

SCIPER: _____

First name: _____ Family name: _____

EXAM
TCP/IP NETWORKING
Duration: 3 hours

Jean-Yves Le Boudec

January 2018

INSTRUCTIONS

1. 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.
2. All problems have the same weight.
3. You may need to make additional assumptions in order to solve some of the questions. If this happens, please describe such assumptions explicitly.
4. Figures are on a separate sheet, for your convenience.
5. No documents, no electronic equipments are allowed.
6. Justify every answer with a short explanation.

PROBLEM 1

Consider the network for problem 1 in the figure sheet. A , B and C are hosts; $BR1$, $BR2$ and $BR3$ are bridges; $R1$ and $R2$ are routers. X is a network box that can be configured in different ways, as explained next. $O1$ to $O6$ are observation points where we observe traffic in both directions of the link. Some selected IP addresses are shown, as well as some selected MAC addresses (denoted with e.g. A , B , $BR1w$, ..., $R1e$, ...). You may need to specify other IP or MAC addresses. 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.

The network masks are $ffff:ffff:ffff:ffff::$ at all interfaces shown on the figure, unless otherwise specified. Addresses such as $fd24:a:b:1::2$ are private IPv6 addresses.

1. Give the uncompressed version (all hexadecimal digits) of the address $fd24:a:b:1::2$.

2. In this question X is a NAT.
 - (a) Give possible values for the complete IP addresses of the interfaces $R1e$ and $R2e$ (i.e. give possible values for the unknown parts of the IP addresses indicated as ?). Shortly justify your answer.

 - (b) A sends one ping message to B . We observe the ping request packets resulting from this activity at observation point $O1$. What are the MAC and IP source and destination addresses in such packets ? What is the Hop Count field, knowing that the HC value is equal to 64 in all IP packets generated by all hosts in this problem ? Put your answers in the table below.

At observation point $O1$:				
MAC source	MAC dest	IP source	IP dest	HC value

- (c) A downloads a huge file from a web server at C using HTTP. A uses the local port 4567. At the same time, B also downloads a file from C , also using HTTP. By coincidence, B uses the same local port number, namely 4567. We observe the IP headers in the packets resulting from this transfer at $O5$ and $O6$, in the direction from A to C and from B to C . Give possible values of the protocol, the source and destination port numbers and the source and destination IP addresses. Give the answers in the tables below.

At observation point O_5, from A and B to C:				
IP source	IP dest	protocol	source port	dest. port

At observation point O_6, from A and B to C:				
IP source	IP dest	protocol	source port	dest. port

3. In this question X is an application layer gateway, acting as web proxy.
- (a) A downloads a huge file from a web server at C using HTTP. A uses the local port 4567. At the same time, B also downloads a file from C , also using HTTP. By coincidence, B uses the same local port number, namely 4567. We observe the IP headers in the packets resulting from this transfer at O_5 and O_6 , in the direction from A to C and from B to C . Give possible values of the protocol, the source and destination port numbers and the source and destination IP addresses. Give the answers in the tables below.

At observation point O_5, from A and B to C:				
IP source	IP dest	protocol	source port	dest. port

At observation point O_6, from A and B to C:				
IP source	IP dest	protocol	source port	dest. port

(b) B sends one ping message to X . Where, among observation points $O2$, $O3$ and $O4$, is the ping request packets resulting from this activity visible ? Justify your answer.

4. In this question we would like that (1) X is configured as a router, (2) the IP address of Xw is unchanged, and (3) the ISP that offers internet service at point $O6$ to N has to inject only one network prefix into its routing tables for the whole of N . IP addresses and masks of interfaces in N (other than Xw) can be modified; no other change is allowed to the configuration of N . Give a possible way to achieve this goal. Give all details of all addresses and masks that may need to be changed by your solution.

SCIPER: _____

First name: _____ Family name: _____

PROBLEM 2

Consider the network for problem 2 in the figure sheet. Unless otherwise specified, there are three ASs, A , B and C . Each AS uses OSPF with Equal Cost Multipath as IGP, and all routers inside each AS uses OSPF. The cost of every link and every directly attached network is 1, except for two links shown on the figure, which have a cost of 10.

