

Cisco CCIE Routing and Switching v5.0 Troubleshooting Practice Labs, Second Edition

★★★★★ 0 Reviews

by [Martin Duggan](#)

Publisher: [Cisco Press](#)

Release Date: *May 2014*

ISBN: 9780133787016

Topic: [Cisco](#)

About This eBook

ePUB is an open, industry-standard format for eBooks. However, support of ePUB and its many features varies across reading devices and applications. Use your device or app settings to customize the presentation to your liking. Settings that you can customize often include font, font size, single or double column, landscape or portrait mode, and figures that you can click or tap to enlarge. For additional information about the settings and features on your reading device or app, visit the device manufacturer's Web site.

Many titles include programming code or configuration examples. To optimize the presentation of these elements, view the eBook in single-column, landscape mode and adjust the font size to the smallest setting. In addition to presenting code and configurations in the reflowable text format, we have included images of the code that mimic the presentation found in the print book; therefore, where the reflowable format may compromise the presentation of the code listing, you will see a “Click here to view code image” link. Click the link to view the print-fidelity code image. To return to the previous page viewed, click the Back button on your device or app.

Cisco CCIE Routing and Switching v5.0 Troubleshooting Practice Labs

Second Edition

Martin James Duggan, CCIE No. 7942

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

<https://t.me/learningnets>

Cisco CCIE Routing and Switching v5.0 Troubleshooting Practice Labs, Second Edition

Martin James Duggan, CCIE No. 7942

Copyright© 2014 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.

ISBN-13: 978-0-13-378633-0

ISBN-10: 0-13-378633-1

Warning and Disclaimer

This book is designed to provide information about the Cisco CCIE Routing and Switching v5.0 lab 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.

<https://t.me/learningnets>

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 book 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 book should not be regarded as affecting the validity of any trademark or service mark.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department atcorpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact international@pearsoned.com.

Feedback Information

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

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

We greatly appreciate your assistance.

<https://t.me/learningnets>

Publisher: Paul Boger

Associate Publisher: Dave Dusthimer

Business Operation Manager, Cisco Press: Jan Cornelssen

Executive Editor: Brett Bartow

Managing Editor: Sandra Schroeder

Senior Project Editor: Tonya Simpson

Copy Editor: Keith Cline

Technical Editor: Neil Shaw

Editorial Assistant: Vanessa Evans

Cover Designer: Mark Shirar

Composition: Jake McFarland

Proofreader: Debbie Williams



Americas Headquarters

Cisco Systems, Inc.

San Jose, CA

Asia Pacific Headquarters

Cisco Systems (USA) Pte. Ltd.

<http://www.cisco.com>

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 Stadium Vision, 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, CCDA, CCDP, CCIE, CCIP, 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, PowerPanels, 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)

<https://t.me/learningnets>

About the Author

Martin James Duggan, CCIE No. 7942, is a network architect for AT&T. He designs network solutions for customers globally and specializes in data center networking and quality of service (QoS). Martin mentors colleagues through their Cisco qualifications and holds regular internal training classes. Before this, Martin was a network architect for IBM performing IP network designs and global network reviews. Martin has been in the industry for 23 years, focusing on Cisco solutions for the past 14 years. Martin is the author of the Cisco Press *CCIE Routing and Switching v5.0 Practice Labs*, First Edition.

About the Technical Reviewer

Neil Shaw, CCIE No. 28866, is a network consultant for AT&T with more than 16 years of IT industry experience. Specializing in data center and enterprise networking for the past 8 years, Neil's experience includes design, implementation, and troubleshooting of IP networks with key skills in security, routing, and switching.

Dedication

I would like to dedicate this publication to my kids: Anna, James, and Jakey. Anna and James, I really value our time together; we can get some snowboarding in now that this project is complete and maybe do some more high-altitude driving lessons. James, based on your GoPro shooting and editing skills, you are bound to win a BAFTA; and Anna, I expect you to be modeling for and running Hollister anytime soon. Jakey, now that you have progressed to solids I think we can begin CCNA classes. Your mom is going to kill me, though, if your first word is *Cisco*. Just give her one of your heart-melting smiles, and we might just get away with it.

Acknowledgments

I want to thank Brett Bartow for once again providing me with this enviable opportunity to write for Cisco Press (and apologies for the slipping deadlines).

Tonya Simpson, thanks for shaping my work with all of your hard work throughout this project.

To Matt Greaves and Ian Risebrough, who kept the lab up and running even when they had more pressing things to do.

To Neil Shaw, who reviewed my work: Thanks for your attention to detail. It's easy to see how you got your number! I give you 6 months in that TT before you look at a new M3.

Contents at a Glance

[Introduction](#)

[Troubleshooting Lab 1 \(The Warmup\)](#)

[Troubleshooting Lab 2 \(Network Down!\)](#)

[Summary](#)

Contents

Introduction

Troubleshooting Lab 1 (The Warmup)

Equipment List

Setting Up the Lab

Lab Topology

Switch Instructions

IP Address Instructions

Pre-Lab Tasks

General Guidelines

Troubleshooting Lab 1

Section 1

Incident 1

Incident 2

Incident 3

Incident 4

Incident 5

Incident 6

<https://t.me/learningnets>

Incident 7

Incident 8

Incident 9

“Ask the Proctor”

Section 1

Incident 1

Incident 2

Incident 3

Incident 4

Incident 5

Incident 6

Incident 7

Incident 8

Incident 9

Troubleshooting Lab 1 Debrief

Section 2

Incident 1

Incident 2

Incident 3

Incident 4 /learningnets

Incident 5

Incident 6

Incident 7

Incident 8

Incident 9

Troubleshooting Lab Wrap-Up

Troubleshooting Lab 2 (Network Down!)

Equipment List

Setting Up the Lab

Lab Topology

Switch Instructions

IP Address Instructions

Pre-Lab Tasks

General Guidelines

Troubleshooting Lab 2

Section 1

Incident 1

Incident 2

Incident 3

<https://t.me/learningnets>

Incident 4

Incident 5

Incident 6

Incident 7

Incident 8

Incident 9

“Ask the Proctor”

Section 1

Incident 1

Incident 2

Incident 3

Incident 4

Incident 5

Incident 6

Incident 7

Incident 8

Incident 9

Troubleshooting Lab 2 Debrief

Section 2

Incident 1 <https://learningnets>

Incident 2

Incident 3

Incident 4

Incident 5

Incident 6

Incident 7

Incident 8

Incident 9

Troubleshooting Lab Wrap-Up

Summary

Are You Ready?

Further Reading

Help and Advice

How Can I Schedule My CCIE Lab Exam?

The Day Before

The Day of the Exam

Pass or Fail, What Next?

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).
- *Italic* indicates arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate an optional element.
- 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 on the first attempt because they are not fully prepared; they generally find that their study plan did not match what was expected of them on the exam. This practice exam has been designed to take you as close as possible to taking the actual lab exam. It can show whether you are ready to schedule your lab or whether you need to reevaluate your study plan.

CCIE EXAM OVERVIEW

The CCIE qualification consists of two separate exams: a 2-hour written exam and an 8-hour hands-on lab exam that includes Troubleshooting, Diagnostics, and Configuration sections. 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 that you are ready to take the lab exam; therefore, you are eligible to schedule the lab exam only after you have passed the written exam. Having 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 is a 5.5-hour hands-on exam in which you are required to configure a series of complex scenarios in strict accordance with the questions; it's tough, but doable. Troubleshooting is now included for 2 hours, and you are also presented with a series of further diagnostic questions for a 30-minute period of the exam. You can find the current lab blueprint content information at the following website:

https://learningnetwork.cisco.com/community/certifications/ccie_routing_switching/lab_exam_v5

The Diagnostics section, which has been added to the lab exam, focuses on the skills required to properly

diagnose network issues. The time for this new section is fixed at 30 minutes (no more, no less). The sequence of delivery of the lab sections is fixed as follows: Troubleshooting, Diagnostics, Configuration. Although the time constraint is an essential component of the Troubleshooting section, the system enables some flexibility by granting candidates the option to extend up to 30 minutes to complete the section. The system does not permit toggling between modules, however. To limit the total exam time to 8 hours, the optional extra time used in the Troubleshooting section is automatically deducted from the time credit allotted for the Configuration section. In contrast, if you spend less than 2 hours in the Troubleshooting section, the Configuration section is credited with the time gained.

SCORING POINT SYSTEM

On the actual exam, a higher number of available points for a question usually indicates that the required solution may take more time to achieve or that multiple lines of configuration might be involved. If you find yourself running short on time, pick questions you think you can answer successfully because each question is unrelated to the next.

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 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 outlined 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 materials 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

Make sure that you can navigate Cisco documentation with confidence, because you will have limited access to [Cisco.com](https://www.cisco.com) when you take the lab exam.

HOME LABS

Although acquiring a personal home lab is ideal, gathering all the equipment you will need can be costly.

CISCO 360 PROGRAM

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

- 1. Assessment:** Students take a diagnostic pre-assessment lab to benchmark their knowledge of various

<https://t.me/learningnets>

networking topics.

2. Planning: Based on the pre-assessment, students create a learning plan that uses a mix of learning components to focus their study.

3. Learning: Students learn by participating in lessons and lectures, reading materials, and working with peers and instructors.

4. Practice: Students use the practice exercises to apply learning on actual network equipment.

5. Mastery: Students measure their understanding by completing assessments of knowledge and skill for various approaches to solving network problems.

6. 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 Cisco 360 Learning 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:

- 3925 series routers - IOS 15.3(T) – Advanced Enterprise Services
- Catalyst 3560X series switches running IOS Version 15.0S – Advanced IP Services

TROUBLESHOOTING INFORMATION AND TECHNIQUES

The Troubleshooting section tests your skills to resolve networking problems using IOS troubleshooting tools and by applying troubleshooting methodologies on any topic that can be found within the v5.0 lab blueprint. 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 fault or faults. The section lasts for 2 hours, with a visible countdown and a 15-minute warning. The questions are symptom based, and the verification will be result based with any stipulated constraints. 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 on to the Configuration section; if you do so, though, you cannot go back to the Troubleshooting section.

Be aware that the routers and switches within the Troubleshooting section of the CCIE lab are virtual devices, and so you will always see interfaces shown as up/up rather than up/down. The topology is formed from a large number of these virtual devices, allowing each incident to be individual and not related to the previous incident. Per-incident constraints are stated within each incident clearly, and no partial credits are awarded. You are required to troubleshoot the incident, provide a solution to the incident by configuration, and then verify the solution.

An incident format is as follows:

Symptoms

- BGP routes are not being discovered on Sw3 in AS 65513.

Validation Test

- Fix the problem so that Sw3 has the following BGP table entry:

[Click here to view code image](https://t.me/learningnets)

<https://t.me/learningnets>

```
Sw3# show ip bgp
```

```
BGP table version is 3, local router ID is 120.100.27.2
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
          r RIB-failure, S Stale, m multipath, b backup-path, x
```

```
best-external, f RT-Filter
```

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

Network	Next Hop	Metric	LocPrf	Weight	Path
r>i120.100.11.0/24	4.4.4.4	0	100	0	65513 i
*> 120.100.20.0/24	0.0.0.0	0		32768	i

Additional Guidelines and Constraints

- You are not permitted to modify the route map applied on Sw3 to its iBGP neighbor.

TROUBLESHOOTING PREPARATION

Build a rock-solid knowledge base by theoretical study and completion of as many practice labs as possible.

Create your own strategy as to how you deal with troubleshooting. This could be starting from Layer 1 and working up, Layer 7 and working down, split-half method, and so on. If you have a tried-and-tested strategy, you will find tackling the incidents much easier than if you use a gung-ho approach and should be able to troubleshoot any scenario presented to you.

Spend time with your colleagues or like-minded friends, have them break or introduce configuration errors into your network that you are not familiar with so that you can troubleshoot it. Each time you complete a practice

lab, ask someone to completely sabotage the configuration 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 that you have a clear and concise problem statement that you can translate to a specific network issue. Identify the symptoms. This is where you will need to call on your hard-earned knowledge and experience. You may be able to call out immediately why a neighbor relationship is failing; if you cannot do so, use the tools available to you in the form of **show** and **debug** commands. When you can clearly see the symptom, you should find the solution apparent. If you have multiple potential solutions, work through them in line with your strategy and go full cycle with implementation, verification, and investigation to find the correct solution. You might find that you have strict guidelines based on 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 will not secure your points. In summary, follow these tips:

- Read the entire question before beginning.
 - Think about the root cause before implementing a solution.
 - If you hit a problem, revert to the initial configuration.
 - Do not simply remove sections of configurations where you believe the issue may lie. (For example, modify access control lists [ACLs] rather than remove them.)
 - Do not add static routes to work around a routing issue unless the constraints state that this is a valid option.
 - Ensure that your validation tests confirm you have successfully resolved the solution.
 - Focus on the problem as it is presented to you and not best practice.
- <https://t.me/learningnets>

- Keep it simple.
- Keep track of time.

CHAPTER OVERVIEW

Each chapter consists of questions related to the v5.0 blueprint within the following sections:

- Layer 2 Technologies
- Layer 3 Technologies
- VPN Technologies
- Infrastructure Security
- Infrastructure Services

Chapter 1, “Troubleshooting Lab 1 (The Warmup),” includes virtual private network (VPN) incidents related to Multiprotocol Label Switching (MPLS); and Chapter 2, “Troubleshooting Lab 2 (Network Down!),” includes VPN incidents related to dynamic multipoint VPN (DMVPN). Each chapter begins with an overview of the required equipment and topology, which for simplicity are identical, and provides general guidelines that you should follow while taking the practice exam. The actual troubleshooting practice lab exam is then presented, which you should take in the advised 2-hour window. An “Ask the Proctor” section is included at the end of each question to provide some clues as to how to answer the question (in case you cannot immediately provide an effective solution for a particular question). A full troubleshooting lab debrief is then provided to explain the desired solution for each question, with associated **show** commands to provide full verification of the working solution. Each chapter then closes with a brief summary offering tips and advice.

The appendixes provide initial and final configurations for each router and switch per chapter to allow you to preconfigure and check your final working configurations at the beginning and end of each troubleshooting

practice lab exam.

Troubleshooting Lab 1 (The Warmup)

The CCIE exam commences with 2 hours of troubleshooting “tickets” followed by 5.5 hours of configuration and a final 30 minutes of additional diagnostic questions. This troubleshooting lab has been timed to last for 2 hours, with a total of 30 points, and to 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. You may, of course, continue until you believe that you have met all the objectives, but you will need to increase your speed as you near your lab exam.

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

If you do not own the seven routers and three switches 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 Cisco 360 Learning 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 real lab will use up to approximately 30 virtual devices in which individual incidents have been created.

EQUIPMENT LIST

You need the following hardware and software components to begin this troubleshooting lab.

- Seven routers loaded with Cisco IOS Software Release 15.3T Advanced Enterprise image and the minimum interface configuration as documented in [Table 1-1](#)

Router	Model	Ethernet I/F
R1	3925	2
R2	3925	1
R3	3925	1
R4	3925	2
R5	3925	2
R6	3925	2
R7	3925	2

Table 1-1 *Hardware Required per Router*

- Three 3560X switches with IOS 15.0S 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 this will make life easier should you load configurations directly from the supplied configurations into your own devices.

