Test 1: NET3012 – IP Architectures & Solutions

Winter 2020

Time: **70** minutes; Test scored out of: 55 Total Marks available: 58
(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 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. Be sure to carefully examine the reference topology provided below.
4. A free tip: find the easy questions and answer them first! Tackle the hard ones afterwards.
5. The interactivity in lectures is just amazing and fantastic! Keep up the good work!
6. 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

Reference Topology

Use the topology below for questions which refer to R1-R6 but do **not** have a topology diagram. Note that this is similar to the topology used throughout the MPLS courseware and slide decks.

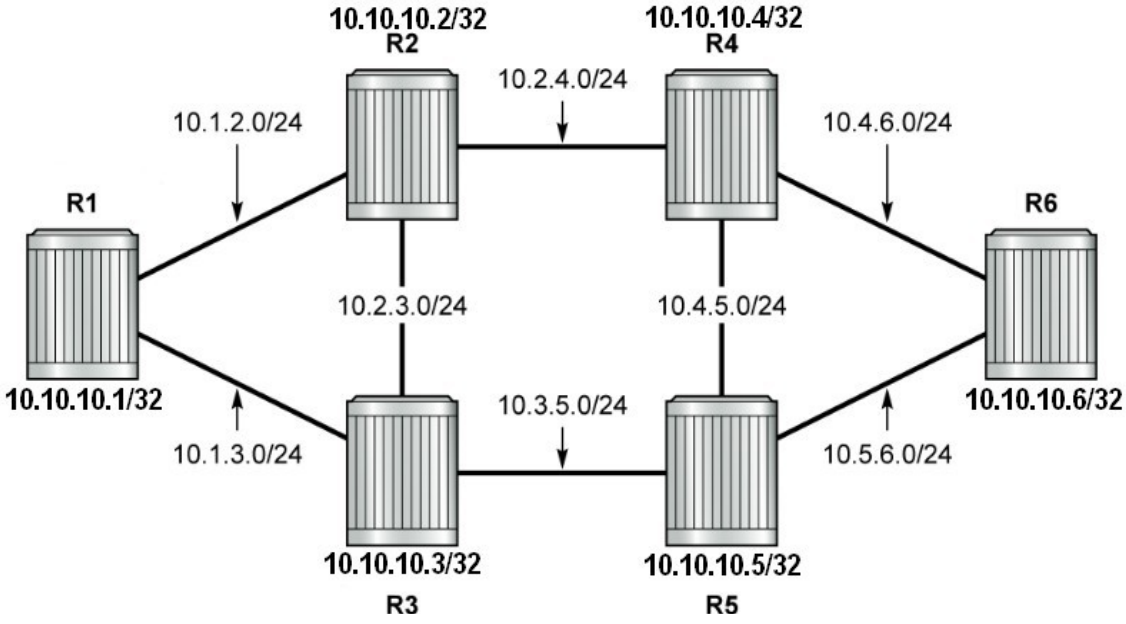

1. **A.** [3 marks; 1 per attribute] **Clearly** identify and **clearly** (but briefly) explain the options for all possible attributes for label distribution, control, and retention modes.

Label distribution: Downstream-unsolicited (link LDP): labels for all known FECs are given out to all neighbours without any request by those neighbours.
Downstream-on-demand (RSVP): labels for a specific FEC is only given out upon request by an upstream router.

Label control: Independent control (not used): labels given out without certainly of being able to complete the LSP to the FEC (i.e. no downstream label)
Ordered control: labels only given out if the router has a downstream label

Label retention: Conservative retention: conserve memory by *discarding unused* labels
Liberal retention: use memory liberally to *save all* labels, used or not

B. [2 marks] For LDP, and for RSVP, **clearly** identify the label handling attributes used.

LDP: Downstream unsolicited; ordered; liberal
RSVP: Downstream on demand; ordered; conservative

2. The diagram below is an edited version of a slide, and some details are either blanked-out or incorrect. Write visibly (!) on the diagram to supply the following information:
A. [1 mark] From the Init message shown, is it possible to determine whether this is

Link or Targeted LDP? If yes, which is it? Both identical, so no way to know

B. [1 mark] First, assume that this is Link LDP. Fix the LDP ID (4 spots)
C. [2 marks] Now assume this is targeted LDP. Provide the destination IP, whether TCP or UDP, and valid/possible port numbers

