Test 1: NET3012 – IP Architectures & Solutions

Winter 2019

Time: 60 minutes; Test scored out of: 48 Total Marks available: 53
(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. The work (and learning!) that's occurring in labs is terrific! Keep up the good work!
4. Be sure to carefully examine the reference topology provided below.
5. 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.

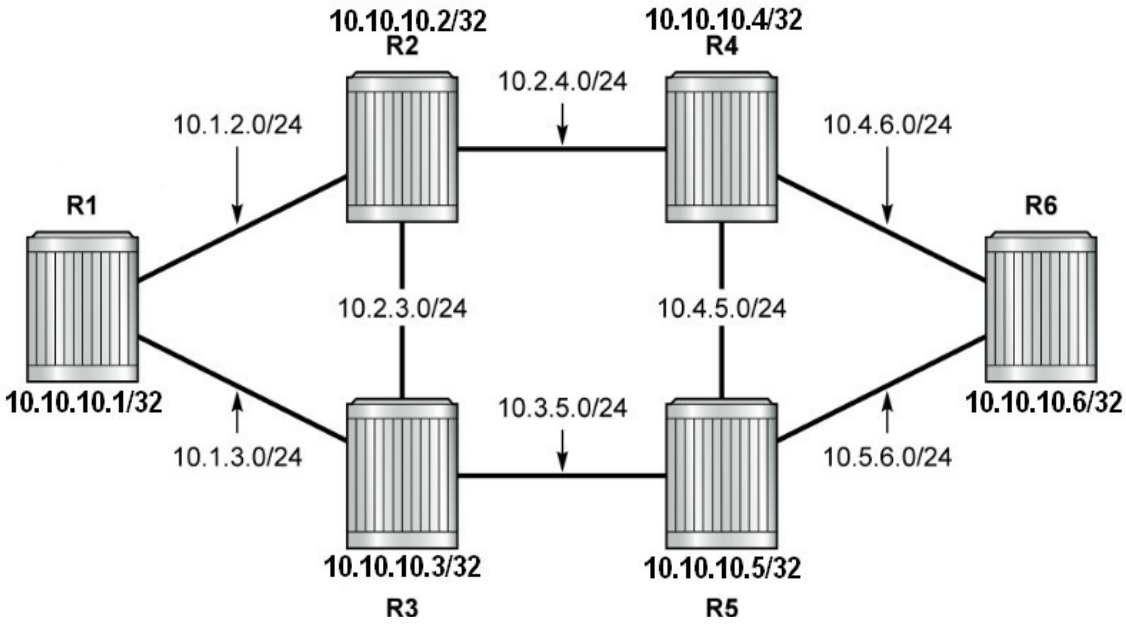

1. A. [2 marks; 1 per triplet] **Clearly** identify at least 6 of the 7 drivers for MPLS [Ref: 1.10]

   – Improved forwarding performance
   – Traffic Engineering applications
   – Building High Available Networks
   – Consolidation of Services over a common infrastructure
   – Delivering Layer 2 and Layer 3 Services
   – Triple Play Solutions
   – Building a BGP Free Core

   B. [2 marks] Based on modules covered and lab work completed so far, identify which drivers involve LDP in one form or another? If none, please state "none".

   [1 mark for first pair; 1 mark for third]

   Improving forwarding performance [lab work: LDP-shortcuts]
   L2 services (t-LDP) [slide decks]
   BGP-free core [post-lab work]

2. [2 marks; Easy question] Draw a **clear** sketch of an MPLS header. Be sure to identify the name and size/length of each field.

   [1 mark for all four fields; 1 mark for all field lengths]

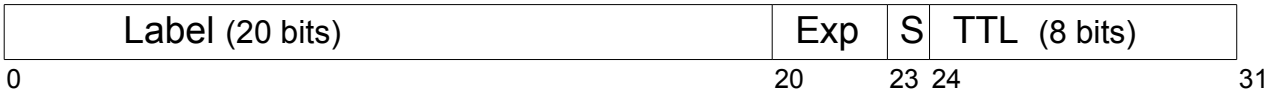

3. [2 marks] Here's a repeat question from one of the weekly quizzes, just using different words! Write three synonyms for the router which:

   – originates an LSP? iLER, head-end, upstream router

   – originates labels for a FEC? eLER, tail-end, downstream router, or "terminate" (router)

4. Think carefully about label distribution for LDP and RSVP in our reference topology.
   A. [1 mark] Is it likely that the LIB matches the LFIB exactly for either protocol? If yes, **clearly** identify which one, or both? If not, state "none".

