



CCIE Routing and Switching v4.0 Troubleshooting Practice Labs

Martin J. Duggan

Cisco Press



CCIE Routing and Switching v4.0 Troubleshooting Practice Labs

Martin Duggan

ciscopress.com

Table of Contents

Chapter 1: Troubleshooting Lab 1 (The Warm-Up).....	11
Chapter 2: Troubleshooting Lab 2 (Network Down!).....	82
Chapter 3: Summary.....	122
Appendix A Practice Lab 1 Router Configurations.....	126
Appendix B Practice Lab 1 Switch Configurations.....	174
Appendix C Practice Lab 2 Router Configurations.....	211
Appendix D Practice Lab 2 Switch Configurations.....	240

About the Author

Martin James Duggan, CCIE No. 7942, is a network architect for AT&T. He designs network solutions for global customers and specializes in data center networking and quality of service. Martin mentors colleagues through their Cisco qualifications and holds regular internal training classes. Previous to this, Martin was a network architect for IBM performing IP network de-signs and global network reviews. Martin has been in the industry for 22 years, focusing on Cisco solutions for the past 12 years. He is the author of the Cisco Press book *CCIE Routing and Switching v4.0 Configuration Practice Labs* and the co-author of *Routing and Switching Practice Labs Volume 1*.

About the Technical Reviewers

Bruno van de Werve, CCIE No. 20066 (Routing and Switching), currently works for the CCIE department in Learning@Cisco as a product manager. Bruno is responsible for managing the content development of the CCIE Routing and Switching written and lab exams. In 2005, he joined Cisco as a Technical Assistance Center engineer supporting architectural issues with classic IOS platforms. After that, Bruno served as a Brussels proctor for more than 2 years.

Luc De Ghein, CCIE No. 1897, has 17 years of experience in networking and is currently employed by Cisco. He is a technical leader in the Technical Assistance Center, focusing on IP routing protocols and Multiprotocol Label Switching (MPLS) technologies. He provides escalation support to Cisco engineers worldwide and teaches others about IP routing, IP multicast, and MPLS technologies.

Dedication

I want to dedicate this publication to my children, Anna and James. You two are my *raison d'être*. The time we get to spend together is extremely valuable to me and will never be enough. I am so proud of you. Anna; your cuddles are priceless. And James, please believe me when I say I really am trying my hardest at tennis; I might hold back a little at Call of Duty, though.

I also want to add a dedication to my lovely Lotte. I know this is not quite as glamorous as your recent TV appearance, but once again you are singled out for your enviable qualities. You make me laugh; you make me happy; you even make me pedal harder. Time spent with you is blissful, and I am extremely privileged to have you in my life.

Mum and Dad, as always, thanks for being there.

Acknowledgments

I want to thank Brett Bartow for again providing me with an opportunity to write for Cisco Press. I must be one of the worst authors in terms of meeting deadlines, so Brett, thank you for your patience and understanding.

To Bruno and Luc, who reviewed the material that comprises this publication, thank you for your suggestions and input. I really value your contributions.


To the guys on the account I work with in Portsmouth: Neil, Ian, Matt, and Andy. Thanks for use of the lab. You are a pleasure to work with, you run the account well, and I know it is always going to be interesting when you need a second opinion on something. I still think about our groundbreaking “Chuckle Brothers” solution to that problem and really grin: from me, to you, to me, to you! We should have implemented it purely for comedy value. Oh, and Richard, thanks for the pizzas while we fixed that issue; that was quality project management.

And thanks to the instructors at Lasham Gliding School. You freely give your time and experience to prevent newbies like me from seeing what it feels like to bail out and pull that enticing rip cord. Thanks for giving me my wings and letting me loose.

To Chris Barney, my CCIE counterpart and Networkers buddy, thanks for standing in for me and keeping me up-to-date about new technology.

To my cricket team (Whitchurch, Hampshire), who still for some unknown and unjustified reason have not put me in as the opening batsman, thanks for having me (and there are plenty more sixes to come).

Icons Used in this Book



Router



Bridge




Hub



DSU/CSU

Catalyst
SwitchMultilayer
SwitchATM
SwitchISDN/Frame Relay
SwitchCommunication
Server

Gateway



Access Server



Token Ring




FDDI



Network Cloud

Frame Relay
Virtual Circuit

Line: Ethernet



Line: Switched Serial



Line: Serial

Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

Boldface indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a show command).

Italics indicate arguments for which you supply actual values.

Vertical bars (|) separate alternative, mutually exclusive elements.

Square brackets [] indicate optional elements.

Braces { } indicate a required choice.

Braces within brackets [{ }] indicate a required choice within an optional element.

Introduction

For more than 10 years, the CCIE program has identified networking professionals with the highest level of expertise. Less than 3 percent of all Cisco certified professionals actually achieve CCIE status. The majority of candidates who take the exam fail at the first attempt because they are not fully prepared. They generally find to their own cost that their study plan did not match what was expected of them in the exam. The troubleshooting labs within this publication have been designed to take you as close as possible to actually taking the Troubleshooting section within the real CCIE lab exam.

Exam Overview

The CCIE qualification consists of two exams: a 2-hour written exam followed by an 8-hour hands-on lab exam that includes a Troubleshooting section. Written exams are computer-based, multiple-choice exams lasting 2 hours and available at hundreds of authorized testing centers worldwide. The written exam is designed to test your theoretical knowledge to ensure you are ready to take the lab exam, and so you are eligible to schedule the lab exam only after you have passed the written exam. Because you have purchased this publication, it is assumed that you have passed the written exam and are ready to practice for the lab exam. The lab exam consists of a 5.5-hour hands-on exam that requires you to configure a series of complex scenarios in strict accordance with the questions. You need to secure a minimum score of 80 percent here; and although it is tough, it is achievable. The lab exam also has a Troubleshooting section that lasts 2 hours, and you also face a series of more questions for a 30-minute period of the exam. Again, you need to secure a minimum score of 80 percent here to be successful. You can find current lab blueprint content information at <https://learningnetwork.cisco.com/docs/DOC-4603>.

Study Roadmap

Taking the lab exam is all about experience. You cannot expect to take it and pass after just completing your written exam and relying on just your theoretical knowledge. You must spend countless hours of rack time configuring features and learning how protocols interact with one another. To be confident enough to schedule your lab exam, review the following points.

Assessing Your Strengths

Using the content blueprint, determine your experience and knowledge in the major topic areas. For areas of strength, practicing for speed should be your focus. For weak areas, you might need training or book study in addition to practice.

Study Materials

Choose lab materials that provide configuration examples and take a hands-on approach. Look for material that is approved or provided by Cisco and its Learning Partners.

Hands-On Practice

Build and practice your lab scenarios on a per-topic basis. Go beyond the basics and practice additional features. Learn the show and debug commands along with each topic. If a protocol has multiple ways of configuring a feature, practice all of them.

Cisco Documentation CD

Make sure you can navigate the Cisco documentation CD with confidence because this is the only resource you are allowed during the lab (or restricted access to the same content on Cisco.com). Make the CD part of your regular study; if you are familiar with it, you can save time during the exam.

Home Labs

Although acquiring a personal home lab is ideal, it can be costly to gather all the equipment you may need.

Cisco 360 Program

The Cisco 360 Learning Program encompasses six stages of activity to support successful learning for students.

- **Assessment:** Students take a diagnostic pre-assessment lab to benchmark their knowledge of various networking topics.
- **Planning:** Based on the pre-assessment, students create a learning plan that uses a mix of learning components to focus their study.
- **Learning:** Students learn by participating in lessons and lectures, reading materials, and working with peers and instructors.
- **Practice:** Students use the practice exercises to apply learning on actual network equipment.
- **Mastery:** Students measure their understanding by completing assessments of knowledge and skill for various approaches to solving network problems.
- **Review:** Students review their work with a mentor or instructor and tune their skills with tips and best practices.

You can find detailed information about the 360 program at https://learningnetwork.cisco.com/community/learning_center/cisco_360/360-rs.

Equipment List and IOS Requirements

The lab exam tests any feature that can be configured on the equipment and the IOS versions indicated here:

- 1841 series routers, IOS 12.4(T)–Advanced Enterprise Services
- 3825 series routers, IOS 12.4(T)–Advanced Enterprise Services
- Catalyst 3560 series switches running IOS Version 12.2S–Advanced IP Services

Troubleshooting Information and Technique

During the Troubleshooting section of the exam, you are presented with a series of trouble tickets (incidents) for preconfigured networks and need to diagnose and resolve the network faults. The section lasts for 2 hours, with a visible countdown and a 15-minute warning before the end. The questions are symptom-based, and the verification will be result based with any stipulated con-straints. As with the Configuration section, the network must be up and running to receive credit. If you finish the Troubleshooting section early, you can proceed to the Configuration section, but you will then not be permitted to go back to the Troubleshooting section.

Be aware that the routers and switches within the Troubleshooting section of the CCIE lab are virtual devices; therefore, you will always see interfaces shown as up/up rather than up/down.

Troubleshooting Preparation

- Build a rock-solid knowledge base by studying the theory and completing as many practice labs as possible.
- Create your own troubleshooting strategy. For instance, you might start from Layer 1 and work up; start at Layer 7 and work down; use the split-half method; and so on. If you have a well-tested strategy, you will find tackling the incidents much easier than if you just follow a gung-ho approach and should be able to troubleshoot any scenario presented.

- Spend time with your colleagues or like-minded friends; have them break or introduce configuration errors into a topology you are not familiar with for you to troubleshoot. Each time you complete a practice lab, ask someone to completely sabotage the configurations and see how long it takes you to fix it, all the time enhancing your strategy and speed.

Troubleshooting Approach

- Define the actual problem from the information provided within the incident. Ensure you have a clear and concise problem statement that you can translate to a specific network issue.
- Identify the symptoms. This is where you must call on your hard-earned knowledge and experience. You may be able to identify immediately why a neighbor relationship is failing. If you cannot, use the tools available to you in the form of show and debug commands. When you can clearly see the symptom, the solution should be apparent.
- If you have multiple solutions, work through them according to your strategy and use the full cycle of implement, verify, and investigate to find the correct solution.
- You may find you have strict guidelines related to an incident. These could be easy fixes that you should not apply; they might mask the issue and not actually address the root cause. Stick to the guidelines; otherwise, you risk not earning points for your solution.

Chapter 1

Troubleshooting Lab 1 (The Warm-Up)

The CCIE exam commences with 2 hours of troubleshooting “tickets” followed by 5.5 hours of configuration and then a final 30 minutes of additional questions. This troubleshooting lab has been timed to last for 2 hours and be representative of the format and difficulty of the tasks you are likely to encounter on your CCIE lab exam.

Ideally, you should work on this exercise for 2 hours and score yourself at that point. Of course, you may continue until you believe you have met all the objectives, but you will need to increase your speed as you near your lab exam.

You now are going to be guided through the equipment requirements and pre-lab tasks in preparation for taking this troubleshooting lab.

If you do not own seven routers and three switches (which have been used to create this troubleshooting lab), consider using the equipment available and additional lab exercises and training facilities available within the CCIE R&S 360 program. You can find detailed information about the 360 program and CCIE R&S exam at the following websites, respectively:

- https://learningnetwork.cisco.com/community/learning_center/cisco_360/360-rs
- https://learningnetwork.cisco.com/community/certifications/ccie_routing_switching

NOTE

The 3825s used in this lab were loaded with c3825-adventerprisek9-mz.124-6.T.bin.

You will likely encounter many more devices within your real CCIE Troubleshooting section of the lab.

Equipment List

For this troubleshooting lab, you need the following hardware and software components:

- Seven routers loaded with Cisco IOS Software Release 12.4 Advanced Enterprise image and the minimum interface configuration, as documented in Table 1-1.

NOTE

The 3560s in this lab were loaded with c3560-ipservicesk9-mz.122-25. SEE.bin.

NOTE

The initial configurations supplied should be used to preconfigure your routers and switches before the lab starts.

If your routers have different interface speeds than those used within this book, adjust the **bandwidth** statements on the relevant interfaces to keep all interface speeds in line. Doing so ensures that you do not get unwanted behavior due to differing Interior Gateway Protocol (IGP) metrics.

Consider asking a like-minded colleague or friend to load the configurations onto your equipment on your behalf; this will ensure that you do not spot any of the potential configuration issues before beginning the exercise.

Table 1-1 Hardware Required per Router

Router	Model	Ethernet I/F	Serial I/F
R1	3825	1	1
R2	3825	1	1
R3	3825	1	1
R4	3825	1	2
R5	3825	—	2
R6	3825	—	2
R7	3825	1	2

Three 3560 switches with IOS 12.2 IP Services.

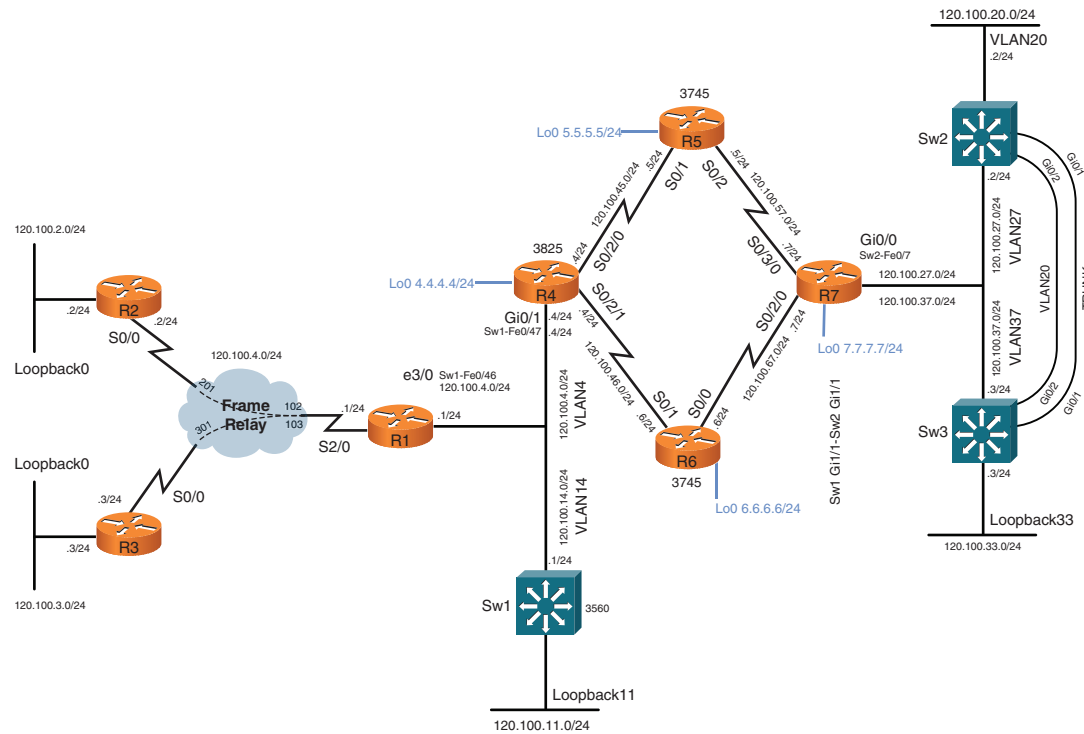
Setting Up the Lab

Feel free to use any combination of routers and switches as long as you fulfill the requirements within the topology diagram, as shown in Figure 1-1. However, it is recommended to use the same model of equipment because doing so will make life easier should you load configurations directly from the supplied configurations into your own devices.

Lab Topology

This troubleshooting exercise uses the topology outlined in Figure 1-1, which you must re-create with your own equipment or by using lab equipment on the CCIE R&S 360 program.