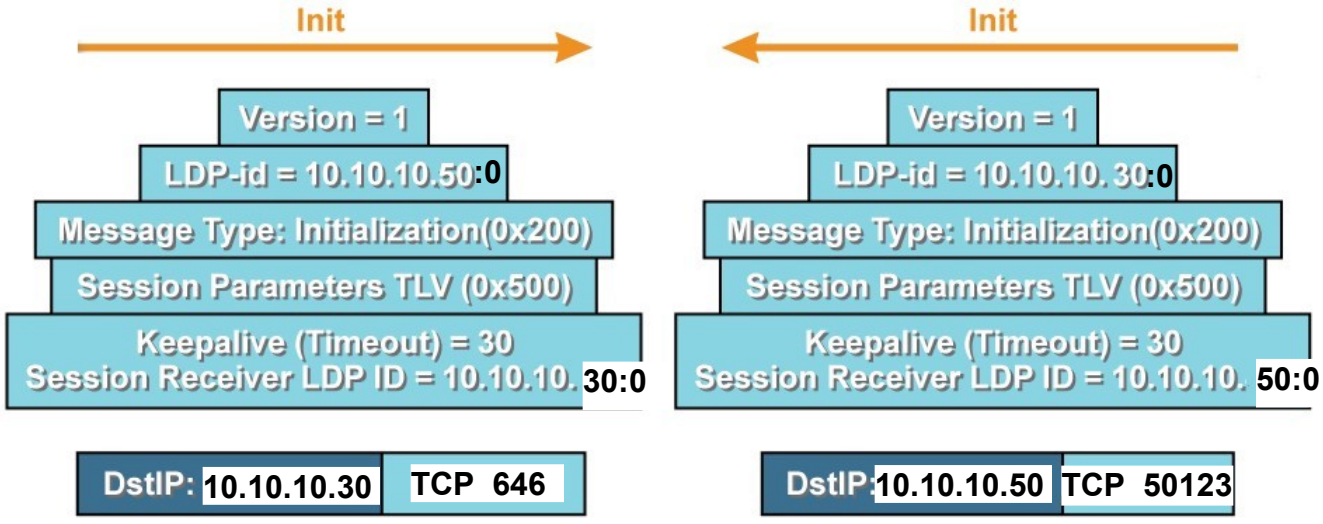

3. [0 marks] Yes/No: Does targeted LDP provide a full mesh of tunnels like link LDP does? [1 mark] **Clearly** explain why or why not.

   No, it does not. All peering sessions are administratively configured so full-mesh is not automatically achieved.

4. [2 marks; 1 per pair] **Clearly** identify at least four (4) of the TE path selection mechanisms available in RSVP-TE. For each one, provide a short 1-2 sentence description.

   1. Follow the IGP: just follow the normal IGP routes for LSP path selection.
   2. TE Metric: use the administratively configured hop metrics, combined with a conventional routing selection
   3. Hop count: Specifies the maximum number of hops (including the head-end) that the LSP may traverse
   4. Strict and Loose hops: nodes that the path should absolutely traverse, either immediately/next or eventually.
   5. Bandwidth: Ensure that the requested bandwidth is available (at least from a "book-keeping" perspective).
   6. Admin groups (aka Link colouring): links tagged with arbitrary attributes by admin, and then either included (manditory) or excluded (must *not* be used).
   7. Shared Risk Link Groups (SRLG): links tagged and then secondary must *not* use any links that are selected/in-use by the primary.

5. [1 mark] In RSVP request messages, the Router Alert option is used. **Clearly** explain the reason and purpose for including this option. Ref: slide 4.18

   This particular Router Alert (there are two!) is in the **IP header**. It identifies that *every* transit router must pass the packet up to the control plane for processing (build PSB) rather than the usual treatment of just forwarding the packet. This is one of the only examples where transit routers modify ("muck with") a packet before forwarding it.

6. [2 marks; 1 per correct line] In RSVP, some messages build up the PSB and RSB, and others clear them out. **Clearly** identify which message(s):

   A. build up the PSB PATH and the RSB RESV

   B. clear them out PATH Tear + RESV Tear

7. [2 marks] Clearly identify the three LDP messages involved in managing the *distribution of labels* (*not sessions*). To get both marks, be sure to indicate what role each one plays.

   Label mapping: advertise a label for a particular FEC
   Label withdraw: invalidate/withdraw a label for a FEC (no longer reachable)
   Label release: acknowledgement of a label withdraw message
   [In 2020, accepted Hello, Init, Notify]

8. Examine the CLI output provided below. It is obtained from the reference topology for R1. Assume that the both the IGP and LDP are fully converged.

```
A:R1# show router ldp bindings active
=============================================================================
Prefix            Op    IngLbl    EgrLbl    EgrIntf/LspId  EgrNextHop
-----------------------------------------------------------------------------
10.10.10.1/32     Pop   131070      --        --             --
10.10.10.2/32     Push    --      131071    1/1/4          10.1.2.2
10.10.10.2/32     Swap  131071    131071    1/1/4          10.1.2.2
10.10.10.3/32     Push    --      131071    1/1/3          10.1.3.3
10.10.10.3/32     Swap  131069    131071    1/1/3          10.1.3.3
10.10.10.4/32     Push    --      131070    1/1/4          10.1.2.2
10.10.10.4/32     Swap  131068    131070    1/1/4          10.1.2.2
10.10.10.5/32     Push    --      131068    1/1/3          10.1.3.3
10.10.10.5/32     Swap  131067    131068    1/1/3          10.1.3.3
10.10.10.6/32     Push    --      131068    1/1/4          10.1.2.2
10.10.10.6/32     Swap  131066    131068    1/1/4          10.1.2.2
-----------------------------------------------------------------------------
```

[2 marks] If ECMP was enabled on all routers, **clearly** explain what changes, if any, would appear in the table above.