   RSVP (which uses conservative label retention)

B. [2 marks] **Clearly** explain *why or why not*?

Conservative label retention *only* keeps labels in-use; if they're in-use, they're in both the LIB and LFIB; if they're not in use, they won't be kept in either the LIB or the LFIB.

5. A student believes they have correctly configured LDP-shortcuts. [Ref: Lab work]
A. [1 mark] **Clearly** identify what is the *best* command to absolutely confirm that user data is being forwarded over the LSP? Be sure to indicate whether to use regular IP commands or OAM commands!

Regular IP "traceroute" since it will be treated the same as the customer traffic.

B. [1 mark] **Clearly** explain *why* your given command is the best.

Ping does **not** tell you anything about *which* hops were used (to return).
Traceroute *does* explicitly *identify* which hops are visible(!) and used (for outbound).

6. [2 marks] I've just paid my ISP a boatload of money to create express-lane LSPs from my house for using Netflix and Google. **Clearly** identify how many LSPs are required and *why* that many. [Be sure to state any required assumptions.]

This answer hinges on the fact that LSPs are unidirectional.
Netflix + Google are both well known content providers that use TCP i.e. bi-directional
[1 mark] Two LSPs are needed for each provider, one for each direction.
[1 mark] A total of **four** LSP are needed.
[Partial marks (1 mark) for clearly identifying unidirectional nature of LSPs]

7. A. [1 mark] **Clearly** define asymmetric routing.

Asymmetric routing describes the situation where the *outbound* path differs from the *return* path.

B. [1 mark] **Clearly** explain why OAM LSP commands have significant potential for asymmetric routing between requests and response(s).

OAM LSP commands use the LSP for *outbound* propagation along the LSP path, but use the IGP for propagation on the *return* path.

C. [2 marks] In a conventional ("normal") network, OAM commands for LDP LSPs are *not* likely to have asymmetric routing, whereas the same commands for RSVP-TE LSPs *are* likely. **Clearly** explain *why* this is the case.

This answer hinges on the fact that LDP always follows the IGP, RSVP may not:
LDP OAM commands follow the LSP … which follows the IGP! Symmetric IGP routing!
RSVP OAM commands follow the LSP … which may be traffic-engineered to a non-IGP path. TE routing outbound but IGP routing return, which may be totally different.

8. [1 mark] **Clearly** and *fully* identify the protocol shown below. [Ref: slide 87, mod 3]