FIGURE 1-1
Lab 1 Topology
Diagram



NOTE

R2, R3, R5, and R6 do not require connectivity to any switches within this troubleshooting lab. R2 and R3 use loopback interfaces to simulate user-connected interfaces and use serial interfaces to connect with the Frame Relay hub router R1. R5 and R6 are P routers within the Multiprotocol Label Switching (MPLS) network and connect to neighboring routers with serial interfaces only.

Switch Instructions

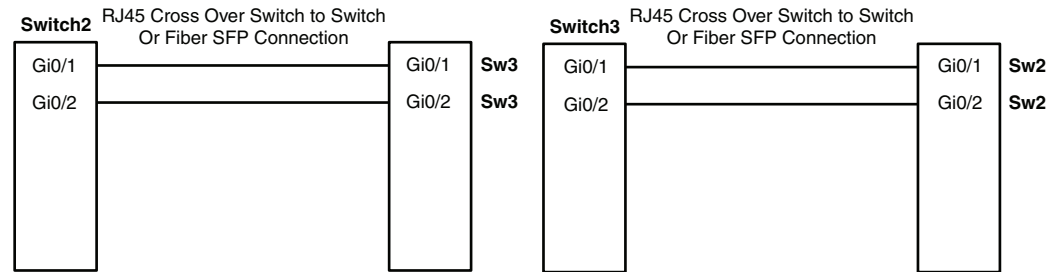
Configure VLAN assignments from the configurations supplied or from Table 1-2.

Table 1-2 VLAN Assignment

VLAN	Switch 1	Switch 2	Switch 3
4	Fa0/46, Fa0/47	—	—
14	Fa0/47	—	—
20	—	Gi0/1, Gi0/2	Gi0/1, Gi0/2
27	—	Gi0/1, Fa0/7	—
37	—	Gi0/1	Gi0/1

Connect your switches together with fiber small form-factor pluggable (SFP) connectors or RJ045 Ethernet crossover cables, as shown in Figure 1-2.

FIGURE 1-2
Switch-to-Switch
Connectivity




Frame Relay Instructions

Configure one of the routers you are going to use in the troubleshooting lab as a Frame Relay switch, or have a dedicated router assigned solely for this task. This troubleshooting lab uses a dedicated router for the Frame Relay switch. A fully meshed environment is configured between all the Frame Relay routers. Pay attention in the lab as to which permanent virtual circuits (PVC) are actually required.

If you are using your own equipment, keep the DCE cables at the frame switch end for simplicity and provide a clock rate to all links from this end.

The Frame Relay connectivity will, after configuration, represent the logical Frame Relay network, as shown in Figure 1-3.

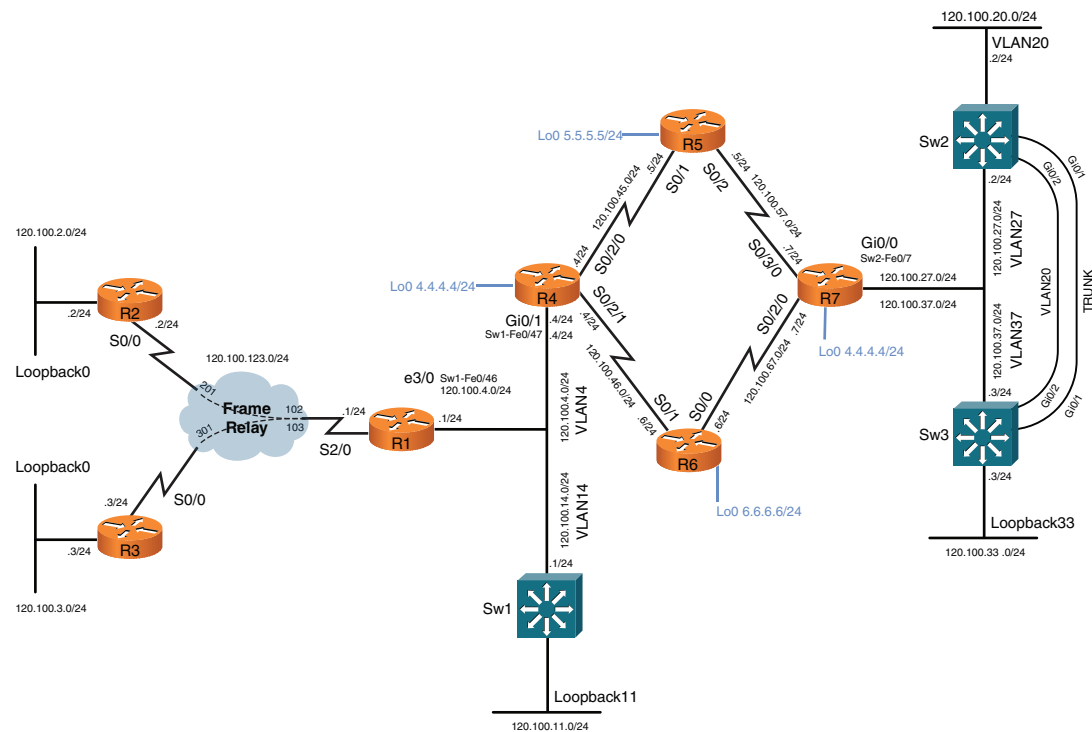
FIGURE 1-3
Frame Relay Logical
Connectivity



IP Address Instructions

For this exercise, you are required to configure your IP addresses, as shown in Figure 1-4, or load the initial router configurations supplied.

FIGURE 1-4
IP Addressing
Diagram



Pre-Lab Tasks

- Build the lab topology as per Figure 1-1 and Figure 1-2.
- Configure your Frame Relay switch router to provide the necessary data link control identifiers (DLCI) as per Figure 1-3.
- Configure the IP addresses on each router, as shown in Figure 1-4. Alternatively, you can load the initial configuration files supplied if your equipment is compatible with that used to create this exercise.

General Guidelines

- Do not configure any static/default routes unless otherwise specified.
- Use only the DLCIs provided in the appropriate figures.
- Your own routing metrics do not need to match exactly those shown in the incidents (unless specifically indicated).
- Tackle questions sequentially. You might find that one trouble ticket needs to be resolved before moving on to the next ticket. (This might not be the case in your real lab exam, however, which will have a much higher number of devices.)
- Get into a comfortable and quiet environment where you can focus for the next 2 hours.
- The incident questions list symptoms, explicit validation tests to confirm you have rectified the incident correctly, and any optional constraints ensure you follow these items correctly to maximize your score.
- When resolving an incident, do not remove any configured feature. You must resolve the misconfiguration instead of removing a whole configuration. (The only exception to this rule is when there is no other choice than removing the faulty configuration to resolve the incident.)
- Have available a Cisco Documentation CD-ROM or access the latest documentation online from <http://www.cisco.com/cisco/web/psa/default.html?mode=prod>.

NOTE

Access just the URL listed here, not the whole Cisco.com website, because if you are permitted to use documentation during your CCIE lab exam, it will be restricted to a limited search function. A well-prepared candidate should not risk losing time during the exam searching for information.

Troubleshooting Lab 1

You now answer questions about network topology and the VPN diagram, as shown in Figures 1-5 and 1-6.

FIGURE 1-5
Lab Topology
Diagram

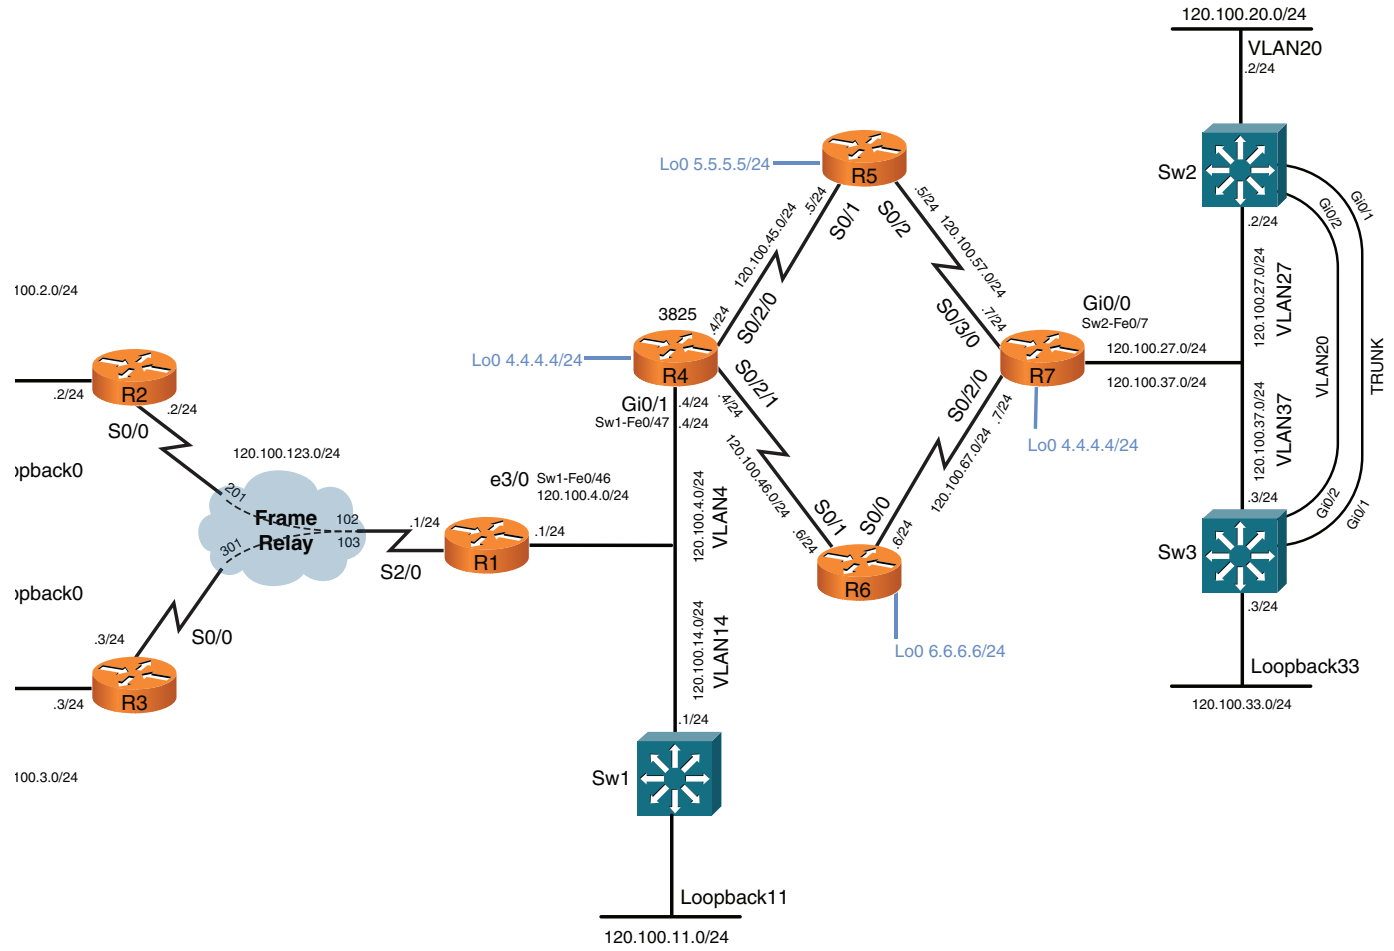
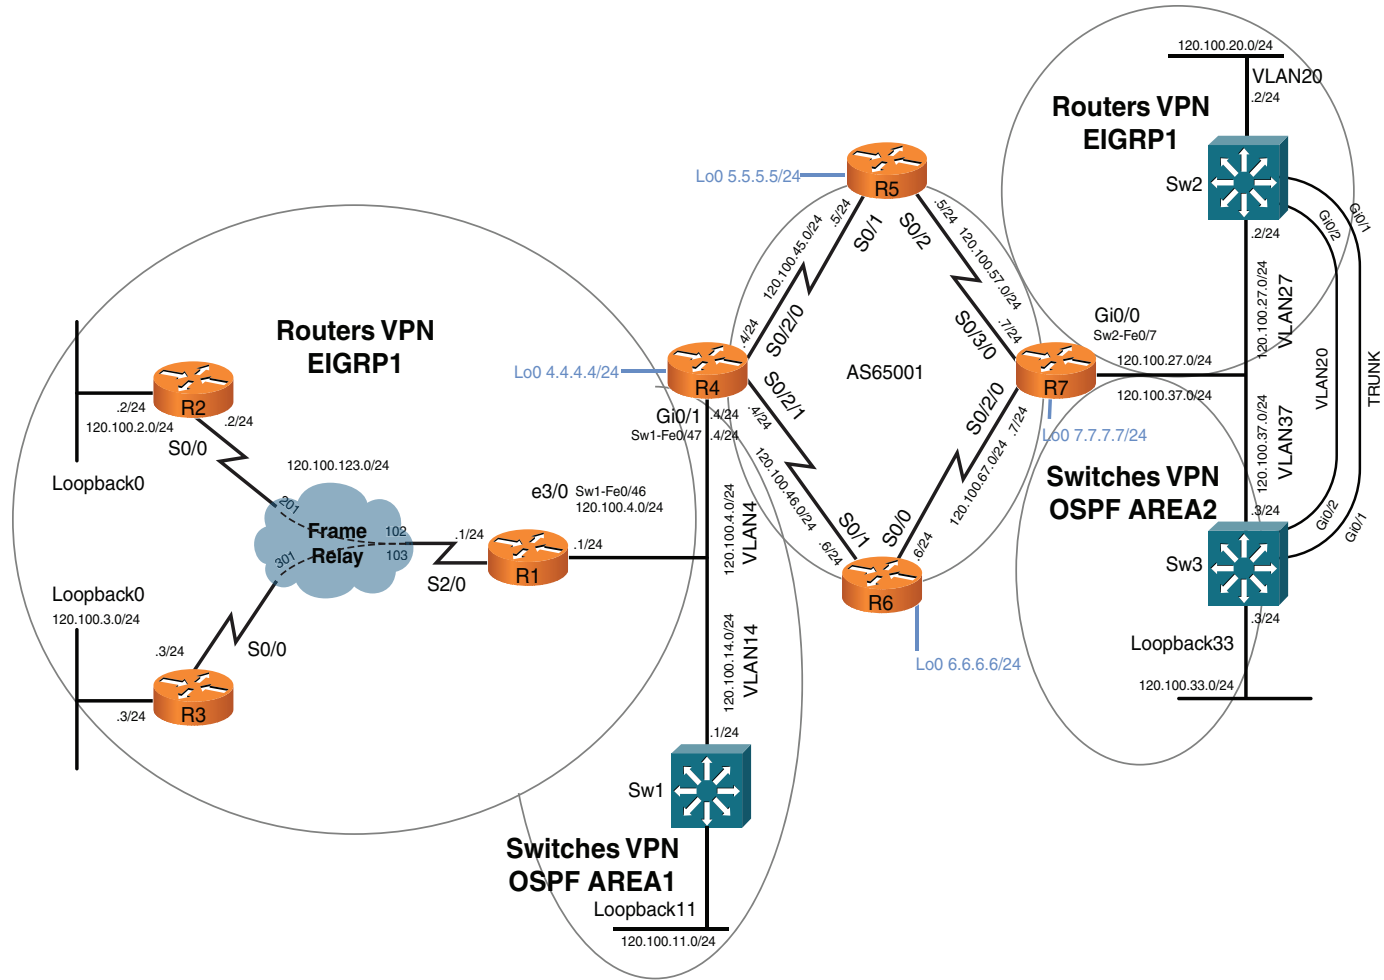



FIGURE 1-6
VPN Diagram



Incident 1

FIGURE 1-7
Incident 1 Diagram



Users on Sw3 VLAN 20 are complaining that they experience poor connectivity to VLAN 20 on Sw2. A variable response is seen when pinging from Sw3 VLAN 20 to Sw2 VLAN 20 interface.

