** and **all** additional sheets of your solution.
3. All problems have the same weight.
4. Briefly justify your answer. For grading, the justification is as important as the solution itself.
5. If you find that you need to make additional assumptions in order to solve some of the questions, please describe such assumptions explicitly.
6. Figures are on a separate sheet, for your convenience.
7. No documents, no electronic equipments are allowed.

PROBLEM 1

Consider the network in the figure sheet. A, B, S1 and S2 are hosts. R1 and R2 are routers. BR1 to BR6 are bridges. The cloud represents the public internet. O1, O2, , ... , O11 are observation points. Machines where both IPv4 and IPv6 addresses appear are dual-stack, others are IPv6 only.

- The network box N1 operates as a NAT for IPv4 and IPv6. The internal port of the NAT is N1n.
- The network boxes N2 operates as a NAT for IPv4 (with internal port N2n) and as a router for IPv6.

All links are full duplex Ethernet. We assume that all machines are correctly configured (unless otherwise specified), proxy ARP is not used and there is no VLAN. A's default gateway is N1 and B's default gateway is N2. The hosts, bridges and routers have been running for some time and their different protocols are in steady-state. There is no other system or interface than shown on the figure.

The figure shows explicit IP addresses. Note that the IPv4 addresses of A and B are identical, as are N1n and N2n's. The symbols A and B also represent the MAC addresses of the hosts A and B. Similarly, the notation N2s, N2n, etc. denotes the MAC addresses of the corresponding interfaces.

1. Explain why it is possible that A and B have same IPv4 addresses.

2. Write S1 and B's IPv6 addresses in non-compressed forms (use hexadecimal digits).

Host	Compressed address	Non-compressed address
S1	3000:a:b:c::20	
B B	3000:a:b:c:2::b1 fe80::1	

3. S1, N1s, N2s and R2n all have the same IPv4 subnet mask. Say which of the following values are valid for this common subnet mask.

   For this table and the next, the usual rules of quizzes apply: (i) if you cross the correct box in one row you obtain the full score for this row; (ii) if you cross zero or two boxes in one row, your score is 0 for this row; (iii) if you cross exactly one box in one row and it is incorrect your score for this row is negative, namely a penalty equal to half the score for this row.

Proposed v4 subnet masks at S1, N1s, N2s and R2n	valid	invalid
255.0.0.0	☐	☐
255.255.0.0	☐	☐
255.255.255.0	☐	☐

4. B and S1 have the same IPv6 subnet mask length. Say which of the following values are valid for this common subnet mask length.

Proposed v6 subnet mask lengths at B and S1	valid	invalid
64	☐	☐
78	☐	☐
80	☐	☐
82	☐	☐

5. The routing table at S1 is as follows.

dest	next-hop
3000:a:b:c:2::/80	3000:a:b:c::12
::	3000:a:b:c::1

There is no other entry, in particular, there is no IPv4 destination in S1's routing table. Now assume that S1 has an IPv6 packet to send to B. Will this work ? If so, to which destination MAC address does S1 send the packet to ? Same question for an IPv4 packet.

6. A sends one UDP packet to S1, S1 sends one to S2, S2 sends one to B and B sends one to A. All are sent using IPv6. We observe the IPv6 packets containing these messages at observation points O3, O5, O6, and O9 (in both directions at every observation point). For each observation point, say whether such packets are observed and if so, give the IPv6 destination address. Give your answer in the table below. Put one line per packet (leave blank if no packet is observed). Justify your answer in the box below the table.

At	IPv6 destination address(es)
O3	
O5	
O6	
O9	
Justification:	

7. B downloads a large web page from S2 using QUIC and IPv6. The local port used by B is 8765. At about the same time, A also downloads a large web page from S2 using QUIC over IPv6. By coincidence, the local port used by A is also 8765. The port used by S2 for HTTP over QUIC is 443. With Wireshark, we observe the packets sent by S2 to B and A at O11, O10, O8, O2, O7 and O1. Write the values of the fields in the tables below.

From S2 to B							
	MAC addresses		IP addresses			port numbers	
At	srce	dest	srce	dest	protocol	srce	dest
O11							
O10							
O8							
O7							

From S2 to A							
	MAC addresses		IP addresses			port numbers	
At	srce	dest	srce	dest	protocol	srce	dest
O11							
O10							
O2							
O1							

8. B now downloads a large web page from S1 using TLS and IPv4. The local port used by B is 4321. At about the same time, A also downloads a large web page from S1 using TLS and IPv4. Again, by coincidence, the local port used by A is also 4321. The port used by S1 for HTTP over TLS is 443. With Wireshark, we observe the packets sent by S1 to B and A at O4, O1 and O7. Write the values of the fields in the tables below.

