SCIPER: ____________________

First name: ________________________ Family name: ________________________

EXAM
TCP/IP NETWORKING
Duration: 3 hours

January 2025

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. Briefly justify your answer. For grading, the justification is as important as the solution itself.
4. If you find that you need to make additional assumptions in order to solve some of the questions, please describe such assumptions explicitly.
5. Figures are on a separate sheet, for your convenience.
6. You can bring and use 4x A4 sheets = 8x AA4 pages of hand-written or type-written notes or the exam booklet that we offer in Moodle (in printed form). You can also use your pocket calculator (i.e. a simple calculator without extra storage or graph plotter).

QUIZ (15PTS)

1. **(1 pt)** A host uses NDP to learn the MAC address of IP host 2001:face:b00c::baba:b0b0:1234. What will be the destination IP address in the IP header of the packet containing this NDP message?

   (a) 2001:face:b00c::baba:b0b0:1234

   (b) ff02::1:ff00:b0b0:1234

   (c) ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff

   (d) ff02::1:ffb0:1234

2. **(1 pt)** An IPv6 host uses SLAAC in order to configure its interface. Which of the following is *always* true at the end of the process?

   (a) The host will be configured with a single link-local, cryptographically generated IPv6 address.

   (b) The host will know the IP address of the DNS server.

   (c) If a router is present in the same subnet and is configured to support SLAAC, the host will be configured with one or several global unicast address.

   (d) The gateway router adds a specific entry with the host's global IP address to its forwarding table.

3. **(1 pt)** A LAN has two WiFi base stations BS1 and BS2, interconnected via a switch and Ethernet cables. Wifi devices $\{A_1, \ldots, A_n\}$ are associated with BS1, and $\{B_1, \ldots, B_m\}$ are associated with BS2. S is a server connected directly to the switch using a separate Ethernet cable. By which mechanism does the switch know where to send packets from S destined to A-type and B-type devices?

   (a) The spanning tree protocol learns all addresses without intervention of the WiFi association data.

   (b) The switch runs a routing protocol to compute shortest paths to all destinations.

   (c) The switch observes all MAC source addresses and associates MAC addresses to ports, while it broadcasts frames to both Ethernet ports if it does not know where the destination is.

   (d) The base stations inform the switch of the presence of A-type and B-type devices (which they know from their association data) using a routing protocol.

4. **(1 pt)** CSMA/CD (collision detection) is a randomized medium-access control protocol that:

   (a) is expected to improve the overall throughput compared to deterministic protocols whenever few users share the same cable and have lots of data to transmit, and always at the same time.

   (b) is expected to improve the overall throughput compared to deterministic protocols whenever many users share the same cable and the probability of transmitting at the same time is very low.

   (c) is used by WiFi, when users are close to each other and there is no "hidden terminal".

   (d) detects collisions via the absence of acknowledgements.

5. **(1 pt)** An application at host A uses a socket with a send-buffer size of 10KB. The application sends one block of 5KB to some host B by doing a single "send" operation on the socket. An application at B attempts to receive the data and does a single "receive" operation on a socket. Is it possible that the data that the application at B gets contains only 2KB out of the 5KB of the message sent by A?

   (a) Yes with UDP and yes with TCP.

   (b) No with UDP and no with TCP.

   (c) Yes with UDP, no with TCP.

   (d) Yes with TCP, no with UDP.

6. **(1 pt)** A TCP sender A detects (via fast retransmit) that the segments 2001:3000 and 3001:4000 were lost. The MSS is 1000 bytes and the usable part of the window (i.e. the maximum amount of bytes that A can send) is 2000 bytes. When A decides to retransmit these two segments:

   (a) A may re-transmit 1 segment with sequence number 3001:4000 and must wait for the acknowledgment of this segment before re-transmitting 2001:3000.

   (b) A may re-transmit 1 segment with sequence number 2001:3000 and must wait for the acknowledgment of this segment before re-transmitting 3001:4000.

   (c) A must re-transmit 2 segments, with sequence numbers 2001:3000 and 3001:4000.

   (d) A may re-transmit 2 segments, with sequence numbers 2001:3000 and 3001:4000 or re-transmit one segment with sequence numbers 2001:4000.

7. **(1 pt)** Say which is true about "Loose Source Routing":

   (a) When a router forwards a packet, it bases its decision on both the source and destination IP addresses.

