
EXAM
TCP/IP NETWORKING

Duration: 4 hours + including technical time for uploading solution.

Jean-Yves Le Boudec

January 2021

INSTRUCTIONS

1. All four problems have the same weight. They are independent and can be done in any order.
2. Briefly justify your answer. For grading, the justification is as important as the solution itself.
3. You may need to make additional assumptions in order to solve some of the questions. If this happens, please describe such assumptions explicitly.

PROBLEM 1

Consider the network in Figure 1. H1 to H6 are hosts. R1 and R2 are routers. B1 to B6 are bridges. O1 to O12 are observation points. All machines are dual-stack.

N3 and N1 operate as NATs for both IPv4 and IPv6: the internal port is B3 [resp. B2] and the external port is the port towards B4 [resp. B6].

N2 operates as a NAT for IPv4 (with internal ports towards observation point O1) 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. There is no other system or interface than shown on the figure.

H1B3, B3H1, B3N3, N3B3, N3B4, B4N3, B4B5, B4B1, B4B6, B1N2, B1B6, B1B4, B1B5, N2B1, N2R1, R1N2, R1H6, H6R1 are MAC addresses.

n4, ..., k2 are ipv4 addresses and are noted in lowercase. N4, ..., K2 are ipv6 addresses and are noted in capital letter.

Some IPv4 and IPv6 addresses are known or partially known, as follow:

1. The IPv4 address t4 is 10.213.141.212.
2. The IPv4 address w1 is 192.168.47.46.
3. The IPv4 address z6 is 8.2.7.11.
4. The IPv4 address s2 is 192.168.47.253.
5. The IPv4 address h1 is 10.213.141.163 .
6. The IPv4 address n4 is 192.168.47.134.
7. The IPv4 address m1 is 8.2.3.7.

1. The IPv6 address H1 is 3390:b:4:5:q::1.
2. The IPv6 address X1 is 3390:b:4:5:e9:74b2::2.
3. The IPv6 address T4 is 3390:b:4:5:e9::1.
4. The IPv6 address V3 is fe80::18d6.
5. The IPv6 address N4 is fe80::ea71.
6. The IPv6 address W1 is fe80::a6c8.

1. Give possible values for the IPv4 address x1 at N3 and for the network mask at this interface.
2. Give possible values for the IPv4 address v3 of H3 and for its network mask.
3. H1 sends one UDP packet to H5, H3 sends one UDP packet to H6 and H2 sends one UDP packet to H4. We look for these UDP packets at observation points O5, O4, O10, O2 and O6, i.e. at the lines between bridges. At which of these observation points will we observe the packets?
4. H1 and H5 send several small UDP packets each to H6, at about the same time, over IPv4, to udp port 510. By chance, the port source numbers are both equal to 16566. We observe these packets at observation points O7, O8 [resp. O12, O9], O1 and O11. What source and destination IP addresses and port numbers do we see ?
5. (continued) At O7, O8 and O1, what source and destination MAC addresses do we see for the packet sent by H1 to H6 ?
6. H3 sends one IP packet to H6 and one to H2, and sets the value of TTL at the source to 48. What is the value of the TTL at destination for each of these packets?
7. The IPv6 network prefix lengths at N1B6 and N2B1 are both equal to 72. The IPv6 address of interface N1B6 is 3390:b:4:5:q::1 where q is a block of 16 bits. Give all the possible values of q (use hexadecimal notation).

8. H1 does a traceroute to fe80::18d6. We look for the traceroute packets at observation points O7, O8, O5, O4, O10, O2, O6, O9 and O12. What are the source and destination IP addresses observed at these points ?