From S1 to B							
	MAC addresses		IP addresses			port numbers	
At	srce	dest	srce	dest	protocol	srce	dest
O4							
O7							

From S1 to A							
	MAC addresses		IP addresses			port numbers	
At	srce	dest	srce	dest	protocol	srce	dest
O4							
O1							

9. S1 now does a traceroute to B's second IPv6 address, namely to `fe80::1`. We observe the traffic resulting from this activity at observation point O7, in both directions. Which packets do we observe?

SCIPER: ____________________

First name: ________________________ Family name: ________________________

PROBLEM 2

Consider the network for problem 2 in the figure sheet. There are four ASs, A, B, C and D with routers A1, A2, B1, C1, C2, C3, R1 and D1. The physical links are shown with plain lines. Each AS uses OSPF with Equal Cost Multipath as IGP, and every router inside each AS uses OSPF. The cost of every link and every directly attached network is 1, except when otherwise specified. C1 injects into OSPF the route `::` (the route with empty prefix), i.e., in C, OSPF will compute shortest path to `::` as if the empty prefix were directly attached to C1.

The figure shows one stub network at router A1 and one at router D1, each shown with its IPv6 network prefix. The lower case symbols such as a1e, b1n, etc. represent IPv6 addresses. H1 is a host, its IPv6 address is also indicated on the figure.

Routers A1, A2, B1, C1, C2 , C3, and D1 use BGP with their external neighbours and as required with their internal neighbours. The router R1 may or may not use BGP, depending on the question. No confederation or route reflector is used.

We assume that the BGP decision process use the following criteria in decreasing order of priority. BGP identifiers are router names such as A1, A2...

1. Highest LOCAL-PREF
2. Shortest AS-PATH
3. Lowest MED, if taken seriously by this network
4. E-BGP is preferred over I-BGP
5. Shortest path to NEXT-HOP, according to IGP
6. Lowest BGP identifier of sender of route is preferred. Here comparison is lexicographic with A<B<C<D and $1 < 2$; for example A1 is preferred over A2, A2 is preferred over B1, etc...

Furthermore, we assume the following.

- Unless otherwise specified, when receiving an E-BGP announcement, every BGP routers tags it with LOCAL-PREF = 0. No other optional BGP attribute (such as MED, etc.) is used in BGP messages.
- No aggregation of route prefixes is performed by BGP.
- The policy in A, B, C and D is such that all available routes are accepted and propagated to neighbouring ASs, as long as the rules of BGP allow.
- Every router redistributes internal OSPF destinations into BGP. Unless otherwise specified, there is no other redistribution.
- Every router performs recursive forwarding-table lookup.
- Equal Cost Multi-Path routing is supported by all routers.
- When writing the AS-PATH attribute received by a router X, do not write the AS that this router is in. For example, if router X in AS A receives a route from router Y in AS B, who received it from AS C, the AS-PATH stored by X in its RIB is `AS-PATH = B C`. However, if X propagates this route over E-BGP to a router in AS D, the message sent by X will have `AS-PATH = A B C`.

1. At time t_1, BGP and OSPF have converged in all ASs.

   (a) At time t_1, what is the best BGP route to `3000:a:b::/48` selected by C1 ? From which BGP peer was it received ? What are its BGP NEXT-HOP and AS-PATH attributes ? Same question for C2 and C3. Give your answers in the table below, with a short justification.

At	From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH
C1		3000:a:b::/48		
C2		3000:a:b::/48		
C3		3000:a:b::/48		

Justification:

   (b) Still at time t_1, what is the best BGP route to `3000:a:b::/48` selected by B1 ? From which BGP peer was it received ? What are its BGP NEXT-HOP and AS-PATH attributes ? Give your answer in the table below, with a short justification.

At	From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH
B1		3000:a:b::/48		

Justification:

(c) Still at time t_1, what is the list of all BGP routes received and stored by C3 in its RIB, with destination = 3000:a:b::/48 ? For every route, indicate from which BGP peer it was received from, the BGP NEXT-HOP, the AS-PATH attributes, and whether it is selected as best. Give your answer in the table below, with a short justification (put as many rows as necessary).