[1 mark] An extra pair of labels to R6, since it is the only *equal cost* FEC in the table.
[1 mark] The labels would egress to R3 (1/1/3 – 10.1.3.3) [R2 labels already exist].

9. **A.**[1 mark] **Clearly** identify at least three (3) protocols which can be used as the L2 transport for MPLS and it's data.

OSI L2! ATM, Frame Relay, Ethernet [and GRE too, but it hasn't been covered yet]

B. [1 mark] Which of these same protocols can be *carried within* MPLS?

All of them, plus TDM

10. [2 marks] Now that we've seen some of both RSVP and LDP, we can ask an important question: Identify which of the following protocols could be used as the IGP for each of the two protocols? Choose all that apply. Ref: slide 1.11, 5.18

LDP: ☐ BGP ☑ EIGRP ☑ ISIS ☑ OSPF ☑ RIP (BGP is *not* an IGP)

RSVP: ☐ BGP ☐ EIGRP ☑ ISIS ☑ OSPF ☐ RIP (Need link-state IGP)

11. **A.** [6 marks] Imagine three routers which implement a Label Switched Path.
Using the diagrams of the routers provided below:

- ☐ clearly identify which *type of PDU* ingresses and which *type of PDU* egresses (i.e. re-draw them at each location where they should be!)
- ☐ clearly illustrate the flow of the frame *through* the router (i.e. which tables are used to process the frame), according to the role labeled for each router. (Ref: Mod 1.39-40)

Sample PDU Types

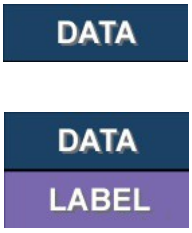

Ref: Slide 1.40 **eLER** (Label op = Pop)

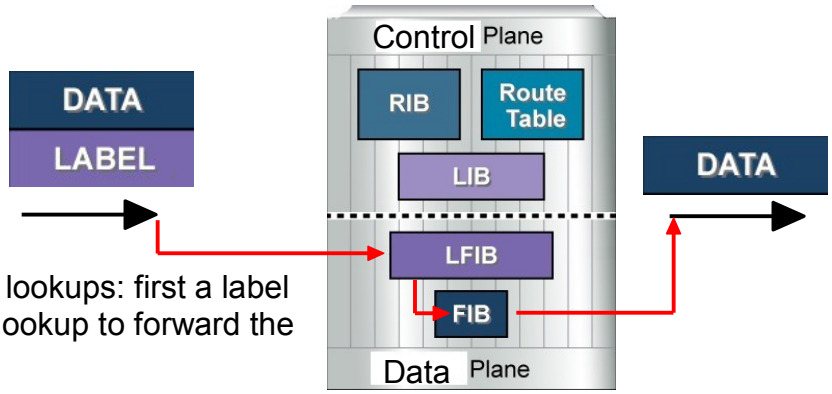

"MPLS egress router performs two lookups: first a label lookup, and then an IP forwarding lookup to forward the unlabeled packet" NRS-II, p. 505

Ref: Slide 1.39 **iLER** (Label op = Push)

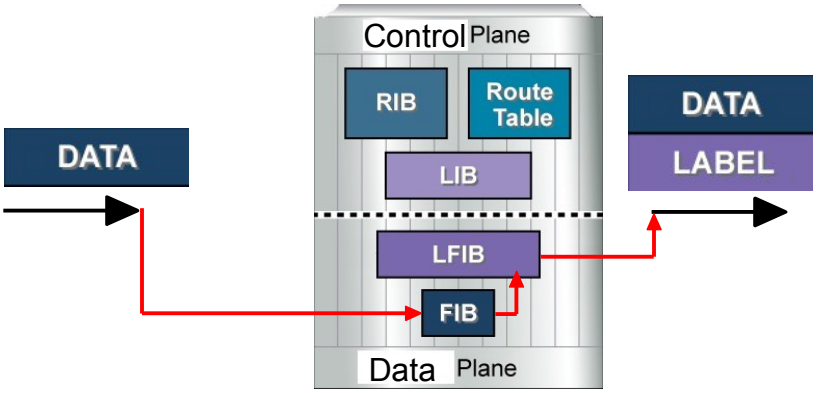

Ref: Slide 1.40 **LSR** (Label op = Swap)