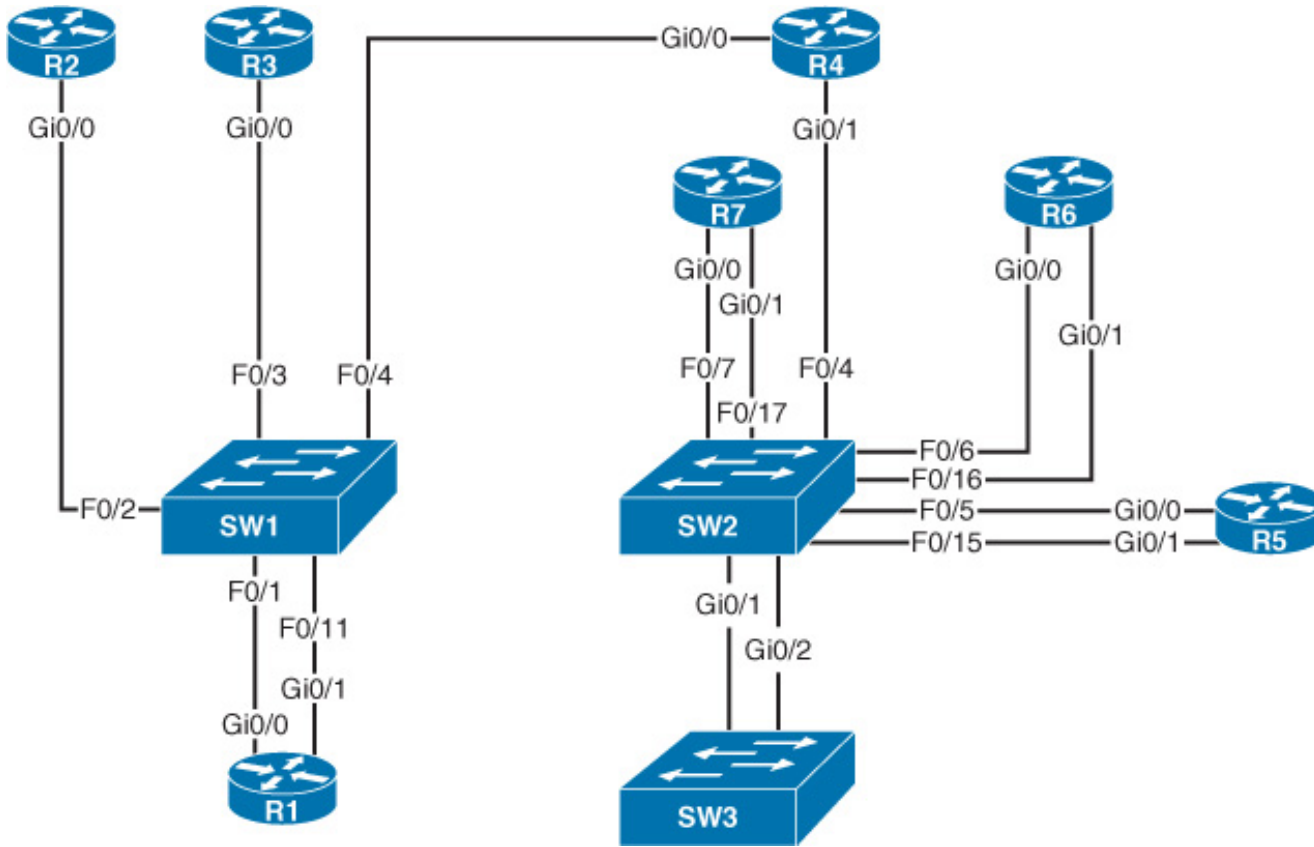


Figure 1-1 Lab 1 Topology Diagram

Note

Use the initial configurations supplied 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 will ensure that you do not get unwanted behavior resulting from differing interior gateway protocol (IGP) metrics.

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

Lab Topology

This troubleshooting exercise uses the topology outlined in [Figure 1-1](#), which you will need to re-create with your own equipment or by using lab equipment on the CCIE R&S 360 program.

Switch Instructions

Configure VLAN assignments from the configurations supplied or from [Table 1-2](#).

VLAN	Switch1	Switch2	Switch3
4	Fa0/4, Fa0/11	—	—
14	Fa0/4	—	—
20	—	Gi0/1, Gi0/2	Gi0/1, Gi0/2
27	—	Gi0/1, Fa0/7	—
45	—	Fa0/4, Fa0/5	—
46	—	Fa0/4, Fa0/6	—
57	—	Fa0/15, Fa0/17	—
67	—	Fa0/16, Fa0/17	—
134	Fa0/1, Fa0/2, Fa0/3	—	—

Table 1-2 VLAN Assignment

Connect your switches together with fiber small form-factor pluggable (SFP) connectors or RJ-45 Ethernet crossover cables, as shown in [Figure 1-2](#).

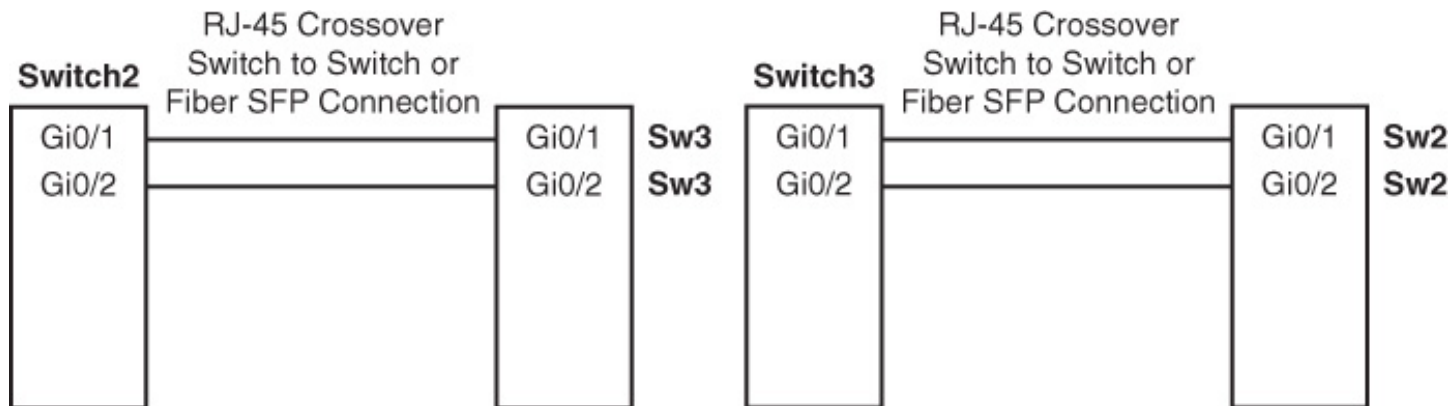


Figure 1-2 Switch-to-Switch Connectivity

Note

For simplicity in this lab, R2 and R3 use loopback interfaces to simulate user-connected interfaces.

IP Address Instructions

For this exercise, you configure your IP addresses as shown in Figure 1-3 or load the initial router configurations supplied.

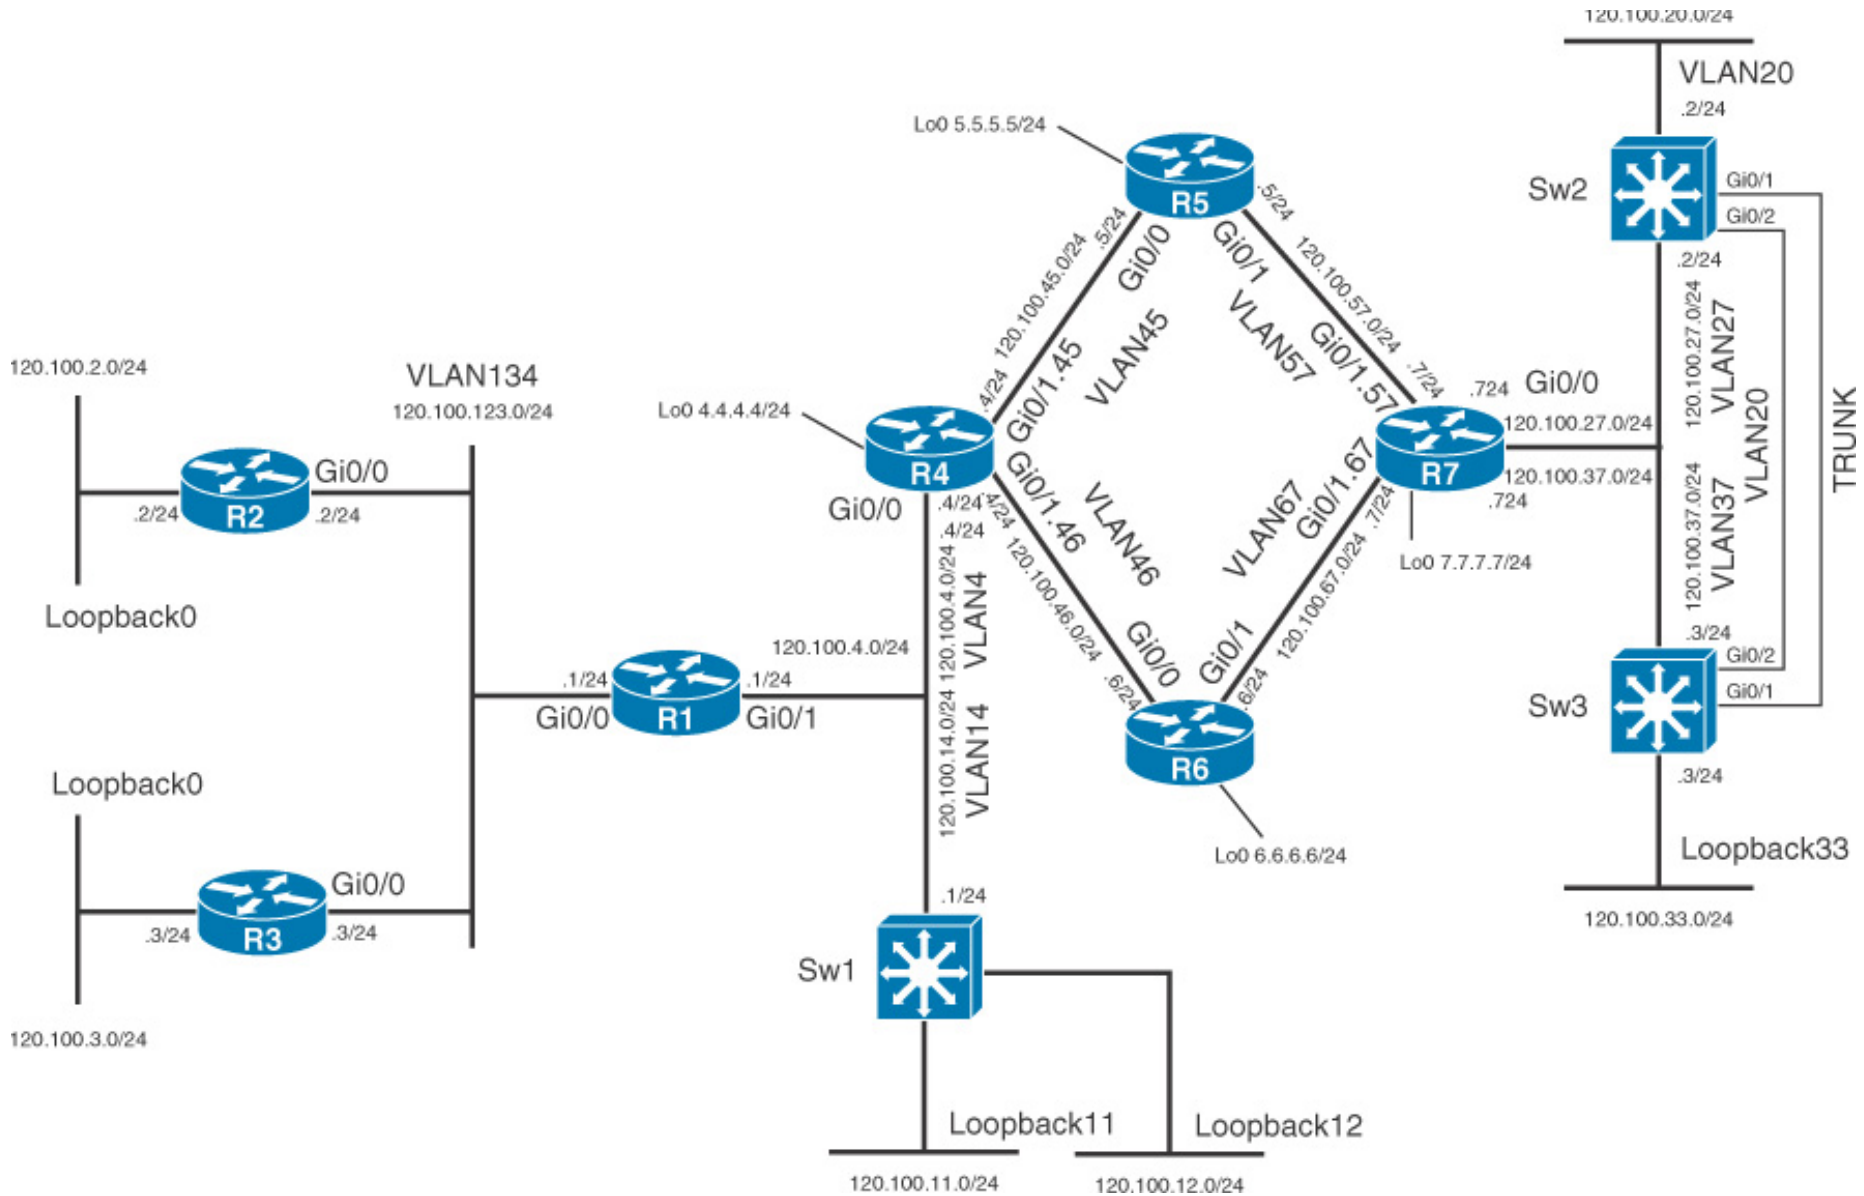


Figure 1-3 IP Addressing Diagram

PRE-LAB TASKS

- Build the lab topology as per [Figure 1-1](#) and [Figure 1-2](#).
- Configure the IP addresses on each router as shown in [Figure 1-3](#). Alternatively, you can load the initial <https://t.me/learningnets>

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.
- Tackle questions sequentially. You might find that one trouble ticket needs to be resolved before moving on to the next ticket. (This will not be the case, however, on your real lab exam, which will have a far greater number of devices with noninterconnected faults.)
- 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 that you follow these items correctly to maximize your score.
- Do not remove any configured feature to resolve an incident. You must resolve the misconfiguration rather than remove 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.)
- Access the latest documentation from <http://www.cisco.com/cisco/web/psa/configure.html>.

Note

Access only the website at the URL shown here, not the whole [Cisco.com](http://www.cisco.com) website (because if you are permitted to use documentation during your CCIE lab exam, your access will be restricted). Well-prepared candidates should not allow themselves to lose time during the exam searching for information.

TROUBLESHOOTING LAB 1

You will now be answering questions in relation to the network topology and virtual private network (VPN) diagram shown in Figures 1-4 and 1-5.