Investigate the issue and rectify it. Confirm by proving a successful ping from Sw3 VLAN 20 interface 120.100.20.3 to Sw2 VLAN 20 interface 120.100.20.2 with a stable response time, as follows:

```
Sw3# ping 120.100.27.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

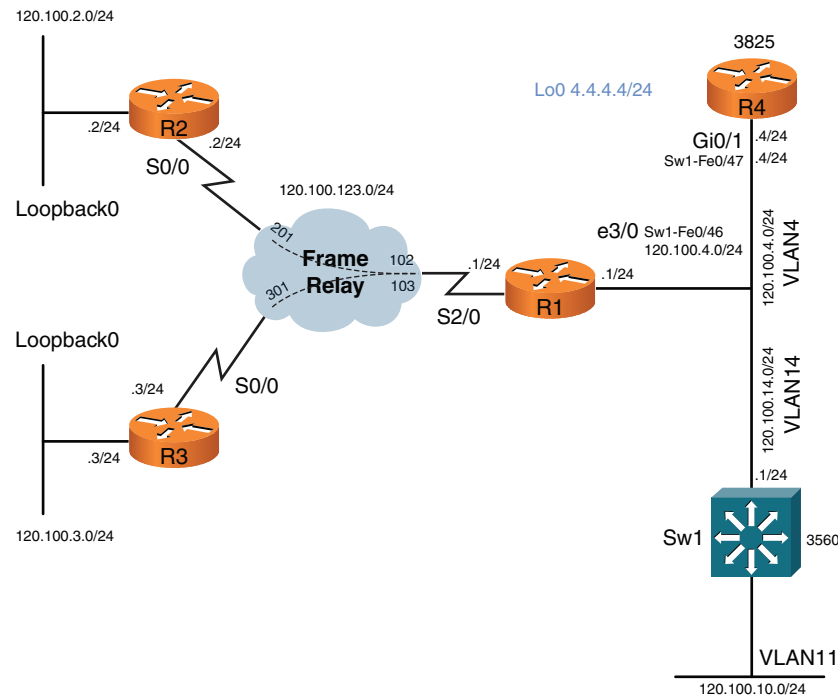
```
Sw3# ping 120.100.27.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
Sw3# ping 120.100.27.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

3 points

Incident 2

FIGURE 1-8
Incident 2 Diagram



Users on R3's 120.100.3.0/24 subnet are complaining that they cannot access resources located on their remote virtual private network (VPN) site on VLAN 20 (switch 2). Initial investigations have led first-line support personnel to think that there is an issue with connectivity up to the provider edge (PE) router R4, which was recently replaced because the subnet 120.100.4.0/24 is not present in R3's routing table.

Investigate the issue and prove connectivity solely to their PE connection (R4's VLAN 4 interface) as follows:

```
R3# ping 120.100.4.4 source 120.100.3.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 120.100.4.4, timeout is 2 seconds:
```

```
Packet sent with a source address of 120.100.3.1
```

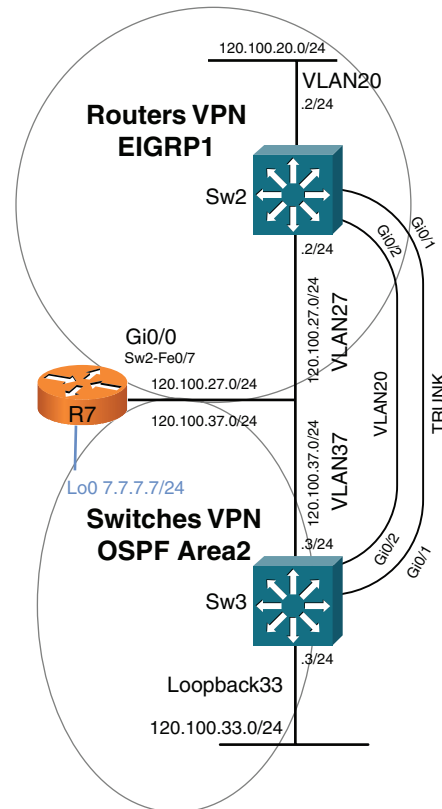
```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

2 points

Incident 3

FIGURE 1-9
Incident 3 Diagram



Users on switch 2 VLAN 20 are complaining that they have no connectivity through their VPN to services located on R2 and R3 (120.100.2.0/24 and 120.100.3.0/24). A traceroute to destination networks shows timeouts immediately, and it appears that an EIGRP neighbor adjacency between Sw2 and R7 is not present.

Investigate the issue by proving connectivity purely between Switch 2 and the local PE router R7. Establishment of an EIGRP adjacency between devices suffices at this point, as follows:

```
Sw2# show ip eigrp neighbors
```

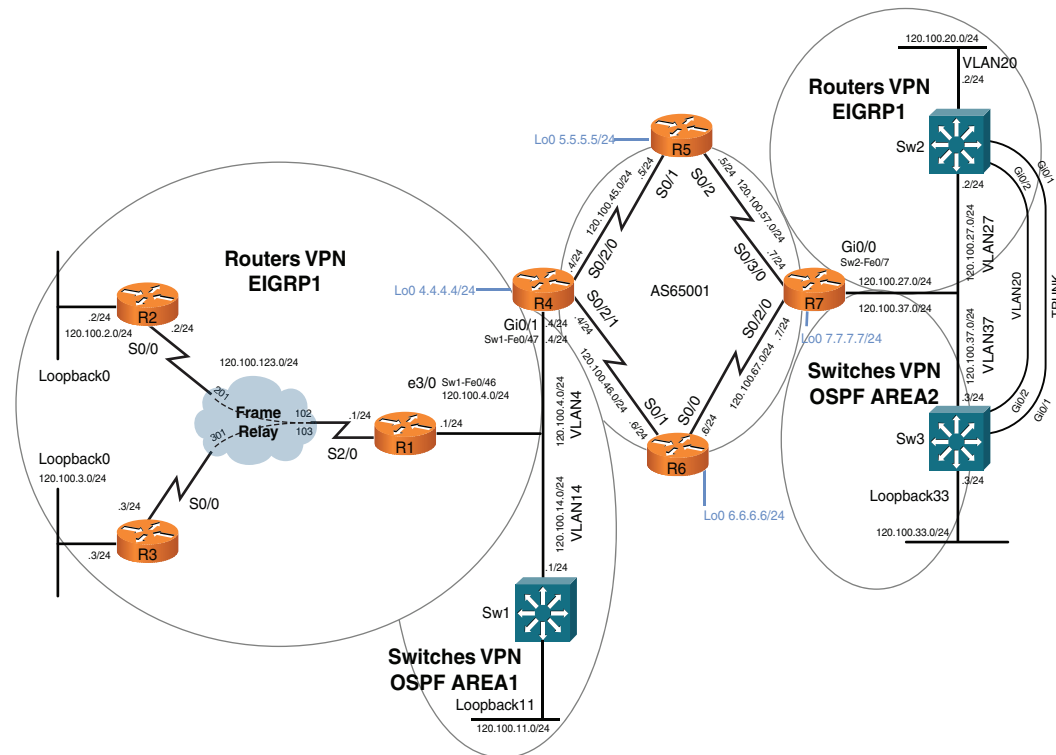
```
EIGRP-IPv4:(1) neighbors for process 1
```

H	Address	Interface	Hold	Uptime	SRTT	RTO	Q	Seq
0	120.100.27.7	Vlan27	12	00:01:11	1	300	0	92

3 points

Incident 4

FIGURE 1-10
Incident 4 Diagram



Users on R3's 120.100.3.0/24 subnet are still complaining that they cannot access resources located on their remote VPN site on VLAN 20 (switch 2). It appears that the MPLS VPN is not functioning correctly.

Investigate the issue and prove connectivity between PE routers via extended ping from R4 VLAN 4 interface and R7 VLAN 27 interface to prove MPLS functionality, as follows:

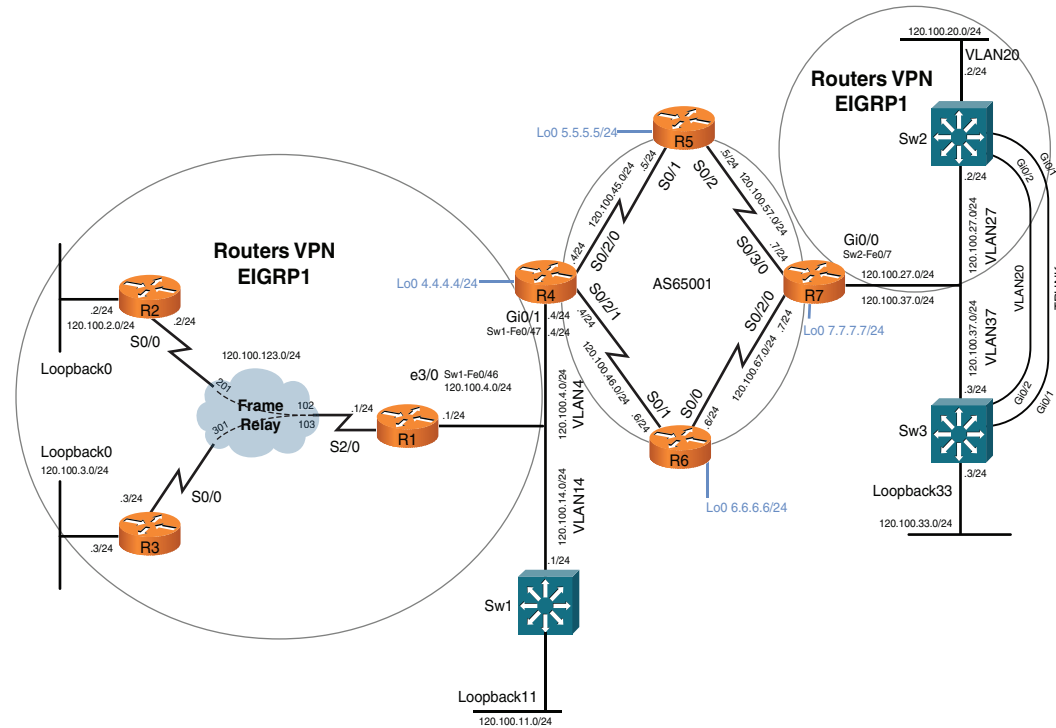
```
R4# ping vrf ROUTERS 120.100.27.7 source 120.100.4.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.27.7, timeout is 2 seconds:
Packet sent with a source address of 120.100.4.4
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

```
R7# ping vrf ROUTERS 120.100.4.4 source 120.100.27.7
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.4.4, timeout is 2 seconds:
Packet sent with a source address of 120.100.27.7
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

6 points

Incident 5

FIGURE 1-11
Incident 5 Diagram



Now that the MPLS issues appear to have been fixed, users on vrf ROUTERS on R2 and R3 are complaining that they cannot see the route to VLAN 27.

Investigate and prove connectivity with a valid route and successful ping from R2 and R3 user subnets to the Sw2 VLAN 27 interface as follows:

```
R2# ping 120.100.27.2 source 120.100.2.2
Type escape sequence to abort
Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:
Packet sent with a source address of 120.100.2.2
!!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 28/31/36 ms

```
R3# ping 120.100.27.2 source 120.100.3.3
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:

Packet sent with a source address of 120.100.3.3

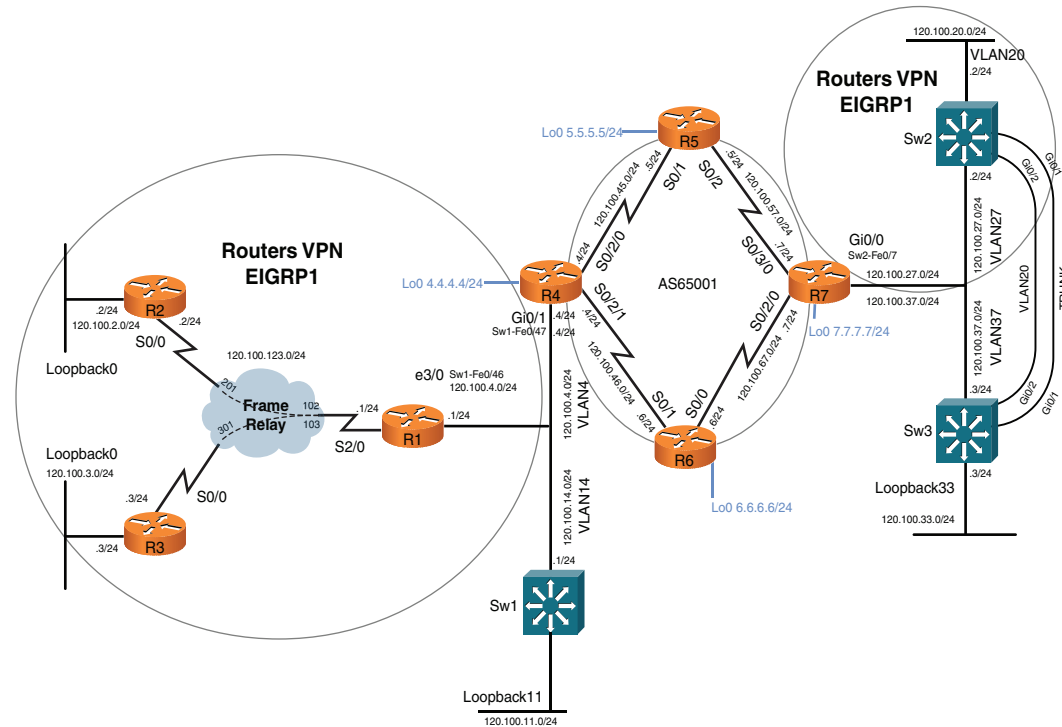
!!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 32/32/36 ms

3 points

Incident 6

FIGURE 1-12
Incident 6 Diagram



Users on vrf ROUTERS on R2 and R3 are complaining they can see only the route to VLAN 27 on Sw2 and not VLAN 20. Investigate the issue and rectify it.

Prove connectivity by extended ping from user subnets on R2 and R3 to Sw2 VLAN 20 as follows:

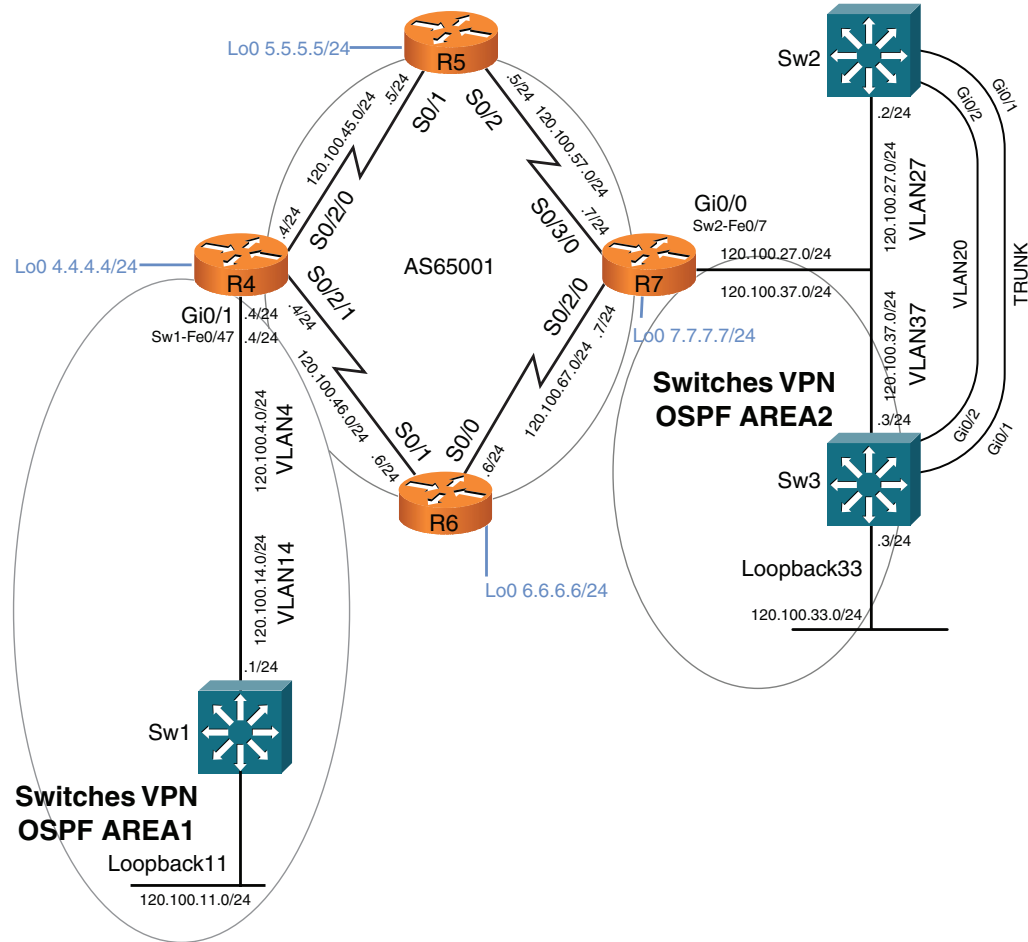
```
R2# ping 120.100.20.2 source 120.100.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.20.2, timeout is 2 seconds:
Packet sent with a source address of 120.100.2.2
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/32/36 ms
```

