Midterm Test: NET3011 – Advanced Switching

Winter 2016

Time: 110 minutes; Test scored out of: 68 Total Marks available: 76
(Allocation of marks is shown beside each question)

Instructions:

1. BEFORE answering any questions, please check that your copy of the test has all pages (as indicated in the footer at the bottom of each page). Please **read all instructions and all questions** carefully, then answer question 0 first!

2. This is a **closed book** test. No textbooks, notes, electronic devices, or any other aids are permitted.

3. If you need extra space, please continue your answer on the last page.

4. If you are uncertain what a question is asking, make reasonable assumptions, write those assumptions down on this test paper, and continue answering the question.

0. What is your:

NAME? Answers

(Continued on next page)

1. [1 mark; Lab] What four commands that will clear a switch, guaranteed?

   delete [flash:]vlan.dat
   erase startup-config
   delete [flash:]multiple-fs
   reload

2. [1 mark; Quiz 2] What are the three technical criteria that define a *managed* switch?

   – Can assign an IP address
   – Can connect remotely
   – Can change settings/configuration

3. [2 marks] **Clearly** define and thus explain the difference between "cut-through" switching and store-and-forward switching.

   – *Store-and-forward* waits to receive the entire frame and checks the CRC **before** forwarding it.
   – *Cut-through* switching starts forwarding the frame as soon as the bits for the destination MAC address have been received.

4. [1 mark] Wireshark can tell whether a frame is Ethernet II or 802.3 format. **Clearly** explain exactly how does Wireshark *always* get it correct?

   The frame format is determined simply & entirely by the "type/length" field that follows the source MAC. Any value < 1536 is a *length* and thus an 802.3 frame; any value >= 1536 is a type field and thus an Ethernet II frame.

5. [2 marks; Quiz 1] What are the five steps in the sequence of operations when an *unmanaged* switch receives a frame?

   1. Receive entire frame [input buffer]
   2. Re-compute CRC
   3. Discard frame if bad CRC; otherwise continue processing
   4. Look in MAC address table for possible match [VLANs ignored since unmanaged]
      – if match, then select single egress port
      – else select flooding as egress method
   5. Queue the frame on selected egress port(s) [output buffer(s)]

6. [1 mark] The fact that routers need to compute new CRCs is old news. **Clearly** explain if / when a *managed* switch ever needs to compute a *new* CRC. If never, simply say "never".

   Managed switches know about VLANs, access ports, and trunk links. VLAN tags may be added or removed when crossing from an access port to a trunk or vice versa. Whenever a tag is added or removed, the CRC must be recalculated.

7. [1 mark] Clearly identify if / when VLAN tags are ever *swapped* by a switch operating only at L2. Tip: consider all operations such as going between access port and trunk, or between trunks.

   As above, tags may either be added *or* removed, but are never swapped.
   (Swapping could occur on a *router*, but not on a L2 switch.)

8. [4 marks] Give at least 2 *distinct* pros or cons for *each* of the two models of VLAN design.

   The possible answers to this question can be found in both textbook and the slidedeck.
   E-to-E: +Grouping, +virtualization, +security, +QoS, +avoiding needless routing,
   +specialized VLANs, -management, -broadcast propagation, -troubleshooting
   Local: +Simpler, +deterministic/predictable, +better link utilization(?),
   +better HA due to matched L2/L3 topology, +finite failure domain, +scalable design

9. [2 marks] In a campus environment, which model for VLAN design is most appropriate for mobile Wifi access? **Clearly** explain *why*.

   The key words in this question are: "mobile" and "wifi". To maintain seamless wifi connectivity while wandering around and moving from WAP to WAP, an *end-to-end* VLAN design is appropriate.

10. [2 marks; 1 mark/pair] Clearly identify at least 4 of the 7 generic VLAN types.

    1. default VLAN
    2. native VLAN
    3. management VLAN
    4. user/data VLAN
    5. wireless VLAN
    5. VoIP VLAN
    7. blackhole or garbage VLAN

11. [2 marks] Draw a **clear** illustration of an 802.1Q tag, showing both the *name* and *length* of each field. Be sure to also show the *tag position* within an Ethernet frame.