The figure shows two stub networks, at routers $A3$ and $A4$, with their IPv6 address prefixes. The lower case symbols such as $b1e$, $b1w$ also represent IPv6 addresses.

Routers $A1$, $A2$, $B1$, $B2$, $C1$ and $C2$ use BGP with their external neighbours and as required with their internal neighbours. Unless otherwise specified, the other routers do not use BGP. We assume that the BGP decision process use the following criteria in decreasing order of priority.

1. Shortest AS-PATH
2. Lowest MED, if taken seriously by this network
3. E-BGP > I-BGP
4. Shortest path to NEXT-HOP, according to IGP
5. Lowest BGP identifier

Furthermore, we assume that no optional BGP attribute (such as MED, LOCAL-PREF etc.) is used in any BGP message and that no aggregation is performed.

1. Each AS is configured to redistribute internal OSPF routes into BGP and to redistribute E-BGP routes into OSPF.
 - (a) When BGP and OSPF have converged in all ASs, which routes are selected by $B1$, $B2$, $A1$ and $A2$? Give the answers in the tables below. Justify your answers.

At B1	DESTINATION NETWORK	NEXT-HOP	AS-PATH
	2001:4::/32		
	2001:3::/32		
At B2	DESTINATION NETWORK	NEXT-HOP	AS-PATH
	2001:baba:bebe/48		
	2001:baba:bebf/48		
Justification:			

At A1	DESTINATION NETWORK	NEXT-HOP	AS-PATH
	2001:4::/32		
	2001:3::/32		
At A2	DESTINATION NETWORK	NEXT-HOP	AS-PATH
	2001:baba:bebe/48		
	2001:baba:bebf/48		
Justification:			

- (b) Recall that each AS is configured to redistribute E-BGP routes into OSPF. When redistributing BGP into OSPF, the OSPF cost of the redistributed route is set to the OSPF cost to the BGP next-hop. When OSPF has converged, what are the routing table entries for the destinations 2001:baba:bebe/48 and 2001:baba:bebf/48 at routers A3 and A4 ? Justify your answer.

At A3	DESTINATION NETWORK	NEXT-HOP	DISTANCE
	2001:baba:bebe/48		
	2001:baba:bebf/48		
At A4	DESTINATION NETWORK	NEXT-HOP	DISTANCE
	2001:baba:bebe/48		
	2001:baba:bebf/48		
Justification:			

- (c) Which path does a packet follow from $A4$ to $2001:baba:bebe::1$ and $2001:baba:bebf::1$? Give the answer in the form of a sequence of routers. Give the answer in the table below.

From $A4$ to	Path
$2001:baba:bebe::1$	
$2001:baba:bebf::1$	
Justification:	

- (d) Assume, in this question only, that $B2$ is compromised and now sends to $A2$ the (bogus) BGP announcement

`DESTINATION = 2001:baba:bebe::/49, AS-PATH = B C, NEXT-HOP=c1w`

Also assume that $A2$ accepts this BGP announcement. When BGP and OSPF have stabilized again in AS A , which path inside AS A does a packet follow from $A4$ to $2001:baba:bebe::1$? Give the answer below in the form of a sequence of routers. Justify your answer.

From $A4$ to	Path
$2001:baba:bebe::1$	
Justification:	

2. The configuration of AS A is now modified: BGP is not redistributed into OSPF at A (but continues to be redistributed at B and C). Explain what can be done in AS A to maintain full connectivity, in particular, we would like all routers in AS A to be able to forward packets to $2001:baba:bebe/48$ and $2001:baba:bebf/48$ (propose only one solution). Note that, in this question, you may make changes to the assumptions that we put at the beginning of Problem 2.
3. Following a re-organization, the three domains A , B and C are now merged into one single large AS. We assume in this question that all routers in the figure use OSPF with the assumptions given at the beginning of Problem 2 (note that all routers in the figure are in the same OSPF domain for the moment). BGP is not used in the rest of this problem.
- (a) Assume that there is a single area, i.e. all routers in the figure are in the same OSPF area. When OSPF has converged, what are the routing table entries for the destinations $2001:baba:bebe/48$ and $2001:baba:bebf/48$ at routers $A3$ and $A4$? Justify your answer.