```
R3# ping 120.100.20.2 source 120.100.3.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.20.2, timeout is 2 seconds:
Packet sent with a source address of 120.100.3.3
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/32/32 ms
```

2 points

Incident 7

FIGURE 1-13
Incident 7 Diagram



Users on Sw3 VLAN30 have reported that they cannot see a route to 120.100.13.0/24 used to access services advertised behind a firewall with a next-hop address of Sw1's Loopback11 interface.

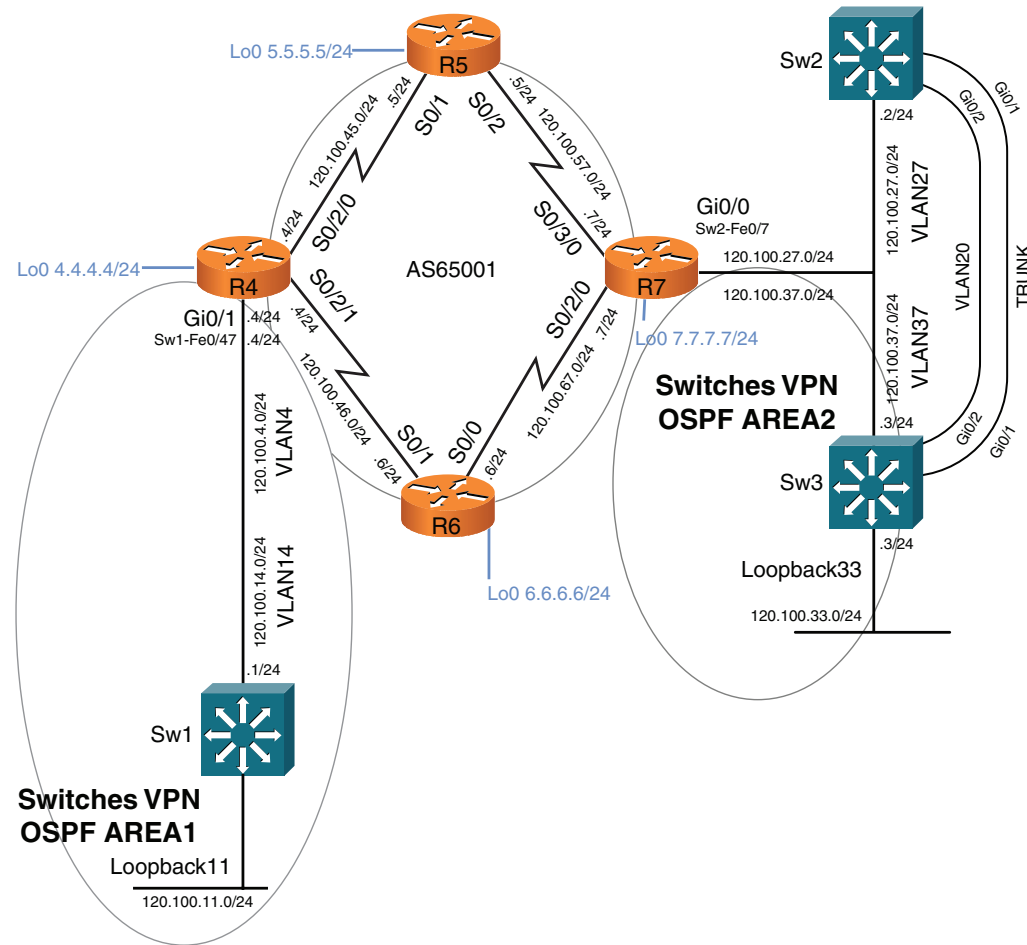
Investigate the issue and ensure that the static route on Sw1 can be seen within Sw3's routing table over OSPF as follows:

```
Sw3# show ip route ospf
      10.0.0.0/26 is subnetted, 2 subnets
O E2   10.20.20.0 [110/20] via 120.100.37.7, 00:20:31, Vlan37
      120.0.0.0/24 is subnetted, 6 subnets
O IA   120.100.12.0 [110/3] via 120.100.37.7, 00:20:36, Vlan37
O E2   120.100.13.0 [110/20] via 120.100.37.7, 00:00:22, Vlan37
O IA   120.100.14.0 [110/2] via 120.100.37.7, 00:20:36, Vlan37
O IA   120.100.11.0 [110/3] via 120.100.37.7, 00:20:36, Vlan37
```

2 points

Incident 8

FIGURE 1-14
Incident 8 Diagram



Users on Sw3 VLAN 37 have reported that they cannot access services hosted on IP address 226.1.1.1 that are hosted on the Sw1 VLAN 14 interface (120.100.14.1). They report that they can no longer even ping this IP address.

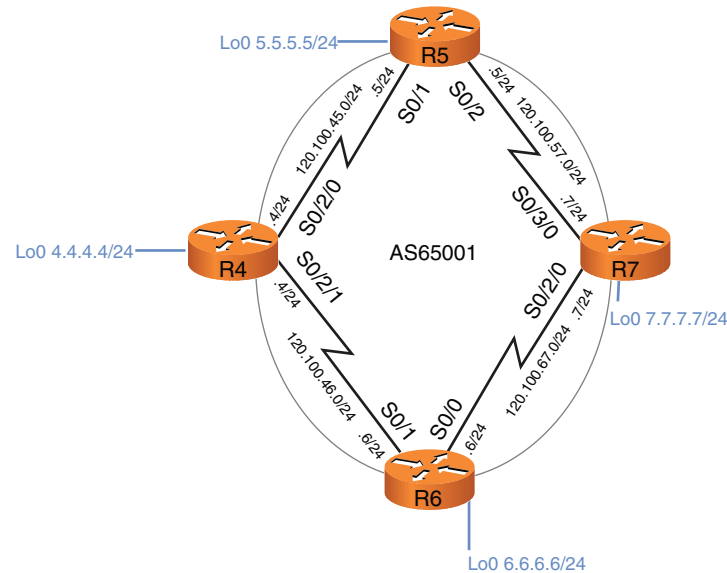
Investigate the issue and perform a successful ping from Sw3 to IP address 226.1.1.1 as follows:

```
Sw3# ping 226.1.1.1
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 226.1.1.1, timeout is 2 seconds
Reply to request 0 from 120.100.14.1, 8 ms
Users have indicated to you that they believe the RP should be PE router R4 (120.100.14.4) and that MDT is used with a group-address of 232.0.0.11.
```

5 points

Incident 9

FIGURE 1-15
Incident 8 Diagram



Your first-line support personnel are complaining that when they telnet to the Loopback0 address of R4 from within the MPLS network that the response is poor. One reported symptom is the screen intermittently locking up.

Investigate the issue with an aim of restoring normal service.

3 points

Ask the Proctor

Section 1

Incident 1

- Q:** I can successfully ping between Sw3 and Sw2 on VLAN 20. I don't believe I have any problems. Do I?
- A:** You should find that your ping response time is variable. This is representative of an underlying issue that you must investigate and rectify. When you receive a constant response time, you can be confident that you have solved the problem.
- Q:** I do not see any neighbor relationships between my PE router and Sw2 and Sw3. Should I be investigating this issue?
- A:** Not yet. You might find that if you solve the first question based on the information provided that connectivity for all the devices on Sw2 and Sw3 will be much more reliable than previously.
- Q:** I am getting log messages on Sw3 that show I have MAC addresses flapping between ports. Is this normal behavior?
- A:** No. This is not normal and is indicative of the problem you are investigating. Think about what this actually means.
- Q:** This looks like a spanning-tree type of issue, but all my cables are good. Do you want me to start isolating ports between Sw2 and Sw3?
- A:** Your cables are likely to be fine, but you can logically disable ports between devices if you want to isolate a potential loop.
- Q:** I have shut down the Gi0/2 interface on Sw3, and everything looks good. Can I leave the topology like this? (After all, I believe I have answered the question.)
- A:** No. This would remove the resiliency between switches for VLAN 20. Find a configuration fix instead.
- Q:** Looking at the configuration I can now see that one port is allowing bridge protocol data units (BPDU) through. May I remove this configuration parameter to stabilize the network?
- A:** ☺

NOTE

In mid 1990's when the large corporate networks began to be connected to less secure public networks (for example, the early Internet), the security-conscious network administrators immediately started to feel the need to secure their internal.

NOTE

Use this section only if you require clues to complete the questions. In the real CCIE lab, the proctor will not enter into any discussions about the questions or answers. Instead, the proctor is present to ensure you do not have problems with the lab environment and to maintain the timing element of the exam.

Incident 2

Q: I have a one-way neighbor relationship between R1 and R3; this cannot be right. Is this a bug?

A: No. Troubleshoot the issue. Remember how EIGRP forms a neighbor relationship and test in this manner over the Frame Relay network.

Q: Do you mean reach ability of multicast of 224.0.0.10?

A: Yes. You should be able to ping this address from each neighbor on the same subnet.

Q: I cannot ping 224.0.0.10 from R1 and R3. I can see that because I am on a point-to-multipoint Frame Relay interface I would need broadcast capability, but it has not been configured. Am I okay to add a **broadcast** statement to R1?

A: ☺

Incident 3

Q: I have another EIGRP one-way neighbor relationship between the PE router R7 and Sw2. If I try to ping 224.0.0.10 again, it fails. But this is an Ethernet network, and broadcast should work by default. Is this correct?

A: Yes. Ethernet is, of course, a broadcast domain. Just think whether there is anything that could potentially stop your devices multicasting to each other.

Q: Ah, okay. I see storm control. Would this stop multicast?

A: ☺

Incident 4

Q: I have just checked the Border Gateway Protocol (BGP) connection between PE routers, and it is up. Could there be an issue with the configuration from PE to CE devices?

A: Potentially. But you could have an MPLS-specific issue.

- Q:** I can see I have an issue with route targets. Should I change the route descriptor and route targets to match?
- A:** Just change either PE router so that the route targets match for import and export on your vrfs. The route descriptors can be unique, and you could lose some vrf-specific configuration if you modified the route descriptors.
- Q:** I do not have LDP neighbors between my PE routers and P routers. The PEs are using LDP and the Ps are using TDP. These should all be the same. Does it matter which protocol I use?
- A:** No.
- Q:** My MP-BGP is up. I have LDP neighbors throughout the MPLS network, but I do not have any vrf-specific routes coming through. Is this a bug?
- A:** No. Use some MPLS-specific troubleshooting commands to aid your problem determination.
- Q:** If I complete an MPLS traceroute, I find that there is no label between PE loopback interfaces and that the mask of these is /24. Shouldn't loopbacks be /32s for MPLS?
- A:** ☺

Incident 5

- Q:** The MPLS network looks good, and I receive BGP routes specific to my vrf, but I don't get these through as EIGRP routes on my CEs. The configuration looks okay. Is something missing?
- A:** Yes. Just remember what EIGRP needs to redistribute routes from other protocols.

Incident 6

- Q:** One PE router has been configured as an EIGRP stub router. May I change this (because this would only allow it to advertise locally connected routes)?
- A:** ☺

Incident 7

- Q:** I have determined that the static route for 120.100.13.0 was not being advertised because of only classful static routes being advertised into OSPF, but it still is not showing up in the remote PE router R7. Is this correct?
- A:** No. There must be something else that stops this route being redistributed into BGP on PE Router R4.
- Q:** There appears to be a route map that stops routes with a TAG value set. Can I remove the TAG value from the static route or change the route map?
- A:** It is quicker and safer to remove the TAG from the route because the route map might be required for additional routes not detailed within the question.

Incident 8

- Q:** Don't I just need Protocol Independent Multicast (PIM) enabled on each interface to run multicast over the network?
- A:** Yes and no. You have an MPLS multicast scenario with MDT and clues as to the solution.
- Q:** I have PIM neighbors all the way along the chain. Is this enough?
- A:** No. This solution also requires PIM neighbors over a tunnel between PE loopback interfaces.
- Q:** I do not have PIM neighbors over my PE loopbacks. I also do not have PIM enabled on these interfaces. Surely, I need this to form a neighbor relationship between loopbacks, right?
- A:** Yes you do. Try enabling PIM on these interfaces.

Incident 9

- Q:** I do get a poor response when I telnet to R5 from R6. Isn't this just the result of CPU loads and traffic rates and considered normal?
- A:** No. Telnet to R5 using a physical address, and note the difference.
- Q:** I can see that control plane policing (CoPP) has been configured on R5 for Telnet traffic directed to the loopback interface. May I modify the packets per second (pps) rate to try to alleviate the problem?
- A:** Yes. Just increase the value until you see an improvement.

Troubleshooting Lab 1 Debrief

The debrief aims to achieve the following:

- Define the problem and identify the symptoms (questioning to the void to identify meaningful symptoms)
- Define the hypothesis and proof-test possible causes
- Design and implement a correct solution (before application of the configuration)
- Verify the resolution within the stipulated guidelines

You should use this section to produce an overall score for the practice lab.

Section 2

Incident 1

Users on Sw3 VLAN 20 are complaining that they experience poor connectivity to VLAN 20 on Sw2. A variable response is seen when pinging from Sw3 VLAN 20 to Sw2 VLAN 20 interface.

Investigate the issue and rectify it. Confirm by proving a successful ping from Sw3 VLAN 20 interface 120.100.20.3 to Sw2 VLAN 20 interface 120.100.20.2 with a stable response time, as follows:

```
Sw3# ping 120.100.27.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
Sw3# ping 120.100.27.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

```
Sw3# ping 120.100.27.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

3 points

So, you're in at the deep end, fault-finding somebody else's network that you are not familiar with and the clock is ticking. You need to depend on your basic troubleshooting skills to be successful. First, it is advisable to confirm what the issue actually is and ping Sw2 VLAN 20 120.100.20.2 from Sw3. Example 1-1 shows that the ping is successful, but indeed with a varied response time. The issue could be because of a number of factors, but you may notice syslog messages on the console of Sw3, which can be examined within the logging buffer. What should be evident is that you have a MAC address flapping between Gi0/1 and Gi0/2 on VLAN 20 (the ports used to connect the switches together, as indicated in Figure 1-1), telltale signs of a classic spanning-tree issue on this VLAN. If you look at the basic spanning-tree topology on Sw3, you will see that all VLANs are forwarding on both Gi0/1 and Gi0/2. (Gi0/1 is a trunk on each side and Gi0/2 is an extension of VLAN 20.) This happens to be identical to Sw2, so you can be confident that you have a spanning-tree loop. If you examine the configurations on Sw3 for Gi0/1, you will find nothing untoward, just a standard trunk configuration. However, Gi0/2 has the command **spanning-tree bpdudfilter enable** configured. By enabling BPDU filtering on Gi0/2, the switch effectively believes it is connected to an end system and not to another switch, as such a loop is formed between devices. The problem is rectified by disabling the BPDU filter, as detailed in Example 1-1, and ping testing shows a stable response time. A correct resolution to this problem is worth 3 points. If you did not fix this issue first, you might encounter numerous connectivity issues with devices connected to each switch.

Example 1-1 *Sw3 and Sw2 Spanning-Tree Issue and Rectification*