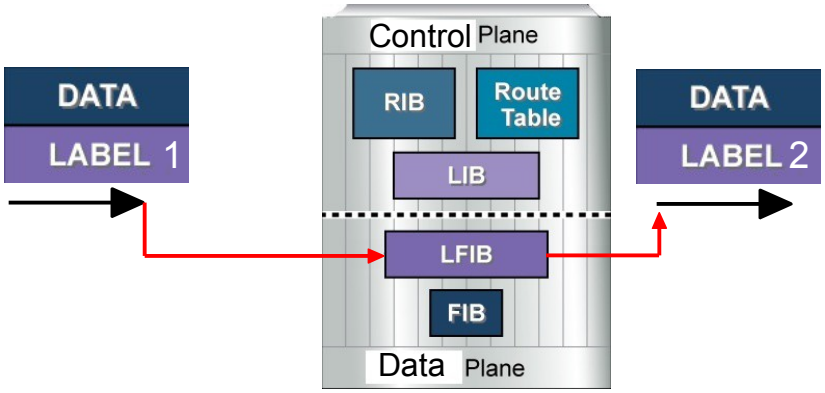

B. [1 mark] Next, clearly label the different **planes** for each of the three types of routers.

C. [1 mark] Lastly, clearly label the type of label operation that's occurring at each router.

12. A. [2 marks] Similarly to the previous question, clearly identify the flow of a ***PATH message*** through a MPLS router. Ref: Module 4.18, 20, 21

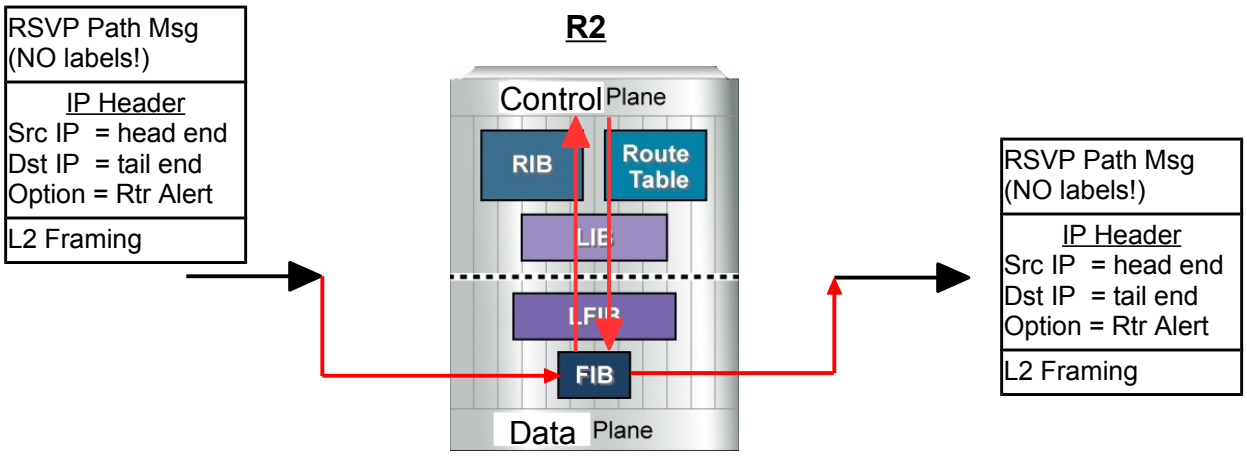

B. [2 marks] Assuming the *PATH message* is sent by R1 to R6, use the space above to:
- ☐ draw a simple sketch of the ingress packet and egress packet in the correct locations;
- ☐ clearly identify actual values of key fields (and any important options) in the IP header, and any label values (if appropriate).

13. [1 mark; Bonus] Consider the characteristics of LDP, and particularly what it produces in terms of LSPs. It could be argued that LDP is not particularly suitable for use in very large networks, with 1000's or 10,000's of nodes. What might be the characteristic of LDP that makes this true?

    Full-mesh, so lots of overhead from 1000's of labels
    Though we may want full mesh within a specific (limited) area, we don't usually want it for the *entire* topology of a *large* (continental-scale?!) provider network.

14. [1 mark] Slide 2.50 describes the purpose of a specific label: it "*indicates to the receiving router that the packet must be passed to the control plane for processing*". Clearly identify the actual value and name for this label.

    The Router Alert (label, not IP header option) has value **1**.