    From: p. 49, Fig 3.5

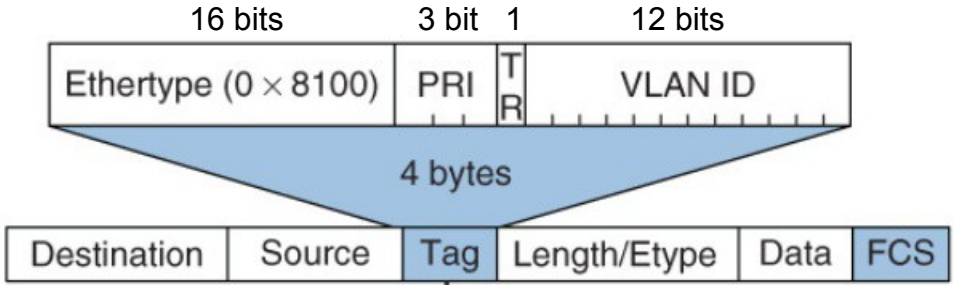

12. [2 marks] Different Cisco switches support different numbers of VLANs. Other than reasons of cost, **clearly** explain the rationale for the differences in the number of supported VLANs.

"According to the intended purpose": designed for the layers: core,distribution, access
An access switch is very unlikely to need to handle as many VLANs as a distribution switch (which possibly uses trunk links!) Thus there is *no need* for a typical access switch to handle very many VLANs, while a distribution switch *must* be capable of it.

13. [2 marks] **Clearly** illustrate VLAN pruning by combining a diagram with a brief explanation.

If switches 3, 5, 6 have *no access* ports in Red VLAN, there is no point whatsoever in them receiving broadcast or flooded traffic for that VLAN.

Pruning ensures that bandwidth (and other switch resources) on the links to switches 3, 5, 6 is not wasted carrying pointless traffic.

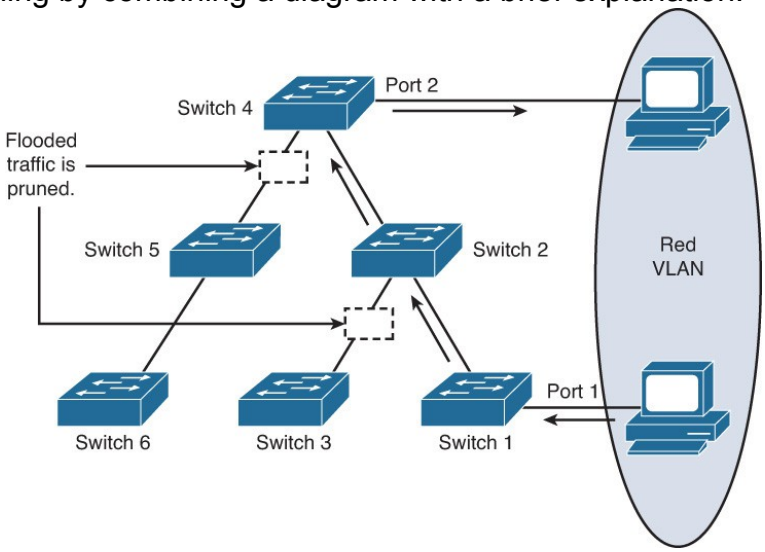

14. [6 marks] **Clearly** describe at least 6 best practices for VLAN and trunking design.
Ref: (Ch 3 slide 21)
– For local VLAN designs, ~1-3 VLANs per access module, limited to few switches
– For local VLAN designs, avoid VTP; use manually allowed VLANs instead
– Avoid using VLAN 1 as "blackhole" for unused ports; use a separate, dedicated VLAN
– Create separate VLANs for each type of traffic (user, mgmt, voice, native, etc)
– Manually configure access mode ports except those specifically intended for trunk link
– Keep all data traffic off VLAN 1; reserve it just for L2 control protocols (DTP, VTP, etc)
– on trunks, turn off Dynamic Trunking Protocol (DTP) and use manual configuration
– on trunks, use IEEE 802.1Q rather than ISL (better QoS support; open standard)
– on trunks: use "switchport nonegotiate" to save time on bootup
– Avoid using Telnet at all costs; enable SSH support on management VLANs

15. VTP has a number of requirements in order to communicate between devices.
A. [1] What must exist/be configured before VTP will even attempt to send a message?

VTP will only attempt to use links operating in *trunk mode*!

B. [2 marks; 1 mark/pair] What are at least four other parameters that must match for successful communication?
Ch 3 slide 23
– VTP version
– VTP domain name
– MD5 hash
– at least one switch in VTP server mode
– VTP password
– VTP revision numbers must be different(!)