The protocol is: t-LDP Key: unicast Hello = Targeted

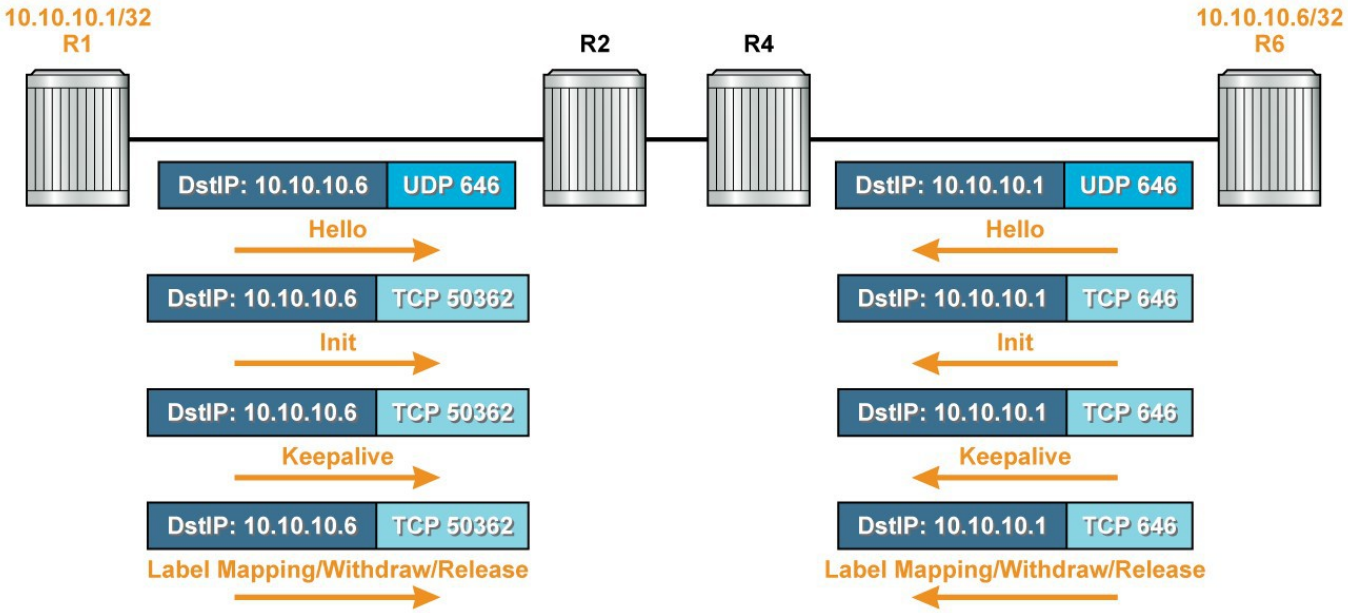

[2 marks] Complete the blank spaces with (possible) values for the L3 and L4 addressing (including protocol!) [1 mark for IP addresses; 1 mark for TCP proto + ports]

[1 mark] What is the protocol's specific *name* for the *L3 address* that you just filled in?

the Transport Address

9. [1 mark; Weekly pop-quiz question; **Bonus**] What are the two controls that need to be flipped "On" to activate traffic engineering for RSVP?

Enable *traffic-engineering* in the IGP
Specify *cspf* within the body of the LSP definition

Do **not** confuse enabling basic RSVP LSPs with enabling **TE** capability for RSVP!

10. This question focuses on label operations and label stacks (compare with next question!)

Study the topology below carefully. I would like to call my grandmother in North Vancouver and ask her what she made for dinner last Sunday, and I'll want to take my turn talking to tell her about the SuperBowl. Fortunately, **A**lgonquin (Customer **A**) has a VoIP connection direct to Vancouver via a VPN service. But note that the connection also provides a VPN service to another customer (How does that affect label operations?!) [Ref: slide 2-7]

Mark up the diagram below to **clearly** indicate:

A. [2 mark] R1, R2, R3 are what **type(s)** / role(s) of MPLS router?
Make sure your answer is *complete*! (Shown above each router)

B. [2 marks] For each router, the full set of **operation**(s) it performs on labels. (below rtr)

C. [2 marks] Draw two diagrams, one between each pair of routers, showing a simplified **datagram** . At a minimum, the datagram must include the payload and a fully identified label stack.

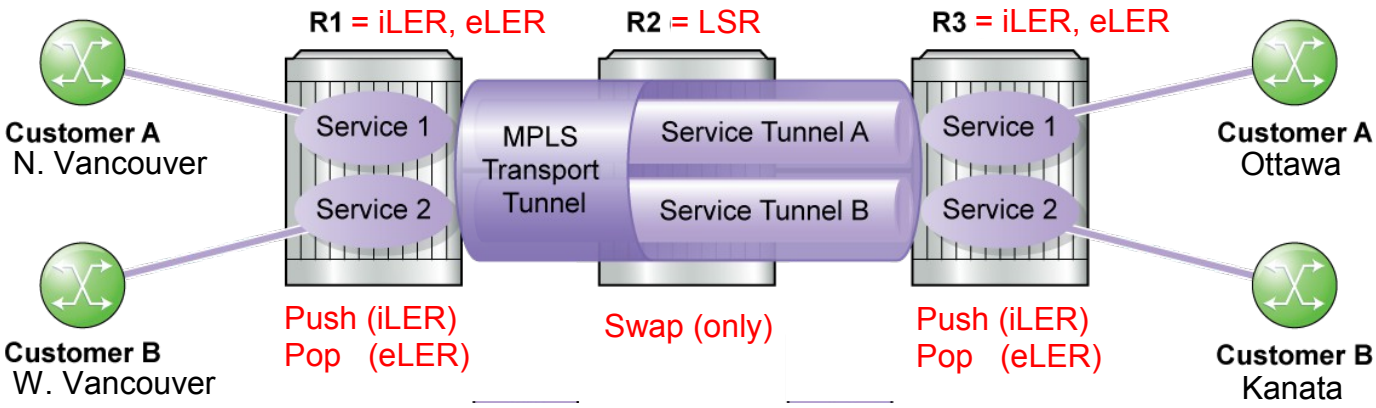

Datagram 1 Datagram 2

This layer exists **only** if L2 VPN is assumed (VPWS, VPLS)