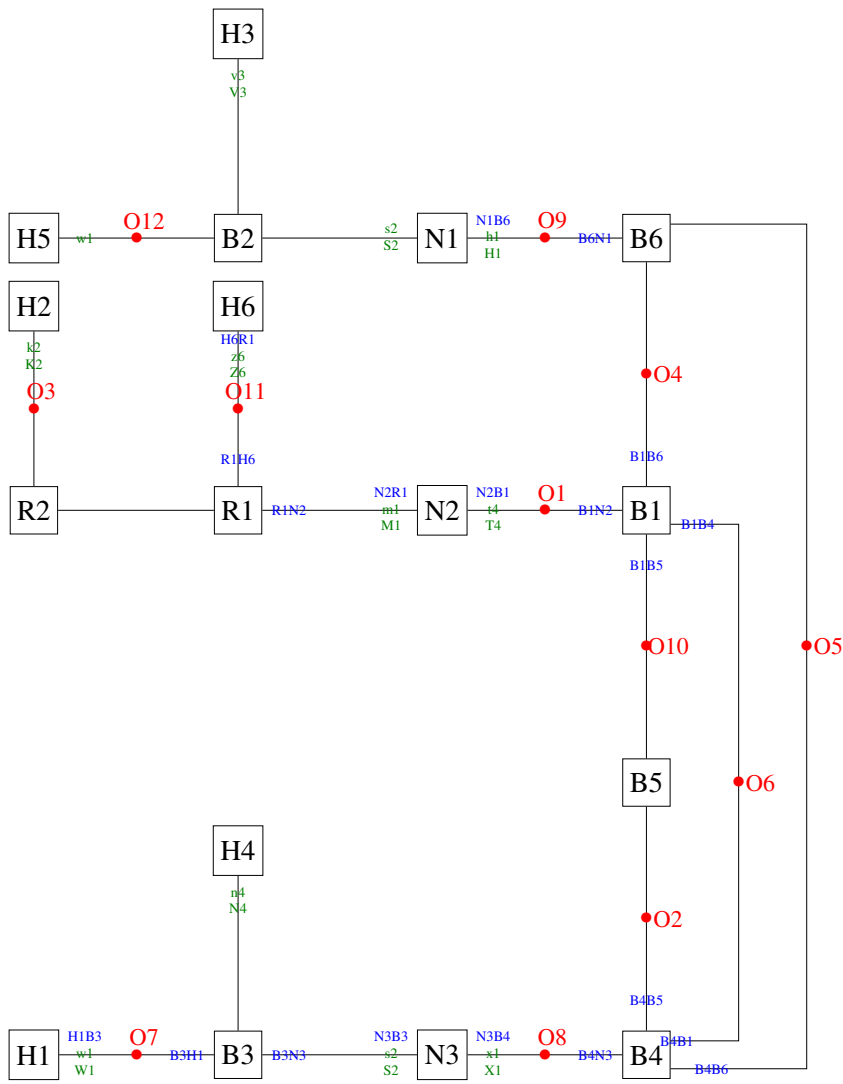


Figure 1: Network in Problem 1

PROBLEM 2

Consider the network in Figure 2. There are three ASs, S , V and R . Every node whose name begins with S belongs to AS S (same with V and R). Each AS uses OSPF with a single area and equal cost multipath. The OSPF cost of every link and every attached network is 1, unless otherwise specified.

The IP addresses at inter-router links are denoted as follows: at $S3$, the IP address of the interface towards $S1$ is $s3s1$, and similarly everywhere else (so for example the IP address at $R3$ of the interface towards $V1$ is $r3v1$, etc.)

The subnetwork $2614:9fca:1f7c::/48$ is attached to node $S5$ and belongs to AS S .

All routers run BGP, unless otherwise specified. The BGP decision processes all 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, unless otherwise specified:

- All routers redistribute internal OSPF routes into BGP.
- Every BGP router injects its BGP routes into its forwarding table.
- The BGP routers do not redistribute external BGP routes into OSPF.
- No optional BGP attribute (such as MED, LOCAL-PREF etc.) is used in any BGP message.
- No aggregation is performed.
- There is no route announcement to $2614:9fca:1f7c::/48$ or $2614:9fca:1f7c:cd56::/64$ other than the ones described in the text and the ones that derive from that.
- The BGP policies accept all routes that are acceptable by the standard operation of BGP.

1. When OSPF and BGP have converged, what are all the BGP routes to destination $2614:9fca:1f7c::/48$ stored by $V4$ in its RIB-In (give the source BGP router, the AS-PATH and NEXT-HOP attribute of the routes)? What are the routes to $2614:9fca:1f7c::/48$ advertised by $V4$ to other BGP routers?
2. Same question with $V5$.
3. Router $V3$ has a packet to forward to destination address $2614:9fca:1f7c:cd56:ff1::9c53$. Which path will this packet take inside AS V ? (Give the path as a sequence of routers, including the first router out of AS V).
4. Due to a configuration error, $R4$ announces to $V5$ the bogus route

$$2614:9fca:1f7c:cd56::/64 \text{ AS-PATH} = R \text{ NEXT-HOP} = r4v5$$

and this route is accepted by all BGP routers in AS V . After BGP and OSPF have stabilized again, router $V1$ has two packets to forward, one with destination address $2614:9fca:1f7c:cd57:b2eb::35c6$, and one with destination address $2614:9fca:1f7c:cd56:ff1::9c53$. Which paths will these packets follow from $V1$ to the exit of AS V ? (Give the path as a sequence of routers, including the first router out of AS V).

5. $R4$ is patched and the configuration error that led to the bogus announcement is removed. Furthermore, we now modify the operation of AS V :
 - router $V1$ does not run BGP (but continues to run OSPF);
 - all other routers in AS V redistribute external BGP routes into OSPF with distance at the point of redistribution equal to 21.