16. [1 mark] Give an example of a VTP command, unique to ver 3, that is **not** saved in the startup config? ("show" commands don't count!)

    Ch 3, slide 26

    `vtp primary vlan` ! only 1 switch can be primary server in VTPv3

    [1 mark; Bonus] Give an example of at least one other VTP command that is **not** saved in the startup config? ("show" commands don't count!)

    (from Lab work) – Just about any VTP command, since it's all stored in vlan.dat!
    – VTP mode, domain name, password

17. A. [1 mark] From which VTP modes can a VLAN be *suspended* ?

    Suspending a VLAN is a *global* action, so only VTP servers can do this.

    B. [1 mark] From which VTP modes can a VLAN be *shutdown* ?

    Shutting down a VLAN is local so VTP servers, clients (and transparent) can do this.

18. The textbook says on p. 8: *The most important aspect to MLS is recognizing that switches can **route** or "switch" frames at wire-rate speeds using specialized hardware*.

    A. [1 mark] **Clearly** identify this type of forwarding, using appropriate terminology.

    This is *Topology Based Switching*, with Cisco's implementation called CEF

    B. [2 marks] **Clearly** explain what specialized hardware is involved, give specific details of any data structures contained in that hardware.

    Cisco implements Topology Based Switching using a *FIB* and *Adjacency table*
    The FIB requires *TCAM* (Ternary Content Addressable Memory)
    The FIB stores the actual routes; the Adjacency Table stores dst + src MAC addresses for all neighbour on all interfaces.

    C. [2 marks] **Clearly** *identify* and *explain* briefly at least two other methods of L3 forwarding.

    Process "switching" is packet forwarding done by software running on the Control Plane
    Route Cache switching processes the 1st packet of any flow in software, caches the result in memory, and then uses that hardware cache to forward all other packets in the same flow.

19. Carefully read the following Wikipedia article (https://en.wikipedia.org/wiki/512K_Day):

> "*While a full IPv4 BGP table as of August 2014, is in excess of 512,000 prefixes, many older routers have a limit of 512k routing table entries. On August 12, outages resulting from full tables hit Ebay, Lastpass and the Microsoft Cloud among others, which TheRegister dubbed 512KDay. A number of Cisco routers commonly in use have TCAM, a form of high-speed content-addressable memory, for storing BGP advertised routes. On impacted routers, the TCAM is default allocated to 512k entries for IPv4 routes, and 512k entries for IPv6 routes. While the reported number of IPv6 advertised routes was only about 20k, the number of advertised IPv4 routes reached the default limit, causing a spillover effect as routers attempted to compensate for the issue by using slow software routing (as opposed to fast hardware routing via TCAM). The main method for dealing with this issue involves operators changing the TCAM allocation to allow more IPv4 entries*"

[2 marks] Give a **clear** explanation, with specific details including example commands, on how you would alter the TCAM allocation on a 3560 switch.

From lab work
Allocation/partitioning of TCAM is done using Switch Database Manager templates
The command is: **`sdm prefer {template-name}`**

20. [2 marks] When using protected ports to achieve PVLAN functionality, it may be necessary to connect two interfaces directly together (see diagram). What configuration command(s) is/are necessary to ensure that neither of those ports gets disabled or put in blocking mode?

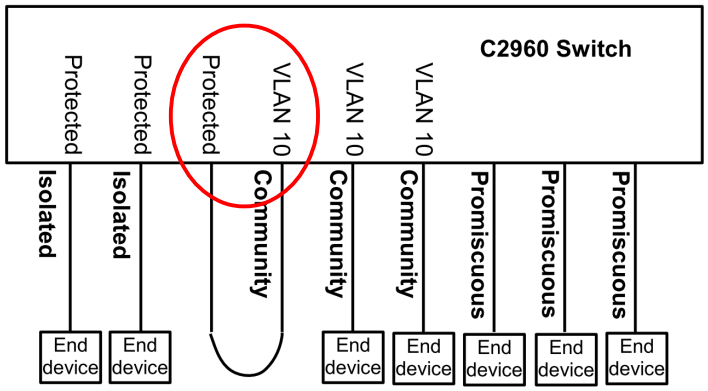

[1] BPDU filtering should be applied [2] directly to the ports.
(If BPDU filtering applied globally, then ports *must* be configured as portfast.)