```
Sw3# ping 120.100.20.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.20.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 9/16/25 ms
Sw3# ping 120.100.20.2
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.20.2, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 8/134/621 ms

Sw3# **show log**

.Mar 11 10:39:10: %SW_MATM-4-MACFLAP_NOTIF: Host 000a.b8b9.1ec6 in vlan 20 is flapping between port Gi0/1 and port Gi0/2

.Mar 11 10:39:25: %SW_MATM-4-MACFLAP_NOTIF: Host 000a.b8b9.1ec6 in vlan 20 is flapping between port Gi0/2 and port Gi0/1

.Mar 11 10:39:40: %SW_MATM-4-MACFLAP_NOTIF: Host 000a.b8b9.1ec6 in vlan 20 is flapping between port Gi0/2 and port Gi0/1

.Mar 11 10:39:55: %SW_MATM-4-MACFLAP_NOTIF: Host 000a.b8b9.1ec6 in vlan 20 is flapping between port Gi0/1 and port Gi0/2

Sw3# **show spanning-tree interface gi0/1**

Vlan	Role	Sts	Cost	Prio.	Nbr	Type
VLAN0001	Root	FWD	4	128.1		P2p
VLAN0007	Desg	FWD	4	128.1		P2p
VLAN0010	Root	FWD	4	128.1		P2p
VLAN0020	Root	FWD	4	128.1		P2p
VLAN0030	Root	FWD	4	128.1		P2p
VLAN0037	Root	FWD	4	128.1		P2p
VLAN0050	Root	FWD	4	128.1		P2p

Sw3# **show spanning-tree interface gi0/2**

Vlan	Role	Sts	Cost	Prio.	Nbr	Type
VLAN0020	Desg	FWD	4	128.2		P2p

Sw2# **show spanning-tree interface gi0/1**

```

Vlan                Role Sts Cost      Prio.Nbr Type
-----
VLAN0001            Desg FWD 4          128.1   P2p
VLAN0010            Desg FWD 4          128.1   P2p
VLAN0020            Desg FWD 4          128.1   P2p
VLAN0027            Desg FWD 4          128.1   P2p
VLAN0030            Desg FWD 4          128.1   P2p
VLAN0037            Desg FWD 4          128.1   P2p
VLAN0050            Desg FWD 4          128.1   P2p

```

```
Sw2# show spanning-tree interface gi0/2
```

```

Vlan                Role Sts Cost      Prio.Nbr Type
-----
VLAN0020            Desg FWD 4          128.2   P2p

```

```
Sw3# show run int gi0/1
```

```

interface GigabitEthernet0/1
  description Link to Sw2
  switchport trunk encapsulation dot1q
  switchport mode trunk
end

```

```
Sw3# show run int gi0/2
```

```

interface GigabitEthernet0/2
  description Link to Sw2
  switchport access vlan 20
  switchport mode access
  spanning-tree bpdufilter enable
end

```

```
Sw3(config)# interface gi0/2
```

```
Sw3(config-if)# no spanning-tree bpdufilter
```

```
Sw3(config-if)# do show spanning-tree interface gi0/1
```



```

Vlan                Role Sts Cost      Prio.Nbr Type
-----
VLAN0001            Root FWD 4         128.1   P2p
VLAN0007            Desg FWD 4         128.1   P2p
VLAN0010            Root FWD 4         128.1   P2p
VLAN0020            Root FWD 4         128.1   P2p
VLAN0030            Root FWD 4         128.1   P2p
VLAN0037            Root FWD 4         128.1   P2p
VLAN0050            Root FWD 4         128.1   P2p

```

```
Sw3(config-if)# do show spanning-tree interface gi0/2
```

```

Vlan                Role Sts Cost      Prio.Nbr Type
-----
VLAN0020            Altn BLK 4         128.2   P2p

```

```
Sw3# ping 120.100.20.2
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 120.100.20.2, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
```

```
Sw3# ping 120.100.20.2
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 120.100.20.2, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
```

Incident 2

Users on R3's 120.100.3.0/24 subnet are complaining that they cannot access resources located on their remote virtual private network (VPN) site on VLAN 20 (switch 2). Initial investigations have led first-line support personnel to think that there is an issue with connectivity up to the provider edge (PE) Router R4, which was recently replaced, because the subnet 120.100.4.0/24 is not present in R3's routing table.

Investigate the issue, and prove connectivity solely to their PE connection (R4's VLAN 4 interface) as follows:

```
R3# ping 120.100.4.4 source 120.100.3.1
```

```
Type escape sequence to abort.  
Sending 5, 100-byte ICMP Echos to 120.100.4.4, timeout is 2 seconds:  
Packet sent with a source address of 120.100.3.1  
!!!!!  
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

2 points

You should begin by familiarizing yourself with the VPN diagram shown as Figure 1-6 to realize you will be working on a vrf-specific configuration of vrf ROUTERS. Using the information provided in the problem statement, you need to be able to prove connectivity from R3 120.100.3.3 (Loopback0 interface) to R4 120.100.4.4 (VLAN 4 interface). A successful extended ping between the two IP addresses will secure your points. In terms of how to troubleshoot this, you must start with routing table entries and IGP neighbor status. Example 1-2 details the routing tables and IGP neighbor status of R3, R1, and R4. Just remember that R4 is a PE router in this scenario, so you need your vrf-specific commands (something easily forgotten in the early hours of the morning or at “headless chicken time” when everyone in the world wants a reason why the network is down). Example 1-2 shows that R3 does not have a route to R1 VLAN 4, yet it has an EIGRP neighbor to R1 over the Frame Relay network. Similarly, R1 does not have a route to R3, but it does not have an EIGRP neighbor with R3 over the Frame Relay. This should seem unusual because a one-way neighbor relationship is not common and is clearly the reason why users would have connectivity issues. R4 has a neighbor with R1 but clearly would not have a route to R3 because of the one-way neighbor relationship between R1 and R3. R1 and R3 are able to ping each other's serial interfaces, so you can confirm that the Frame Relay network is operational. However, it is worth attempting to ping 224.0.0.10 to ensure EIGRP is able to function correctly over Frame Relay to troubleshoot the neighbor issue. The result of which is that neither R1 nor R3 can successfully ping 224.0.0.10. You now should understand that you are looking at a multicast/broadcast issue over Frame Relay.

Example 1-2 R3 and R1 and R4 Route Verification

```

R3# show ip route | include 120.100.4.
R3# show ip eigrp neighbor
IP-EIGRP neighbors for process 1
H   Address                Interface   Hold Uptime   SRTT   RTO   Q   Seq Type
      (sec)                (ms)                Cnt Num
0   120.100.123.1          Se0/0      136 00:00:43   1   5000   1   0

R1# show ip route | include 120.100.3.
R1# show ip eigrp neighbor
IP-EIGRP neighbors for process 1
H   Address                Interface   Hold Uptime   SRTT   RTO   Q   Seq Type
      (sec)                (ms)                Cnt Num
1   120.100.4.4            Et3/0      14 3d22h       1    200   0   8
2   120.100.123.2         Se2/0      175 4d22h      435  2610   0   8

R4# show ip route vrf ROUTERS | include 120.100.3.
R4# show ip eigrp vrf ROUTERS neighbor
IP-EIGRP neighbors for process 1
H   Address                Interface   Hold Uptime   SRTT   RTO   Q   Seq
      (sec)                (ms)                Cnt Num
0   120.100.4.1            Gi0/1.4    12 3d22h       1    300   0   94

R1# ping 120.100.123.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.123.3, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms
R1# ping 224.10.10.10

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 224.10.10.10, timeout is 2 seconds:
.

```

```
R3# ping 120.100.123.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 120.100.123.1, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

```
R3# ping 224.0.0.10
```

```
Type escape sequence to abort.
```

```
Sending 1, 100-byte ICMP Echos to 224.0.0.10, timeout is 2 seconds:
```

```
.
```

Inspection of the configuration on the Frame Relay hub R1 in Example 1-3 shows a valid configuration, but R3 is missing the broadcast statement in the **frame-relay map** to R1. When configured, the multicast ping functions and the neighbor relationship is now two-way between R1 and R3, permitting valid routing updates. Testing is detailed in Example 1-3 with an extended ping between R3 and R4, which also proves the routing between R1 and R4. The question indicates that Router R4 was recently replaced, and this clearly had no relevance to the issue. This might have led you to focus on R4 as the root cause of the problem. Because just as in real life, you may be presented with misleading or irrelevant information when dealing with faults. A correct resolution to this problem, as detailed in Example 1-3, is worth 2 points.

Example 1-3 R1 and R3 Frame-Relay Testing and Configuration

```
R1# show run int s2/0
```

```
interface Serial2/0
```

```
ip address 120.100.123.1 255.255.255.0
```

```
encapsulation frame-relay
```

```
frame-relay map ip 120.100.123.2 102 broadcast
```

```
frame-relay map ip 120.100.123.3 103 broadcast
```

```
end
```

```
R3# show run int s0/0
interface Serial0/0
 ip address 120.100.123.3 255.255.255.0
 encapsulation frame-relay
 clock rate 128000
 frame-relay map ip 120.100.123.1 301
end
```

```
R3# conf t
R3(config)# int s0/0
R3(config-if)# frame-relay map ip 120.100.123.1 301 broadcast
R3(config-if)# ^Z
R3# ping 224.0.0.10
```

Type escape sequence to abort.

Sending 1, 100-byte ICMP Echos to 224.0.0.10, timeout is 2 seconds:

```
Reply to request 0 from 120.100.123.1, 52 ms
```

```
R3# show ip eigrp neighbor
```

```
IP-EIGRP neighbors for process 1
```

H	Address	Interface	Hold Uptime (sec)	SRTT (ms)	RTO	Q Cnt	Seq	Type
0	120.100.123.1	Se0/0	126 00:02:42	1	5000	1	0	

```
R3# show ip route | include 120.100.4.
```

```
D 120.100.4.0 [90/2195456] via 120.100.123.1, 00:00:14, Serial0/0
```

```
R3#
```

```
R1# ping 224.0.0.10
```

Type escape sequence to abort.

```

Sending 1, 100-byte ICMP Echos to 224.0.0.10, timeout is 2 seconds:

Reply to request 0 from 120.100.4.4, 8 ms
Reply to request 0 from 120.100.123.2, 40 ms
Reply to request 0 from 120.100.123.3, 32 ms
R1# show ip eigrp neighbor
IP-EIGRP neighbors for process 1
H   Address                Interface    Hold Uptime    SRTT  RTO  Q  Seq Type
                               (sec)        (ms)          Cnt Num
0   120.100.123.3           Se2/0       129 00:00:52 1421 5000 0 158
1   120.100.4.4             Et3/0       14 3d22h      1   200 0 8
2   120.100.123.2           Se2/0       129 4d23h     352 2112 0 8
R1# sh ip route | include 120.100.3.
D    120.100.3.0 [90/2195456] via 120.100.123.3, 00:01:02, Serial2/0

R3# ping 120.100.4.4 source 120.100.3.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.4.4, timeout is 2 seconds:
Packet sent with a source address of 120.100.3.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms

```

Incident 3

Users on switch 2 VLAN 20 are complaining that they have no connectivity through their VPN to services located on R2 and R3 (120.100.2.0/24 and 120.100.3.0/24). A traceroute to destination networks shows timeouts immediately, and it appears that an EIGRP neighbor adjacency between Sw2 and R7 is not present.

Investigate the issue by proving connectivity purely between Switch 2 and the local PE router R7. Establishment of an EIGRP adjacency between devices suffices at this point, as follows:

```
Sw2# show ip eigrp neighbors
EIGRP-IPv4:(1) neighbors for process 1
H   Address                Interface          Hold Uptime   SRTT   RTO   Q   Seq
0   120.100.27.7            Vlan27            12           00:01:11  1     300   0  92
```

3 points

The ticket dictates that you just need an EIGRP adjacency at this point, so initial troubleshooting should be checking the adjacency between Sw2 and R7 over the VLAN 27 interface, as indicated in Figure 1-6. Example 1-4 shows that the PE Router R7 has an adjacency with Sw2, but Sw2 does not have a neighbor relationship with R7. If you get a one-way neighbor relationship like this, myriad issues could be to blame. So, you must get back to basics first and prove connectivity. Bear in mind, however, that this could again be multicast related, as shown in the previous question. Example 1-4 shows a ping from Sw2 to R7 over the locally connected interface. This ping is successful, so a multicast ping can be attempted. If successful, this proves that multicast packets are received at R7 and replied to by the EIGRP process. The ping fails. So, you can now be certain you are looking for a multicast-related issue between Sw2 and R7.

Example 1-4 R7 and S2 EIGRP Neighbor Testing

```
R7# show ip eigrp vrf ROUTERS neighbors
IP-EIGRP neighbors for process 1
H   Address                Interface          Hold Uptime   SRTT   RTO   Q   Seq
                                (sec)           (ms)           Cnt Num
0   120.100.27.2            Gi0/0.27          12 00:01:11   1  5000  2  0
```

```
Sw2# show ip eigrp neighbors
EIGRP-IPv4:(1) neighbors for process 1
Sw2# ping 120.100.27.7
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.27.7, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

```
Sw2# ping 224.0.0.10
```

Type escape sequence to abort.

Sending 1, 100-byte ICMP Echos to 224.0.0.10, timeout is 2 seconds:

```
Reply to request 0 from 120.100.27.7, 8 ms
```

R7# ping vrf ROUTERS 120.100.27.2

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:

```
!!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/8 ms

R7# ping vrf ROUTERS 224.0.0.10

Type escape sequence to abort.

Sending 1, 100-byte ICMP Echos to 224.0.0.10, timeout is 2 seconds:

```
.
```

It could be that there is an access control list (ACL) in place denying multicast traffic or EIGRP traffic in one direction, so it is time for a methodical verification of configurations between relevant interfaces and routing processes on R7 and Sw2. Example 1-5 details the resultant configurations on each device. All interface and EIGRP configuration appears to be as normal until you inspect the interface where R7 connects to Sw2 (Fast Ethernet 0/7). You will see the command **storm-control multicast level 0.00**. This effectively disable multicast on this port. This could be a valid configuration for a host port but not with this level for a port connecting to an EIGRP router. When removed, you will see the neighbor relationship is immediately formed between Sw2 and R7. Basic troubleshooting steps resolve the issue. Focus on the problem to provide the necessary clues to save your time and secure your points on the exam. As correct resolution to this problem, as detailed in Example 1-5, is worth 3 points.

