Homework 2: L2 and L3

Due: 11:59 PM, Oct 18, 2018

Contents

1 Spanning Tree **1**

2 IP Routing **2**

3 Spanning Tree Protocol **2**

1 Spanning Tree

Consider the network below. The three switches run the spanning tree protocol.

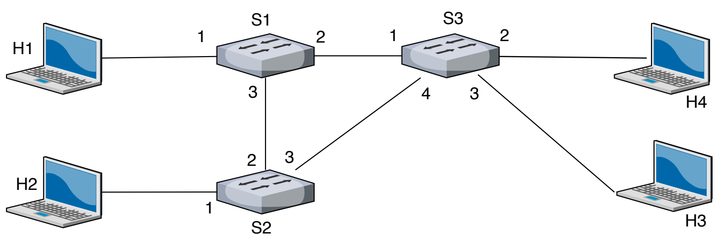

1. For each switch port, write its final state after the STP converges. Denote each port by its switch and port number (e.g., S1.1, S2.3, etc), and choose between three states: Root Port, Designated Port, or Discarding Port.

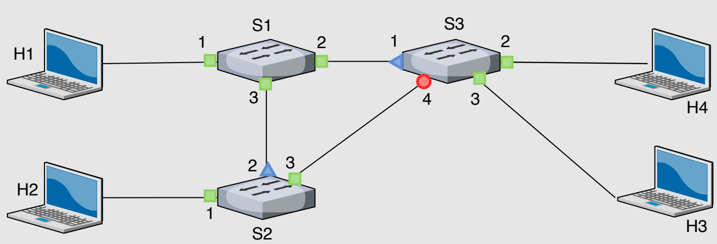

S1.1: Designated; S1.2: Designated; S1.3: Designated

S2.1: Designated; S2.2: Root; S2.3: Designated

S3.1: Root; S3.2: Designated; S3.3: Designated; S3.4: Discarding

The root is S1, because it has the lowest ID. The link between S3 and S2 is the one that forms a loop. The two switches are at the same distance, and thus the one that has the lowest id wins, the other disables its port.

2. Now assume every link is a 10Gbps link, and that switches are can switch at full bandwidth between all its ports (i.e., switches are not the bottlenecks). H1 wants to send to H4, and H2 wants to send to H3. Assuming that flows share links fairly, what is the bandwidth that each pair of nodes gets?

Each node will get 5Gbps, because their flows will be sharing the link between S1 and S3. This is bad, because if S2 were using its direct link to S3, they would each be getting 10Gbps.

3. Is there a way to renumber the switches in this network to improve the bandwidth in this case? How?

Yes. We have to make S3 be the root. So if we renumber them such that S3->S1, and either (S1->S3, S2->S2), or (S1->S2, S2->S3), each flow will use 10Gbps.

2 IP Routing

1. List two advantages of a distance vector protocol over a link state protocol.

(a) DV generates less messages per node
(b) DV is a simpler protocol

2. List two advantages of a link state protocol over a distance vector protocol.

(a) LS is more robust, will not form loops
(b) LS is more resilient to nodes spreading incorrect information (this can affect a link, but not entire paths).
(c) LS convergence is more deterministic (no count-to-infinity problems).
(d) DV is a simpler protocol

3 Spanning Tree Protocol

For the topology below, consider what happens when H1 sends a packet P to H2. All H1 knows beforehand are the IP addresses of H2 and of R1 (which is its default gateway). Before each transmission, assume all caches are empty (ARP caches, learning switch tables). The box below the router is a learning Ethernet switch.

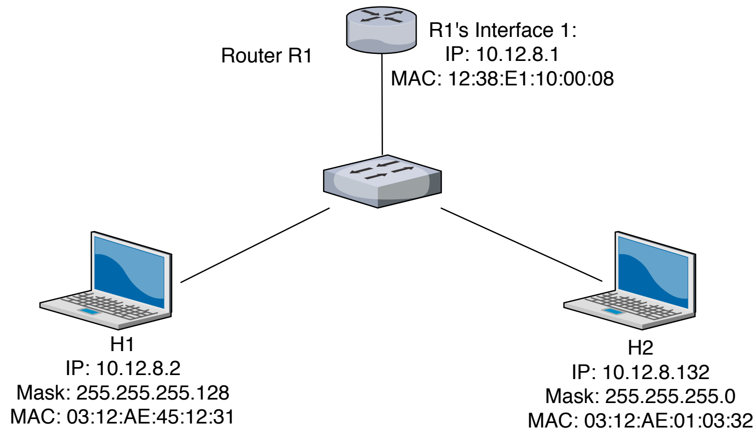

1. If H1's network mask is 255.255.255.128, list the following values. Assume that H2 has not responded to P, but assume all ARP queries and responses necessary for the communication have been made.

   (a) H1's entry(ies) in its ARP table (IP and MAC):

   10.12.8.1 -> 12:38:E1:10:00:08

   (b) Packet P destination IP address:

   10.12.8.132 (H2)

   (c) Packet P Ethernet destination MAC address:

   12:38:E1:10:00:08 (R1)

   (d) The switch's entry(ies) in the MAC learning table (MAC and Port):

   03:12:AE:45:12:31, P1; 12:38:E1:10:00:08, P2; 03:12:AE:01:03:32, P3

In this scenario, H1 determines that H2 is *not* on its subnet, because H2.ip & H1.mask != H1.ip & H1.mask. Thus, H1 has to send the packet first to the router via the LAN. The destination IP address of the packet is still H2 (this is how the router knows what to do!), but it is sent to the gateway. H1 has to issue an ARP query asking for the MAC address of the Router's IP address. This is why the router's address is the one in the ARP table in H1, and why the Ethernet address in the packet is the router's. The switch learns all three MAC addresses: it gets two packets from H1 (the ARP query and the actual packet P; it gets three packets from the router: the response to the ARP query from H1, an ARP query for the MAC address of H2, and packet P, forwarded to H2; it gets one packet from H2, the response to the router's ARP query.

2. Now assume H1's network mask is set to 255.255.255.0, and that all tables and caches have been cleared. Again, H1 sends an IP packet to H2, and H2 has not responded to this IP packet. What are the answers to the same questions:

   (a) H1's entry(ies) in its ARP table (IP and MAC):

   10.12.8.132 -> 03:12:AE:01:03:32

   (b) Packet P destination IP address:

   10.12.8.132 (H2)

   (c) Packet P Ethernet destination MAC address:

   03:12:AE:01:03:32 (H2)

   (d) The switch's entry(ies) in the MAC learning table (MAC and Port):

   03:12:AE:45:12:31,P1; 03:12:AE:01:03:32, P3

   With the change in the mask, now H1 thinks H2 is in its own subnet, and can send a packet directly to H2 over the LAN. So it issues an ARP request for H2's IP address, and send a packet with IP destination set to H2's IP, and with the L2 address of H2. The switch learns the MAC/port mappings for the two hosts, but not of the router, which is not involved in this exchange.

3. Given the MAC addresses, what (if anything) can you say about H1, H2, and H3?

   H1 and H2 share the same most significant bit of the MAC address, which means that they are from the same manufacturer.

Please let us know if you find any mistakes, inconsistencies, or confusing language in this or any other CS168 document by filling out the anonymous feedback form:

`https://piazza.com/brown/fall2018/csci1680`.