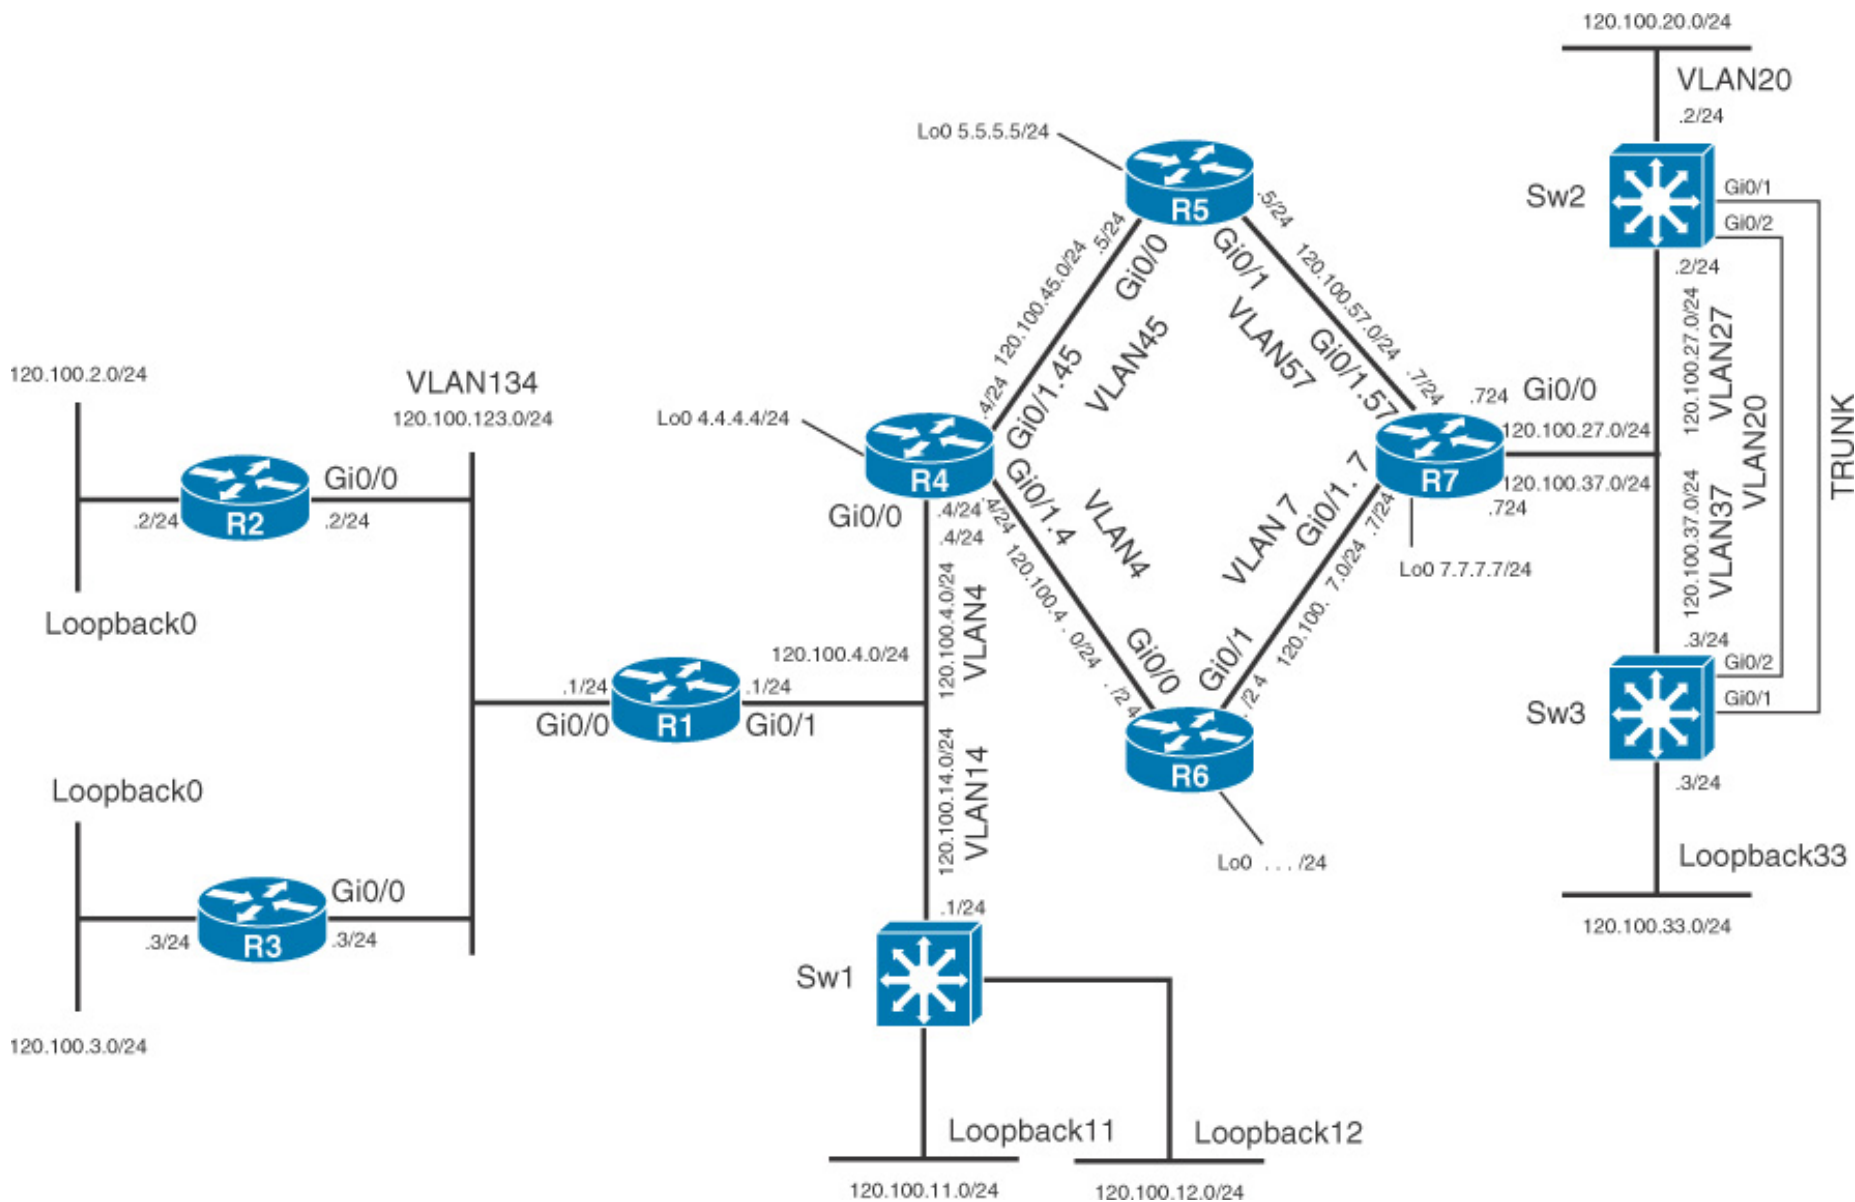


Figure 1-4 Lab Topology Diagram

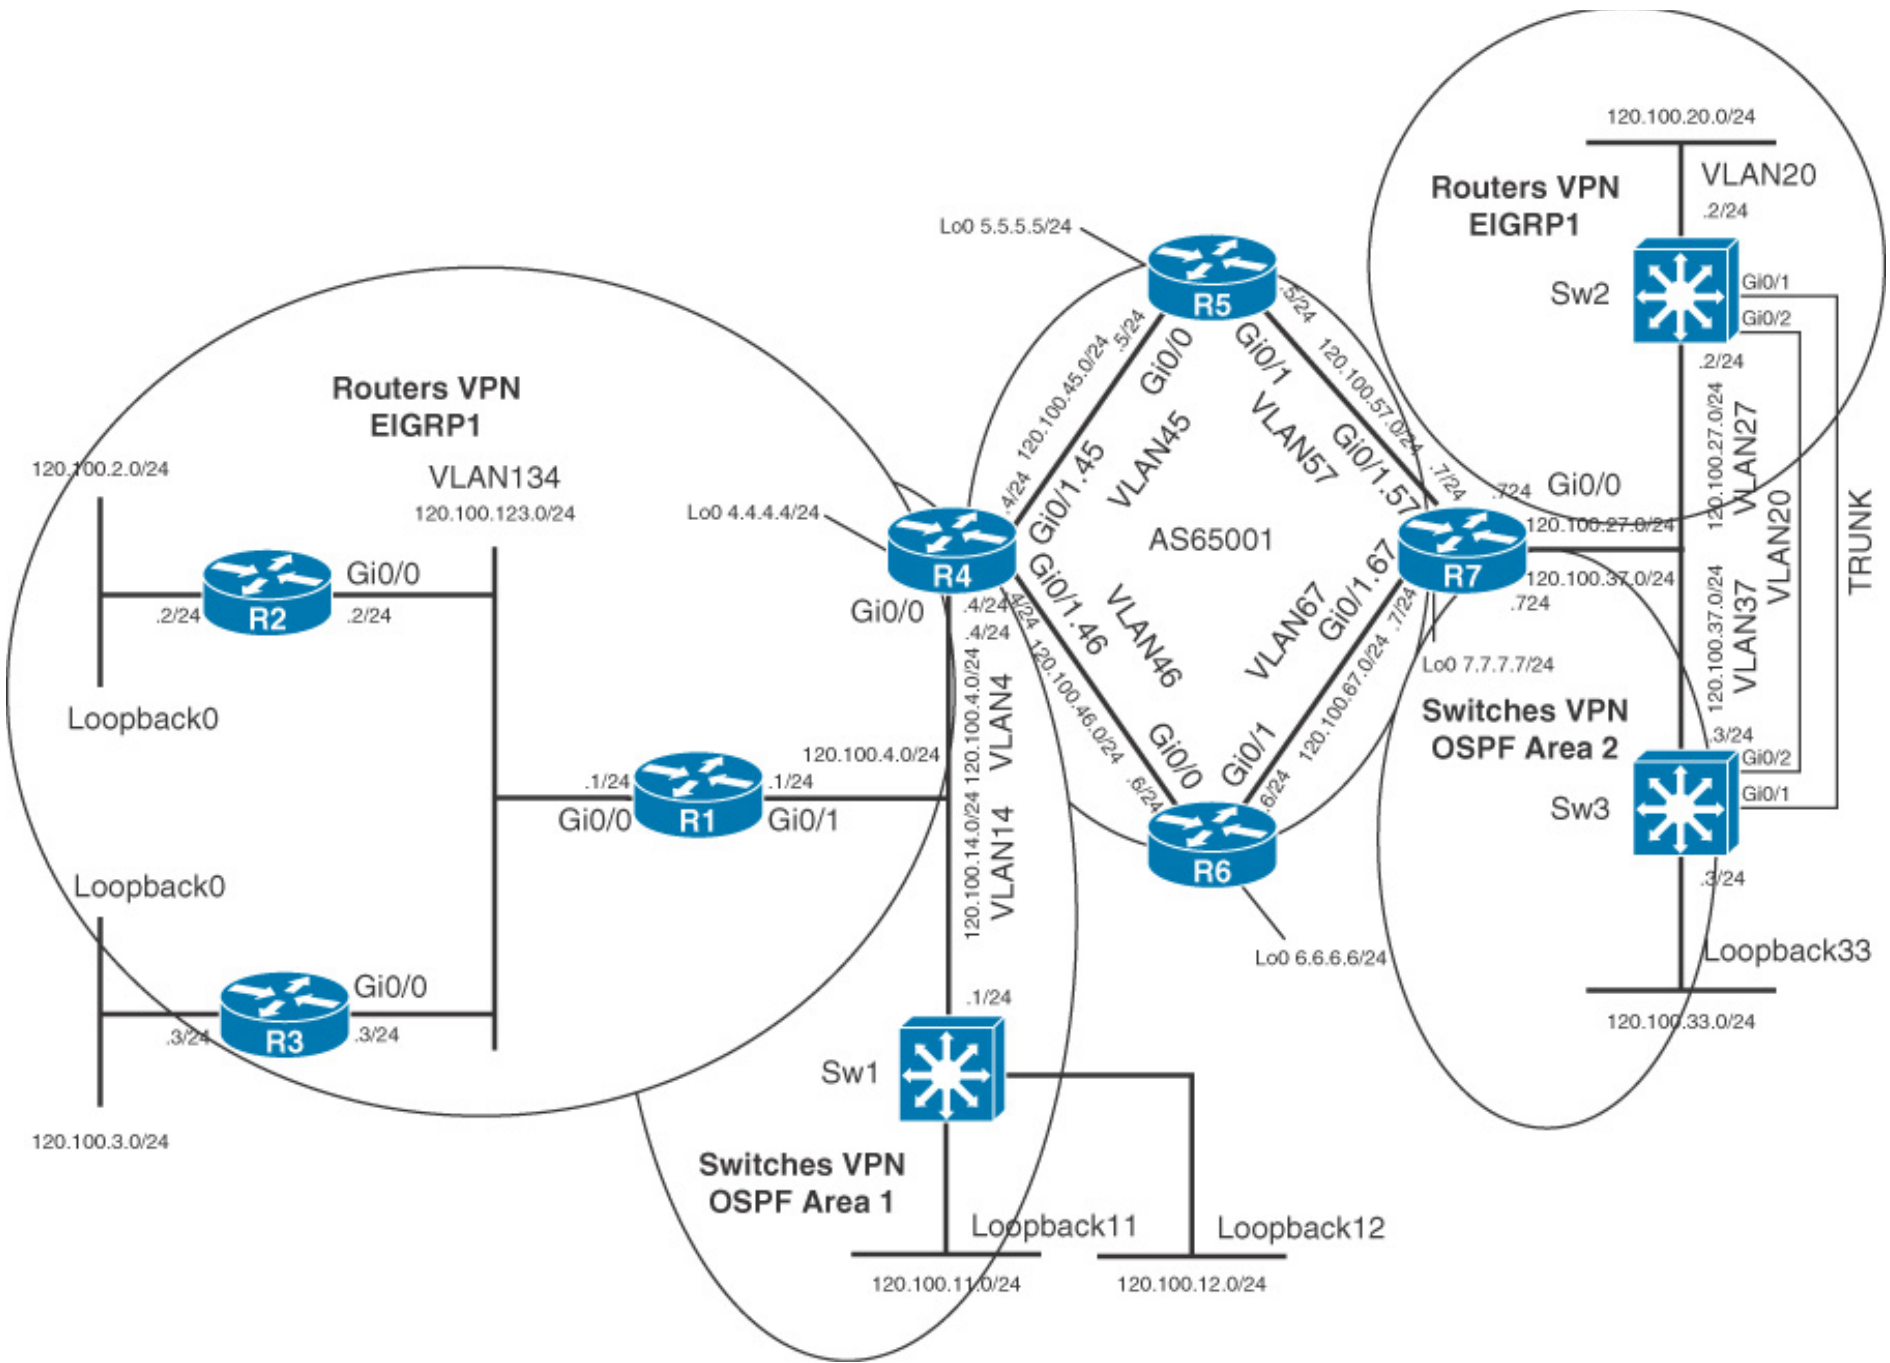


Figure 1-5 VPN Diagram

SECTION 1

Incident 1




Figure 1-6 Incident 1 Diagram

- Users on Sw3 VLAN20 are complaining that they experience poor connectivity to VLAN20 on Sw2. A variable response is seen when pinging from Sw3 VLAN20 to the Sw2 VLAN20 interface.
<https://t.me/learningnets>

- Investigate the issue and fix it. Confirm by proving a successful ping from Sw3 VLAN20 interface 120.100.20.3 to Sw2 VLAN20 interface 120.100.20.2 with a stable response time as follows:

[Click here to view code image](#)

```
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 = 1/5/9 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
```

```
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
```

3 points

Incident 2

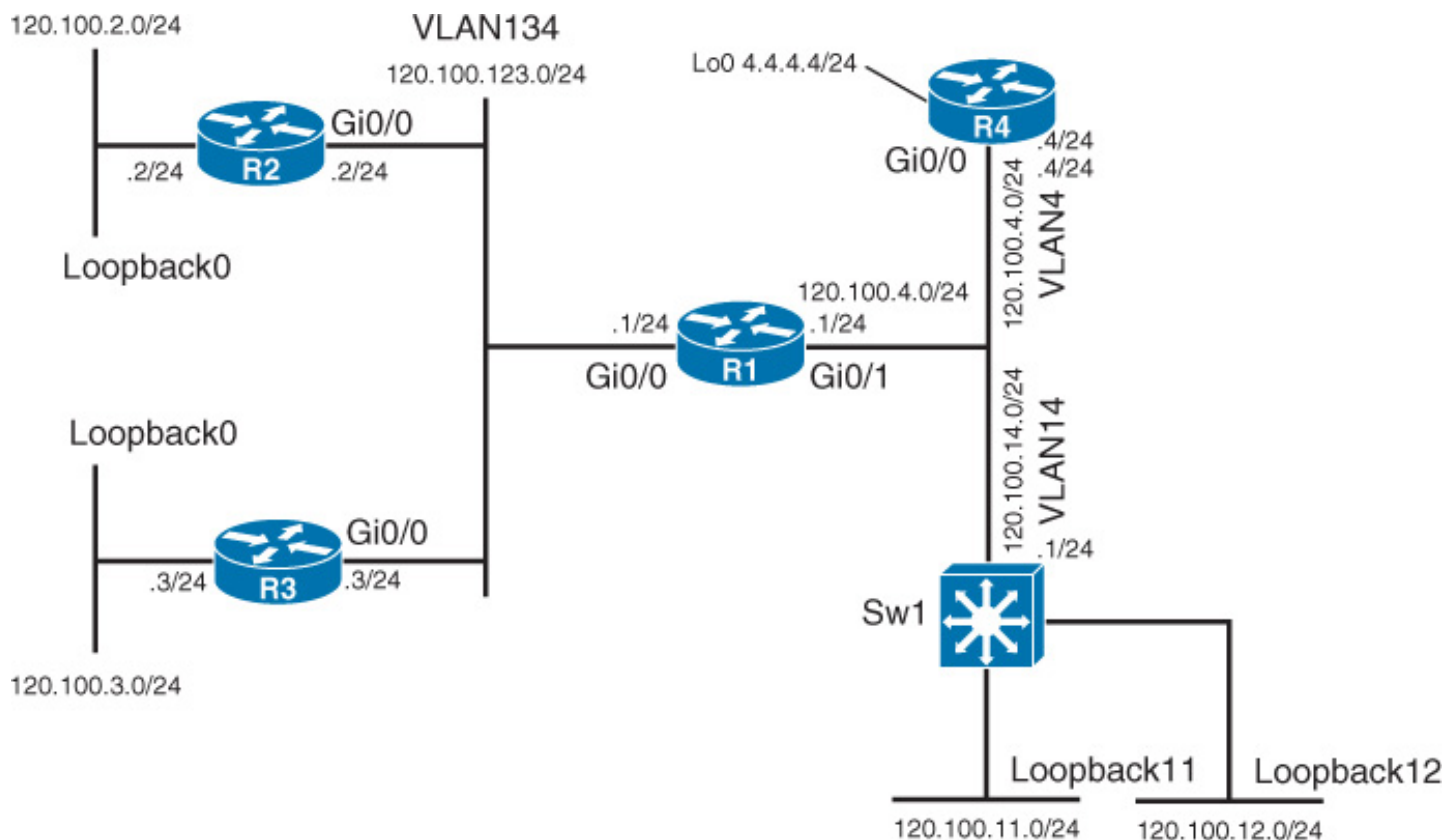


Figure 1-7 Incident 2 Diagram

- Users on R3's 120.100.3.0/24 subnet are complaining that they cannot access resources located on their remote VPN site on VLAN20 (Switch2). Initial investigations have led first-line support personnel to believe that an issue exists with connectivity up to the provider edge (PE) router R4, which was recently replaced.
- Investigate the issue and prove connectivity solely to their PE connection (R4's VLAN4 interface) as follows:

[Click here to view code image](#)

```
R3# ping 120.100.4.4 source 120.100.3.3
```

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 120.100.3.3

!!!!