   (b) When a router forwards a packet, it bases its decision on the source and destination IP addresses as well as the source and destination port numbers.

   (c) The source writes in the packet header the complete sequence of all intermediate hops from source to destination.

   (d) The source writes in the packet header a sequence of some intermediate nodes that the packet must visit on its way to the final destination.

8. **(1 pt)** In a WiFi network with a single channel and a single base station, how many data frames can be transmitted concurrently?

   (a) At most one in total.

   (b) At most one from a mobile device to base station and one from base station to a mobile device.

   (c) At most one per mobile device to base station and one from base station.

   (d) At most one per mobile device to base station and one per mobile device from base station.

9. **(1 pt)** Alice's host is connected to a switch and uses MACSEC, with encryption and authentication. Eve uses a repeater on the cables between the Alice's host and the switch S, and can therefore launch a woman-in-the-middle attack. Say what Eve will be able to do without being detected:

   (a) She can correctly read the IP packets sent by Alice.

   (b) She can modify only the source IP address of the packets sent by the Alice.

   (c) She can correctly read Alice's MAC address.

   (d) She can modify the source MAC address of the frames sent by Alice.

10. **(1 pt)** Hosts A, B are in the same VLAN but are connected to different layer-2 switches, S_1 and S_2, respectively, which are connected via a trunk link. C and D are also connected to S_1 and S_2, but they are in a different VLAN. A sends an Ethernet frame with MAC destination address ff:ff:ff:ff:ff:ff. Which host(s) will receive this frame?

    (a) A and B.

    (b) B.

    (c) C and D.

    (d) D.

11. **(1 pt)** Host H belongs to a big enterprise network. H receives an IP packet from a system A, with source IP address A = fe80::f8b8:c271:904a:80c5 and Hop Limit = 32. What can H infer?

    (a) A is in the same private network as H, but not necessarily in the same subnet as an enterprise network may have several routers internally.
    (b) A is not inside the enterprise network.
    (c) The Hop Limit set by A when it sent this packet was 32.
    (d) A has used stateful DHCP to configure its IP address.

12. **(1 pt)** If a server uses SYN cookies, then

    (a) the server has to keep per connection state information after receiving a SYN packet.
    (b) the server does not keep any connection-state information after receiving a SYN packet; instead, it encodes the state in the SEQ Number field of the SYNACK, using a cryptographic function.
    (c) the server does not keep any connection-state information after receiving a SYN packet; instead, it encodes the state in the ACK Number field of the SYNACK, using a cryptographic function.
    (d) the server is vulnerable to DoS attackers that send multiple bogus SYN packets to connect to it.

13. **(1 pt)** A domain network is composed of border and backbone routers. Each border router connects to external domains and at least one backbone router. The backbone routers are not directly connected to any external networks and form a full mesh with each other. All routers run OSPF, and all border routers run BGP. OSPF is redistributed into BGP, but BGP is not redistributed into OSPF. All BGP routes are injected into the forwarding table at every router that runs BGP. In order to make sure that all transit traffic is properly routed through this domain, the network operator envisions activating the following functions:
    i. Backbone routers also run BGP.
    ii. All routers support recursive table lookup in their forwarding table.

    Which of the following is true?

    (a) The configuration does not work even if both functions are activated, because some function is missing for transiting packets through the backbone.
    (b) The configuration works if (i) alone is activated; (ii) may be activated but is not necessary.
    (c) The configuration works if (ii) alone is activated; (i) may be activated, but is not necessary.
    (d) The configuration works if both are activated.

14. **(2 pts)** You are the operator of a private sensor network with 500 sensors, which are cheap devices with limited computational resources, and two monitors A and B (A is the main monitor and B is backup), which should receive a measurement of 2 bytes from each sensor of every 1 sec. The topology is as follows: the sensors are grouped in 5 non-overlapping groups; each group forms a subnet with its own gateway router; the 5 gateway routers are then connected to a backbone router, where the 2 monitors are also connected to. Which is the best approach to follow?