▲ payload - Customer
♣ IP Header - Customer Src IP, DST IP; TTL = XXX (OR TTL = **-1** if L3 VPN)
♫ L2 Framing - Customer
Service Label = **S1**
Transport Label = **T1**
L2 framing – Provider (different hop-by-hop)

▲ payload - Customer
♣ IP Header - Customer Src IP, DST IP; TTL = XXX (OR TTL = **-1** if L3 VPN)
♫ L2 Framing - Customer
Service Label = **S1**
Transport Label = **T2**
L2 framing – Provider (different hop-by-hop)

11. This question focuses on the datagrams and potential changes to them as a VPN payload.

Demonstrate your knowledge of at least two VPN services (L1 or L2, and L3) by sketching three datagrams, one at: (1) ingress, (2) in transit, and (3) egress of the service.

- For all **three** diagrams in each set, you should include a "layer" for: (i) customer L2 framing; (ii) customer L3 header; (iii) customer data
- For the transit datagram, you must include a properly identified label stack, in the correct layer relative to the customer layers.
- Markings to indicate when layers are identical [a common symbol, whether a letter, digit, shape, etc] and when they are different [different symbols] as they transit across.

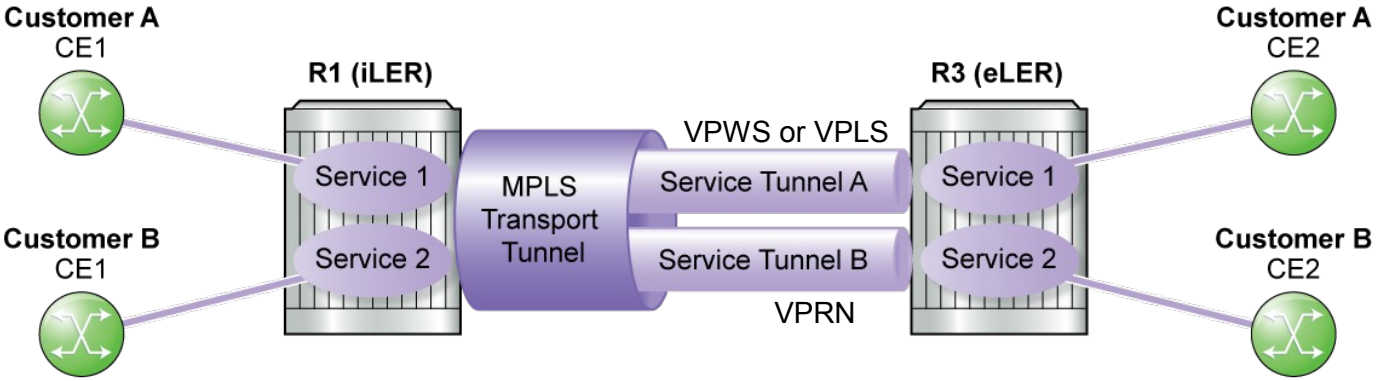

A. [3 marks] Sketch the three datagrams for **Service A**, a *VPWS* or *VPLS*. Remember:
- customer layers
- label stack
- mark layers to indicate when identical, when different

▲ payload - Customer
♣ IP Header - Customer Src IP, DST IP; TTL = XXX
♫ L2 Framing -Customer

Must show layers in correct order, identifying both L2 and L3 at a minimum.

▲ payload - Customer
♣ IP Header - Customer Src IP, DST IP; TTL = XXX
♫ L2 Framing -Customer
Service Label **A**
Transport Label **X**
L2 framing – Provider (different hop-by-hop)

▲ payload - Customer
♣ IP Header - Customer Src IP, DST IP; TTL = XXX
♫ L2 Framing -Customer

Every single bit of customer datagram is delivered unchanged to eLER.

B. [3 marks] Sketch the three datagrams for **Service B**, a *VPRN*. Remember:
- customer layers
- label stack
- mark layers to indicate when identical, when different

▲ payload - Customer
♣ IP Header - Customer Src IP, DST IP; **TTL = XXX**
♫ L2 Framing -Customer

▲ payload - Customer
♥ IP Header - Customer Src IP, DST IP; **TTL = -1**
Service Label **B**
Transport Label **X**
• L2 framing – Provider (different hop-by-hop)

Different Service, same Transport!