15. Carefully examine the diagram below. Assume that no other protocols or databases are implemented on the routers.

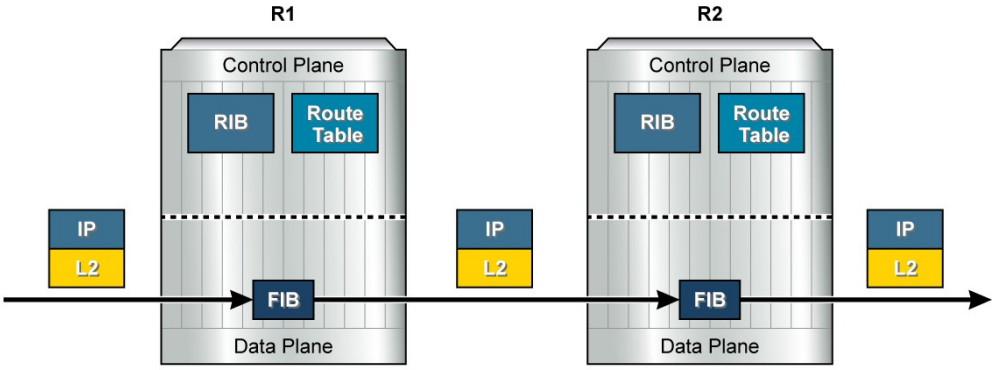

A. [1 mark] **Clearly** explain what R1 would do if a frame containing MPLS arrived on the ingress interface?

Drop the packet; there's just no way to process it

B. [1 mark] **Clearly** explain the mechanics of how R1 reaches it's decision on how to handle the frame. Your explanation should include relevant fields and data that the router would check (**Circle them** in the diagram!)

Check the Ethertype field in the L2 (Ethernet) header

16. [2 marks] Clearly identify (by "name") which routers, if any, decrement the TTL* …

A. on one or more labels (for at least 1 VPN type) LSR

B. in the customer's IP header (for at least 1 VPN type) iLER, eLER

(*If you have any doubts about how to answer:
– Copying a TTL value is not the same as decrementing it.
– If it makes any difference, state whether your answer applies to pipe mode or uniform mode.) Ref: slides 2.17-2.19

17. [1 mark] [Following on from the above question.] In circumstances where an MPLS router decrements a label's TTL, clearly define which (or both) label(s) are changed:

A. the TTL on the top label – Yes/No? Yes

B. the TTL on the bottom label – Yes/No? No

18. [1 mark] Slide 3.56 below shows the datagram for a LSP-ping, i.e. MPLS Echo Request. The question is, to which field/header shown in the diagram does the TTL apply?

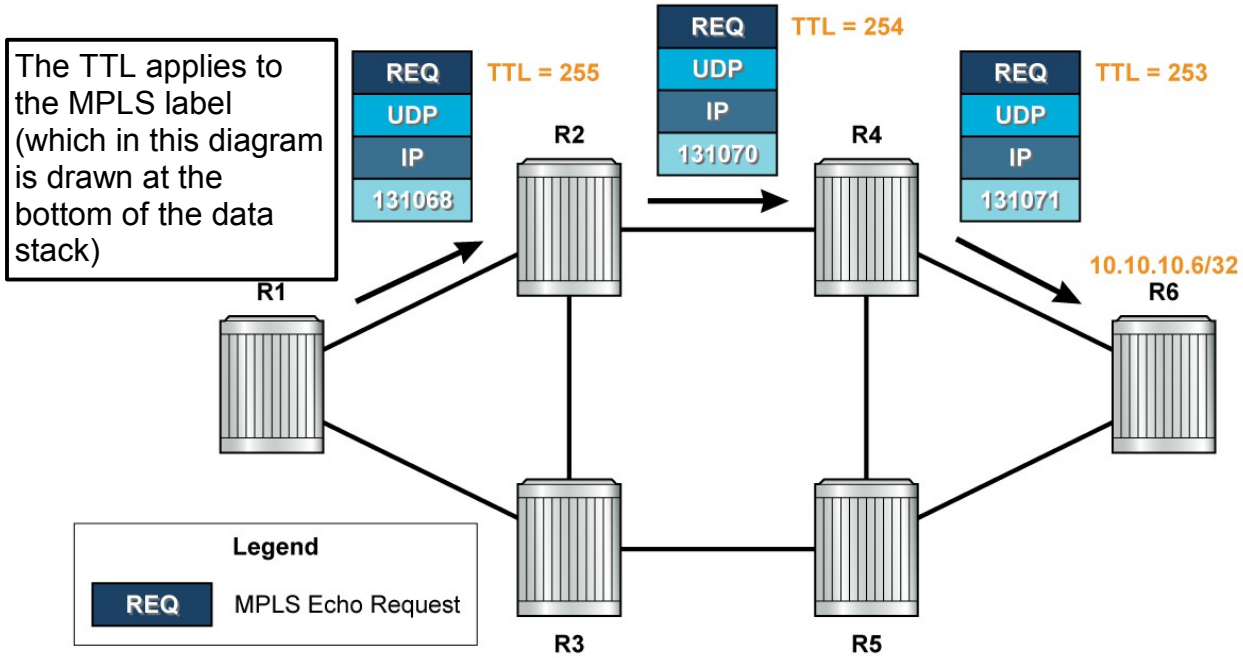

19. One of the drivers of MPLS is solving hyper-aggregation issues.
[0 marks] Yes/No: Can LDP assist with solving the hyper-aggregation problem?
[1 mark] Explain clearly why or why not.

No, not at all. LDP always follows the IGP, so it's no better at avoiding hyper-aggregation than the IGP that it's following.