    (a) Use TCP between each sensor and each monitor and a good congestion control algorithm to avoid congestion in the network.
    (b) Use TCP and enable any-source multicasting via BIER, because the number of sensors is large and you want to reduce the state stored in the routers.
    (c) Use UDP and enable any-source multicasting via BIER, because the number of sensors is large, hence it is better to reduce the state that PIM would need to store in the routers because of them.
    (d) Use UDP and enable any-source multicasting via PIM, because the number of monitors is small, hence the state stored in the routers because of them is small and maintaining a multicast flow overlay for this would be an overkill.

PROBLEM 1 (20PTS)

Consider the network for Problem 1 in the figure sheet. H0, H1, H2, H3, H4, H5, and H6 are hosts. R0, R1, R2, and R3 are routers. S0, S1, S2, and S4 are switches. N1 is an IPv4 NAT. O1, O2, and O3 are observation points. All machines are dual-stack. All necessary IPv4 and IPv6 addresses are shown in the figure, as well as necessary MAC addresses (denoted with e.g., !!!!!!!)

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 hosts, switches and routers have been running for some time, their different protocols have converged and the forwarding tables of all routers and switches are in their final state.** There is no other system or interface than those shown on the figure.

Question 1 (!!!!! pts):

1. **(!!!! pts)** Write the two IPv6 addresses of H0 in uncompressed format.

2. **(!!!!! pts)** Select whether the following IPv4 network masks lengths for hosts H3, H4, H5, and H6 are valid.

Network Mask Length					
H3	H4	H5	H6	Valid	Invalid
20	22	24	20		
16	24	24	16		
24	26	26	24		

3. **(!!!!! pts)** Assume that the IPv6 prefix lengths for hosts H1 to H6 are 56 bits. Which values of X, Y, and K within the IPv6 addresses of H1, H2, and H4 are valid?

Proposed Values of X, Y, and K	Valid	Invalid
$X = 1, Y = 1, K = 100$		
$X = 2, Y = 2, K = 100$		
$X = 1, Y = 11, K = aaa$		
$X = 1, Y = 10, K = aaa$		

Question 2 (!!!! pts): Which of the three addresses of H6 can H3 use to reach it. Similarly, which of H4's three addresses can H5 use to reach H4?

Question 3 (!!!! pts):

1. **(!!!!! pts)** H0 uploads a large file to www.largefiles.com that is hosted in H1 using HTTP and IPv4. H5 and H6 also upload similar files using the same protocols. We observe the packet headers at the MAC, IP and transport layer at points O1 and O3 for H0, and at O1 and O2 for H5, H6.

   Give possible values in the tables below. You can use given variable names when the specific value is not known (e.g., ”X” in 2002:1:X:1::9 or ”Nw”). Assume that each client application is using one local port between 50000 and 50100 (any).

Direction from H0 to H1						
	MAC addresses		IPv4 addresses		port numbers	
At	src	dst	src	dst	src	dst
O3						
O1						

Direction from H5 to H1						
	MAC addresses		IPv4 addresses		port numbers	
At	src	dst	src	dst	src	dst
O2						
O1						

Direction from H6 to H1						
	MAC addresses		IPv4 addresses		port numbers	
At	src	dst	src	dst	src	dst
O2						
O1						

2. **(!!!!! pts)** The organization operating the NAT box and network to its right-hand side has been running an IPv4 file-hosting server, similar to H1, on H4 for internal use. After setting up the server, does the organization need to take any further steps so that H0 can access the server?

   If not, what are the packets observed at O2 and O3 when H0 uploads a file to H4 using TLS (fill-in the table below)?

   If yes, explain why and describe exactly what you would do without changing the topology.

Direction from H0 to H4						
	MAC addresses		IPv4 addresses		port numbers	
At	src	dst	src	dst	src	dst
O3						
O2						

Question 4 (!!!! pts): Suppose that suddenly H6 reboots and its caches become empty. Suppose that H3 is compromised by an adversary, while H6 reboots. Briefly describe how the adversary could prevent H6 from sending traffic to H2 and what could the network administrator do to avoid such attacks.

SCIPER: ____________________

First name: ________________________ Family name: ________________________

PROBLEM 2 (20PTS)

In the following questions, we assume that the BGP decision process uses the following criteria in decreasing order of priority.