▲ payload - Customer
♦ IP Header - Customer Src IP, DST IP; **TTL = -2**
♪ New L2 Framing – Cust.

Customer framing is stripped & replaced; TTL is -2 (in pipe mode, possibly more in uniform mode)

12. [2 marks] Consider the lab work you've done with LDP-shortcuts. Carefully study the additional results presented below from testing LDP-shortcuts on the lab topology.

```
*A:R2# configure router ldp no shortcut-local-ttl-propagate
*A:R2# traceroute 10.10.10.10
traceroute to 10.10.10.10, 30 hops max, 40 byte packets
 1  10.10.10.10 (10.10.10.10)    2.24 ms  2.22 ms  2.14 ms

*A:R2# show router arp
===============================================================
ARP Table (Router: Base)
===============================================================
IP Address      MAC Address       Expiry    Type   Interface
---------------------------------------------------------------
10.10.10.2      02:05:ff:00:00:00 00h00m00s Oth    system
10.2.6.2        52:54:00:a9:28:fe 00h00m00s Oth[I] first
10.2.6.6        52:54:00:a4:83:c9 03h55m33s Dyn[I] first
---------------------------------------------------------------
No. of ARP Entries: 3
===============================================================

*A:R10# show router arp
===============================================================
ARP Table (Router: Base)
===============================================================
IP Address      MAC Address       Expiry    Type   Interface
---------------------------------------------------------------
10.10.10.10     02:0d:ff:00:00:00 00h00m00s Oth    system
10.6.10.6       52:54:00:97:b7:29 03h56m34s Dyn[I] second
10.6.10.10      52:54:00:b9:26:12 00h00m00s Oth[I] second
---------------------------------------------------------------
No. of ARP Entries: 3
===============================================================
*A:R10#
```

The Question: **Clearly** explain exactly what part(s) of the traffic get transported over an LDP-shortcut. You may choose freely between using words, a simple diagram(s), or a combination of the two. Be sure to reference **all** layers of the traffic, including L2, L3, and above. If you feel there's any ambiguity, make reasonable assumptions, **write those assumptions down**, and then continue answering the question.

Neither ARP table shows the system interface from the other end, so the original L2 framing is *not* transported via the LDP-shortcut. We can safely assume that LDP-shortcuts only transport L3 and above.

(No assumptions are necessary for answering this question.)

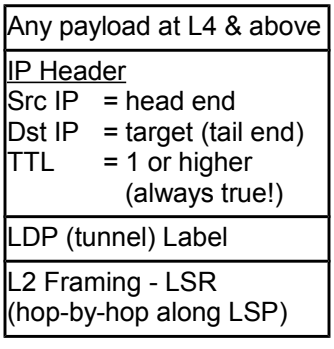

13. [5 marks; 1 each part] Carefully study the topology and label information presented here.

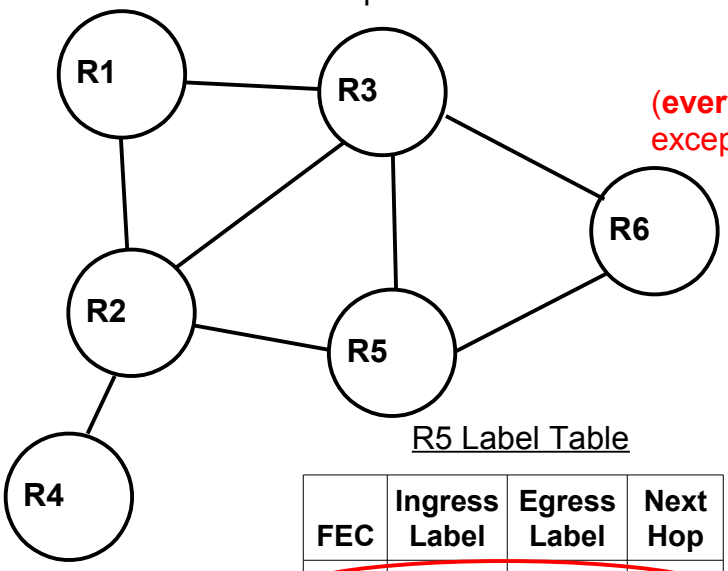

R5 Label Table

FEC	Ingress Label	Egress Label	Next Hop
X	–	220	R6
Y	–	130	R3
Z	–	160	R6
Q	–	330	R2