20. Slide 1.38 is titled "Building Tables: MPLS Control Plane". It contains the sentences:

"*A selection process might be performed on the LIB when constructing the LFIB. Thus, the LIB might contain some redundant entries, those are not actually used on the data plane (LFIB) at a given time.*"

[1 mark] **Clearly** explain the reason for "redundant entries" in the LIB, with specific reference to protocols studied in this course.
NB: the questions says "redundant" and not "incorrect".
LDP: it has liberal label retention which means keeping (redundant) labels

[1 mark] What CLI command(s) allow you to see "redundant entries" in the LIB?
[Ref: both lecture material, as well as lab work]

Show the LIB (not the LFIB): show router ldp bindings

21. [2 marks] Carefully examine the diagram below. Show that you understand the label distribution process: Beside each label, clearly indicate which router *originated* the label.

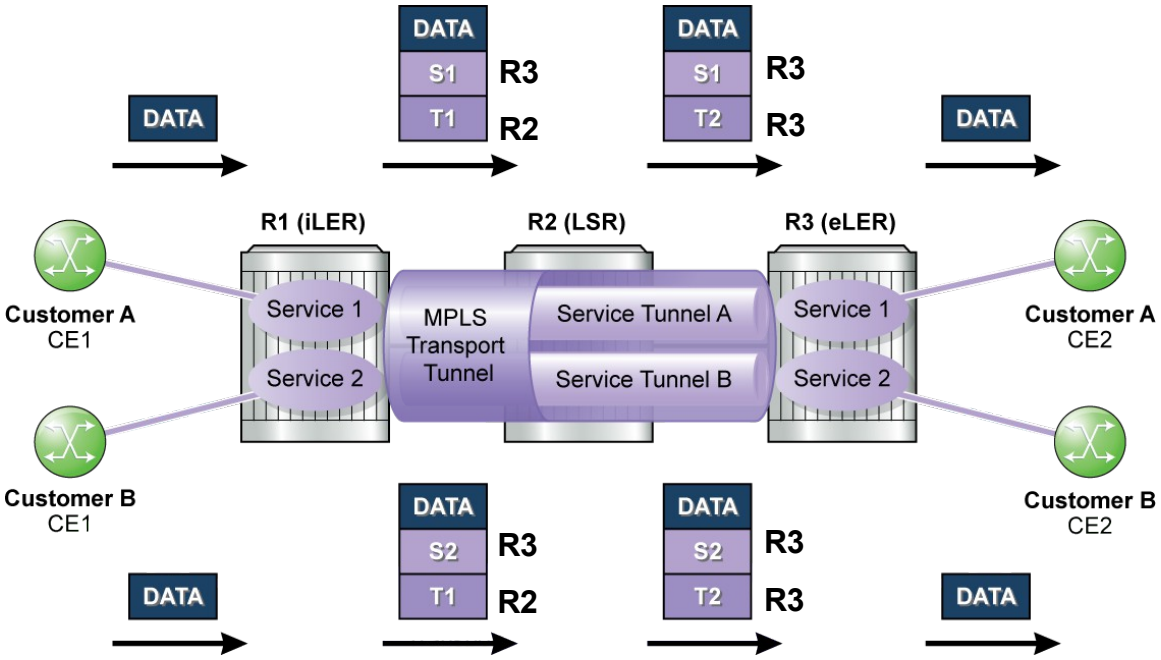

22. [1 mark] Again, show that you understand the label distribution process. Carefully examine the diagram below, which depicts the packet for an ICMP Echo Request flowing through an MPLS tunnel from R1 to R6. In the diagram, clearly indicate what exact label values would be used for the Echo Reply. (For any unknown values, write the word "unknown".)