1. Highest LOCAL-PREF.
2. Shortest AS-PATH.
3. E-BGP is preferred over I-BGP.
4. Shortest path to NEXT-HOP, according to IGP.
5. Lowest BGP identifier of sender of route is preferred; the 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, **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 all ASs is that all available routes are accepted and propagated to neighbouring ASs, as long as the rules of BGP allow it.
- Every router redistributes internal OSPF destinations into BGP.
- Every router performs recursive forwarding-table lookup.
- No confederation or route reflector is used.
- Besides what is shown on each figure, there are no other stub networks.

In the following questions, you will need to fill in the local RIBs (route information bases) of various routers. To help you, we have added all destination prefixes for you to fill in the rest of the information:

- If you think a prefix is not in the local database, cross the line.
- If you think a prefix has several entries, write them one on top of the other as shown in the example below. If you run out of space, you can add extra lines at the end of the table.

At Example :			
Destination Network	From BGP Peer	AS-PATH	Best route ?
X.X.X.X/Y	Z1 Z2	ASN ASM ASM	yes no
Justification:			

- If you think a prefix is missing, you can add it at the end of the table.

Question 1 (X pts): Consider the network for Problem 2 in the figure sheet. There are six ASs (AS1 to AS6), with border routers (A1, A2, A3, B1, C1, etc). In each domain, there is an **I-BGP mesh** that is not shown in the figure. **Consider the situation after BGP has converged**. Justify each answer.

1. **(2pts)** List all the routes in the local database of Router $D1$.

At $D1$:			
Destination Network	From BGP Peer	AS-PATH	Best route ?
10.20.0.0/14 10.24.0.0/13 10.32.0.0/14 10.36.0.0/14 10.40.0.0/13 10.56.0.0/13 10.64.0.0/13			
Justification:			

2. **(2pts)** List all the routes in the local database of Router $B2$.

At B_2 :			
Destination Network	From BGP Peer	AS-PATH	Best route ?
10.20.0.0/14			
10.24.0.0/13			
10.32.0.0/14			
10.36.0.0/14			
10.40.0.0/13			
10.56.0.0/13			
10.64.0.0/13			
Justification:			

3. Consider the Router $B2$, what is the Next-Hop for the following routes? Explain.
   (a) 10.20.0.0/14
   (b) 10.56.0.0/13

4. **(2pts)** Suppose that there exists a malicious AS, which is customer of AS3 (i.e. a stub domain served by AS3), starts sending bogus announcements for prefixes 10.36.0.0/15 and 10.38.0.0/15.
   (a) How do these announcements change the local RIB of Router $B2$?
   (b) What is the AS-path followed by packets going from 10.20.0.1 to 10.35.1.1?
   (c) What is the AS-path followed by packets going from 10.20.0.1 to 10.39.1.1?

Question 2 (X pts): Consider again the same network as before (Problem 2 in the figure sheet without the malicious AS). **For the purposes of this question, aggregation is used by BGP and whenever possible ASs also aggregate internally originated prefixes with prefixes that are learnt externally.** Consider the situation after BGP has converged. **Justify each answer**.

1. **(2pts)** List all the routes in the local database of Router $D1$.

At $D1$:			
Destination Network	From BGP Peer	AS-PATH	Best route ?
10.20.0.0/14			
10.24.0.0/13			
10.32.0.0/14			
10.36.0.0/14			
10.40.0.0/13			
10.56.0.0/13			
10.64.0.0/13			
Justification:			

2. **(3pts)** List all the routes in the local database of Router $B2$.

At $B2$:			
Destination Network	From BGP Peer	AS-PATH	Best route ?
10.20.0.0/14			
10.24.0.0/13			
10.32.0.0/14			
10.36.0.0/14			
10.40.0.0/13			
10.56.0.0/13			
10.64.0.0/13			
Justification:			

3. **(2pts)** As previously, suppose now that there exists a malicious AS, which is customer of AS3 (i.e. a stub domain served by AS3), starts sending bogus announcements for prefixes 10.36.0.0/15 and 10.38.0.0/15.
   (a) How do these announcements change the local RIB of Router $B2$?
   (b) What is the AS-path followed by packets going from 10.20.0.1 to 10.35.1.1?
   (c) What is the AS-path followed by packets going from 10.20.0.1 to 10.39.1.1?

SCIPER: ____________________

First name: ________________________ Family name: ________________________

PROBLEM 3 (20PTS)

Consider the network for Problem 3 on the figure sheet.