At C3 :				
From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH	Best route ?
	3000:a:b::/48 3000:a:b::/48 ...			
Justification:				

(d) Still at time t_1, H1 sends a packet to `3000:a:b:c:d:e:f::1`. Which path will this packet follow? At every router except the final router A1, say by which protocol (BGP or OSPF, or both) the next-hop is obtained. Justify your answer.

2. In this question we assume that the network is restarted, with the following change in the configuration:

   - R1 also runs BGP.

   (a) At time t_2, BGP and OSPF have converged in all ASs and H1 sends a packet to `3000:a:b:c:d:e:f::1`. Which path will this packet follow? At every router except the final router A1, say by which protocol (BGP or OSPF, or both) the next-hop is obtained. Justify your answer.

   (b) At time $t_3 > t_2$, C1 is re-started. Unfortunately, there is a configuration error in the new configuration file at C1: the OSPF interface database of C1 is configured with the prefix `3000:a:b:c::/64` (i.e the OPSF software believes that the network `3000:a:b:c::/64` is directly attached at C1). The default prefix `::` continues to be injected into OSPF by C1. At time $t_4 > t_3$, OSPF and BGP have converged again and R1 has a packet to send to `3000:a:b:c:d:e:f::1`. Which path will this packet follow? At every router except the final router A1, say by which protocol (BGP or OSPF, or both) the next-hop is obtained. Justify your answer.

3. In this question we assume that the network is restarted, with the following change in the configuration:

   - Inside AS A: routers tag all incoming E-BGP announcements received from AS C with LOCAL-PREF = 75 and from AS-D with LOCAL-PREF = 50. Other incoming E-BGP announcements are tagged with LOCAL-PREF = 0, same as before.
   - Inside AS C: routers tag all incoming E-BGP announcements received from AS A with LOCAL-PREF = 100 and from AS-D with LOCAL-PREF = 50. Other incoming E-BGP announcements are tagged with LOCAL-PREF = 0, same as before.

   The rest of the configuration is as in the previous question.
   At time t_3, BGP and OSPF have converged in all ASs. What is the best route to `4000:c:d::/48` selected by B1 ? Justify your answer.

SCIPER: ____________________

First name: ________________________ Family name: ________________________

PROBLEM 3

Consider the network for Problem 3 on the figure sheet.

- R1, R2 and R3 are routers. H1, H2, H3, S1, S2 and S3 are hosts. The links are full duplex with same rate in both directions.
- S1 [resp. S2, S3] sends one large flow of data to H1 [resp. H2, H3] (there are thus 3 unidirectional flows in total).
- There is no other system and no other flow than shown on the figure. There is no other capacity constraint than the link capacities shown on the figure. We also neglect the impact of the acknowledgement flows in the reverse direction.
- We neglect all overheads and assume that the link capacities can be fully utilized at bottlenecks.
- For $i = 1, 2, 3$ we call x_i the rate of the flow from Si to Hi .

1. Which of the following allocations, in Mb/s, are Pareto-efficient ? Justify your answer.
   (a) $x_1 = 1,\ x_2 = 2,\ x_3 = 6$
   (b) $x_1 = 0,\ x_2 = x_3 = 4.5$
   (c) $x_1 = x_2 = x_3 = 3$
   (d) $x_1 = 7,\ x_2 = x_3 = 1$

2. Assume the rates x_1, x_2, x_3 of the three flows are allocated according to max-min fairness. What are all the possible allocations ? Which ones are Pareto-efficient ?

3. Assume the rates x_1, x_2, x_3 of the three flows are allocated according to proportional fairness. What are all the possible allocations ? Which ones are Pareto-efficient ?

4. In this question we assume that the three flows use TCP Reno with ECN. Queuing at all routers is FIFO with RED enabled. The round trip times are:

   - 40 ms for flow 1,
   - 10 ms for flows 2 and 3.

   These numbers include all processing times. The MSS is the same for all flows and is equal to 1250 Bytes = 10^4 bits. We assume that the offered window is very large. Compute the rates of flows 1, 2 and 3.

5. In this question we continue with the same assumptions as in the previous question, except for the following modifications:
   - R3 now acts as an application layer gateway for flow 1.
   - The connection S1 - R3 has the same RTT as flows 2 and 3, namely 10 ms.
   - The RTT of the connection R3 - H1 is 122 ms.

   What is the rate achieved by flow 1 (from S1 to H1)? What are the rates achieved by flows 2 and 3?