At A3	DESTINATION NETWORK	NEXT-HOP	DISTANCE
	2001:baba:bebe/48		
	2001:baba:bebf/48		
At A4	DESTINATION NETWORK	NEXT-HOP	DISTANCE
	2001:baba:bebe/48		
	2001:baba:bebf/48		
Justification:			

(b) We would now like to configure this OSPF network with 3 areas, such that:

- i. $B1$ and $B2$ are in the backbone area
- ii. $A1, A2, A3, A4$ are in a common area
- iii. $C1, C2, C3, C4$ are in a common area
- iv. no router is a single point of failure

Give a possible allocation of routers to areas that satisfy these constraints. Justify your answer.

SCIPER: _____

First name: _____ Family name: _____

PROBLEM 3

Consider the network for problem 3 on the figure sheet.

- A , B and C are routers. The capacity of the links between them is 480 Mb/s. The links are full duplex with same rate in both directions.
 - There are 4 unidirectional flows, as shown on the figure. There is no other system and no other flow than shown on the figure. There is no other capacity constraint than the two 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.
1. Assume the rates x_1, x_2, x_3, x_4 of the four flows are allocated according to max-min fairness. Compute the values of x_1, x_2, x_3, x_4 .

2. Which of the following allocations, in Mb/s, are Pareto-efficient ? Justify your answer.

- (a) $x_1 = 160, x_2 = 320, x_3 = x_4 = 160$
- (b) $x_1 = 240, x_2 = 240, x_3 = x_4 = 120$
- (c) $x_1 = 240, x_2 = 240, x_3 = x_4 = 100$
- (d) $x_1 = 240, x_2 = 240, x_3 = 140, x_4 = 100$

3. Assume the rates x_1, x_2, x_3, x_4 of the four flows are allocated according to proportional fairness. Compute the values of x_1, x_2, x_3, x_4 . Is this a Pareto-efficient allocation ?

4. In this question flow 1 is using UDP and sends at a constant rate equal to 80 Mb/s. Flows 2, 3, 4 use TCP Reno with ECN. Queuing at all routers is FIFO with RED enabled. The round trip times are:
- 300 ms for flows 2 and 3,
 - 100 ms for flow 4.

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 2, 3 and 4.

5. Assume now that the flows 2, 3, 4 use TCP Cubic with ECN instead of TCP Reno with ECN; the rest is as in the previous question. Can you guess how the allocation of rates would differ ? We don't ask you to compute the rates; simply put an *X* in the correct boxes in the table below, with a short justification.

the rate achieved with TCP Cubic and ECN is ...	flow 2	flow 3	flow 4
higher			
same			
lower			
Justification:			

SCIPER: _____

First name: _____ Family name: _____

PROBLEM 4

1. $R1$ is a router and uses RIP, a distance-vector routing protocol that implements the distributed Bellman-Ford algorithm. All link costs are equal to 1. At time t_0 , the routing table at $R1$ contains

Destination Network	Next-Hop	Distance
10.0.0.0 / 8	10.1.1.1	5
10.1.0.0 / 16	10.1.2.1	2

The router $R2$ is a neighbour of $R1$ and has IP address 10.1.2.1. At $t_1 > t_0$, $R1$ receives from $R2$ the distance-vector message

Destination = 10.0.0.0/8, distance = 3

Destination = 10.1.0.0/16, distance = 3

No other message is received between t_0 and t_1 . Just after processing this message, what is the state of the routing table at $R1$? Give your solution by filling the empty cells in the table below.