- The figure illustrates a data center network in which two clusters (one for file storage services and one for web services) are connected through a link.
- T0, T1, T2, T3, A0, and A1 are routers, connected via point-to-point links.
- f_0, f_1, f_2, f_3, and f_4 are unidirectional flows (as indicated by the arrows). There is no other system and no other flow than those shown on the figure.
- The capacity of each link is shown in the figure and they are all in Gb/s. The links are full duplex with the same rate in both directions. There is no other capacity constraint.
- We neglect the impact of the acknowledgement flows in the reverse direction.
- We also neglect all overheads and assume that the link capacities can be fully utilized at bottlenecks.

Question 1 (4 pts): Assume the rates are allocated by some central bandwidth manager according to max-min fairness. What are all the possible rate allocations? Justify.

Question 2 (5 pts): Assume the rates are allocated by some central bandwidth manager, which of the following rate allocations are Pareto-efficient? Justify.

1. $f_0 = 1, f_1 = 7, f_2 = 2, f_3 = 0.5, f_4 = 0.5$
2. $f_0 = 1.5, f_1 = 5, f_2 = 5, f_3 = 0.7, f_4 = 0.3$
3. $f_0 = 1, f_1 = 10, f_2 = 0, f_3 = 0.3, f_4 = 0.7$

Question 3 (4 pts): Assume the rates are allocated by some central bandwidth manager as the following: $f_0 = 1, f_1 = 3, f_2 = 7, f_3 = 0.3, f_4 = 0.7$ Gb/s. Proof that this rate allocation is not propotionally fair.

Question 4 (4 pts): In this question, **flow F0 is using UDP** and sends data at a constant rate of 0.75 Gb/s. Assume that all other flows use **TCP Reno** and are long-running TCP flows. Also assume that all switches use RED queuing with supporting ECN, and the marking probability is $q \ll 0.01$ for all flows. The round trip times (RTTs) for flows is the following: $RTT_{f_1} = 250ms, RTT_{f_2} = 100ms, RTTf_3 = 450ms, RTT_{f_4} = 150ms$; these RTTs include all processing times. All flows use the same MSS and the offered window is very large.

What are the rates attained by each flow in the long run?

Question 5 (4 pts): In this question we have the same setting as in question 4 but we enable **TCP Cubic instead of TCP Reno**. What will happen to the rates of flows F1 and F2? Justify.

Question 6 (4 pts): Consider the same setting as in Question 4. But all TCP flows use **DCTCP** as their congestion control algorithm, and they use the same apporoach in TCP Cubic for the additive increase phase. An external flow f_5 is added to the network by entering the network at A0 and communicating with S3. This external flow uses **TCP Cubic** as its congestion control algorithm and it has the same RTT as f_4.

Is this setting fair? why? if no, how can you obtain fairness. If yes, under which circumstances would it become unfair?

SCIPER: ____________________

First name: ________________________ Family name: ________________________

PROBLEM 4 (25PTS)

Question 1 (10 points):

Consider the 3-area OSPF network of the figure for Problem 4 in the figure sheet. All lines represent physical links and the associated numbers/weights are the OSPF costs. All boxes are routers; Area 0 is backbone, whereas Areas 1 and 2 are not. Routers $A0, A1$ (resp. $A3, A4$) are border routers between Areas 0 and 1 (resp. Areas 0 and 2). All depicted subnets (of the form 10.0.X.0/24) are stub networks. All network interfaces have been correctly configured and OSPF has converged.

Suppose that as the network administrator you want to allow for a video-streaming server with IP address 10.0.16.1 to serve many customers in the stub networks of Area 2. So, you decide to enable IP multicasting via BIER at all routers and run the mutlicast flow overlay function on your SDN controller (because your network also supports SDN).

1. What would be the BIER index forwarding table at router $A2$?

Destination BFER	Forwarding Bit Mask	Next Hop

2. After some time, you observe that no customer receives the streamed video, and you realize that out of all routers $X0$–$X2$, only $X2$ has enough resources to actually serve as a BIER router. Can you provide a fix to this problem without upgrading your routers? Briefly explain what you plan to do and make your assumptions about interfaces/addresses clear.