6. In this question we continue with the same assumptions as in the previous question, except for the following modifications:
   - The link R3-H1 is wireless and has a loss ratio equal to $\ell = 0.0001$

   What is the rate achieved by flow 1 (from S1 to H1)? What are the rates achieved by flows 2 and 3?

7. In this question we continue with the same assumptions as in the previous question, except for the following modifications:

   - the loss ratio of the link R3-H1 is now equal to $\ell = 0.01$

   What is the rate achieved by flow 1 (from S1 to H1)? What are the rates achieved by flows 2 and 3?

SCIPER: ____________________

First name: ________________________ Family name: ________________________

PROBLEM 4

1. Lisa uses a media server to stream music in her house (see on the figure sheet). The audio server uses IP multicast, with source specific multicast. It uses the multicast address `232.1.2.3`. The music stream is unidirectional, from the media server to whomever is listening.

   (a) Lisa and Bart receive the music stream, each on their own separate device. Homer also would like to receive the audio stream. Say what happens at the IP layer when Homer decides to receive the audio stream. In particular, among the following machines: the media server, Lisa, Bart and Homer's machines, say which have to send specific IP layer control messages for Homer to be be able to receive the audio stream. (Application layer messages are not considered here).

   (b) We assume in this question that Lisa's home network is a single bridged LAN. Lisa does a packet capture at the media server and at her device; she observes only the packets that contain the audio stream. Which addresses does she see in the packets sent by the media server ? Put the answer in the table below (see on the figure sheet for device addresses).

At media server:			
MAC source	MAC dest	IP source	IP dest

At Lisa's device:			
MAC source	MAC dest	IP source	IP dest

   (c) Assume that Lisa's home network contains several routers and Lisa is more than one hop away from the media server. Is there any change to the answer of the previous question ?

2. Consider the network for Problem 4, Question 2 on the figure sheet. All routers are BIER routers; R1 to R8 are BFERs (i.e., BIER Egress routers). The IP forwarding table of R10 is shown (only the lines that contain destination IP address = R1 to R8 are shown).
The BIER network uses a bit-string of 8 bits, where position 1 is the least significant bit and position 8 is the most significant bit. Fill in the values of the forwarding bit masks in the Bier Index Forwarding table at R10. Use hexadecimal notation.

Destination BFER	Forwarding Bit Mask	Next-Hop
1		R9
2		R9
3		R12
4		R12
5		R5
6		R6
7		R12
8		R12

3. Consider the network for Problem 4, Question 3 in the figure sheet. Borducom is an Internet Service Provider (ISP) that has an IPv4 backbone and provides IPv6 connectivity by means of 6rd. Borducom has a box R which is a 6rd relay router; it connects to IPV4 on one side, and to IPv6 on the other side. The IPv6 side of R6 is connected to Tornado Networks, another ISP, which provides only IPv6 connectivity and has an IPv6-only backbone. Tornado networks has a box N that acts as a NAT64. The block of addresses 100.64.0.0/16 is reserved by Tornado Networks for the IPv4 side of the NAT64. The prefix 3000:a:b:ffff::/96 is reserved by Tornado Networks for the IPv6 side of the NAT64.
You are a customer of Borducom and have a box M that connects your home network to Borducom. The box M acts as a 6rd router, and receives from Borducom the address shown on the figure.
The prefix 2bbb:bbbb::/32 is reserved by Borducom for 6rd. The IPv6 prefix allocated to a customer network (such as our home network) is the concatenation of this prefix and the customer's IPv4 address provided by Borducom. This makes a 64-bit prefix, which can be used freely by the home network. The address 3.88.88.88 is reserved by Borducom to represent the IPv6 internet and is allocated to Borducom's 6rd relay router R.
The hosts and routers have been running for some time and their different protocols are in steady-state.

   (a) Your home network is IPv6 only. Give a possible value for A's IPv6 address.

   (b) The host A downloads a file from S6, using HTTP over TLS. We observe the IP headers of the packets resulting from this transfer, in the direction from S6 to A, at observation points O1, O2, O4 and O5. For each observation point, say whether there is a packet, and if so, give the IPv4 or/and IPv6 source and destination addresses. Write your answer in the table below.

Packets from S6 to A	
At	IPv4 or/and IPv6 source and destination addresses
O5	
O4	
O2	
O1	
Justification:	

(c) S6 downloads a file from S4, using HTTP over TLS. We observe the IP headers of the packets resulting from this transfer, in the direction from S4 to S6, at observation points O2, O3, O4 and O5. For each observation point, say whether there is a packet, and if so, give the IPv4 or/and IPv6 source and destination addresses. Write your answer in the table below.