Destination Network	Next-Hop	Distance
10.0.0.0 / 8		
10.1.0.0 / 16		

2. 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 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

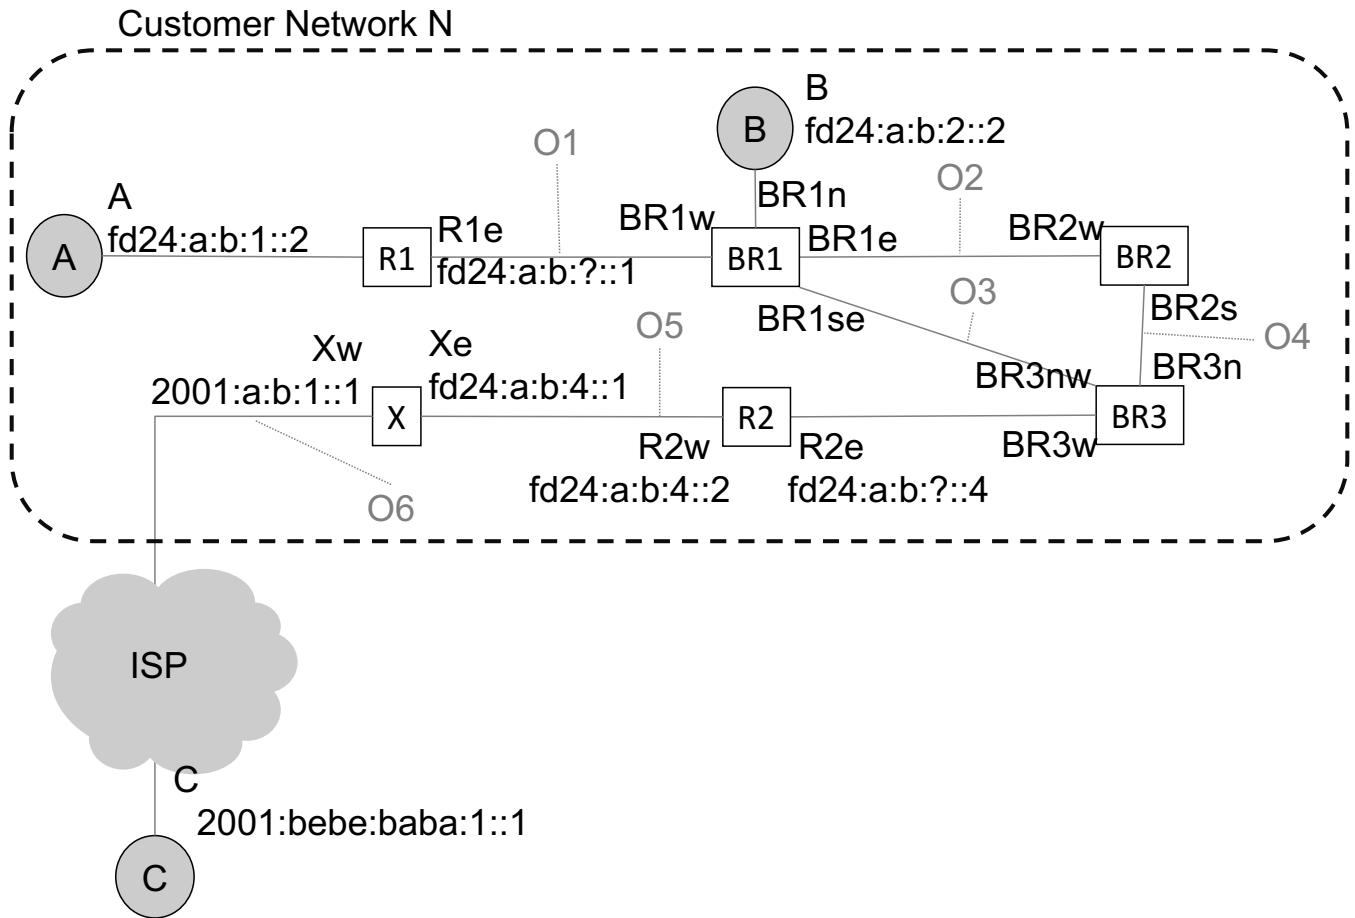
(c) Assume that Lisa's home network contains several routers and Lisa is more than one hop away from the media server. Is there any change to the answer of the previous question ?