Question 2 (8 points):

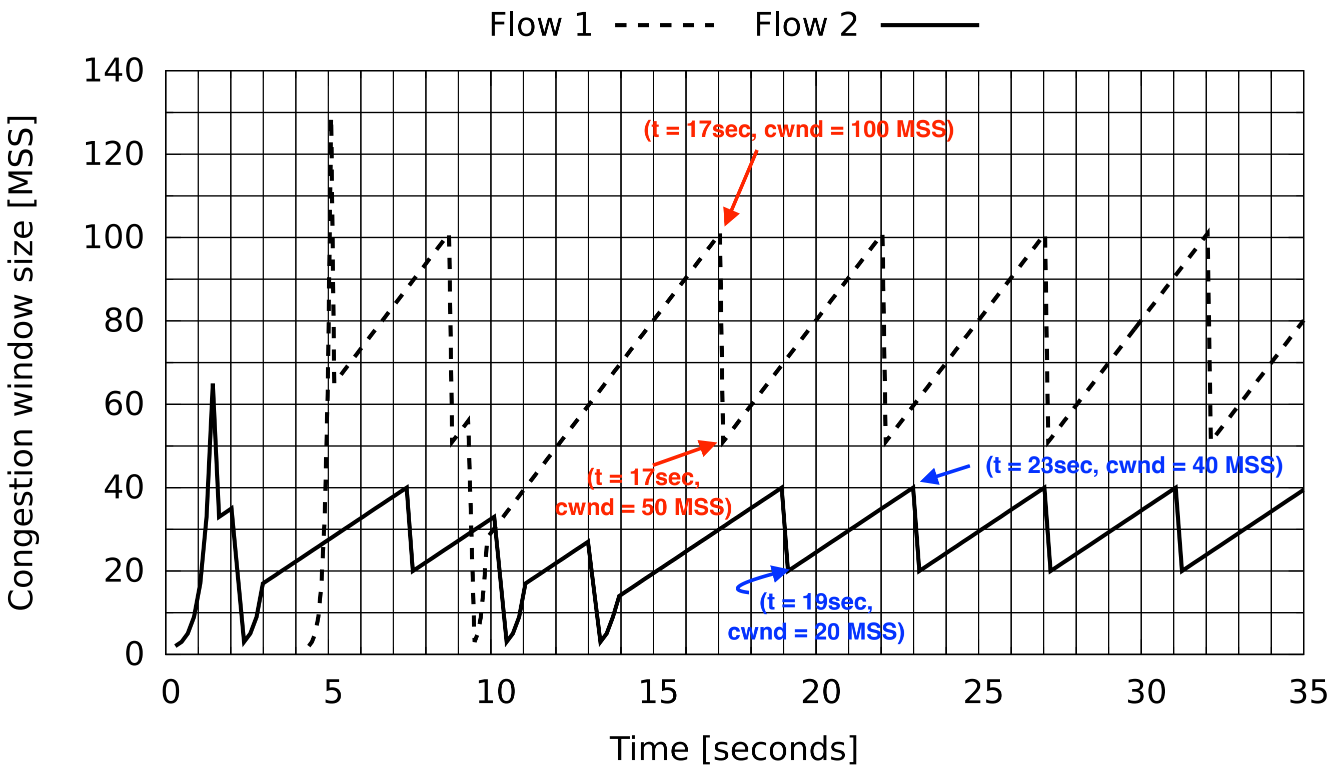

Figure 1: Congestion window of the sender as a function of time.

Consider a server sending a large file to Client 1 and Client 2 using TCP Reno. We record the sender's TCP congestion window size for both Flow 1 and Flow 2 and plot it in Figure 1 of Problem 4.
For simplicity, in the following questions, assume that:

- there is **no** packet reordering and every time packet loss is detected, **only one** packet has been dropped;
- the duration of Reno's fast recovery phase is negligible (i.e. the **multiplicative factor is** 0.5);
- the additive term of the additive increase phase is **exactly equal to** (instead of a bit less than) 1.

1. Identify all timeout events (for both flows) and circle them on the graph. Briefly mention two possible reasons why timeouts can happen in reality.

2. Further to the simplifying assumptions above, assume that the RTT of both flows, Flow 1 and Flow 2, are constant. What are their values? Briefly justify your answer.