Figure 2: Network in Problem 2

- The rest is unchanged. After BGP and OSPF have stabilized: Router $V1$ has a packet to forward to destination address $2614 : 9fca : 1f7c : cd56 : ff1 :: 9c53$. Which path will this packet take inside AS V ? (Give the path as a sequence of routers, including the first router out of AS V). Compare to the paths that would have been followed if the configuration were the same as in Questions 1, 2 and 3.
6. Router $V3$ now has packets to send to a large number of destination addresses such as $2614 : 9fca : 1f7c : cd56 : n : k : z : h$, where n, k, z and h each represents a sequence of four hex digits. Which paths will these packets take? (Give the path as a sequence of routers, including the first router out of AS V). Compare to the paths that would have been followed if the configuration were the same as in Questions 1, 2 and 3.

PROBLEM 3

Consider the network in Figure 3. $H0$, $H6$ and $H3$ are downloading very large contents from $H4$, simultaneously. The intermediate boxes are routers, unless otherwise specified. The line rates are shown in the figure with $m = 200$, $u = 1064$, and $K = 40432$, in Mb/s. The links are full duplex Ethernets. Let p [resp. w, s] be the rate at which $H0$ [resp. $H6, H3$] downloads contents from $H4$. All other flows (including ACKs) can be neglected.




Figure 3: Network in Problem 3

1. Assume the rates are allocated by some central bandwidth manager. Which of the following rate allocations are Pareto-efficient?
 - (a) $p = 399, w = 133, s = 100$.
 - (b) $p = 798, w = 266, s = 200$.
 - (c) $p = 532, w = 532, s = 200$.
 - (d) $p = 798, w = 266, s = 100$.
2. Assume the rates are allocated by some central bandwidth manager according to max-min fairness. What are all the possible rate allocations?
3. Assume the rates are allocated by some central bandwidth manager according to proportional fairness. What are all the possible rate allocations?
4. $H0$, $H6$ and $H3$ are now using TCP-Reno to perform their transfer from $H4$. The RTTs are: $\text{RTT}_{H0} = 113$, $\text{RTT}_{H6} = 15$ and $\text{RTT}_{H3} = 5$, in ms. All have same MSS equal to 1344, in bits. ECN is enabled in all hosts and all routers. What is the rate obtained by every source ?
5. $R3$ is now working as a web tunnel (web proxy) for $H0$, also using TCP Reno with ECN and with same MSS as all hosts. The RTT between $R3$ and $H4$ is the same as between $H6$ and $H4$. In this question only, the link between $H0$ and $R3$ is wireless with a loss rate equal to 7.32 %, and the RTT between $H0$ and $R3$ is equal to 4.26 ms. (there is no congestion loss on this link, only wireless transmission losses). What is the rate achieved by $H0$?

PROBLEM 4

1. The Agostinho Neto University (UAN) offers a repository of media files (lyrics, recorded audios and videos) to the public. The repository is hosted by a number of media servers. The UAN campus has access to IPv6 and to IPv4; the media servers use only IPv4 addresses. The media servers use old hardware and software, partly written in Cobol; migrating them to IPv6 would require expensive upgrades and is not an option. There are n media servers, and each of them has a DNS name of its own. The renderer application on the user device accesses the media files by using URLs, such as `https://musica-hoji-ya-henda.uan.ao/2340993.flac` where `musica-hoji-ya-henda.uan.ao` is the DNS name of the media server that hosts the file of interest.
You joined the UAN Network Team as engineer and you are asked to design a solution for allowing external users that have only IPv6 connectivity to access the repository. Your solution should not modify the UAN media servers; it should also not modify the IPv6 configuration of external users. Please sketch at least two different solutions (one paragraph each).
2. The network of the National Security Agency of Borduria (NSAB) uses addresses in the block 133.133.0.0/16 and is organized with subnet prefixes of length 24 bits. For example, the subnet 133.133.99.0/24 is used by the management team, the subnet 133.133.1.0/24 is used for the accounting department and the subnet 133.133.42.0/24 contains the information servers used by all NSAB departments. NSAB decided to deploy the SuperSecureIx operating system in all of their computers. This is a very secure operating system designed by the NSAB engineers. With SuperSecureIX, the IP address of a computer is assigned to a user and cannot be modified (DHCP is not used). Unfortunately, the NSAB engineers did not take the TCP/IP Networking course and, instead, read random things found on forums. As a result, they believe that, in class B networks, the network mask is 255.255.0.0 (Class B is an obsolete classification of IP address blocks) and since the NSAB addresses are in class B, they hardcoded this subnet mask in SuperSecureIx. As this is a very secure operating system, no one dares modify it.
Immediately after deploying SuperSecureIX, many network problems appeared. Can you propose a solution ? Your solution should not involve any change to end-user devices but can involve the network infrastructure (which, fortunately, is not based on SuperSecureIX and can be modified).