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

2 points

Incident 3

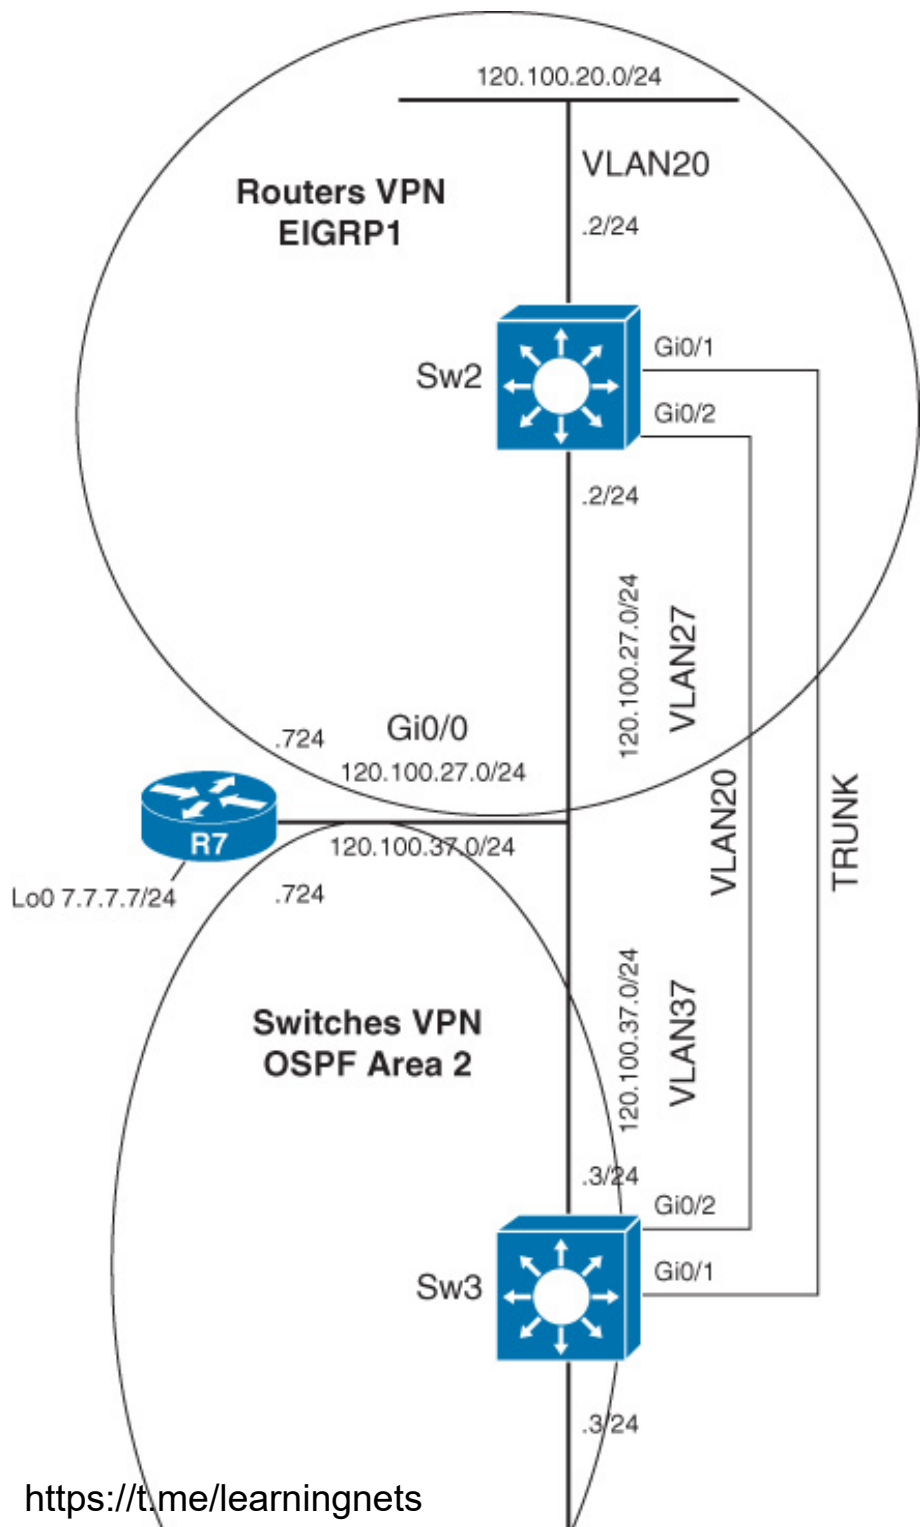




Figure 1-8 Incident 3 Diagram

- Users on Switch2 VLAN20 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 Switch2 and the local PE router R7. Establishment of an EIGRP adjacency between devices will suffice at this point in time, as follows:

[Click here to view code image](#)

```
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

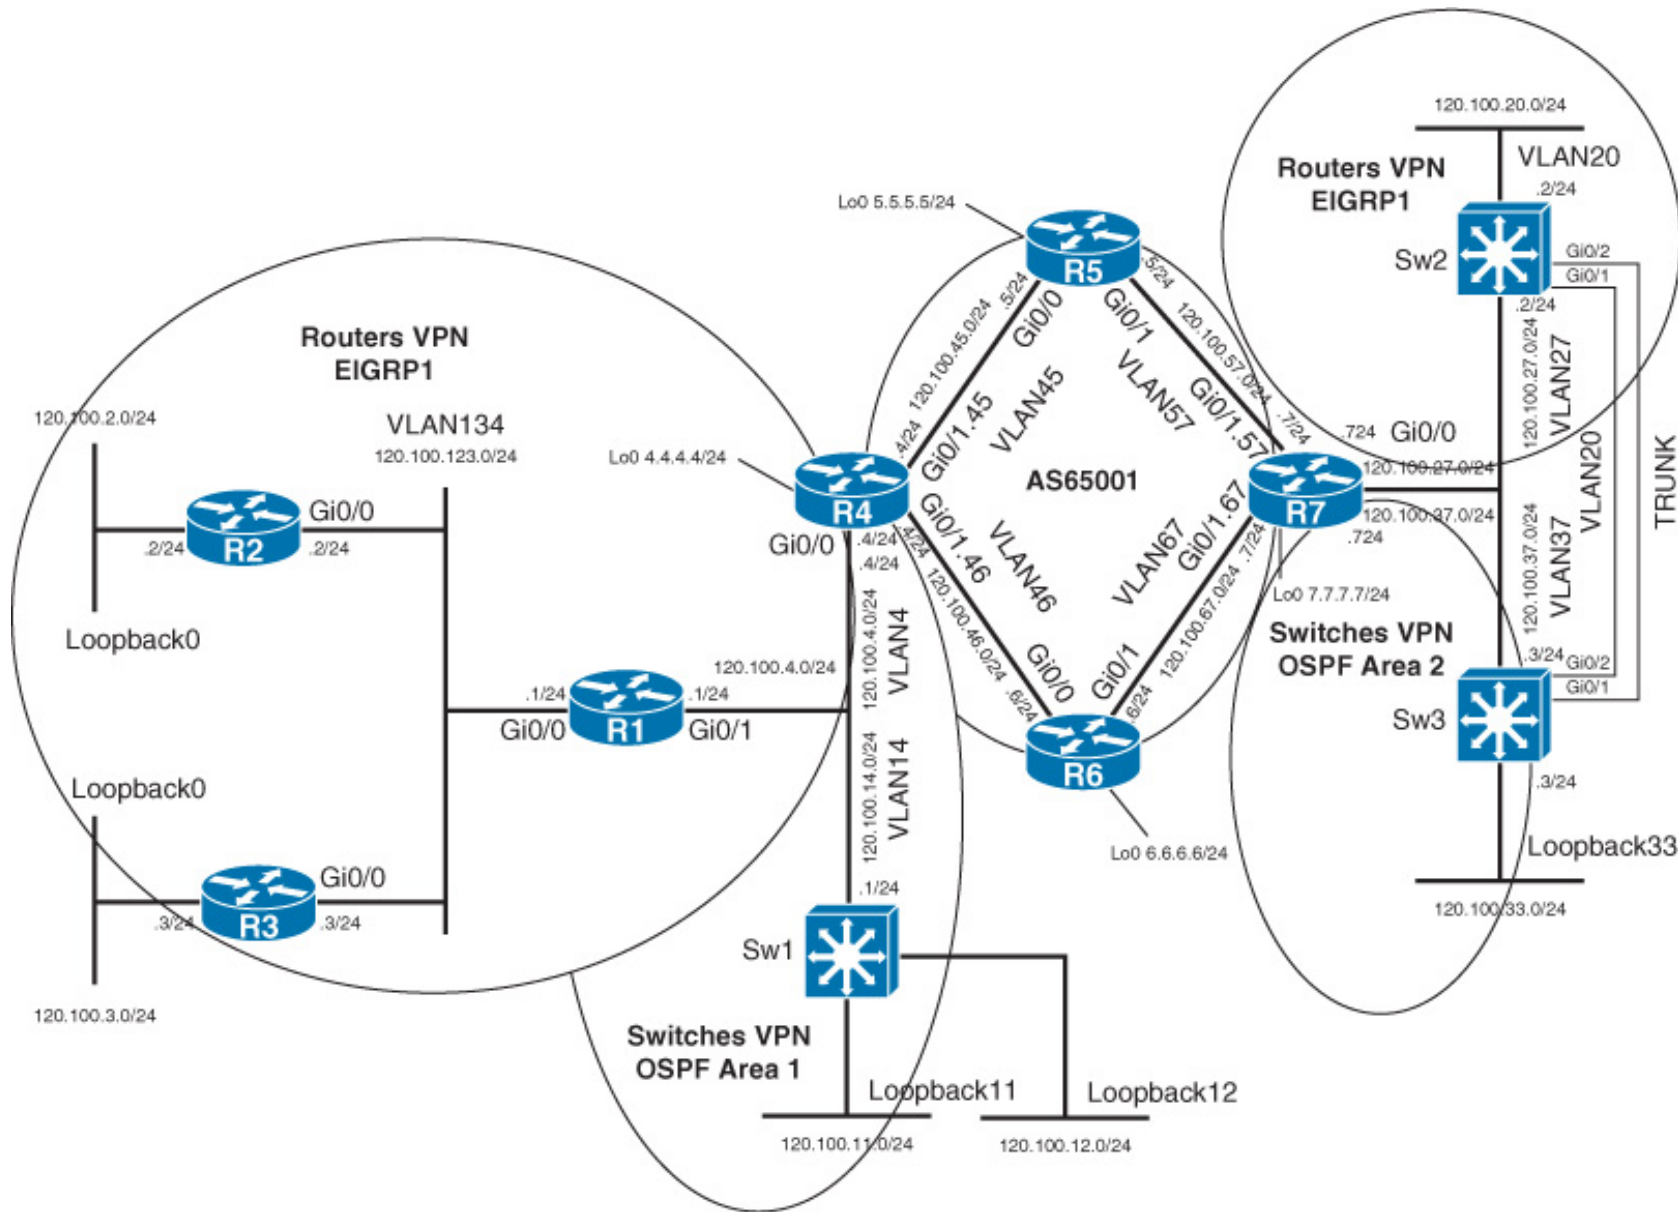


Figure 1-9 Incident 4 Diagram

■ 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 VLAN20 (Switch2). It appears that the MPLS VPN is not functioning correctly.

<https://t.me/learningnets>

- Investigate the issue and prove connectivity between PE routers via an extended ping from the R4 VLAN4 interface and R7 VLAN27 interface to prove MPLS functionality, as follows:

[Click here to view code image](#)