3. Now, suppose that the congestion window of Flow 1 after T = 17secs continues to follow the same pattern (i.e. its behavior stabilizes). If the MSS of Flow 1 is 1250 bytes (i.e., 10000 bits) and the offered window from Client 1 is very large, what is the long-term loss rate (i.e. fraction of segments being lost) that Flow 1 experiences?
Hint: Compute first the average throughput from the information in the figure. Or if this is too hard, focus on one cycle of the sawtooth curve and compute the loss rate directly.

4. Now, suppose that the congestion window of Flow 2 also stabilizes after T = 17secs, Flow 2 also uses MSS = 1250 bytes, and Client 2's offered window is very large. Do both flows achieve the same long-term throughput after time T? If yes, prove your argument. If no, explain which flow will performs better and why this happens.

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.

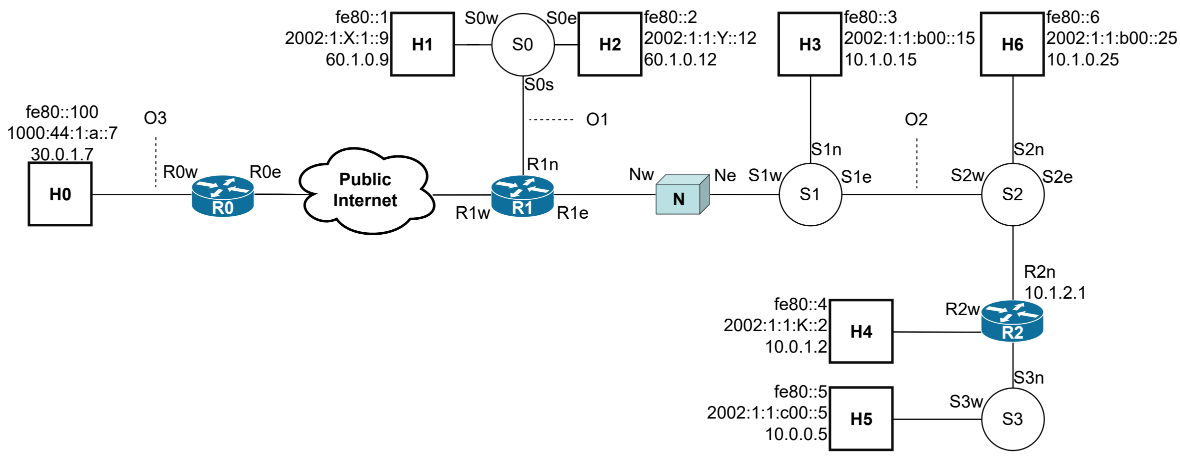

Problem 1.