21. I read somewhere that: "To achieve the optimal traffic distribution in a LAG, always bundle an even number of links."

A. [2 marks] **Clearly** describe Cisco's two basic load balancing algorithms for a LAG.

The load balancing algorithms are: XOR and hash
An XOR is used when there are *two* input params (i.e. source and dest addresses)
A hash is used when there is *one* input param (i.e. source or dest address)

B. [1 mark] **Clearly** explain why the above comment is, or isn't, correct.

As highlighted in class, the above statement is *not* correct; optimal distribution requires the number of links to be *a power of two* (i.e. 2, 4, or 8). Six links (an even number) results in two of the links having *double* the traffic of the other four!

C. [1 mark] How is it even possible for a network administrator to determine if a chosen load balancing option is effective?

Two possible answers:
– the command show interface {i/f} gives counts of traffic in/out on an interface
– tools such as NetFlow or NBAR assist with capturing and analyzing all traffic
(Please note that Wireshark isn't a particularly good tool for this application.)

D. [2 marks] **Clearly** explain why it would be desirable to have different load-balancing choices configured for each end of a LAG. Use a specific example in your answer.

Consider a bunch of hosts (left side) accessing a server or gateway (right side)
a LAG between a pair of switches is used to join the hosts to the server

– using source address (MAC or IP) results in good distribution to the server, but *no* distribution from the server (always the *same* source address!)
– likewise, using destination address (whether MAC or IP) gives good results from the server, but *no* distribution to the server (always the *same* destination address!)
– a good solution would source address on the host side (left) and destination address on the server side.

22. [2 marks] The laptop on the very right side is pinging it's gateway, RTA. Identify the VLAN tags on an ICMP Echo Req and Reply throughout the PVLAN. [1=trunk tags; 1= all others]

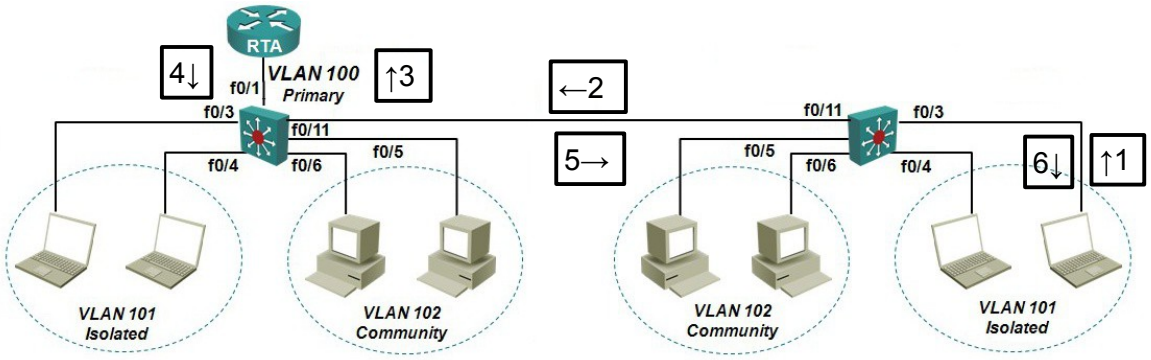

1-Echo Req from PC: **none** 2-Echo Req on trunk: 101 3-Echo Req to RTA: **none**

4-Echo Reply from RTA: **none** 5-Echo Reply on trunk: 100 6-Echo Reply to PC: **none**

23. A. [4 marks] You've memorized the criteria in the STP election processes. Draw a network topology diagram which **clearly** illustrates *each* of the four possible election criteria being used to determine the Root port.