```
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

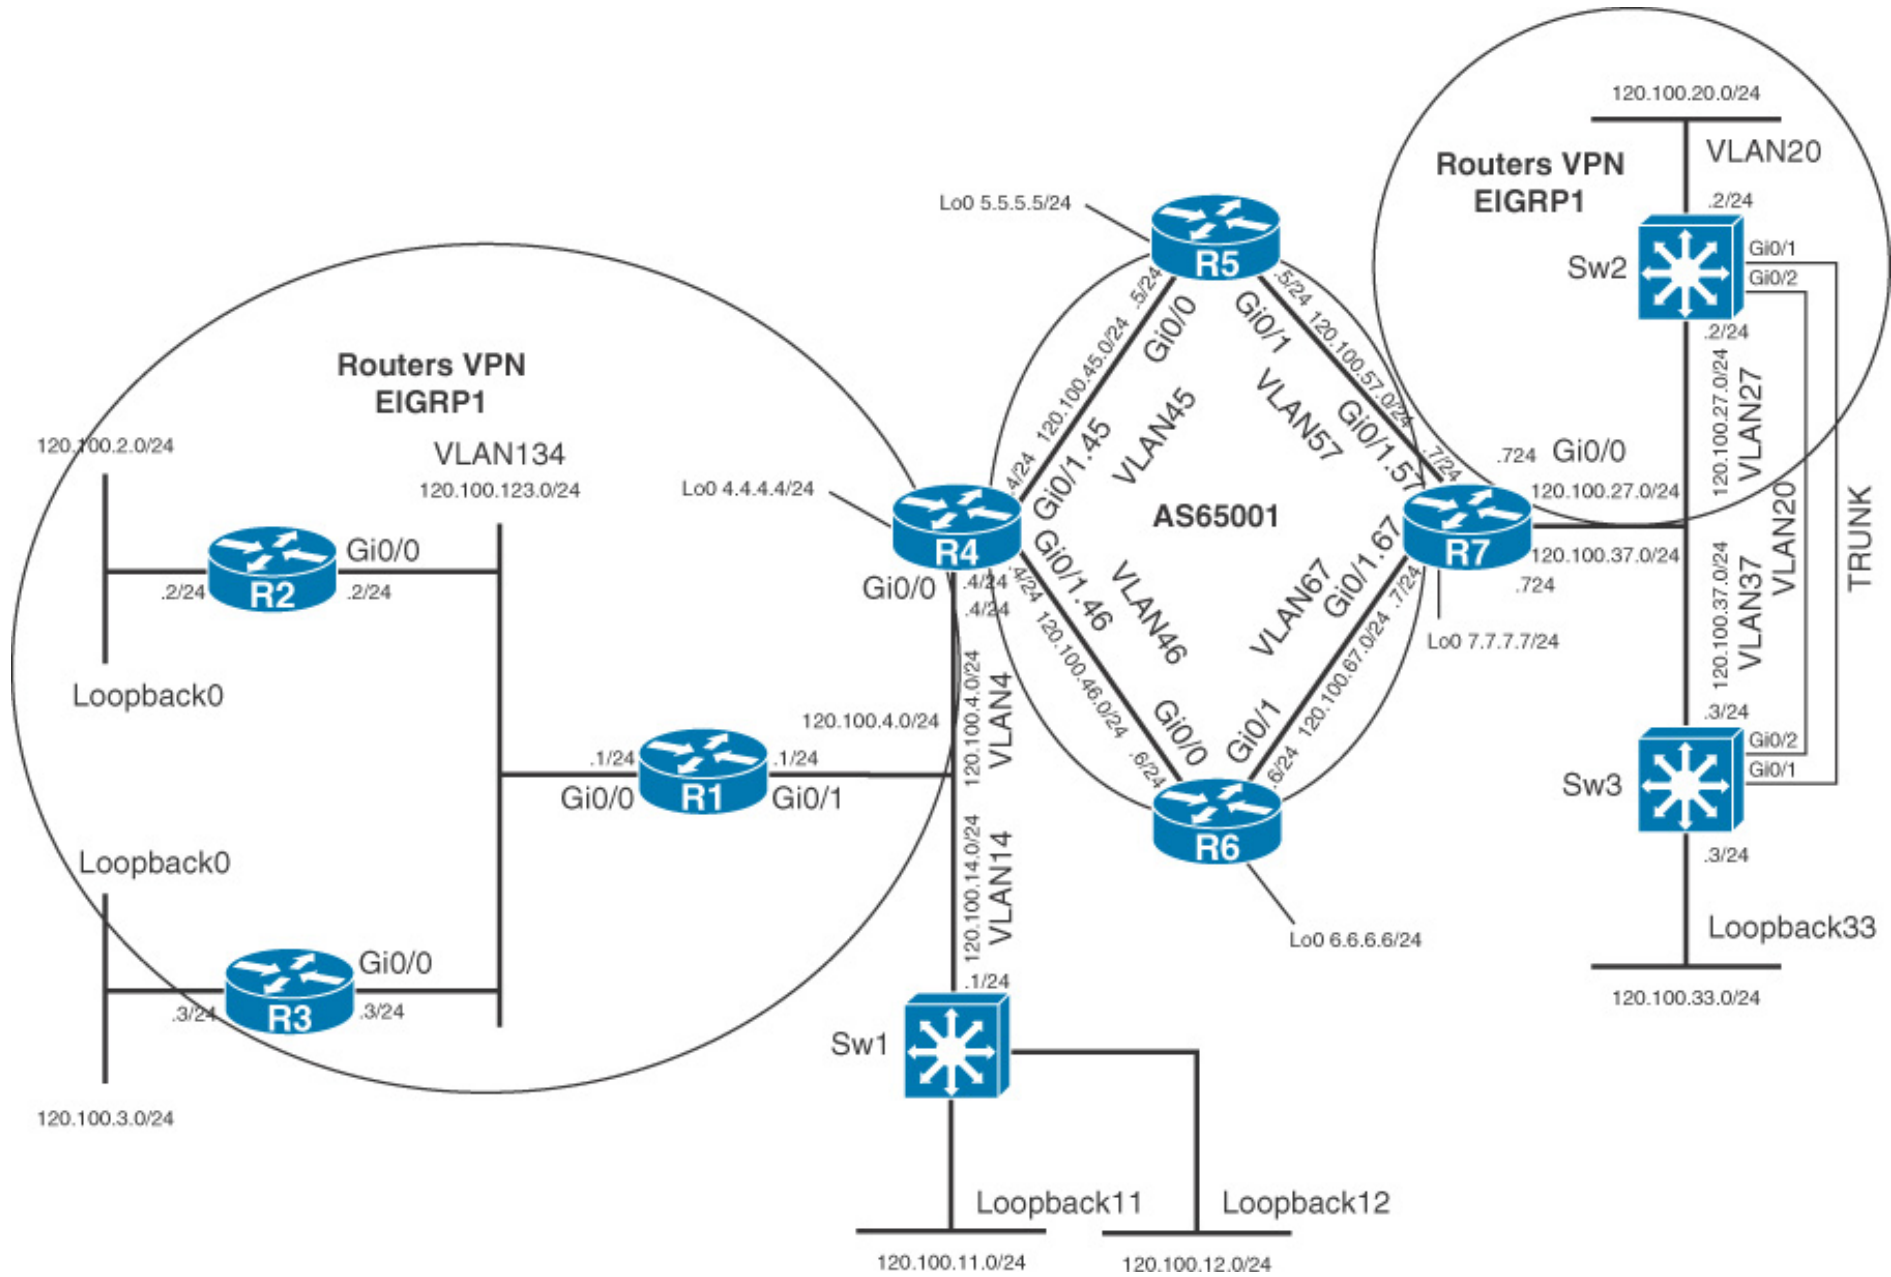


Figure 1-10 Incident 5 Diagram

- Now that the MPLS issues appear to have been fixed, users on virtual routing and forwarding (VRF) routers <https://t.me/learningnets>

on R2 and R3 are complaining that they cannot see the route to VLAN27.

■ Investigate and prove connectivity with a valid route and successful ping from R2 and R3 user subnets to the Sw2 VLAN27 interface, as follows:

[Click here to view code image](#)

```
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

- Users on VRF routers on R2 and R3 are complaining that they can only see the route to VLAN 27 on Sw2 and not VLAN20. Investigate the issue and fix it.
- Prove connectivity by an extended ping from user subnets on R2 and R3 to Sw2 VLAN20, as follows:

[Click here to view code image](#)