Push (all 4 rows)

Ingress label	Egress label	Next Hop
R1		
240	350 **invalid**	R2 **FEC A**
110	140	R3 **FEC Q**
330	POP	IP **FEC Y**
140	130	R2 (**FEC Z**)
R2		
170	330	R1 **FEC Y**
140	POP	IP **FEC B**
130	150	R4 **FEC Z**
330	110	R1 **FEC Q**
R3		
160	130	R2 **FEC Z**
110	POP	IP **FEC C**
220	POP	IP **FEC X**
130	170	R2 **FEC Y**
200	240	R1 **FEC A**
140	200	R6 **FEC Q**
R4		
150	POP	IP **FEC Z**
210	POP	IP **FEC D**
160	200	R3
R6		
170	160	R5
220	220	R3 **FEC X**
350	350	R5
200	POP	IP **FEC Q**
160	160	R3 **FEC Z**

Swap (**every** row except Pop)

Impossible (no link to R3!)

(Part **D**)

(Part **D**)

A. Pop is shown several times in the tables. **Clearly** indicate several examples of where each of (i) **push,** and (ii) **swap** would be in the tables.

B. How many labels did R1 distribute **to** neighbours?

four (240, 110, 330, 140)

C. Ignoring R3 (which is clearly very busy), which router is advertising (i.e. origin for) the most FECs? (Not to be confused with originating the most LSPs!)

R4 – it has **two** POPs (eLER for two FECs)

D. The diagram does *not* show a complete LIB for R5. For how many LSPs is it an LSR or eLER? **Clearly** explain your answer.

Two: on R6, we see two labels given out by R5, 160 (first row) and 350 (third row). R6 will use these when sending datagrams to R5 (either for swap as an LSR, or pop as eLER)

E. **[Bonus]** The numbers used were simply for teaching purposes BUT if they were real, what is the proper name for the label range they are taken from? [Ref. Slide 2-11]

The label range for "Static LSPs" (32 – 1, 023)

14. [3 marks; 1 per attribute] **Clearly** identify and (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

15. Professor Anderson is testing MPLS networks, using experimental Linux software to imitate Nokia's MPLS functionality. Carefully examine the data on the previous page. It is a snapshot in time of the LIBs from 6 routers. Then answer the following questions:

A. [2 marks] The very last line for R6 shows an identical value in both label columns. Is this value identical *randomly* or is it *necessarily* identical? **Clearly** explain and justify your answer.
[1 mark for correct answer; 1 mark for justification]
Matching ingress and egress label values are completely *random*; the ingress label is *handed out* by a router (i.e. R6) and the egress label is *received* by that router (R6)
Only in the case of two labels being handed out *by the same router* are they necessarily identical.

B. [2 marks; Tough Question] We have no information about *when* these labels were *distributed* but we can still consider the other two attributes of control and retention. **Clearly** identify any *evidence* that indicates what attributes are revealed by the label tables.

For any marks, **must** correctly use evidence from the label table to support the attribute!
1. Generally, the tables support Ordered control and Conservative retention.
2. Very first row in R1: a sign of "liberal retention" (hanging onto an unused label)
   Egress label 350-to-R2 remains even though R2 no longer uses that label.
   (I will accept an argument that this might be a sign of Independent Mode, but less likely since R1 is offering label 240 only **because** it (still) has label 350 to R2).
3. Note that last row for R4 is handled in part C below)

C. [2 marks; Tough **Bonus** ; skip it if time is running out?] At least one label is invalid because it's completely impossible. Which one(s)? **Clearly** explain why?

The row "160 200 R3" on R4 because R3 is **not** a (direct) next-hop from R4!
(This one needs some time & careful attention to detail, thus the 2 mark weight)

16. A. [0 marks] Yes/No: Is Professor Anderson pleased, even delighted, with the work that's being done in labs?
    B. [1 mark] Justify your answer by quoting written proof!
    "The work (and learning!) that's occurring in labs is terrific! Keep up the good work!"
    If you **read the instructions** on the cover page, the answer is in item #3!
    Could this be considered an example of TL;DR = F? ...

Extra Work