Packets from S4 to S6	
At	IPv4 or/and IPv6 source and destination addresses
O2	
O3	
O4	
O5	
Justification:	

4. Consider the network for Problem 4, Question 4 in the figure sheet. A,B are LERs; the ports a,b of A and B are MPLS ports; the ports c,d,e,f of A and B are non-MPLS Ethernet ports. M,N,P and Q are LSRs. R is a non-MPLS router.
The host H sends a ping packet to the destination address 10.1.1.201; the destination replies and the ping reply is received by H. We observe the ping request and reply packets at observation points O1 to O6 (if they appear at such points). Give the MPLS labels and IP destination addresses (if any) that we see in such packets. Leave the space blank if there is no such header or no such packet.

Ping packet from H			
At	Outermost MPLS label	Next MPLS label	IP destination address
O1			
O2			
O3			
O4			
O5			
O6			

Ping packet to H			
At	Outermost MPLS label	Next MPLS label	IP destination address
O1			
O2			
O3			
O4			
O5			
O6			

TCP IP EXAM - FIGURES

For your convenience, you can separate this sheet from the main document. Do not write your solution on this sheet, use only the main document. You do not need to return this sheet.

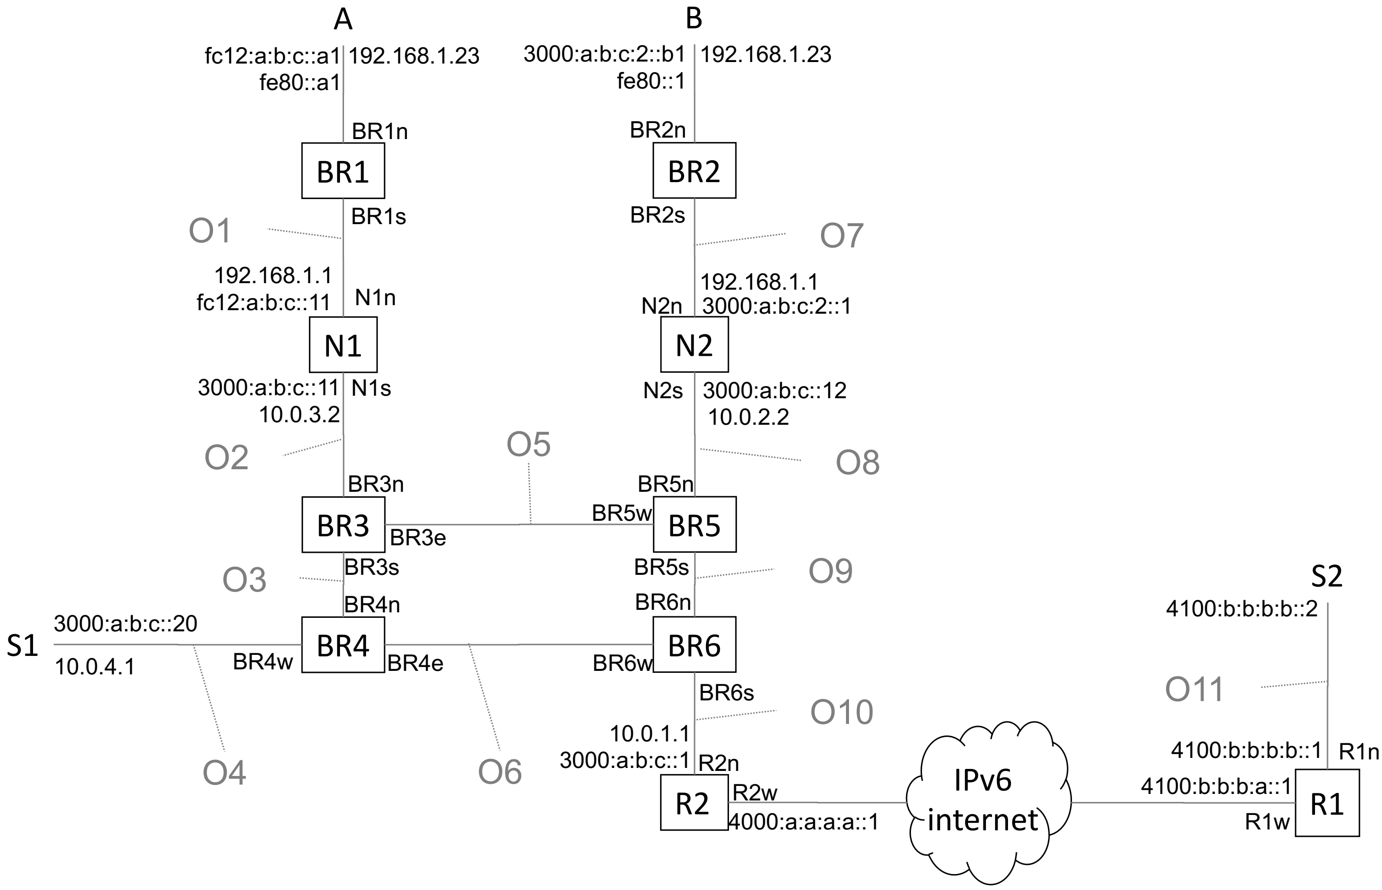

Problem 1.

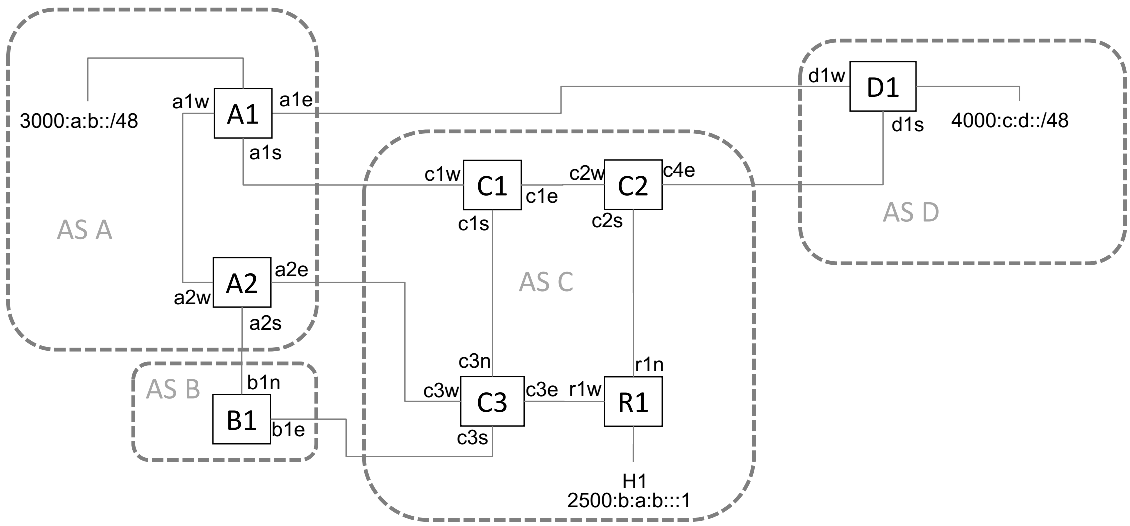

Problem 2.

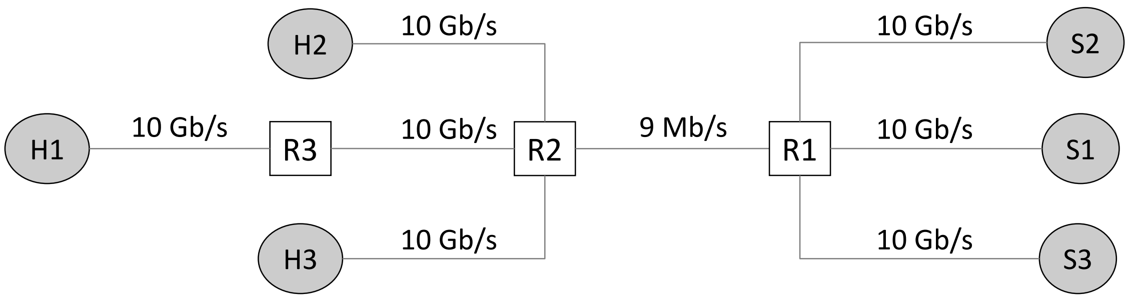

Problem 3.