```
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

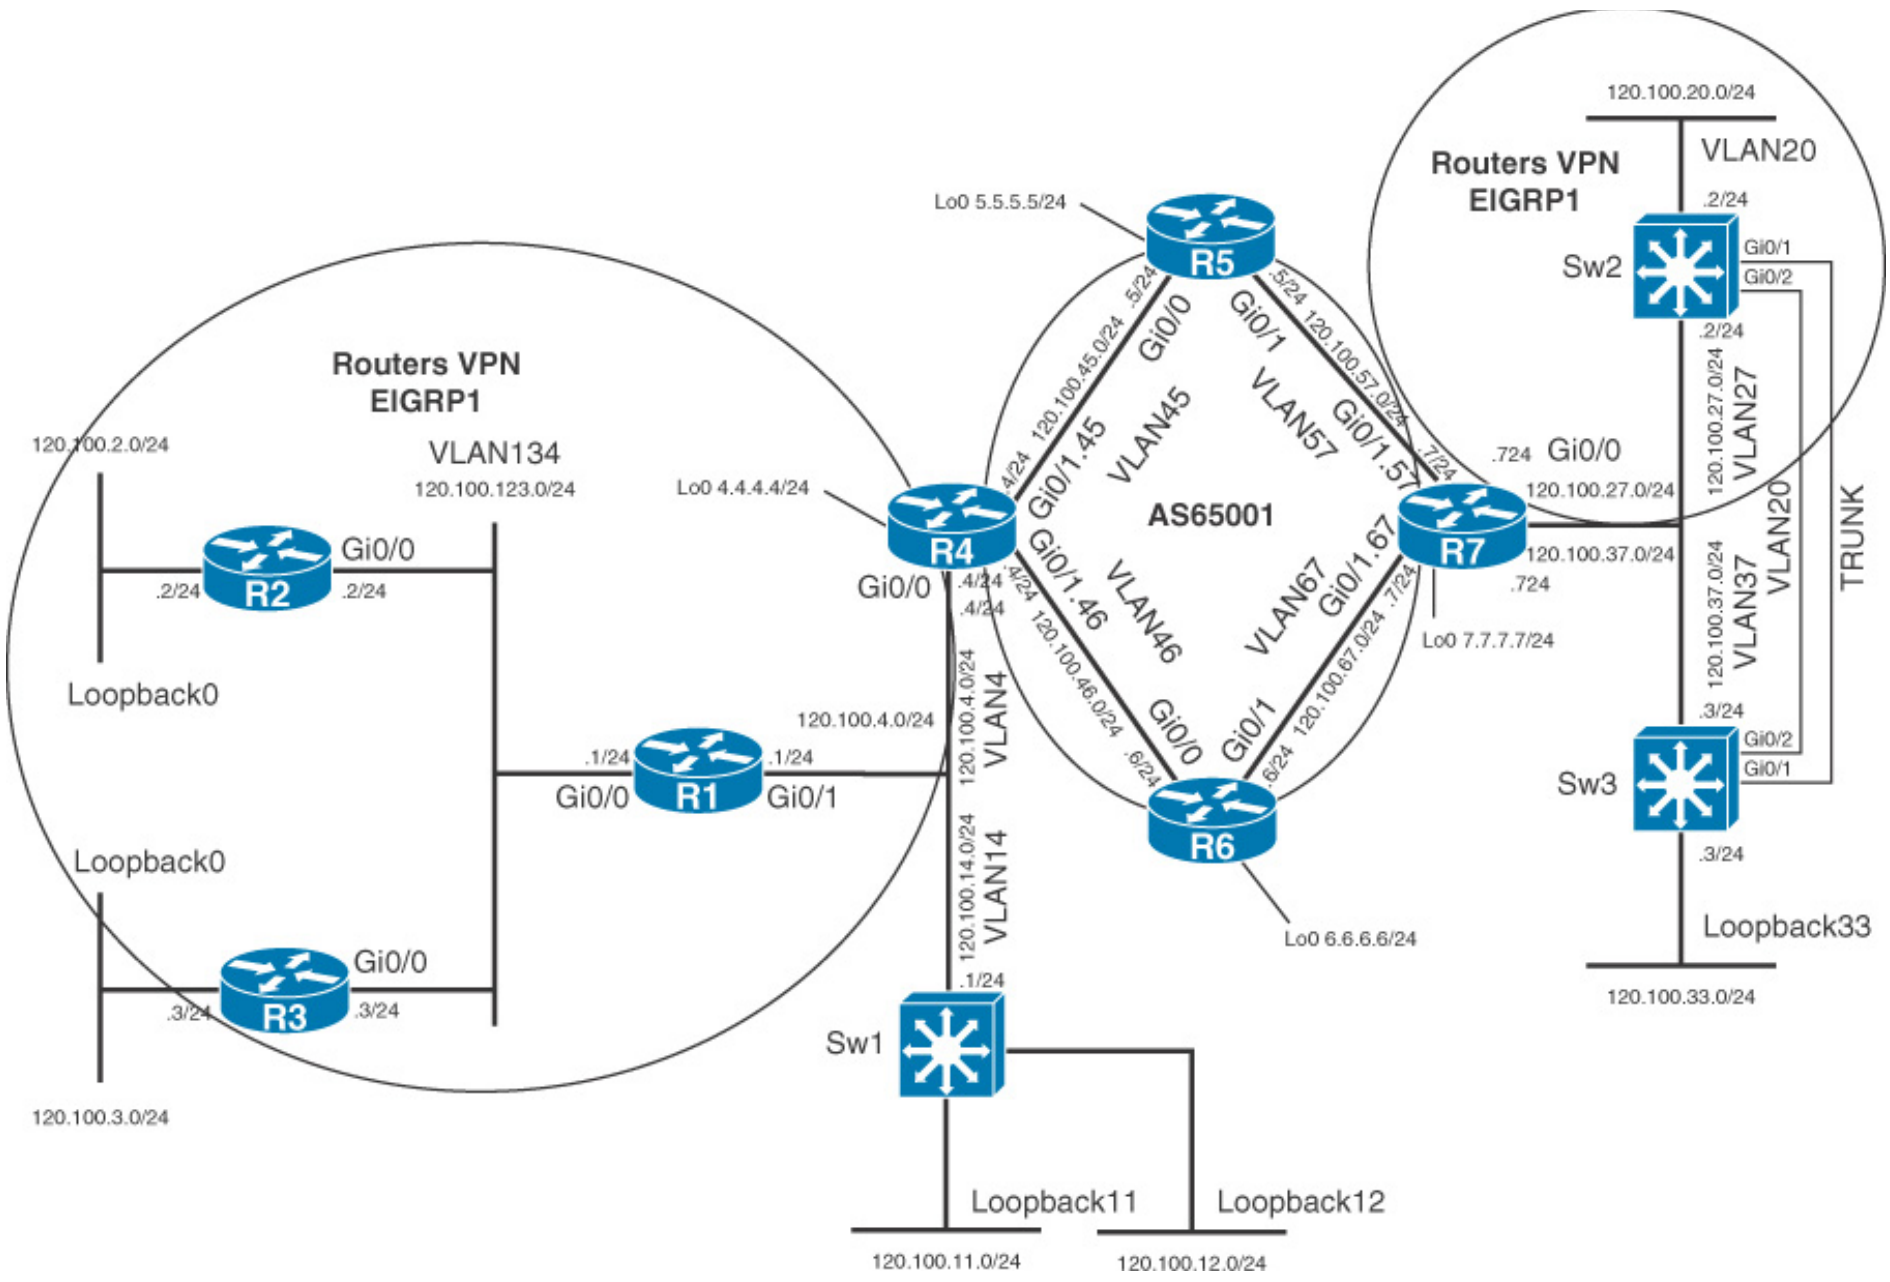


Figure 1-11 Incident 6 Diagram

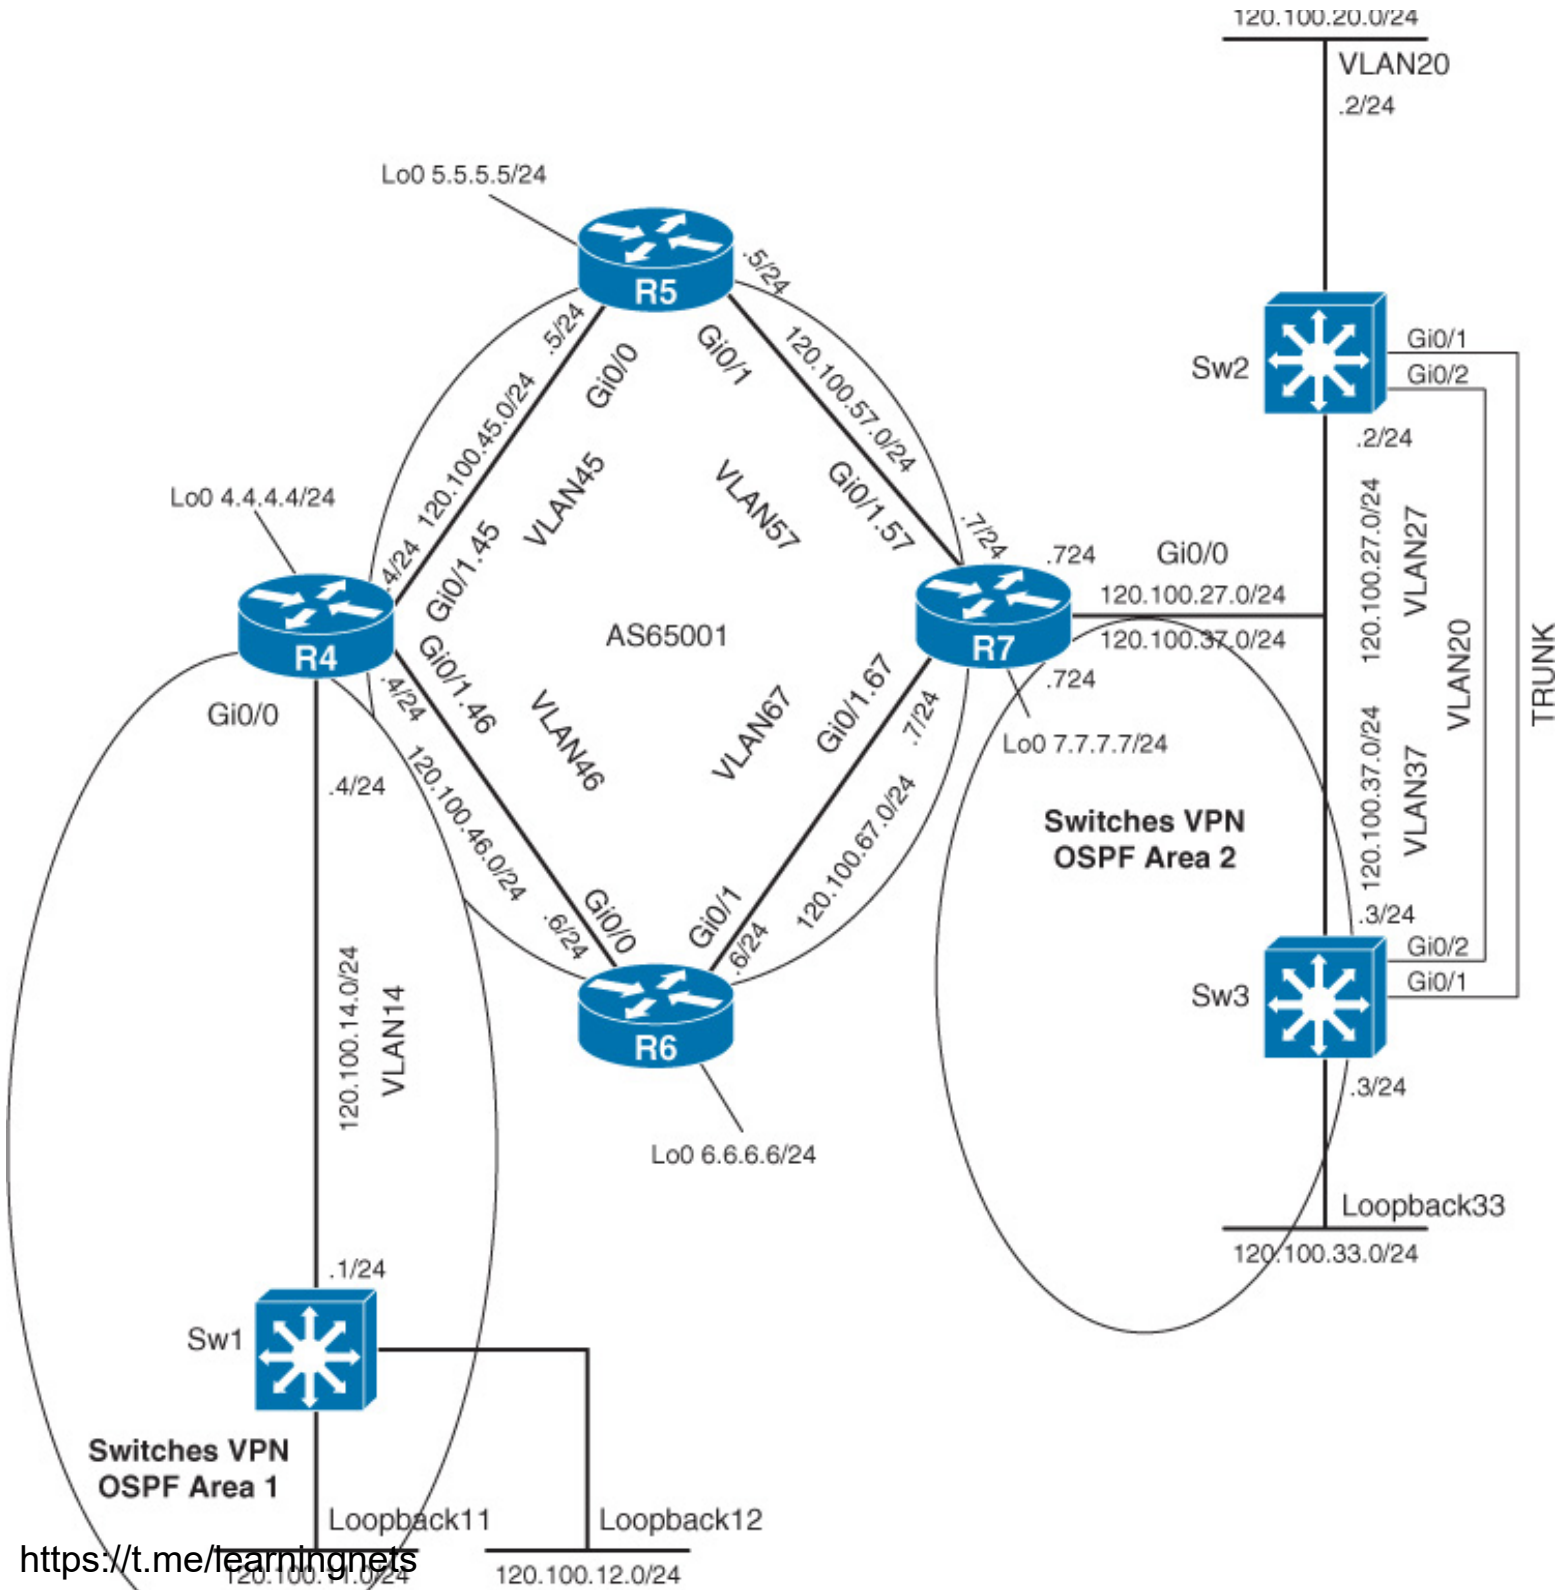
Incident 7

- Users on Sw3 VLAN30 have reported that they do not have visibility of 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:

[Click here to view code image](#)

```
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

- Users on Sw3 VLAN37 have reported that they cannot access services hosted on IP address 226.1.1.1, which are hosted on the Sw1 VLAN14 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:

[Click here to view code image](#)

```
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
```

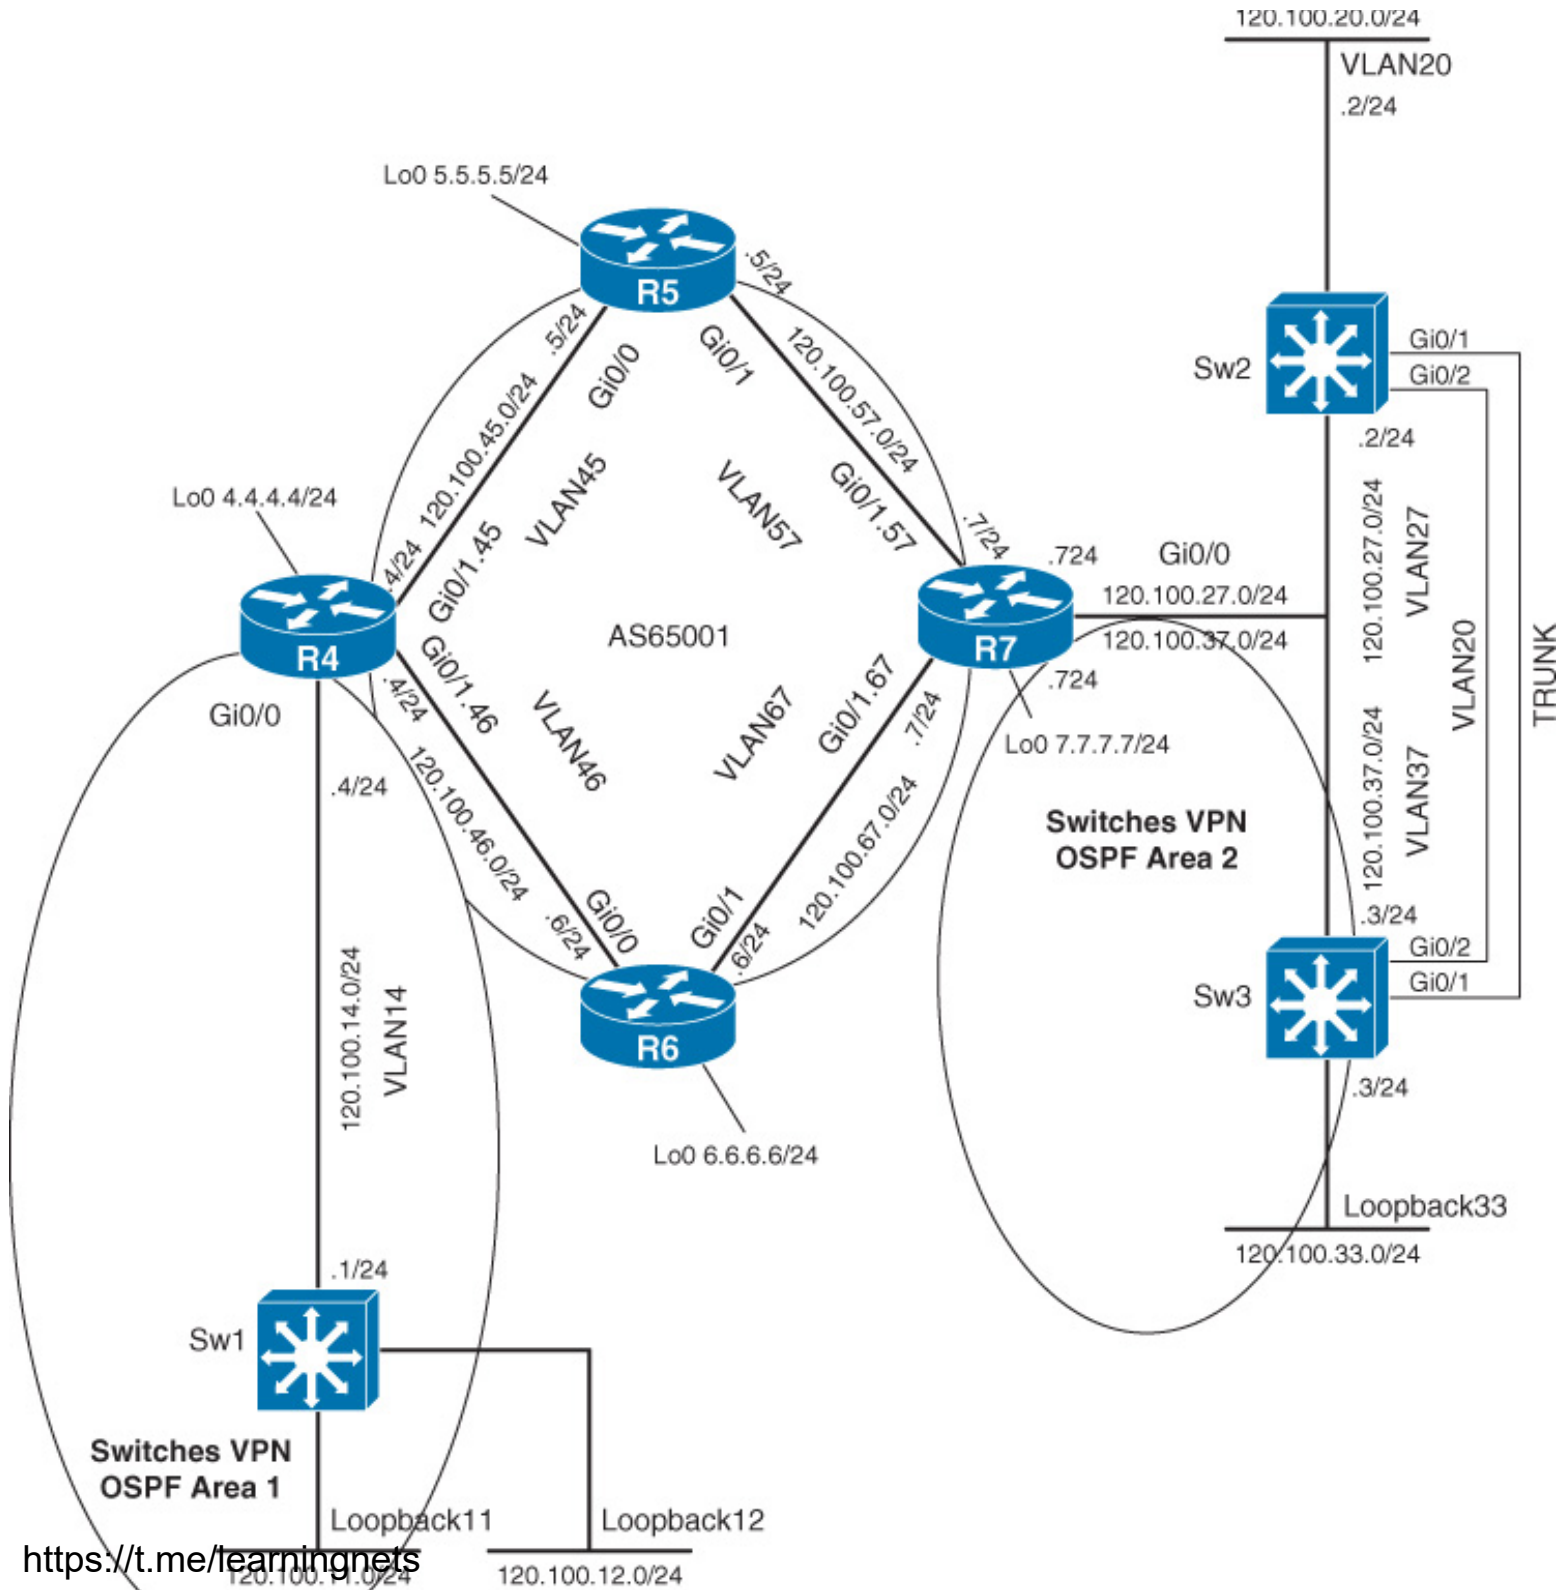


Figure 1-13 *Incident 8 Diagram*

- 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

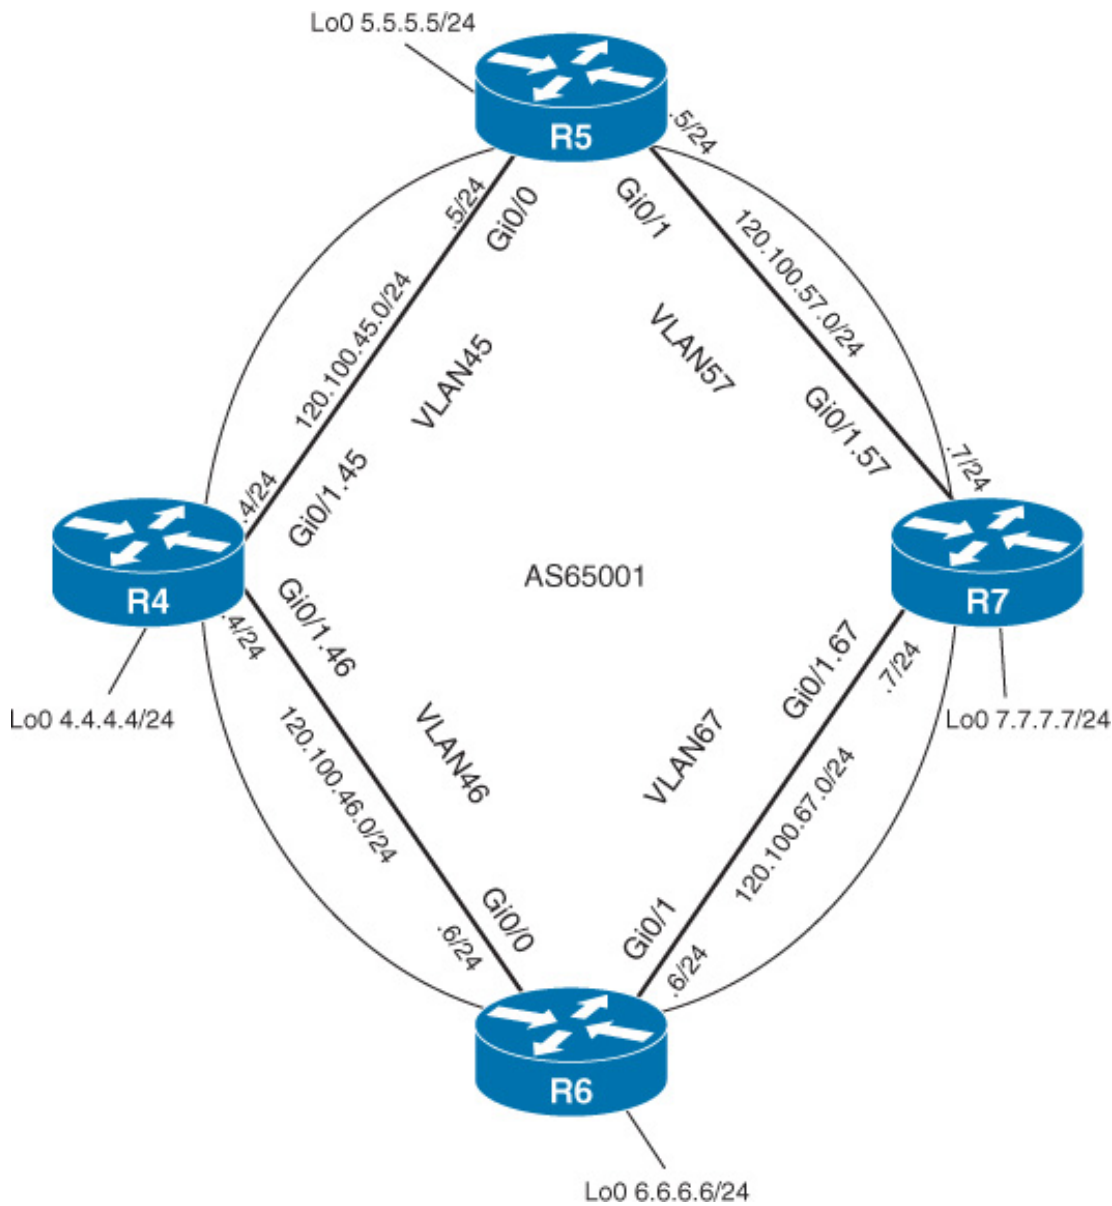


Figure 1-14 Incident 9 Diagram

- Your first-line support personnel are complaining that when they telnet to the Loopback address of R5 from within the MPLS network the response is very poor. One reported symptom is the screen intermittently
- <https://t.me/learningnets>

locking up.

- Investigate the issue with an aim of restoring normal service.

3 points

“ASK THE PROCTOR”

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 will be present to ensure that you do not have problems with the lab environment and to maintain the timing element of the exam.

Section 1

Incident 1

Q . I can successfully ping between Sw3 and Sw2 on VLAN20. I don't believe I have any problems.

A . You should find that your ping response time is variable. This is representative of an underlying issue that you must investigate and rectify. Once you receive a constant response time, you can be confident that you have solved the problem.

Q . I don't see any neighbor relationships between my PE router and Sw2 and Sw3. Should I be investigating

<https://t.me/learningnets>

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 isn't normal and indicates 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, as I believe I have answered the question?

A . No, this would remove the resiliency between switches for VLAN20. Find a configuration fix instead.

Q . Looking at the configuration, I can now see that one port is allowing BPDUs through. Can I remove this

configuration parameter to stabilize the network?

A . Yes.

Incident 2

Q . I have a route to 120.100.0.0/16 from R1 on R4, so this should work, right?

A . No, troubleshoot the issue. Trace the connectivity from R4 to see what happens.

Q . I can see when tracing the route from R4 that there is a connectivity issue on R1. Is this down to the configured summary route of 120.100.0.0/16 on R1? I can see that there is also a summary route configured on R3. Can I remove one summary route?

A . Yes.

Incident 3

Q . I have an 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 if there is anything that could potentially stop your devices multicasting to each other.

Q . Ah, okay, I see storm control. This would stop multicast, wouldn't it?

A . Yes.

Incident 4

Q . I have just checked the BGP connection between PE routers, which 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 modify the route
descriptors.

Q . I don't have Label Distribution Protocol (LDP) neighbors between my PE routers and P routers. The PEs

are using LDP, and the Ps are using Tag Distribution Protocol (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.

<https://t.me/learningnets>

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 . Yes.

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. Can I change this, as this would only allow it
.....
to advertise locally connected routes?

A . Yes.

Incident 7

Q . I have determined that the static route for 120.100.13.0 was not being advertised due to only classful
.....
static routes being advertised into OSPF, but it still isn't 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.

<https://t.me/learningnets>

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 would be 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 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. Surely, this is enough, right?

A . No, this solution also requires PIM neighbors over a tunnel between PE loopback interfaces.

Q . I don't have PIM neighbors over my PE loopbacks. I also don't have PIM enabled on these interfaces.

Surely, I would need this to form a neighbor relationship between loopbacks, right?

A . Yes, you would. 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 down to CPU loads and traffic rates

and considered normal?

<https://t.me/learningnets>

A . No, telnet to R5 using a physical address and note the difference.

.....
Q . I can see that CoPP has been configured on R5 for telnet traffic directed to the loopback interface. Can I modify the 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 items:

- Definition of the problem and identification of the symptoms (questioning to the void to identify meaningful symptoms)
- Definition of hypothesis and proof-testing of possible causes
- Design and implementation of a final solution (prior to application of the configuration)
- Verification of 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 VLAN20 are complaining that they experience poor connectivity to VLAN20 on Sw2. A variable response is seen when pinging from Sw3 VLAN20 to the Sw2 VLAN 20 interface.
- Investigate the issue and fix it. Confirm by proving a successful ping from the Sw3 VLAN20 interface 120.100.20.3 to the Sw2 VLAN20 interface 120.100.20.2 with a stable response time, as follows:

<https://t.me/learningnets>

[Click here to view code image](#)