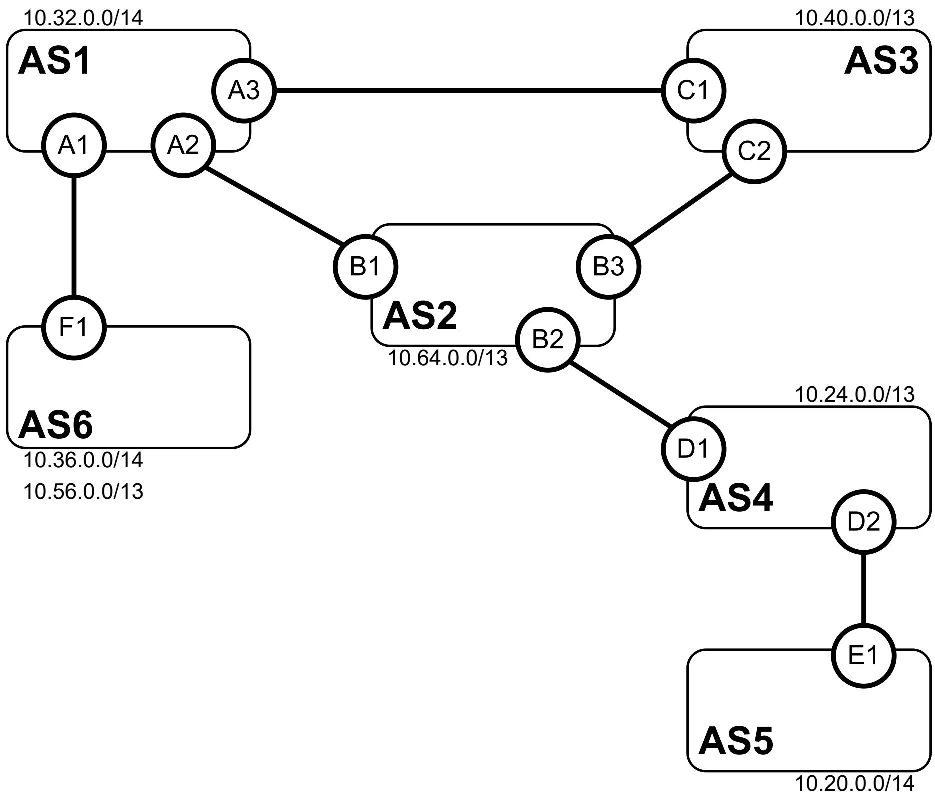

Problem 2

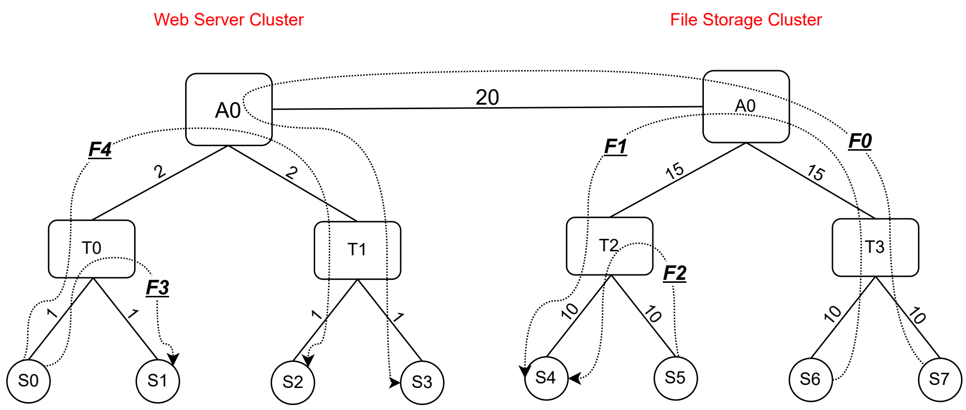

Problem 3

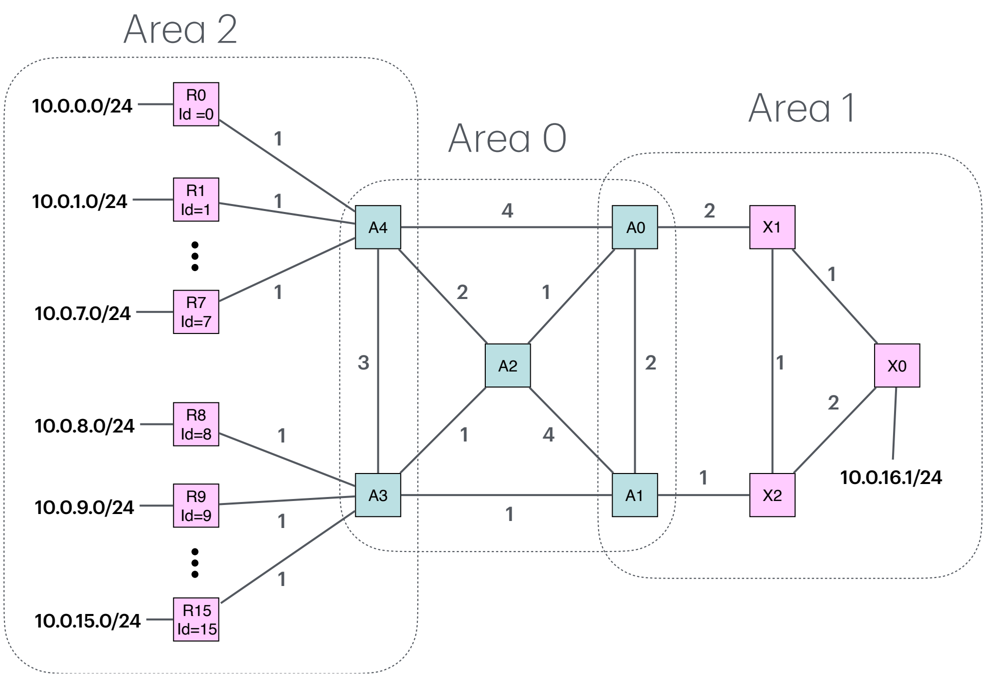

Problem 4