Example 1-5 *R7 and Sw2 Interface and EIGRP Configuration and Verification*

```
R7# show run int gi0/0.27
interface GigabitEthernet0/0.27
  encapsulation dot1q 27
  ip vrf forwarding ROUTERS
  ip address 120.100.27.7 255.255.255.0
end
```

```
R7# show run | section address-family ipv4 vrf ROUTERS
address-family ipv4 vrf ROUTERS
  redistribute bgp 65001
  network 120.100.27.0 0.0.0.255
  eigrp stub connected summary
  no auto-summary
  autonomous-system 1
```

```
Sw2# show run int vlan 27
interface Vlan27
  ip address 120.100.27.2 255.255.255.0
end
```

```
Sw2# show run | begin eigrp 1
router eigrp 1
  no auto-summary
  network 120.100.20.0 0.0.0.255
  network 120.100.27.0 0.0.0.255
!
```

```
Sw2# show run int fast 0/7
interface FastEthernet0/7
  description Link to R7 G0/0
  switchport trunk encapsulation dot1q
  switchport mode trunk
```

```

speed 100
duplex full
storm-control multicast level 0.00
end

```

```

Sw2(config)# int fast 0/7
Sw2(config-if)# no storm-control multicast level 0.00
Sw2(config-if)#
.Feb 28 10:14:36: %DUAL-5-NBRCHANGE: EIGRP-IPv4:(1) 1: Neighbor 120.100.27.7 (Vlan27) is up: new adjacency

```

Incident 4

Users on R3's 120.100.3.0/24 subnet are still complaining that they cannot access resources located on their remote VPN site on VLAN 20 (switch 2). It appears that the MPLS VPN is not functioning correctly.

Investigate the issue and prove connectivity between PE routers via extended ping from R4 VLAN 4 interface and R7 VLAN 27 interface to prove MPLS functionality, as follows:

```

R4# ping vrf ROUTERS 120.100.27.7 source 120.100.4.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.27.7, timeout is 2 seconds:
Packet sent with a source address of 120.100.4.4
!!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

```

```

R7# ping vrf ROUTERS 120.100.4.4 source 120.100.27.7
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.4.4, timeout is 2 seconds:
Packet sent with a source address of 120.100.27.7
!!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

```

6 points

Ouch! You have just been asked to confirm a whole MPLS network is functioning correctly from PE to PE. This is a great deal of testing and configuration verification to run through. The quickest approach is to look at the vrf-specific routing tables at each PE router to see what routes are present and to determine whether you need to work out to the CE devices or back into the MPLS network. Example 1-6 details the vrf routing tables for the PE routers. As you can see, there are no routes learned via BGP between PE routers, and therefore there is no end-to-end connectivity throughout this VPN.

Example 1-6 *R4 and R7 PE VRF Route Verification*

R4# **show ip route vrf ROUTERS**

```

    120.0.0.0/24 is subnetted, 4 subnets
D       120.100.14.0
        [90/284160] via 120.100.4.1, 00:59:31, GigabitEthernet0/1.4
C       120.100.4.0 is directly connected, GigabitEthernet0/1.4
D       120.100.3.0
        [90/2198016] via 120.100.4.1, 00:59:31, GigabitEthernet0/1.4
D       120.100.123.0
        [90/2172416] via 120.100.4.1, 00:59:31, GigabitEthernet0/1.4

```

R7# **show ip route vrf ROUTERS**

```

    120.0.0.0/24 is subnetted, 2 subnets
C       120.100.27.0 is directly connected, GigabitEthernet0/0.27

```

So, what are the steps to resolve a potential MPLS issue? Dive into MPLS traceroute, check LDP neighbors, and check per-vrf redistribution: The list is long and varied. Getting back to basics will solve this and any problem. Prove your connectivity, and then build up the layers of detail until you have enough information to solve the issue. You may or may not have a huge amount of experience with MPLS, but you should know that the PE routers need to peer with each other over MP-BGP to exchange vrf-specific routing entries detailed within extended community values by use of route descriptors and route targets. To peer

with each other, the MPLS-specific network should be capable of running an IGP to transport the PE loopback addresses and run MPLS on each connecting interface throughout the network to transport labels. This is pretty high level, but that just about sums up what happens and should give you a basis for troubleshooting. Start by checking an extended ping from R4 to R7 via their peering loopback interfaces. Then, if successful, check the BGP neighbor relationship. Example 1-7 shows the extended ping between PEs, which is successful, and verification that the BGP session between them is indeed established and active.

Example 1-7 *R4 and R7 PE Testing*

```
R4# ping 7.7.7.7 source 4.4.4.4
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 7.7.7.7, timeout is 2 seconds:
```

```
Packet sent with a source address of 4.4.4.4
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

```
R4# show ip bgp neighbors 7.7.7.7 | include BGP state
```

```
BGP state = Established, up for 00:01:05
```

```
R7# ping 4.4.4.4 source 7.7.7.7
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
```

```
Packet sent with a source address of 7.7.7.7
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

```
R7# show ip bgp neighbors 4.4.4.4 | include BGP state
```

```
BGP state = Established, up for 00:02:02
```

A quick test to determine whether you are advertising vrf-specific routes between PEs is to use the `show ip bgp vpnv4 vrf X` command. This tells you what you are advertising to the remote PE and what you are learning from the remote PE (based on the next hop of the route). Example 1-8 shows the output from this command from each PE, and you can see that you are only advertising at each edge and not receiving any routes from the remote PEs. You can therefore begin to look into the MPLS functionality as to why you are not receiving routes from each PE. Also note that in the output the RD is shown as 1:1000 for this vrf on R4 and 2:100 on R7. This could be fine being different based on RT import and export values, but it is worth checking. Example 1-8 also shows the vrf-specific RTs, and you can see they are different on each PE router. By changing either PE to import and export the same RT per vrf, you can ensure your vrf routes are advertised successfully over MP-BGP. You should be able to spot that the second vrf is also inconsistent and that you can rectify this here or at a later time when investigating connectivity problems for that specific vrf.

Example 1-8 R4 and R7 PE Testing and Configuration

```
R4# show ip bgp vpnv4 vrf ROUTERS
BGP table version is 23, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 1:1000 (default for vrf ROUTERS)					
*> 120.100.2.0/24	120.100.4.1	2300416		32768	?
*> 120.100.3.0/24	120.100.4.1	2198016		32768	?
*> 120.100.4.0/24	0.0.0.0	0			32768 ?
*> 120.100.14.0/24	120.100.4.1	284160		32768	?
*> 120.100.123.0/24	120.100.4.1	2172416		32768	?

```
R7# show ip bgp vpnv4 vrf ROUTERS
BGP table version is 416, local router ID is 7.7.7.7
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
```

Origin codes: i - IGP, e - EGP, ? - incomplete

```

      Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 2:100 (default for vrf ROUTERS)
*> 120.100.27.0/24  0.0.0.0          0          32768 ?

```

R4# show run | section ip vrf

```

ip vrf ROUTERS
  rd 1:1000
  route-target export 1:1000
  route-target import 2:1000
ip vrf SWITCHES
  rd 1:2000
  route-target export 1:2000
  route-target import 2:2000
  mdt default 232.0.0.11

```

R7# show run | section ip vrf

```

ip vrf ROUTERS
  rd 2:100
  route-target export 2:100
  route-target import 1:100
ip vrf SWITCHES
  rd 2:200
  route-target export 2:200
  route-target import 1:200
  mdt default 232.0.0.11

```

R4(config)# ip vrf ROUTERS

R4(config-vrf)# no route-target export 1:1000

R4(config-vrf)# no route-target import 2:1000

R4(config-vrf)# route-target export 1:100

```
R4(config-vrf)# route-target import 2:100
R4(config-vrf)# ip vrf SWITCHES
R4(config-vrf)# no route-target export 1:2000
R4(config-vrf)# no route-target import 2:2000
R4(config-vrf)# route-target export 1:200
R4(config-vrf)# route-target import 2:200
```

```
R4# show ip bgp vpnv4 vrf ROUTERS
```

```
BGP table version is 44, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 1:1000 (default for vrf ROUTERS)					
*> 120.100.2.0/24	120.100.4.1	2300416		32768	?
*> 120.100.3.0/24	120.100.4.1	2198016		32768	?
*> 120.100.4.0/24	0.0.0.0	0			32768 ?
*> 120.100.14.0/24	120.100.4.1	284160		32768	?
*>i120.100.27.0/24	7.7.7.7	0	100		0 ?
*> 120.100.123.0/24	120.100.4.1	2172416		32768	?

```
R7# show ip bgp vpnv4 vrf ROUTERS
```

```
BGP table version is 431, local router ID is 7.7.7.7
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 2:100 (default for vrf ROUTERS)					
*>i120.100.2.0/24	4.4.4.4	2300416	100		0 ?
*>i120.100.3.0/24	4.4.4.4	2198016	100		0 ?

```
*>i120.100.4.0/24 4.4.4.4 0 100 0 ?
*>i120.100.14.0/24 4.4.4.4 284160 100 0 ?
*> 120.100.27.0/24 0.0.0.0 0 32768 ?
*>i120.100.123.0/24 4.4.4.4 2172416 100 0 ?
```

Example 1-8 shows that you have the routes available between PE routers after configuring symmetrical RT **import** and **export** statements. To receive your points on the question, though, you must perform an extended ping from the vrf interfaces on each PE router to prove end-to-end connectivity. Example 1-9 shows the extended ping, which happens to fail. This would indicate a core MPLS-related issue. If you check for an LDP neighbor relationship over MPLS-configured interfaces, you will see that both PE routers are running LDP and do not have any neighbors when they should in fact be neighbored with R4 and R5 P routers. R4 and R5, of course, will therefore not have a neighbor relationship with LDP, but you should notice that they are using TDP as the label distribution protocol. You could either change your P routers or your PE routers to run the same protocol, but Example 1-9 shows configuring the P routers to run LDP with the global command **mpls label protocol ldp**. The neighbor relationship is now formed across the MPLS network, but the extended ping still fails.

Example 1-9 R4 and R7 PE Testing and Configuration

```
R4# ping vrf ROUTERS 120.100.27.4 source 120.100.4.4
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 120.100.27.4, timeout is 2 seconds:
```

```
Packet sent with a source address of 120.100.4.4
```

```
.....
```

```
Success rate is 0 percent (0/5)
```

```
R7# ping vrf ROUTERS 120.100.4.4 source 120.100.27.7
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 120.100.4.4, timeout is 2 seconds
```

```
Packet sent with a source address of 120.100.27.7
```

```
.....
```


Success rate is 0 percent (0/5)

R4# **show mpls ldp neighbors**

R4# **show mpls interfaces**

Interface	IP	Tunnel	Operational
Serial0/2/0	Yes (ldp)	No	Yes
Serial0/2/1	Yes (ldp)	No	Yes

R7# **show mpls ldp neighbors**

R7# **show mpls interfaces**

Interface	IP	Tunnel	Operational
Serial0/2/0	Yes (ldp)	No	Yes
Serial0/3/0	Yes (ldp)	No	Yes

R5# **show mpls ldp neighbors**

R5# **show mpls interfaces**

Interface	IP	Tunnel	Operational
Serial0/1	Yes (tdp)	No	Yes
Serial0/2	Yes (tdp)	No	Yes

R6# **show mpls ldp neighbors**

R6# **show mpls interfaces**

Interface	IP	Tunnel	Operational
Serial0/0	Yes (tdp)	No	Yes
Serial0/1	Yes (tdp)	No	Yes

R5(config)# **mpls label protocol ldp**

```

*Mar  7 05:23:25.844: %LDP-5-NBRCHG: LDP Neighbor 7.7.7.7:0 (1) is UP
*Mar  7 05:23:29.800: %LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (2) is UP
R6(config)# mpls label protocol ldp
*Mar  7 06:04:39.372: %LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (1) is UP
*Mar  7 06:04:40.256: %LDP-5-NBRCHG: LDP Neighbor 7.7.7.7:0 (2) is UP

R4# ping vrf ROUTERS 120.100.27.4 source 120.100.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.27.4, timeout is 2 seconds:
Packet sent with a source address of 120.100.4.4
.....
Success rate is 0 percent (0/5)

R7# ping vrf ROUTERS 120.100.4.4 source 120.100.27.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.4.4, timeout is 2 seconds
Packet sent with a source address of 120.100.27.7
.....
Success rate is 0 percent (0/5)

```

So, you now must be into the realms of an MPLS issue, where life can be complicated unless you work with MPLS on a daily basis. A good tool to check the LSP is an MPLS traceroute. Example 1-10 shows an MPLS traceroute between PE loopback addresses. As you can see, this fails from either direction because R5 and R6 P routers inform each PE router that no corresponding label exists. This is a classic case of incorrect initial configuration. You know your BGP session is working correctly, and you know you have LDP established correctly between your devices. If you look at your MPLS router loopback0 interfaces, you will see they have a mask configured of /24, and you should know by theory and through practical experience that your peering and LDP router ID loopbacks should be configured as /32 to operate correctly within MPLS. If you check the routing table, you see the routes come through as /32 host routes, but this is simply because of the way Open Shortest Path First (OSPF) Protocol advertises loopback networks (unless you adjust the OSPF network type of the interface). Example 1-10 shows that

when the loopbacks are re-addressed to /32 hosts the MPLS traceroute functions correctly and the end-to-end vrf-specific ping is successful between PE routers. A correct resolution to this problem, as detailed in Examples 1-6 through 1-10, is worth 6 points.

Example 1-10 *R4 and R7 PE Testing and Configuration*

```
R4# traceroute mpls ipv4 7.7.7.7/32
```

```
Tracing MPLS Label Switched Path to 7.7.7.7/32, timeout is 2 seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
       'L' - labeled output interface, 'B' - unlabeled output interface,
       'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
       'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label ent
       'P' - no rx intf label prot, 'p' - premature termination of LSP,
       'R' - transit router, 'I' - unknown upstream index,
       'X' - unknown return code, 'x' - return code 0
```

```
Type escape sequence to abort.
```

```
 0 120.100.46.4 MRU 1500 [Labels: 16 Exp: 0]
B 1 120.100.46.6 MRU 1504 [No Label] 0 ms
B 2 120.100.46.6 MRU 1504 [No Label] 0 ms
B 3 120.100.46.6 MRU 1504 [No Label] 0 ms
B 4 120.100.46.6 MRU 1504 [No Label] 0 ms
B 5 120.100.46.6 MRU 1504 [No Label] 4 ms
B 6 120.100.46.6 MRU 1504 [No Label] 0 ms
B 7 120.100.46.6 MRU 1504 [No Label] 0 ms
B 8 120.100.46.6 MRU 1504 [No Label] 0 ms
B 9 120.100.46.6 MRU 1504 [No Label] 4 ms
B 10 120.100.46.6 MRU 1504 [No Label] 0 ms
B 11 120.100.46.6 MRU 1504 [No Label] 0 ms
B 12 120.100.46.6 MRU 1504 [No Label] 0 ms
B 13 120.100.46.6 MRU 1504 [No Label] 4 ms
```

```
B 14 120.100.46.6 MRU 1504 [No Label] 0 ms
B 15 120.100.46.6 MRU 1504 [No Label] 0 ms
B 16 120.100.46.6 MRU 1504 [No Label] 0 ms
B 17 120.100.46.6 MRU 1504 [No Label] 4 ms
B 18 120.100.46.6 MRU 1504 [No Label] 0 ms
B 19 120.100.46.6 MRU 1504 [No Label] 0 ms
B 20 120.100.46.6 MRU 1504 [No Label] 0 ms
B 21 120.100.46.6 MRU 1504 [No Label] 4 ms
B 22 120.100.46.6 MRU 1504 [No Label] 0 ms
B 23 120.100.46.6 MRU 1504 [No Label] 0 ms
B 24 120.100.46.6 MRU 1504 [No Label] 0 ms
B 25 120.100.46.6 MRU 1504 [No Label] 4 ms
B 26 120.100.46.6 MRU 1504 [No Label] 0 ms
B 27 120.100.46.6 MRU 1504 [No Label] 0 ms
B 28 120.100.46.6 MRU 1504 [No Label] 0 ms
B 29 120.100.46.6 MRU 1504 [No Label] 4 ms
B 30 120.100.46.6 MRU 1504 [No Label] 0 ms
```