```
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 = 1/5/9 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
```

```
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 = 1/5/9 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

<https://t.me/learningnets>

clock is ticking. You will really need to depend on your basic troubleshooting skills to be successful. First, you should confirm what the issue actually is and ping Sw2 VLAN20 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 due to a number of factors, but you may notice syslog messages on the console of Sw3, which you can examine within the logging buffer. What should be evident is that you have a MAC address flapping between Gio/1 and Gio/2 on VLAN20 (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 Gio/1 and Gio/2. (Gio/1 is a trunk each side and Gio/2 is an extension of VLAN20.) 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 Gio/1, you will find nothing untoward, just a standard trunk configuration, but Gio/2 has the command **spanning-tree bpdupfilter enable** configured. With BPDU filtering on Gio/2 enabled, the switch effectively believes it is connected to an end system and not to another switch, as such a loop is formed between devices. You can fix this problem by disabling the BPDU filter, as detailed in [Example 1-1](#), and ping testing shows a stable response time. If you solved this problem successfully, you have scored 3 points. If you did not fix this issue, you might first find you encounter numerous connectivity issues for devices connected to each switch.

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

[Click here to view code image](#)

```
Sw3# ping 120.100.20.2
```

Type escape sequence to abort.
<https://t.me/learningnets>

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 Gio/1 and port Gio/2

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

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

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

Sw3# show spanning-tree interface gi0/1

```

-----
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 bpdupfilter enable
end
```

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

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

```
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
```

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 VPN site on VLAN20 (Switch2). Initial investigations have led first-line support personnel to believe that an issue exists with connectivity up to the PE router R4, which was recently replaced.
- Investigate the issue and prove connectivity solely to their PE connection (R4's VLAN4 interface), as follows:

[Click here to view code image](#)

```
R3# ping 120.100.4.4 source 120.100.3.3
```

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 120.100.3.3

<https://t.me/learningnets>

!!!!

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

2 points

Begin by familiarizing yourself with the VPN diagram in [Figure 1-7](#) to understand that you will be working on a VRF-specific configuration of VRF routers. Using the information provided in the problem statement, you will need to be able to prove connectivity from R3 120.100.3.3 (Loopback0 interface) to R4 120.100.4.4 (VLAN4 interface). A successful extended ping between the two IP addresses will secure your points. In terms of how to troubleshoot this, well, you would have to 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. Remember that R4 is a PE router in this scenario, so you will 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 have a route to R1 VLAN4 so must have an EIGRP neighbor to R1. Similarly, R4 does have a route to 120.100.0.0/16 but not a specific route to 120.100.3.0/24, so a traceroute to 120.100.3.3 from R4 would certainly provide information as to where the issue lies. [Example 1-2](#) shows that the traceroute from R4 stops at R1 (120.100.4.1), a ping from R1 to 120.100.3.3 fails, and the routing table clearly shows why as the route is via Null0.

Example 1-2 *R3 and R1 and R4 Route Verification and Testing*

[Click here to view code image](#)

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

```
D 120.100.4.0/24
```

<https://t.me/learningnets>

```
R3# ping 120.100.4.4
```

Type escape sequence to abort.

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

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 124/131/144 ms

```
R3# ping 120.100.4.4 source 120.100.3.3
```

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 120.100.3.3

.....

Success rate is 0 percent (0/5)

```
R4# show ip route vrf ROUTERS | include 120.100.3.
```

```
R4# show ip route vrf ROUTERS | include 120.100.0.0
```

```
D 120.100.0.0/16
```

```
R4# show ip route vrf ROUTERS 120.100.0.0
```

Routing Table: ROUTERS

Routing entry for 120.100.0.0/16

Known via "eigrp 1", distance 90, metric 15360, type internal

Redistributing via eigrp 1

Last update from 120.100.4.1 on GigabitEthernet0/0.4, 00:43:32 ago

Routing Descriptor Blocks:

```
* 120.100.4.1, from 120.100.4.1, 00:43:32 ago, via GigabitEthernet0/0.4
Route metric is 15360, traffic share count is 1
Total delay is 20 microseconds, minimum bandwidth is 1000000 Kbit
Reliability 255/255, minimum MTU 1500 bytes
Loading 1/255, Hops 1
```

```
R4# traceroute vrf ROUTERS 120.100.3.3
```

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

```
Tracing the route to 120.100.3.3
```

```
VRF info: (vrf in name/id, vrf out name/id)
```

```
1 120.100.4.1 76 msec 152 msec 128 msec
```

```
2 120.100.4.1 !H * !H
```

```
R1# show ip route | include 120.100.3.
```

```
R1# show ip route | include 120.100.0.0
```

```
D 120.100.0.0/16 is a summary, 00:42:20, Null0
```

```
R1# show ip route 120.100.3.3
```

```
Routing entry for 120.100.0.0/16
```

```
Known via "eigrp 1", distance 5, metric 10240, type internal
```

```
Redistributing via eigrp 1
```

```
Routing Descriptor Blocks:
```

```
* directly connected, via Null0
```

```
Route metric is 10240, traffic share count is 1
```

Total delay is 10 microseconds, minimum bandwidth is 1000000 Kbit

Reliability 255/255, minimum MTU 1500 bytes

Loading 1/255, Hops 0

```
R1# ping 120.100.3.3
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.3.3, timeout is 2 seconds:

.....

Success rate is 0 percent (0/5)

The summary route on R1 to 120.100.0.0/16 is responsible for the issue. Inspection of the configuration on R1 and R3 as shown in [Example 1-3](#) details that there are, in fact, two summaries configured: one on R3's Ethernet interface toward R1, and one on R1's Ethernet interface toward R4. Because R1 does not have a specific route to 120.100.3.0/24, the local route to Null0 created by the summary route is followed, resulting in the black holing of traffic directed to network 120.100.3.0/24. By removing either one of these summaries, you can resolve the problem. If you have rectified the ticket, as detailed in [Example 1-3](#), you have scored 2 points.

Example 1-3 *R1 and R3 Summary Route Verification and Testing*

[Click here to view code image](#)

```
R3# show run | section eigrp
```

router eigrp CCIE
<https://t.me/learningnets>

```
!  
address-family ipv4 unicast autonomous-system 1  
!  
af-interface GigabitEthernet0/0  
summary-address 120.100.0.0 255.255.0.0  
exit-af-interface  
!  
topology base  
exit-af-topology  
network 120.100.3.0 0.0.0.255  
network 120.100.123.0 0.0.0.255  
exit-address-family
```

```
R1# show run | section eigrp
```

```
router eigrp CCIE
```

```
!  
address-family ipv4 unicast autonomous-system 1  
!  
af-interface GigabitEthernet0/1  
summary-address 120.100.0.0 255.255.0.0  
exit-af-interface  
!
```

```
topology base
```

```
exit-af-topology
network 120.100.4.0 0.0.0.255
network 120.100.123.0 0.0.0.255
exit-address-family
```

```
R1(config)# router eigrp CCIE
R1(config-router)# address-family ipv4 unicast autonomous-system 1
R1(config-router-af)# af-interface GigabitEthernet0/1
R1(config-router-af-interface)# no summary-address 120.100.0.0 255.255.0.0
```

```
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 Switch2 VLAN20 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 Switch2 and the local PE router R7.

Establishment of an EIGRP adjacency between devices will suffice at this point in time, as follows:

[Click here to view code image](#)

```
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 in time, so initial troubleshooting should be checking the adjacency between Sw2 and R7 over the VLAN27 interface, as indicated in [Figure 1-6](#). [Example 1-4](#) shows that the PE router R7 has an adjacency with Sw2, yet Sw2 does not have a neighbor relationship with R7. If you get a one-way neighbor relationship like this, a number of issues could be to blame, but you need to 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 will prove that multicast packets are received at R7 and replied to by the EIGRP process. The ping fails, so you can now be certain that you are looking for a multicast-related issue between Sw2 and R7.

Example 1-4 *R7 and Sw2 EIGRP Neighbor Testing*

[Click here to view code image](#)

```
R7# show ip eigrp vrf ROUTERS neighbors
```

```
EIGRP-IPv4 VRF(ROUTERS) Address-Family Neighbors for AS(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 an ACL is 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 (FastEthernet 0/7). You will see the command **storm-control multicast level 0.00**. This effectively disables 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. Once removed, you will see the neighbor relationship is immediately formed between Sw2 and R7. Basic troubleshooting steps will resolve the issue. Home in on the problem to provide the necessary clues to save your time and secure your points on the exam. If you have resolved the ticket, as detailed in [Example 1-5](#), you have scored 3 points.

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

Click here to view code image
<https://t.me/learningnets>

```
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 | begin eigrp
```

```
router eigrp JAKE
```

```
!
```

```
address-family ipv4 unicast vrf ROUTERS autonomous-system 1
```

```
!
```

```
topology base
```

```
redistribute bgp 65001
```

```
exit-af-topology
```

```
network 120.100.27.0 0.0.0.255
```

```
eigrp stub connected summary
```

```
exit-address-family
```

```
Sw2# show run int vlan 27
```

```
interface Vlan27
```

```
ip address 120.100.27.2 255.255.255.0
```

```
https://t.me/learningnets
```

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 Trunk Link to R7 Go/o
```

```
switchport trunk encapsulation dot1q
```

```
switchport mode trunk
```

```
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 VLAN20 (Switch2). It appears that the MPLS VPN is not functioning correctly.
- Investigate the issue and prove connectivity between PE routers via an extended ping from the R4 VLAN4 interface and R7 VLAN27 interface to prove MPLS functionality, as follows:

[Click here to view code image](#)

```
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

<https://t.me/learningnets>

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 shown here, no routes are learned via BGP between PE routers, and therefore there will be no end-to-end connectivity throughout this VPN.

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

[Click here to view code image](#)

```
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, 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 will 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 the 10,000-foot view, 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, and 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*

[Click here to view code image](#)

```
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:  
https://t.me/learningnets
```

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 see whether you are advertising VRF-specific routes between PEs is the **show ip bgp vpnv4 vrf X** command. This will tell 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. Note, as well, that in the output the RD is shown as 1:1000 for this VRF on R4 and 2:100 on R7; now 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 will be able to ensure that your VRF routes are advertised successfully over MP-BGP. You should be able to spot that the second VRF is also inconsistent and can fix 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

[Click here to view code image](#)

```
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
```

Note

Be aware that if you remove the RDs on a PE router the VRF-specific configuration will be removed from the address family within BGP. Therefore, you should copy the relevant configuration and reapply it after changing the RD and import export values. For verification, consider viewing the startup configuration if you lose configuration.