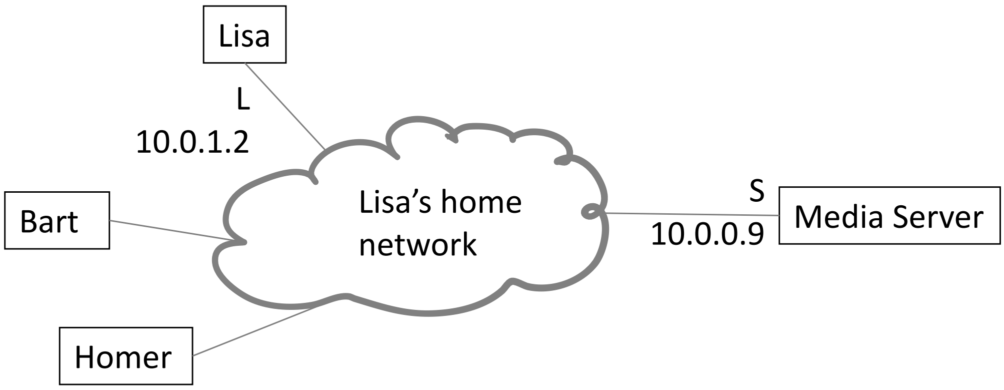

Problem 4, Question 1. L and S are MAC addresses.

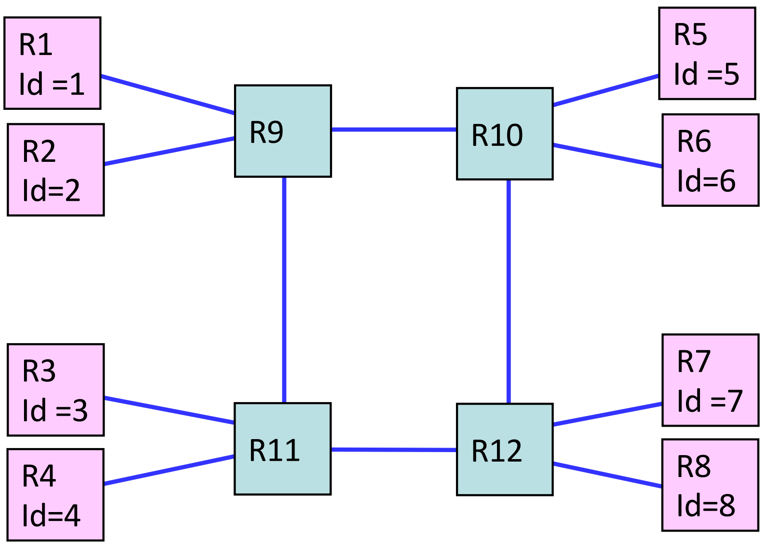

IP Forwarding Table at R10

Dest. IP addr	Next-Hop
R1	R9
R2	R9
R3	R12
R4	R12
R5	R5 (onlink)
R6	R6 (onlink)
R7	R12
R8	R12

Problem 4, Question 2.

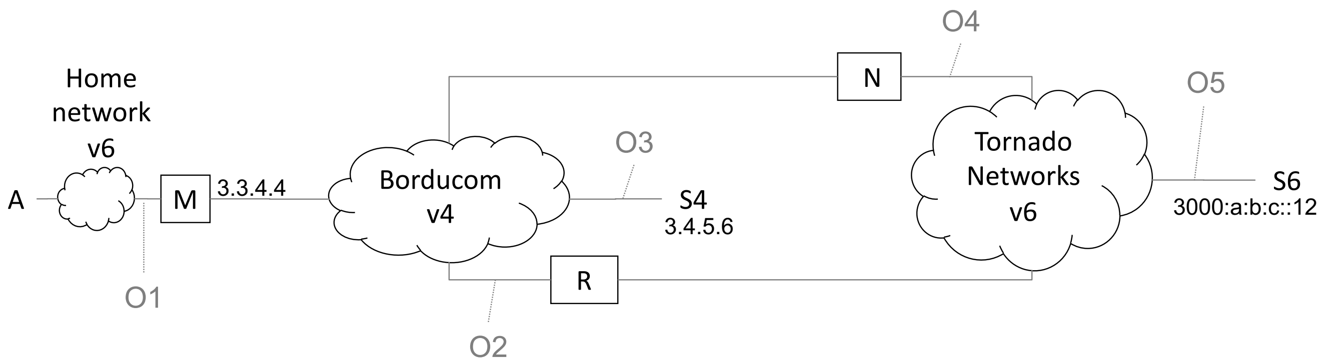

Problem 4, Question 3.