```
R7# traceroute mpls ipv4 4.4.4.4/32
```

```
Tracing MPLS Label Switched Path to 4.4.4.4/32, timeout is 2 seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
       'L' - labeled output interface, 'B' - unlabeled output interface,
       'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
       'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
       'P' - no rx intf label prot, 'p' - premature termination of LSP,
       'R' - transit router, 'I' - unknown upstream index,
       'X' - unknown return code, 'x' - return code 0
```

```
Type escape sequence to abort.
```

```
0 120.100.67.7 MRU 1500 [Labels: 18 Exp: 0]
B 1 120.100.67.6 MRU 1504 [No Label] 0 ms
```

```
B 2 120.100.67.6 MRU 1504 [No Label] 0 ms
B 3 120.100.67.6 MRU 1504 [No Label] 0 ms
B 4 120.100.67.6 MRU 1504 [No Label] 0 ms
B 5 120.100.67.6 MRU 1504 [No Label] 4 ms
B 6 120.100.67.6 MRU 1504 [No Label] 0 ms
B 7 120.100.67.6 MRU 1504 [No Label] 0 ms
B 8 120.100.67.6 MRU 1504 [No Label] 0 ms
B 9 120.100.67.6 MRU 1504 [No Label] 4 ms
B 10 120.100.67.6 MRU 1504 [No Label] 0 ms
B 11 120.100.67.6 MRU 1504 [No Label] 0 ms
B 12 120.100.67.6 MRU 1504 [No Label] 0 ms
B 13 120.100.67.6 MRU 1504 [No Label] 4 ms
B 14 120.100.67.6 MRU 1504 [No Label] 0 ms
B 15 120.100.67.6 MRU 1504 [No Label] 0 ms
B 16 120.100.67.6 MRU 1504 [No Label] 0 ms
B 17 120.100.67.6 MRU 1504 [No Label] 4 ms
B 18 120.100.67.6 MRU 1504 [No Label] 0 ms
B 19 120.100.67.6 MRU 1504 [No Label] 0 ms
B 20 120.100.67.6 MRU 1504 [No Label] 0 ms
B 21 120.100.67.6 MRU 1504 [No Label] 4 ms
B 22 120.100.67.6 MRU 1504 [No Label] 0 ms
B 23 120.100.67.6 MRU 1504 [No Label] 0 ms
B 24 120.100.67.6 MRU 1504 [No Label] 0 ms
B 25 120.100.67.6 MRU 1504 [No Label] 4 ms
B 26 120.100.67.6 MRU 1504 [No Label] 0 ms
B 27 120.100.67.6 MRU 1504 [No Label] 0 ms
B 28 120.100.67.6 MRU 1504 [No Label] 0 ms
B 29 120.100.67.6 MRU 1504 [No Label] 4 ms
B 30 120.100.67.6 MRU 1504 [No Label] 0 ms
R7#
```

```
R4# show run int lo0
```

```
interface Loopback0
 ip address 4.4.4.4 255.255.255.0
 ip pim sparse-mode
end
```

```
R7# show run int lo0
interface Loopback0
 ip address 7.7.7.7 255.255.255.0
 ip pim sparse-mode
end
```

```
R5# show run int lo0
interface Loopback0
 ip address 5.5.5.5 255.255.255.0
end
```

```
R6# show run int lo0
interface Loopback0
 ip address 6.6.6.6 255.255.255.0
end
```

```
R4(config)# int lo0
R4(config-if)# ip add 4.4.4.4 255.255.255.255
```

```
R5(config)# int lo0
R5(config-if)# ip add 5.5.5.5 255.255.255.255
```

```
R6(config)# int lo0
R6(config-if)# ip add 6.6.6.6 255.255.255.255
```

```
R7(config)# int lo0
R7(config-if)# ip add 7.7.7.7 255.255.255.255
```

```
R4# traceroute mpls ipv4 7.7.7.7/32
Tracing MPLS Label Switched Path to 7.7.7.7/32, timeout is 2 seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
       'L' - labeled output interface, 'B' - unlabeled output interface,
       'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
       'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
       'P' - no rx intf label prot, 'p' - premature termination of LSP,
       'R' - transit router, 'I' - unknown upstream index,
       'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
 0 120.100.46.4 MRU 1500 [Labels: 16 Exp: 0]
 I 1 120.100.46.6 MRU 1504 [Labels: implicit-null Exp: 0] 0 ms
! 2 120.100.67.7 1 ms
```

```
R7# traceroute mpls ipv4 4.4.4.4/32
Tracing MPLS Label Switched Path to 4.4.4.4/32, timeout is 2 seconds
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
       'L' - labeled output interface, 'B' - unlabeled output interface,
       'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
       'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
       'P' - no rx intf label prot, 'p' - premature termination of LSP,
       'R' - transit router, 'I' - unknown upstream index,
       'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
 0 120.100.67.7 MRU 1500 [Labels: 18 Exp: 0]
 I 1 120.100.67.6 MRU 1504 [Labels: implicit-null Exp: 0] 0 ms
! 2 120.100.46.4 1 ms
```

```
R7#
```

```
R4# ping vrf ROUTERS 120.100.27.7 source 120.100.4.4
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 120.100.27.7, timeout is 2 seconds:
```

```
Packet sent with a source address of 120.100.4.4
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

```
R7# ping vrf ROUTERS 120.100.4.4 source 120.100.27.7
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 120.100.4.4, timeout is 2 seconds:
```

```
Packet sent with a source address of 120.100.27.7
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Incident 5

Now that the MPLS issues appear to have been fixed, users on vrf ROUTERS on R2 and R3 are complaining that cannot see the route to VLAN 27.

Investigate and prove connectivity with a valid route and successful ping from R2 and R3 user subnets to the Sw2 VLAN 27 interface as follows:

```
R2# ping 120.100.27.2 source 120.100.2.2
```

```
Type escape sequence to abort
```

```
Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:
```

```
Packet sent with a source address of 120.100.2.2
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/31/36 ms
```

```
R3# ping 120.100.27.2 source 120.100.3.3
```



```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:
Packet sent with a source address of 120.100.3.3
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/32/36 ms
```

3 points

You have just proven a vrf ping from 120.100.4.0/24 to 120.100.27.0/24 over the MPLS network, so you can be confident that this is not an MPLS issue and instead appears to be an issue with PE router R4 not propagating the BGP-learned route of 120.100.27.0/24 over EIGRP to R2 and R3 within the vrf. Example 1-11 shows that the route is not present in R1, so there is no point in checking R2 and R3 further downstream. The redistribution configuration is examined on R4, and you should be able to tell (even if you are not an MPLS guru) that there is a default metric missing on the redistribution of BGP into EIGRP on the appropriate vrf. It is worth checking the configuration of the remote PE also at this point, either for verification of the default metric or just to see whether this is also missing here (which it is). Example 1-11 shows the configuration of the EIGRP default metric in both PE routers and the resulting router verification and extended ping from R2 and R3 to Sw2's VLAN 27 interface. A correct resolution to this problem is worth 3 points.

Example 1-11 *R4 and R7 PE VRF Route Verification and Default Metric Configuration*

R4# **show ip route vrf ROUTERS**

```

    120.0.0.0/24 is subnetted, 6 subnets
B       120.100.27.0 [200/0] via 7.7.7.7, 6d19h
D       120.100.14.0 [90/284160] via 120.100.4.1, 1w0d, GigabitEthernet0/1.4
C       120.100.4.0 is directly connected, GigabitEthernet0/1.4
D       120.100.2.0 [90/2300416] via 120.100.4.1, 6d20h, GigabitEthernet0/1.4
D       120.100.3.0 [90/2198016] via 120.100.4.1, 1w0d, GigabitEthernet0/1.4
D       120.100.123.0 [90/2172416] via 120.100.4.1, 1w0d, GigabitEthernet0/1.4
```

R1# **show ip route**

```

10.0.0.0/26 is subnetted, 1 subnets
```

```
C      10.10.10.0 is directly connected, Ethernet0/0
      120.0.0.0/24 is subnetted, 5 subnets
C      120.100.14.0 is directly connected, Ethernet3/0.14
C      120.100.4.0 is directly connected, Ethernet3/0
D      120.100.2.0 [90/2297856] via 120.100.123.2, 6d20h, Serial2/0
D      120.100.3.0 [90/2195456] via 120.100.123.3, 1w2d, Serial2/0
C      120.100.123.0 is directly connected, Serial2/0
```

```
R4# show run | section address-family ipv4 vrf ROUTERS
address-family ipv4 vrf ROUTERS
  redistribute bgp 65001
  network 120.100.4.0 0.0.0.255
  no auto-summary
  autonomous-system 1
```

```
R4(config)# router eigrp 10
R4(config-router)# address-family ipv4 vrf ROUTERS
R4(config-router-af)# default-metric 1000 1 1 1 1500
```

```
R7# show run | section address-family ipv4 vrf ROUTERS
address-family ipv4 vrf ROUTERS
  redistribute bgp 65001
  network 120.100.27.0 0.0.0.255
  eigrp stub connected summary
  no auto-summary
  autonomous-system 1
```

```
R7(config)# router eigrp 1
R7(config-router)# address-family ipv4 vrf ROUTERS
R7(config-router-af)# default-metric 1000 1 1 1 1500
```

```
R1# show ip route eigrp
```

```
120.0.0.0/24 is subnetted, 6 subnets
D    120.100.27.0 [90/284160] via 120.100.4.4, 00:00:17, Ethernet3/0
D    120.100.2.0 [90/2297856] via 120.100.123.2, 6d20h, Serial2/0
D    120.100.3.0 [90/2195456] via 120.100.123.3, 1w2d, Serial2/0
```

```
R2# ping 120.100.27.2 source 120.100.2.2
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:

Packet sent with a source address of 120.100.2.2

```
!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 28/31/36 ms

```
R3# ping 120.100.27.2 source 120.100.3.3
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.27.2, timeout is 2 seconds:

Packet sent with a source address of 120.100.3.3

```
!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 32/32/36 ms

Incident 6

Users on vrf ROUTERS on R2 and R3 are complaining they can see only the route to VLAN 27 on Sw2 and not VLAN 20. Investigate the issue and rectify it.

Prove connectivity by extended ping from user subnets on R2 and R3 to Sw2 VLAN 20 as follows:

```
R2# ping 120.100.20.2 source 120.100.2.2
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.20.2, timeout is 2 seconds:

Packet sent with a source address of 120.100.2.2

```

!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/32/36 ms

R3# ping 120.100.20.2 source 120.100.3.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.20.2, timeout is 2 seconds:
Packet sent with a source address of 120.100.3.3
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/32/32 ms

```

2 points

You have proven the vrf-specific MPLS network previously, so you can be confident that this is not an MPLS issue and that it is a good idea to use the split-half method of troubleshooting and divide your network into two to home in on the issue. Because you are confident of the MPLS network, you can focus between Sw2 and the PE Router R7. Example 1-12 begins by looking at the PE to see whether the route is being learned over EIGRP from Sw2. As you can see, the route is indeed present. So, in theory, it should be redistributed into BGP and over MPLS to PE Router R4, but it is not. This should tell you that it is not being redistributed correctly at PE R7. Inspection of the configuration on R7 will reveal that EIGRP has been configured as an EIGRP stub router, not a good choice for a PE router that is required to advertise its CE devices' routes over the MPLS network. After the stub-specific configuration has been removed from R7, the route is propagated correctly, and an end-to-end ping from R2 and R3 to VLAN 20 on Sw2 will secure your points. A correct resolution to this problem, as detailed in Example 1-12, is worth 2 points.

Example 1-12 *R4 and R7 PE VRF Route Verification, Configuration and Testing on R2 and R3*

```

R7# show ip route vrf ROUTERS

      120.0.0.0/24 is subnetted, 7 subnets
C       120.100.27.0 is directly connected, GigabitEthernet0/0.27
D       120.100.20.0 [90/28416] via 120.100.27.2, 6d19h, GigabitEthernet0/0.27
B       120.100.14.0 [200/284160] via 4.4.4.4, 6d18h
B       120.100.4.0 [200/0] via 4.4.4.4, 6d18h
B       120.100.2.0 [200/2300416] via 4.4.4.4, 6d18h
B       120.100.3.0 [200/2198016] via 4.4.4.4, 6d18h

```

```
B      120.100.123.0 [200/2172416] via 4.4.4.4, 6d18h
```

```
R4# show ip route vrf ROUTERS bgp
```

```
120.0.0.0/24 is subnetted, 5 subnets
```

```
B      120.100.27.0 [200/0] via 7.7.7.7, 00:04:12
```

```
R7# show run | section address-family ipv4 vrf ROUTERS
```

```
address-family ipv4 vrf ROUTERS
```

```
network 120.100.27.0 0.0.0.255
```

```
no auto-summary
```

```
autonomous-system 1
```

```
eigrp stub connected summary
```

```
exit-address-family
```

```
R7(config)# router eigrp 1
```

```
R7(config-router)# address-family ipv4 vrf ROUTERS
```

```
R7(config-router-af)# no eigrp stub
```

```
R4# show ip route vrf ROUTERS bgp
```

```
120.0.0.0/24 is subnetted, 7 subnets
```

```
B      120.100.27.0 [200/0] via 7.7.7.7, 6d18h
```

```
B      120.100.20.0 [200/28416] via 7.7.7.7, 00:00:40
```

```
R2# ping 120.100.20.2 source 120.100.2.2
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 120.100.20.2, timeout is 2 seconds:
```

```
Packet sent with a source address of 120.100.2.2
```

```

!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/32/36 ms

R3# ping 120.100.20.2 source 120.100.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 120.100.20.2, timeout is 2 seconds:
Packet sent with a source address of 120.100.3.3
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/32/32 ms

```

Incident 7

Users on Sw3 VLAN30 have reported that they cannot see a route to 120.100.13.0/24 used to access services advertised behind a firewall with a next-hop address of Sw1's Loopback11 interface.

Investigate the issue and ensure that the static route on Sw1 can be seen within Sw3's routing table over OSPF as follows:

```

Sw3# show ip route ospf
 10.0.0.0/26 is subnetted, 2 subnets
O E2   10.20.20.0 [110/20] via 120.100.37.7, 00:20:31, Vlan37
 120.0.0.0/24 is subnetted, 6 subnets
O IA   120.100.12.0 [110/3] via 120.100.37.7, 00:20:36, Vlan37
O E2   120.100.13.0 [110/20] via 120.100.37.7, 00:00:22, Vlan37
O IA   120.100.14.0 [110/2] via 120.100.37.7, 00:20:36, Vlan37
O IA   120.100.11.0 [110/3] via 120.100.37.7, 00:20:36, Vlan37

```

3 points

Okay, so now you are looking at a potential OSPF issue on a different VPN over the MPLS network. You know the MPLS OSPF network, the MP-BGP, and that MPLS is functioning correctly, so any problems are likely to be between the PE and CE devices or with PE route redistribution. There could also (in theory) be an issue with RTs again, but you should have spotted that inconsistency in a previous question (and as detailed in Example 1-8). Again, using the split-half method, it is wise to see whether the route is indeed configured on Sw1 and is being received by the local PE router (120.100.13.0 255.255.255.0 to

Sw1's Loopback11 interface of 120.100.11.1). Example 1-13 shows that the route is indeed present on Sw1 and being redistributed into OSPF but it is not seen on PE router R4 in the vrf-specific routing table. There is, of course, a neighbor relationship between the two devices, but you should spot that the **redistribute** command on Sw1 is missing the **subnets** command and hence only classful subnets would by default be redistributed into OSPF. Once the **subnets** command is configured on Sw1, the route is propagated and present in the PE router R4 but not present in the remote PE router R7. However, PE R7 is receiving routes from PE R4 for networks 120.100.10.0/24, 120.100.11.0/24, and 120.100.14.0/24, so the issue must now be with the specific route to 120.100.13.0/24 as it is redistributed into BGP on PE router R4.