– Assume all settings are at default

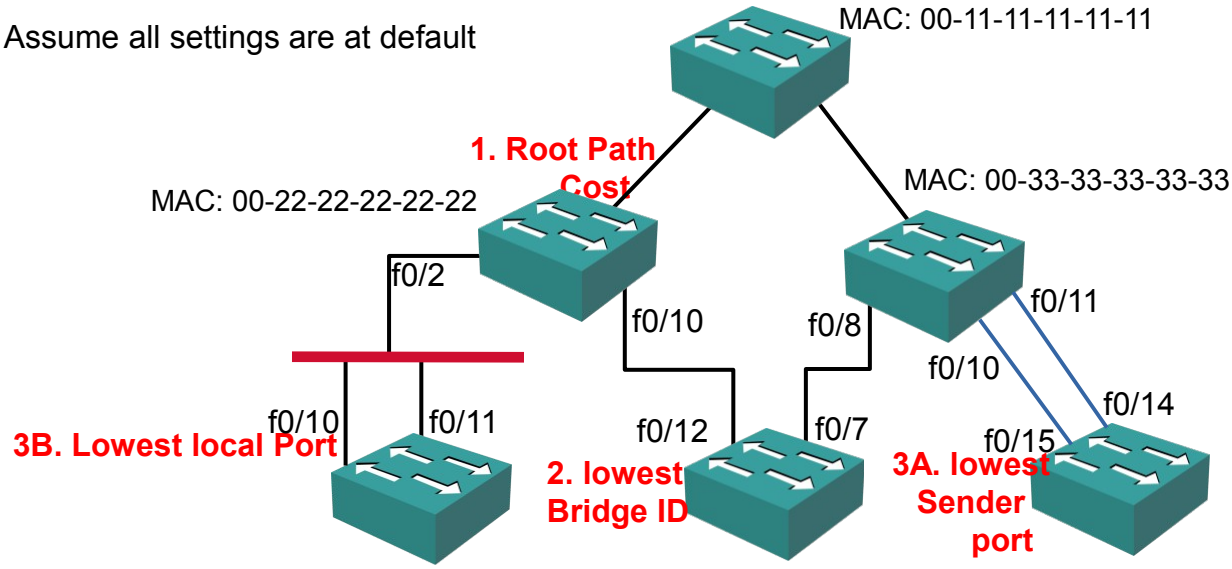

B. [1 mark] Your topology diagram certainly indicates a shared-media segment (i.e. non point-to-point). Give an example of Ethernet technology that creates this type of segment.

– an unmanaged (cheap $10) switch
– will also accept bus-topology, such as (very old) 10-Base2 or 10-Base5 Ethernet

24. A. [7 marks] In the wireshark capture below, circle the field(s) which provide the answer to the following questions about STP.