3. The web servers that serve the url `https://sovkom.coolstuff` are distributed across several sites. The managers of sovkom would like to spread the load between the different sites. How can that be done without modifying anything on user devices ?

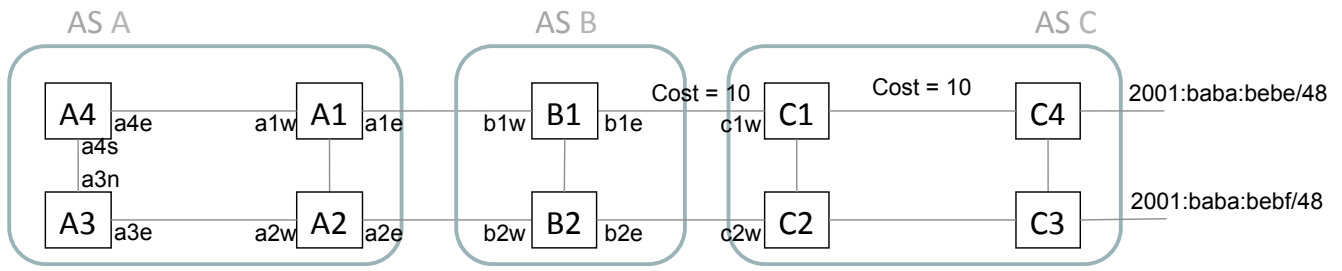
4. Joe has a dual stack IPv4 / IPv6 machine and has both IPv4 and IPv6 access. Joe connects to `https://lca.epfl.ch` and `https://infoscience.epfl.ch`. The web server `lca.epfl.ch` uses IPv6 only whereas the web server `infoscience.epfl.ch` uses IPv4 only. By which mechanism does Joe's machine know whether IP4 or IPv6 should be used ?

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.



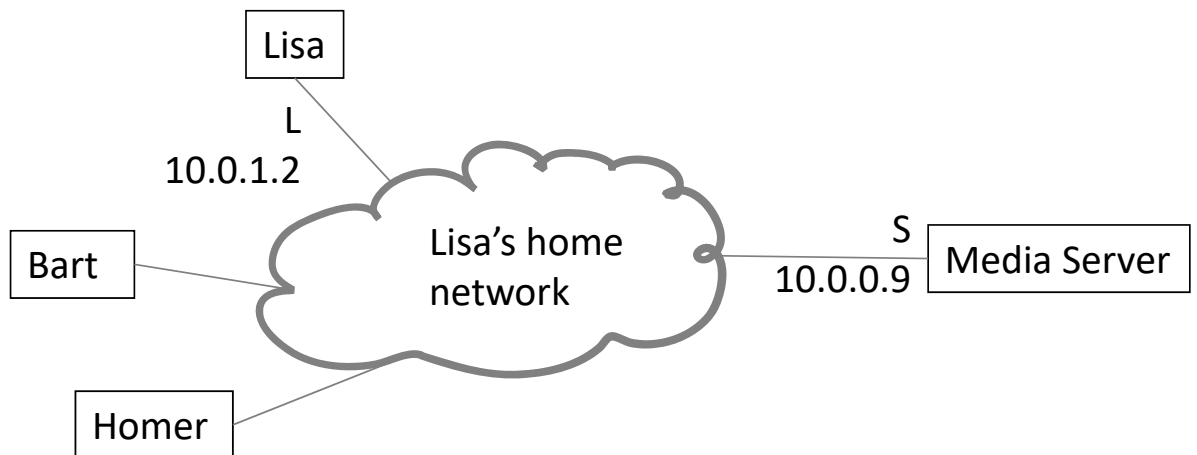
Problem 1



Problem 2



Problem 3



Problem 4, Question 2. L and S are MAC addresses.