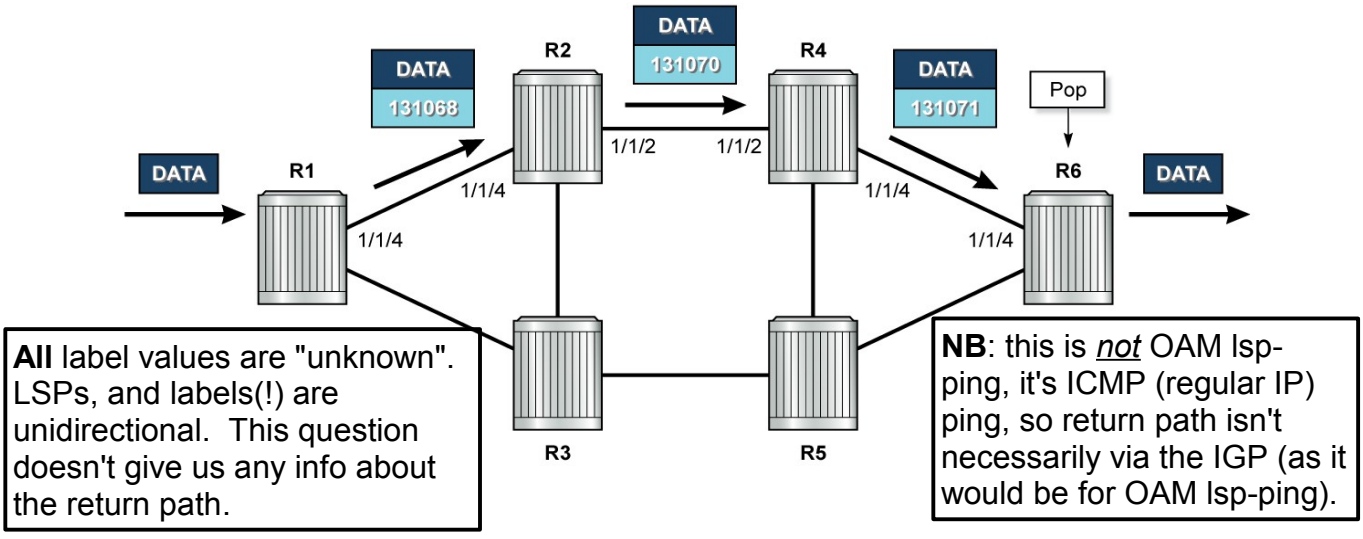

23. [1 mark] Clearly indicate which type of label space is implemented for each of: (a) frame mode, and (b) cell mode MPLS. (Be clear about which is which!)

Frame mode = per platform; cell mode = per interface

24. [4 marks] Carefully examine the diagram below and then mark it with answers. Ref:slide 2.9
    A. For each and every layer throughout the diagram, clearly label it with a "C" or "SP" to indicate whether it derives from the Customer (C) or the service provider (SP).
    B. Using symbols (such as triangle, square, circle, asterisk, etc), clearly indicate fields that are identical (same symbol) or different (different symbols). Don't forget about TTL!
    C. Clearly indicate on the diagram which is the Bottom label and which is the Top label.
    D. Demonstrate your knowledge about the numeric ranges of label values. Assuming all labels are dynamically allocated, write valid/*possible* label values on the diagram, in the correct location.

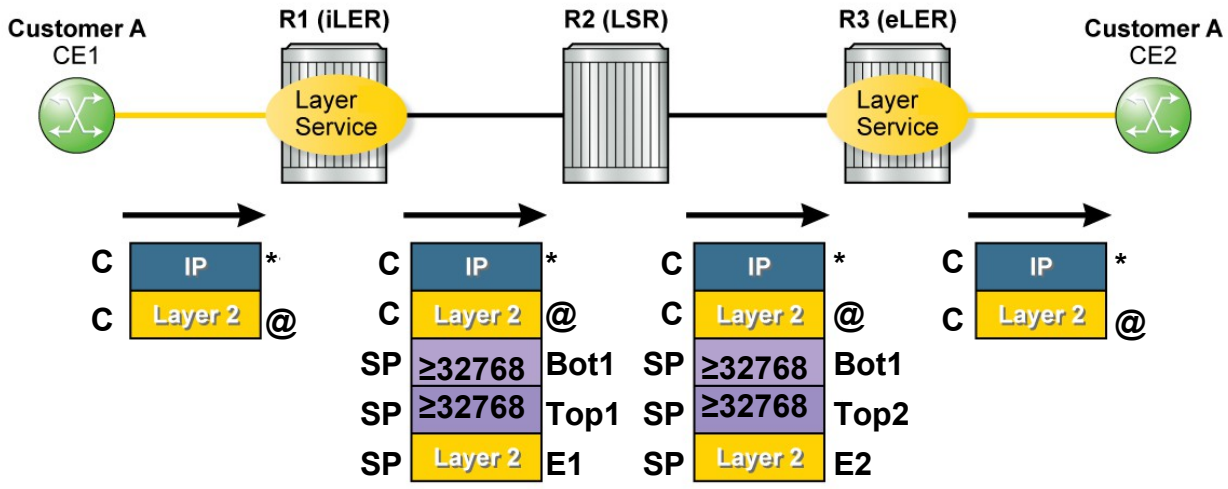

25. [1 mark] Carefully examine the diagram below. For each label, write beside the diagram to **clearly** indicate the numeric value of the "S" field. Ref: slide 2.15

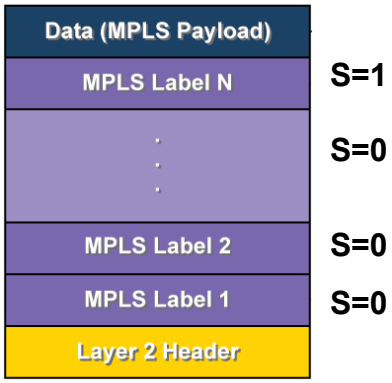

26. [1 mark] RFC7274 discusses uses of labels. It states

    "*The special-purpose label value of [xxx] is only used in [control plane] signaling, never in the data plane. Could it (and should it) be used in the data plane? If so, how and for what purpose?*"

    There is only one possible label being discussed. Clearly identify the actual value and name for this label. Ref: Slide 2.44-2.46

    The only label signalled but never used for data is Implicit Null, with a value of **3**

27. [2 marks] As we go through the course material, the later sections fill in gaps from earlier sections. Carefully examine the two figures below (ref: slides 2.31, 2.32). We now know actual names for the messages shown.