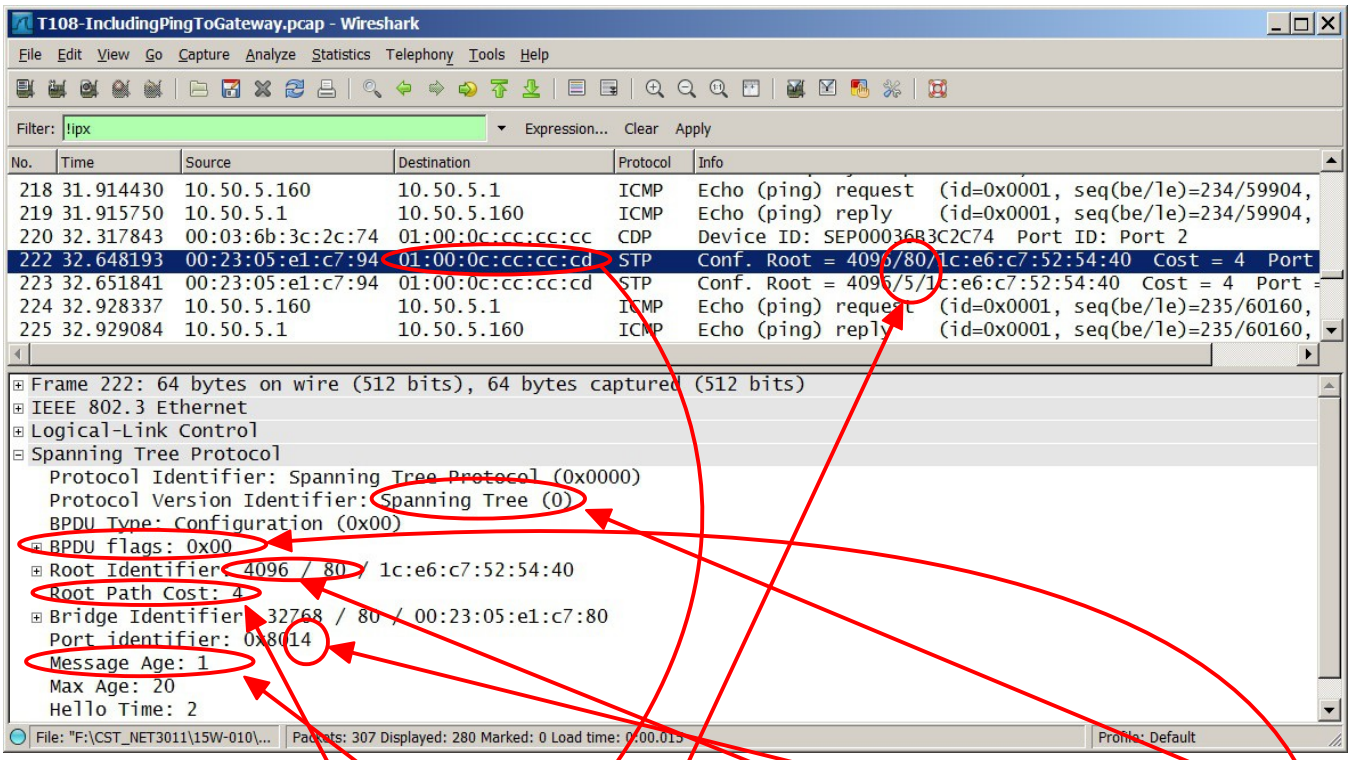

1. What is the destination MAC addr?

2. What version (STP, RSTP, MST)? **STP (0)**

3. How far from Root Bridge? **2 (msg age=1)**

4. What interface? **Port Identifier=0x14**

5. What speed link(s)? **1Gbps (cost=4)**

6. Any non-default values? **Root BID=4096; VLAN=80**

7. What VLANs exist on this link? **5, 80**

8. Whether this is this a TCA or TC BPDU? **No**

B. [1 mark] Is the interface identified above on the sending switch or the receiving switch?

It's a port sending switch (port ID is advertised; port cost is always local)

C. [1 mark] Based on your knowledge of Cisco 2960 and 3560 equipment, give the actual, exact interface.

Port ID = port priority + port number = 0x8014; so port = 0x14 => port 20 (i.e. **gi0/18**)

D. [1 mark] What would be the cost advertised by the switch with MAC 1C:E6:C7:52:54:40?

It's identified as the Root Bridge by the BPDU; the Root *always* advertises a **cost = 0**

25. Examine the following output carefully. Be sure to **use it** as the basis for all your answers!

```
DLS1# show spanning-tree
VLAN0001
 Spanning tree enabled protocol ieee
 Root ID Priority  32769
 Address     1c17.d3d2.df00
 Cost        ???
 Port        11 (FastEthernet0/4)
 Hello Time  2sec  Max Age 20sec Forward Delay 15sec

 Bridge ID Priority 32769  (priority 32768 sys-id-ext 1)
 Address     7010.5c16.2580
 Hello Time  2sec  Max Age 20sec Forward Delay 15sec
 Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- ---------
Fa0/1               Desg FWD 10        128.3    P2p
Fa0/2               Desg FWD 20        128.4    P2p
Fa0/3               Altn BLK 30        128.5    P2p
Fa0/4               Root FWD 40        128.11   P2p
```

A. [1 mark] In the output above, which value is changed in response to receiving an STP TC BPDU? (This is not a trick question; *one* of the shown values *is* modified.)

TC BPDU's change the aging time (**`Aging Time 300 sec`**) to help reconvergence

B. [1 mark] Assuming the switch above is directly connected to the Root via a single 1Gbps link, what is the value of the Root path cost?

From lab work: Output shows **configured** costs; Fa0/4 is root port, config'd cost = 40 (so default link cost is thus irrelevant); with single link ("direct") connection, **cost = 40**

C. [1 mark] Can configuration made only on this switch affect the value of the Root path cost? If yes, give the (reasonably correct) command ; if no, simply say "no".

Yes, definitely: configure the **port cost on Fa0/3 or Fa0/4** which lead to the root

D. [2 marks] Using the available information, draw a sketch of where this switch is situated in the topology. Be as specific as possible but if sufficient information isn't available, you may label interfaces simply as "towards Root", "to peer Sw", or "to lower level switches".

Diagram should look reasonably close to diagram to the right.