Example 1-13 Sw1 and R4 PE vrf Route Verification

```
Sw1# show ip route static
      120.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
S      120.100.13.0/24 is directly connected, Loopback11
Sw1# show run begin router ospf 1
router ospf 1
  redistribute static
  network 120.100.11.0 0.0.0.255 area 1
  network 120.100.12.0 0.0.0.255 area 1
  network 120.100.14.0 0.0.0.255 area 1

R4# show ip route vrf SWITCHES 120.100.13.0
% Subnet not in table
R4#
R4# show ip ospf neighbor | include 120.100.14.1
120.100.14.1    1    FULL/DR    00:00:39    120.100.14.1    GigabitEthernet0/1.14
R4#

Sw1(config)# router ospf 1
Sw1(config-router)# redistribute static subnets

R4# show ip route vrf SWITCHES 120.100.13.0
Routing entry for 120.100.13.0/24
```

```

Known via "ospf 2", distance 110, metric 20
Tag 999, type extern 2, forward metric 1
Redistributing via bgp 65001
Last update from 120.100.14.1 on GigabitEthernet0/1.14, 00:00:53 ago
Routing Descriptor Blocks:
* 120.100.14.1, from 120.100.14.1, 00:00:53 ago, via GigabitEthernet0/1.14
  Route metric is 20, traffic share count is 1
  Route tag 999

```

```

R7# show ip bgp vpnv4 vrf SWITCHES
BGP table version is 446, local router ID is 7.7.7.7
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```

```

      Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 2:200 (default for vrf SWITCHES)
*>i120.100.11.0/24 4.4.4.4          2      100      0 ?
*>i120.100.12.0/24 4.4.4.4          2      100      0 ?
*>i120.100.14.0/24 4.4.4.4          0      100      0 ?
*> 120.100.33.1/32 120.100.37.3     2      32768    ?
*> 120.100.37.0/24 0.0.0.0          0      32768    ?

```

When you inspect the redistribution configuration on PE Router R4, as detailed in Example 1-14, you can see that the route map labeled OSPF-BGP is used for redistributing OSPF into BGP within the BGP address family. The route map is denying routes received with an assigned TAG value of 999 and forwarding all other routes. Inspection of the static route on Sw1 shows a TAG value of 999 has been assigned, so you must either remove this TAG value or adjust the route map. The quickest and safest way is to remove the TAG value on Sw1 (because the route map might be required in certain circumstances). Example 1-14 details the modification of the static route on Sw1 and the resulting route propagation onto PE router R7 and through the MPLS network to Sw3. You are not requested to confirm connectivity by a ping test in the question. A correct resolution to this problem is worth 3 points.

Example 1-14 *Sw1 Static Route Propagation and Verification*

```

R4# show run | section address-family ipv4 vrf SWITCHES
address-family ipv4 vrf SWITCHES
  redistribute ospf 2 vrf SWITCHES route-map OSPF-BGP
  no synchronization

R4# show run | section route-map OSPF-BGP
redistribute ospf 2 vrf SWITCHES route-map OSPF-BGP
route-map OSPF-BGP deny 10
  match tag 999
route-map OSPF-BGP permit 20

Sw1# show run | include ip route 120.100.13.0
ip route 120.100.13.0 255.255.255.0 Loopback11 tag 999

Sw1(config)# no ip route 120.100.13.0 255.255.255.0 Loopback11 tag 999
Sw1(config)# ip route 120.100.13.0 255.255.255.0 Loopback11

R7# show ip bgp vpnv4 vrf SWITCHES | include 120.100.13.0
*>i120.100.13.0/24 4.4.4.4                20    100    0 ?

Sw3# show ip route ospf
  10.0.0.0/26 is subnetted, 2 subnets
O E2   10.20.20.0 [110/20] via 120.100.37.7, 00:20:31, Vlan37
  120.0.0.0/24 is subnetted, 6 subnets
O IA   120.100.12.0 [110/3] via 120.100.37.7, 00:20:36, Vlan37
O E2   120.100.13.0 [110/20] via 120.100.37.7, 00:00:22, Vlan37
O IA   120.100.14.0 [110/2] via 120.100.37.7, 00:20:36, Vlan37
O IA   120.100.11.0 [110/3] via 120.100.37.7, 00:20:36, Vlan37

```

Incident 8

Users on Sw3 VLAN 37 have reported that they cannot access services hosted on IP address 226.1.1.1 that are hosted on the Sw1 VLAN 14 interface (120.100.14.1). They report that they can no longer even ping this IP address.

Investigate the issue and perform a successful ping from Sw3 to IP address 226.1.1.1 as follows:

```
Sw3# ping 226.1.1.1
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 226.1.1.1, timeout is 2 seconds
Reply to request 0 from 120.100.14.1, 8 ms
```

Users have indicated to you that they believe the RP should be PE router R4 (120.100.14.4) and that MDT is used with a group address of 232.0.0.11.

5 points

Just when you think you are nearing completion of the troubleshooting lab you are presented with a multicast failure, not just a normal one but one over an MPLS network. You have sufficient information to know that Sw1 VLAN 14 is likely to have or should have an IGMP **join** command for 226.1.1.1 and that Multicast Distribution Tree (MDT) is in use with a group address of 232.0.0.11. Start simple with your troubleshooting rather than diving right in. To begin, attempt a multicast ping on the local PE R4 to the multicast destination to see whether the multicast is functioning locally. This may prove PIM adjacency and the IGMP join group on Sw1. Example 1-15 shows the local test is successful. So, you can now move to the next step, which is a ping from PE router R7 initiated from the correct vrf. This test fails, so you now know you are looking at an MPLS core-related multicast issue. You may or may not be familiar with MDT, but you should know that for multicast to function correctly you need a multicast protocol enabled on your interfaces and an RP configured. If you check your PIM neighbor relationships throughout the network, you will see that you have the correct neighbors throughout the MPLS chain. The command **show ip pim mdt bgp** shows you the next hop for the configured MDT as the remote PE on each PE router, so you can assume that this configuration is valid and working. What you should realize is that MDT forms a PIM neighbor relationship between the two PE routers. Example 1-15 shows that this neighbor relationship has not formed. You need to be vrf specific when validating. This PE-to-PE neighbor relationship is formed between PE loopback addresses in the form of a tunnel, so an issue must exist in relation to these. The example shows that PIM has not been enabled on the Loopback0 interfaces of each PE router, either. When enabled, the neighbor relationship is formed, and a successful ping is now achievable to 226.1.1.1 from Sw3. The question could have been harder. For example, required configuration items such as the RP address on R4 itself or the command **ip pim ssm default**

could have been removed from each MPLS router (items that are possibly not well known but still required for the multicast solution to function effectively). A correct resolution to this problem, as detailed in Example 1-15, is worth 5 points.

Example 1-15 Multicast Testing and Configuration

```
R4# ping vrf SWITCHES 226.1.1.1
```

Type escape sequence to abort.

Sending 1, 100-byte ICMP Echos to 226.1.1.1, timeout is 2 seconds:

```
Reply to request 0 from 120.100.14.1, 1 ms
```

```
R7# ping 226.2.2.2
```

Type escape sequence to abort.

Sending 1, 100-byte ICMP Echos to 226.2.2.2, timeout is 2 seconds:

.

```
R4# show ip pim neighbor
```

PIM Neighbor Table

Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
S - State Refresh Capable

Neighbor Address	Interface	Uptime/Expires	Ver	DR Prio/Mode
120.100.45.5	Serial0/2/0	1w3d/00:01:43	v2	1 / S
120.100.46.6	Serial0/2/1	1w3d/00:01:22	v2	1 / S

```
R5# show ip pim neighbor
```

PIM Neighbor Table

Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
S - State Refresh Capable

Neighbor	Interface	Uptime/Expires	Ver	DR
----------	-----------	----------------	-----	----

```

Address                                     Prio/Mode
120.100.45.4   Serial0/1           1w3d/00:01:40   v2   1 / S
120.100.57.7   Serial0/2           1w6d/00:01:36   v2   1 / S

```

R6# **show ip pim neighbor**

PIM Neighbor Table

Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
 S - State Refresh Capable

```

Neighbor      Interface      Uptime/Expires   Ver   DR
Address                                     Prio/Mode
120.100.67.7   Serial0/0       1w6d/00:01:38   v2   1 / S
120.100.46.4   Serial0/1       1w3d/00:01:23   v2   1 / S

```

R7# **show ip pim neighbor**

PIM Neighbor Table

Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
 S - State Refresh Capable

```

Neighbor      Interface      Uptime/Expires   Ver   DR
Address                                     Prio/Mode
120.100.67.6   Serial0/2/0     1w6d/00:01:18   v2   1 / S
120.100.57.5   Serial0/3/0     1w6d/00:01:29   v2   1 / S

```

R4# **show ip pim mdt bgp**

```

Peer (Route Distinguisher + IPv4)      Next Hop
MDT group 232.0.0.11
  2:2:200:7.7.7.7                       7.7.7.7
  2:2:2000:7.7.7.7                      7.7.7.7

```

R7# **show ip pim mdt bgp**

```

Peer (Route Distinguisher + IPv4)      Next Hop
MDT group 232.0.0.11
  2:1:2000:4.4.4.4                       4.4.4.4

```

```
R4# show ip pim vrf SWITCHES neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor      Interface                Uptime/Expires   Ver   DR
Address
120.100.14.1  GigabitEthernet0/1.14    1w3d/00:01:44   v2    1 / S
R4#
```

```
R7# show ip pim vrf SWITCHES neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor      Interface                Uptime/Expires   Ver   DR
Address
120.100.37.3  GigabitEthernet0/0.37    1w1d/00:01:31   v2    1 / S
```

```
R4# show run int lo0
Building configuration...
interface Loopback0
 ip address 4.4.4.4 255.255.255.255
end
```

```
R7# show run int lo0
interface Loopback0
 ip address 7.7.7.7 255.255.255.255
end
```

```
R4(config)# int lo0
R4(config-if)# ip pim sparse-mode
```

```
R7(config)# int lo0
R7(config-if)# ip pim sparse-mode
```

```
R4# show ip pim vrf SWITCHES neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor      Interface          Uptime/Expires   Ver   DR
Address
120.100.14.1  GigabitEthernet0/1.14  1w3d/00:01:31   v2    1 / S
7.7.7.7       Tunnel0                00:00:11/00:01:33 v2    1 / DR S
```

```
R7# sh ip pim vrf SWITCHES neigh
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      S - State Refresh Capable
Neighbor      Interface          Uptime/Expires   Ver   DR
Address
120.100.37.3  GigabitEthernet0/0.37  1w1d/00:01:18   v2    1 / S
4.4.4.4       Tunnel0                00:00:08/00:01:35 v2    1 / S
```

```
Sw3# ping 226.1.1.1
```

Type escape sequence to abort.

Sending 1, 100-byte ICMP Echos to 226.1.1.1, timeout is 2 seconds:

```
Reply to request 0 from 120.100.14.1, 8 ms
```

Incident 9

Your first-line support personnel are complaining that when they telnet to the Loopback0 address of R4 from within the MPLS network that the response is very poor. One reported symptom is the screen intermittently locking up.

Investigate the issue with an aim of restoring normal service.

3 points

You have to consider what is normal service and why it is normal (to a loopback address). When you have specific information such as this, it is worth initiating a Telnet session to a physical address to compare the response seen to that of the loopback interface. In theory, if you have a good response to the physical you should also have a good response to the loopback interface, unless of course there is something specific to the loopback such as a local policy route map or something similar. Example 1-16 shows there is nothing untoward with respect to the loopback interface on R4, but closer inspection of the configuration shows that control plane policing (CoPP) has been configured for Telnet traffic sent to host 4.4.4.4 (Loopback 0). If the command **show policy-map control-plane** is issued and viewed, you can see that some packets are being dropped because of the unrealistic police settings of dropping anything over two packets per second. The police rate can be adjusted, via trial and error, to a relatively low value to ensure that access to R4 via its loopback interface does not result in dropped packets or intermittent screen lockups. Example 1-16 shows that a value of 400 pps with a burst rate of 200 is sufficient for normal service. A correct resolution to this problem (altering the police rate on R4) is worth 3 points.

Example 1-16 *CoPP Testing and Configuration*

```
R4# show run
-
interface Loopback0
  ip address 4.4.4.4 255.255.255.255
  ip pim sparse-mode
-
class-map match-all management
  match access-group 120
-
```

```
policy-map management
class management
  police rate 1 pps burst 2 packets
  conform-action transmit
  exceed-action drop
policy-map control-plane
-
access-list 120 permit tcp any host 4.4.4.4 eq telnet
-
control-plane
service-policy input management

R4# show policy-map control-plane
Control Plane

Service-policy input: management

Class-map: management (match-all)
  60 packets, 2738 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
Match: access-group 120
police:
  rate 1 pps, burst 2 packets
  conformed 24 packets; actions:
    transmit
  exceeded 36 packets; actions:
    drop
  conformed 0 pps, exceed 0 pps

Class-map: class-default (match-any)
  419 packets, 32405 bytes
  5 minute offered rate 1000 bps, drop rate 0 bps
```


Match: any

```
R4(config)# policy-map management
R4(config-pmap)# class management
R4(config-pmap-c)# police rate 400 pps burst 200 packets
R4(config-pmap-c-police)# conform-action transmit
R4(config-pmap-c-police)# exceed-action drop
```

```
R4# show policy-map control-plane
```

Control Plane

Service-policy input: management

```
Class-map: management (match-all)
  348 packets, 15631 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: access-group 120
  police:
    rate 400 pps, burst 200 packets
    conformed 216 packets; actions:
      transmit
    exceeded 0 packets; actions:
      drop
    conformed 0 pps, exceed 0 pps
```

```
Class-map: class-default (match-any)
  1530 packets, 116823 bytes
  5 minute offered rate 1000 bps, drop rate 0 bps
  Match: any
```

Troubleshooting Lab Wrap-Up

So, how did it go? Did you fix the tickets or did you run out of time? If you scored more than 24 points from a potential 30, well done. If you accomplished this within 2 hours or less, you are well on your way to being prepared for tickets that you will likely face during the Troubleshooting section of the real exam.

The questions here might have seemed a little vague (or even misleading), but they closely echo real-world scenarios (in which you are given just limited information and have to work out what is really happening yourself to get to root cause). In the real world, however, you might not have the constraints assigned in these scenarios; they are included here to just ensure you correctly resolve the issues.

CCIE Routing and Switching v4.0 Troubleshooting Practice Labs

Martin Duggan

Technical Editor: **Luc De Ghein**

Copyright © 2012 Pearson Education, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

First Printing October 2011

ISBN-10: 0-13-271155-9

ISBN-13: 978-0-13-271155-5

Warning and Disclaimer

This book is designed to provide information about the CCIE Routing and Switching v4.0 written exam. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc

Trademark Acknowledgments

All terms mentioned in this ebook that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc. cannot attest to the accuracy of this information. Use of a term in this ebook should not be regarded as affecting the validity of any trademark or service mark.

Feedback Information

At Cisco Press, our goal is to create in-depth technical ebooks of the highest quality and value. Each ebook is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members of the professional technical community.

Reader feedback is a natural continuation of this process. If you have any comments on how we could improve the quality of this ebook, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please be sure to include the ebook title and ISBN in your message.

We greatly appreciate your assistance.

Corporate and Government Sales

The publisher offers excellent discounts on this ebook when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact: **U.S. Corporate and Government Sales** 1-800-382-3419 corpsales@pearsontechgroup.com.

For sales outside the United States please contact: **International Sales** international@pearsoned.com



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CDA, CDDP, CCIE, COIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork, Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanel, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)