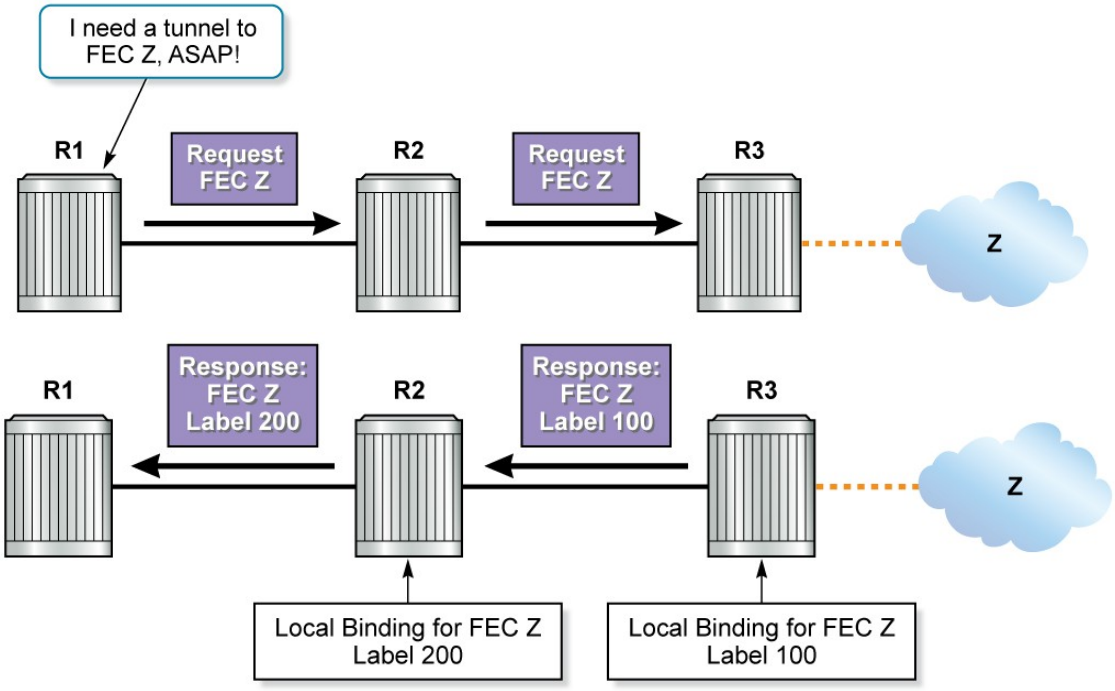

A. What is the protocol depicted above? RSVP (or RSVP-TE)

B. What are the actual, exact message names? Top = PATH, Bottom = RESV

28. [1 mark] In this question, you must compare and contrast forwarding via the IGP vs forwarding via a LSP. For data going from R1 to R6 (see cover page), clearly describe the nodes where FEC lookup is done for IGP forwarding, and for LSP forwarding.

IGP: FEC lookup at each and every hop (slides 1.25, 1.30)
LSP: FEC lookup *only* at iLER (slide 1.31)

29. [1 mark] For MPLS, frame mode differs from cell mode in at least three ways: how labels are inserted into the datagram, the label space, and the number of LDP sessions when there's multiple links. **Clearly** explain, *including a diagram*, any one of these three.
Ref: slide 2.20; slide 2.39, slide 3.21

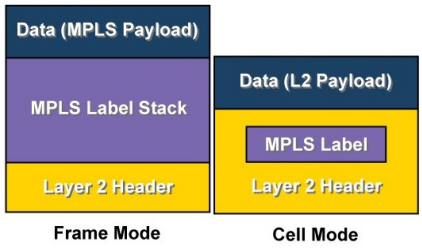

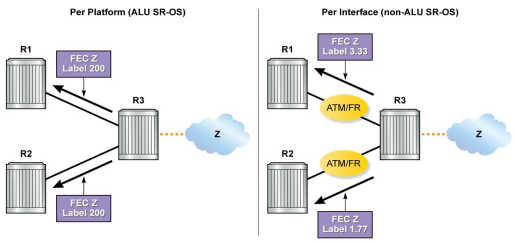

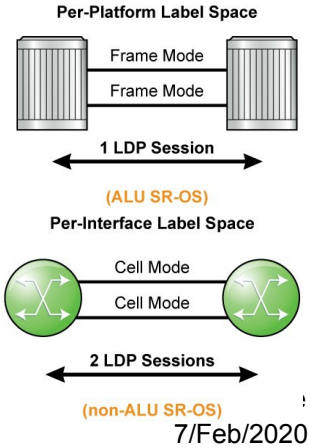

cell: labels *in* existing header single label set vs label set per i/f

Extra Work