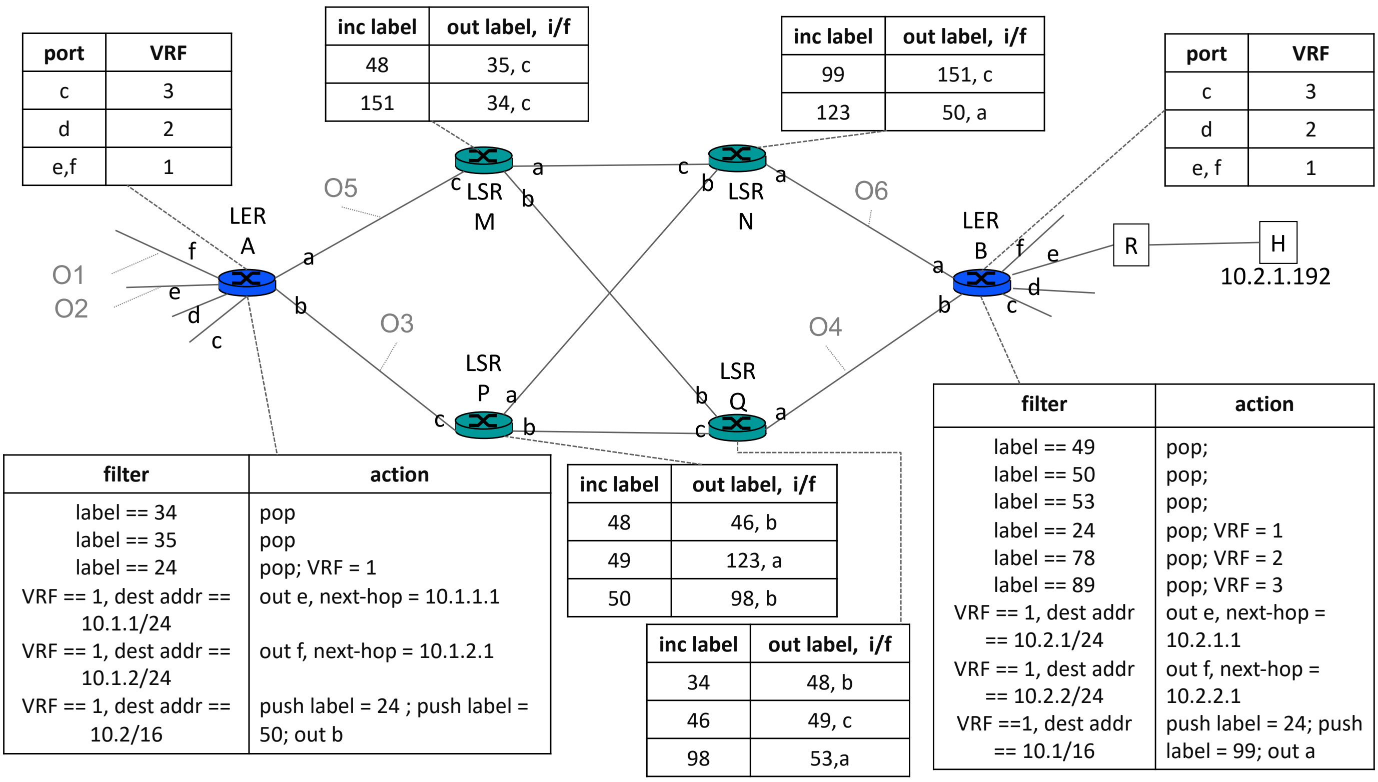

port	VRF
c	3
d	2
e,f	1

inc label	out label, i/f
48	35, c
151	34, c

inc label	out label, i/f
99	151, c
123	50, a

port	VRF
c	3
d	2
e, f	1

filter	action
label == 34	pop
label == 35	pop
label == 24	pop; VRF = 1
VRF == 1, dest addr == 10.1.1/24	out e, next-hop = 10.1.1.1
VRF == 1, dest addr == 10.1.2/24	out f, next-hop = 10.1.2.1
VRF == 1, dest addr == 10.2/16	push label = 24 ; push label = 50; out b

inc label	out label, i/f
48	46, b
49	123, a
50	98, b

inc label	out label, i/f
34	48, b
46	49, c
98	53,a

filter	action
label == 49	pop;
label == 50	pop;
label == 53	pop;
label == 24	pop; VRF = 1
label == 78	pop; VRF = 2
label == 89	pop; VRF = 3
VRF == 1, dest addr == 10.2.1/24	out e, next-hop = 10.2.1.1
VRF == 1, dest addr == 10.2.2/24	out f, next-hop = 10.2.2.1
VRF ==1, dest addr == 10.1/16	push label = 24; push label = 99; out a

Problem 4, Question 4.