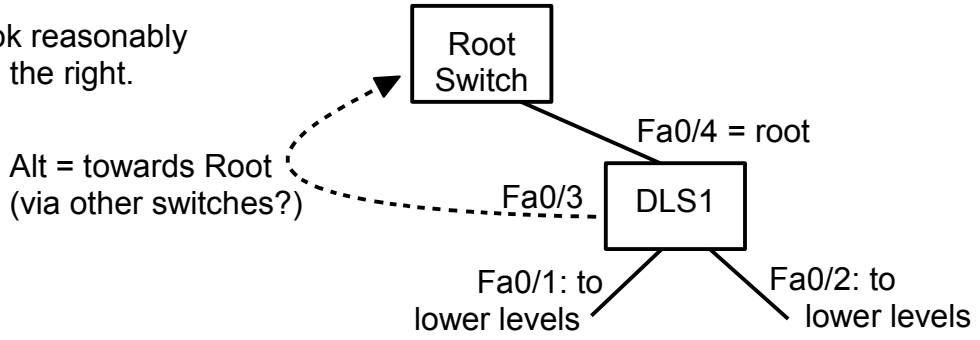

26. [3 marks] In the diagram below, **clearly** mark where to apply each STP optimization:

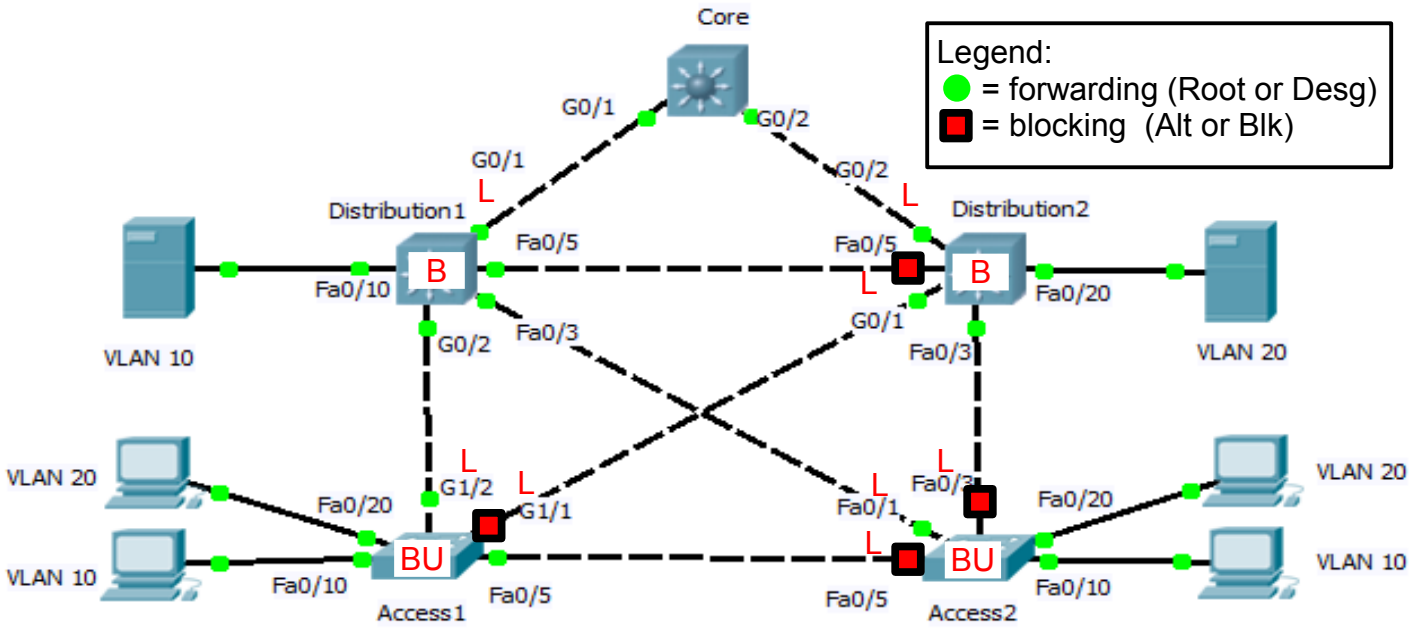

– Mark everywhere to apply Backbone fast with a "B"
– Mark everywhere to apply Loopguard with an "L" – every root or alt/blocking port (8)
– Mark everywhere to apply Uplink fast with a "U" – leaf-node switches (x2)

Backbone fast: per switch; every switch in topology (i.e. 4x)
Loopguard: per port; every non-designated (root, alt) port (i.e. 8x)
Uplink fast: per switch; every leaf switch (i.e. 2x – Access1, Access2)

Extra Work