[Click here to view code image](#)

```
router bgp 65001
  address-family ipv4 vrf SWITCHES
    redistribute ospf 2 vrf SWITCHES route-map OSPF-BGP
  no synchronization
  exit-address-family
!
  address-family ipv4 vrf ROUTERS
    redistribute connected
    redistribute eigrp 1
    no synchronization
  exit-address-family
```

[Click here to view code image](#)

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 for this question, though, you need to be able to 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 R5 and R6 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*

[Click here to view code image](#)

```
R4# ping vrf ROUTERS 120.100.27.4 source 120.100.4.4
https://t.me/learningnets
```

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
GigabitEtherneto/1.45		Yes (ldp)	No Yes
GigabitEtherneto/1.46		Yes (ldp)	No Yes

```
R7# show mpls ldp neighbors
```

```
R7# show mpls interfaces
```

Interface	IP	Tunnel	Operational
GigabitEthernet0/1.57	Yes	(ldp)	No Yes
GigabitEthernet0/1.67	Yes	(ldp)	No Yes

```
R5# show mpls ldp neighbors
```

```
R5# show mpls interfaces
```

Interface	IP	Tunnel	Operational
GigabitEthernet0/0	Yes	(tdp)	No Yes
GigabitEthernet0/1	Yes	(tdp)	No Yes

```
R6# show mpls ldp neighbors
```

```
R6# show mpls interfaces
```

Interface	IP	Tunnel	Operational
GigabitEthernet0/0	Yes	(tdp)	No Yes
GigabitEthernet0/1		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 (and 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 <https://t.me/learningnets>

MPLS traceroute between PE loopback addresses. As shown, this fails from both directions because R5 and R6 P routers are informing each PE router that no corresponding label exists. This is a classic case of incorrect initial configuration (as you know your BGP session is working correctly and you know that you have LDP established correctly between your devices). If you look at your MPLS router loopback interfaces, you will see that 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 checked the routing table, you would see the routes come through as /32 host routes, but this is simply due to the way OSPF advertises loopback networks (unless you adjust the OSPF network type of the interface). [Example 1-10](#) shows that when the loopbacks are readdressed to /32 hosts the MPLS traceroute functions correctly and the end-to-end VRF-specific ping is successful between PE routers. If you have solved this ticket successfully, as detailed in [Examples 1-6](#) through [1-10](#), you have scored 6 points.

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

[Click here to view code image](#)

```
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,
```

<https://t.me/learningnets>

'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.

o 120.100.67.7 MRU 1500 [Labels: 18 Exp: o]

B 1 120.100.67.6 MRU 1504 [No Label] o ms

B 2 120.100.67.6 MRU 1504 [No Label] o ms

B 3 120.100.67.6 MRU 1504 [No Label] o ms

B 4 120.100.67.6 MRU 1504 [No Label] o ms

B 5 120.100.67.6 MRU 1504 [No Label] 4 ms

B 6 120.100.67.6 MRU 1504 [No Label] o ms

B 7 120.100.67.6 MRU 1504 [No Label] o ms

B 8 120.100.67.6 MRU 1504 [No Label] o ms

B 9 120.100.67.6 MRU 1504 [No Label] 4 ms

B 10 120.100.67.6 MRU 1504 [No Label] o ms

B 11 120.100.67.6 MRU 1504 [No Label] o ms

B 12 120.100.67.6 MRU 1504 [No Label] o ms

B 13 120.100.67.6 MRU 1504 [No Label] 4 ms

B 14 120.100.67.6 MRU 1504 [No Label] o ms

B 15 120.100.67.6 MRU 1504 [No Label] o ms

B 16 120.100.67.6 MRU 1504 [No Label] o ms

B 17 120.100.67.6 MRU 1504 [No Label] 4 ms

B 18 120.100.67.6 MRU 1504 [No Label] o ms

B 19 120.100.67.6 MRU 1504 [No Label] o ms

B 20 120.100.67.6 MRU 1504 [No Label] o 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 they cannot see the route to VLAN27.
- Investigate and prove connectivity with a valid route and successful ping from R2 and R3 user subnets to the Sw2 VLAN27 interface, as follows:

[Click here to view code image](#)

```
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 proved 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. Instead, it 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 a default metric is 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 simply 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 VLAN27 interface. If you solved this problem successfully, you have scored 3 points.

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

[Click here to view code image](#)

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

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

```
B    120.100.27.0 [200/0] via 7.7.7.7, 6d19h
```

```
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, 1wod, GigabitEthernet0/1.4
D 120.100.123.0 [90/2172416] via 120.100.4.1, 1wod, 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, 4 subnets

C 120.100.4.0 is directly connected, GigabitEthernet0/1
D 120.100.2.0 [90/2297856] via 120.100.123.2, 6d20h, GigabitEthernet0/0
D 120.100.3.0 [90/2195456] via 120.100.123.3, 1w2d, GigabitEthernet0/0
C 120.100.123.0 is directly connected, GigabitEthernet0/0
```

```
R4# show run | section eigrp
```

```
router eigrp JAKE
!
address-family ipv4 unicast vrf ROUTERS autonomous-system 1
!
topology base
 redistribute bgp 65001
 exit-af-topology
 network 120.100.4.0 0.0.0.255
```

```
eigrp stub connected summary
```

```
exit-address-family
```

```
R4(config)# router eigrp JAKE
```

```
R4(config-router)# address-family ipv4 unicast vrf ROUTERS autonomous-system 1
```

```
R4(config-router-af)# topology base
```

```
R4(config-router-af-topology)# default-metric 1000 1 1 1 1500
```

```
R7# show run | section eigrp
```

```
router eigrp JAKE
```

```
!
```

```
address-family ipv4 unicast vrf ROUTERS autonomous-system 1
```

```
!
```

```
topology base
```

```
redistribute bgp 65001
```

```
exit-af-topology
```

```
network 120.100.27.0 0.0.0.255
```

```
eigrp stub connected summary
```

```
exit-address-family
```

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

```
R7(config-router)# address-family ipv4 unicast vrf ROUTERS autonomous-system 1
```

```
R7(config-router-af)# topology base
```

```
R7(config-router-af-topology)# default-metric 1000 1 1 1 1500
```

```
R1# show ip route eigrp | include 120.100.27.0
```

```
D 120.100.27.0 [90/284160] via 120.100.4.4, 00:00:17, GigabitEthernet0/1
```

```
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 only see the route to VLAN27 on Sw2 and not VLAN20.
- Investigate the issue and fix it. Prove connectivity by an extended ping from user subnets on R2 and R3 to Sw2 VLAN20, as follows:

[Click here to view code image](#)

```
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

<https://t.me/learningnets>

You proved the VRF-specific MPLS network previously, so you can be confident that this is not an MPLS issue. 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 shown, the route is indeed present, so in theory it should be redistributed into BGP and over MPLS to PE router R4, but it isn't. 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 VLAN20 on Sw2 will secure your points. If you have resolved this issue successfully, as detailed in [Example 1-12](#), you have scored 2 points.

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

[Click here to view code image](#)

```
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 eigrp
```

```
router eigrp JAKE
```

```
!
```

```
address-family ipv4 unicast vrf ROUTERS autonomous-system 1
```

```
!
```

```
topology base
```

```
redistribute bgp 65001
```

```
exit-af-topology
```

```
network 120.100.27.0 0.0.0.255
```

```
eigrp stub connected summary
```

```
exit-address-family
```

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

```
R7(config-router)# address-family ipv4 unicast vrf ROUTERS autonomous-system 1
```

```
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 do not have visibility of 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:

[Click here to view code image](#)

```
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 that the MPLS OSPF network, the MP-BGP, and MPLS are 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 determine whether the route is indeed configured on Sw1 and <https://t.me/learningnets>

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, yet 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 redistribution 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. PE R7 is, however, 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*

[Click here to view code image](#)

```
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 a TAG value assigned 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 will need to 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 for certain circumstances (and it is the quickest option). [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. If you have successfully resolved this issue, you have scored 3 points.

[Click here to view code image](#)

```
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
```

<https://t.me/learningnets> 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 VLAN37 have reported that they cannot access services hosted on IP address 226.1.1.1 that are hosted on the Sw1 VLAN14 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:

[Click here to view code image](#)

```
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
<https://t.me/learningnets>

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 VLAN14 is likely to have or should have an IGMP join command for 226.1.1.1 and Multicast Distribution Tree is in use with a group address of 232.0.0.11.

Start simple with your troubleshooting instead of diving right in. Attempt a multicast ping on the local PE R4 to the multicast destination to see whether the multicast is functioning locally to begin with. This would 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: 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, and 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 there must be an issue in relation to these. The example shows that PIM has also not been enabled on the Loopback0 interfaces of each PE router. Once 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, with 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 required all the same for the <https://t.me/learningnets>

multicast solution to function effectively. If you have solved the tickets, as detailed within [Example 1-15](#), you have scored 5 points.

Example 1-15 *Multicast Testing and Configuration*

[Click here to view code image](#)

```
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
120.100.45.5	GigabitEthernet0/1.45	1w3d/00:01:43	v2	1 / S
120.100.46.6	GigabitEthernet0/1.46	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 Address	Interface	Uptime/Expires	Ver	DR
120.100.45.4	GigabitEthernet0/0	1w3d/00:01:40	v2	1 / S
120.100.57.7	GigabitEthernet0/1	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 Address	Interface	Uptime/Expires	Ver	DR
120.100.67.7	GigabitEthernet0/1	1w6d/00:01:38	v2	1 / S
120.100.46.4	GigabitEthernet0/0	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 Address	Interface	Uptime/Expires	Ver	DR
120.100.67.6	GigabitEthernet0/1.67	1w6d/00:01:18	v2	1 / S
120.100.57.5	GigabitEthernet0/1.57	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 Address	Interface	Uptime/Expires	Ver	DR
120.100.14.1	GigabitEthernet0/0.14	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 Address	Interface	Uptime/Expires	Ver	DR
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 Address	Interface	Uptime/Expires	Ver	DR
120.100.14.1	GigabitEthernet0/0.14	1w3d/00:01:31	v2	1 / S
7.7.7.7	Tunnel0	00:00:11/00:01:33	v2	1 /

DR S

```
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 Address	Interface	Uptime/Expires	Ver	DR
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 R5 from within the MPLS network 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 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 <https://t.me/learningnets>

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 the like.

Example 1-16 shows that 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 (Loopback0). If the command **show policy-map control-plane** is issued and viewed, you can see that indeed some packets are being dropped to the unrealistic police settings of anything over two packets per second being dropped. The police rate can be adjusted 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 by trial.

Example 1-16 shows that a value of 400pps with a burst rate of 200 is sufficient for normal service. If you have solved the issue successfully by altering the police rate on R4 to a value where a telnet session is reliable, you have scored 3 points.

Example 1-16 *CoPP Testing and Configuration*

[Click here to view code image](#)

```
R4# show run
```

```
-
```

```
interface Loopback0
```

```
ip address 4.4.4.4 255.255.255.255
```

<https://t.me/learningnets>

```
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
```

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 the time frame of 2 hours or less, you are well on your way to being prepared for tickets that you are likely to face during the Troubleshooting section of the real exam.

The questions might have seemed a little vague or, in fact, misleading, but they closely resemble real-world scenarios in which you are given only limited information and have to work out what is really happening yourself to get to the root cause.

Troubleshooting Lab 2 (Network Down!)

EQUIPMENT LIST

You need the following hardware and software components to commence this troubleshooting lab:

- Seven routers loaded with Cisco IOS Software Release 15.3T Advanced Enterprise image and the minimum interface configuration as documented in [Table 2-1](#)

Router	Model	Ethernet I/F
R1	3925	2
R2	3925	1
R3	3925	1
R4	3925	2
R5	3925	2
R6	3925	2
R7	3925	2

Table 2-1 *Hardware Required per Router*

- Three 3560X switches with IOS 15.0S 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 shown in [Figure 2-1](#). It is, however, recommended to use the same model of equipment because this will make life easier should you load configurations directly from the supplied configurations into your own devices.

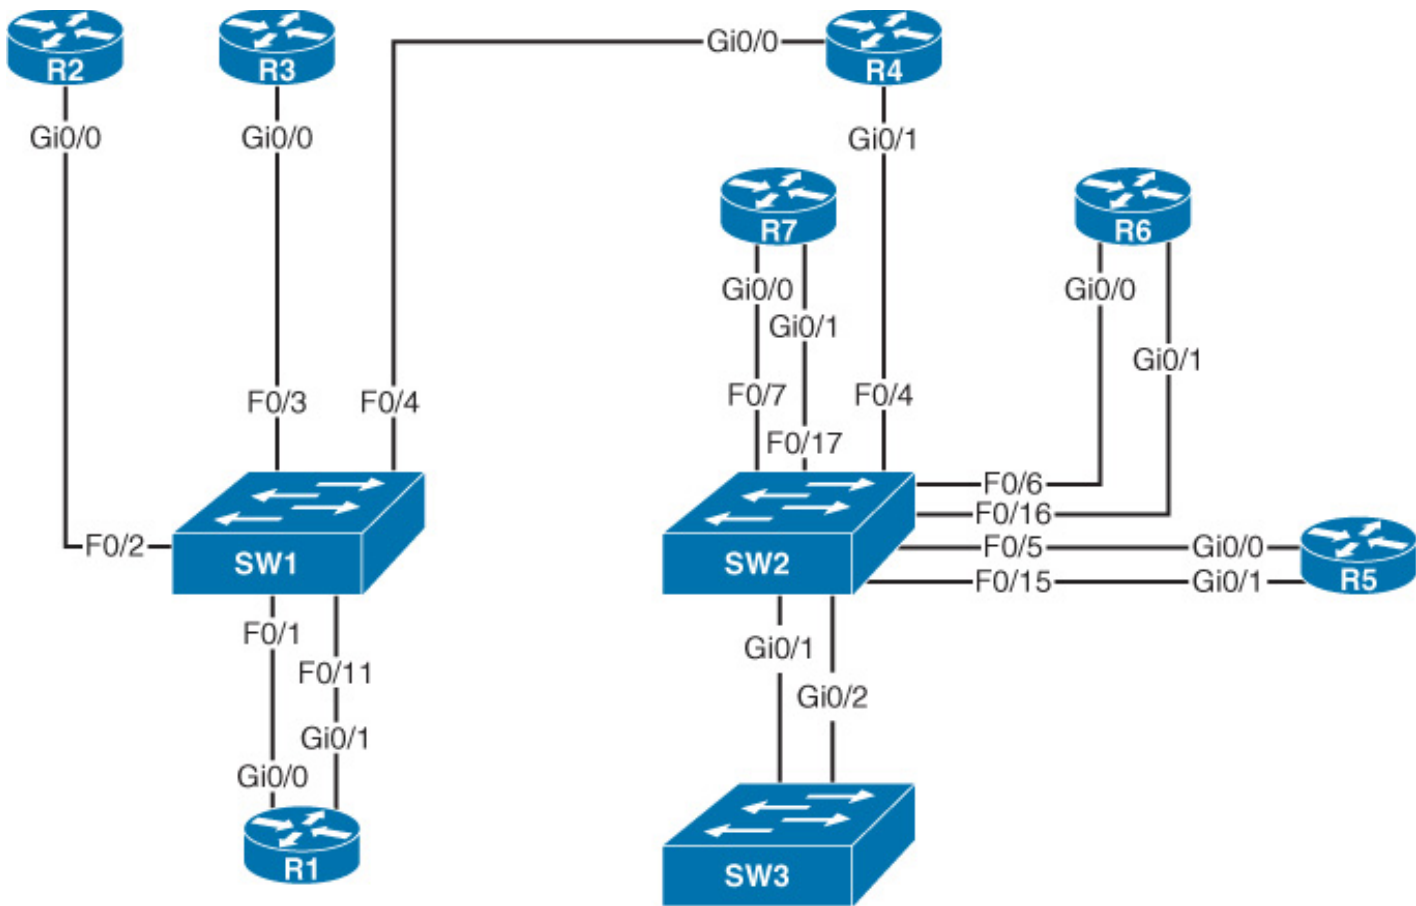


Figure 2-1 Lab 1 Topology Diagram

Note

Use the initial configurations supplied 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 will ensure that you do not get unwanted behavior resulting from differing interior gateway protocol (IGP) metrics.

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

Lab Topology

This troubleshooting exercise uses the topology outlined in [Figure 2-1](#), which you will need to re-create with your own equipment or by using lab equipment on the CCIE R&S 360 program.

Switch Instructions

Configure VLAN assignments from the configurations supplied or from [Table 2-2](#).

VLAN	Switch1	Switch2	Switch3
4	Fa0/4, Fa0/11	—	—
14	Fa0/4	—	—
20	—	Gi0/1, Gi0/2	Gi0/1, Gi0/2
27	—	Gi0/1, Fa0/7	—
45	—	Fa0/4, Fa0/5	—
46	—	Fa0/4, Fa0/6	—
57	—	Fa0/15, Fa0/17	—
67	—	Fa0/16, Fa0/17	—
134	Fa0/1, Fa0/2, Fa0/3	—	—

Table 2-2 VLAN Assignment

Connect your switches with fiber small form-factor pluggable (SFP) connectors or RJ-45 Ethernet crossover cables, as shown in [Figure 2-2](#).

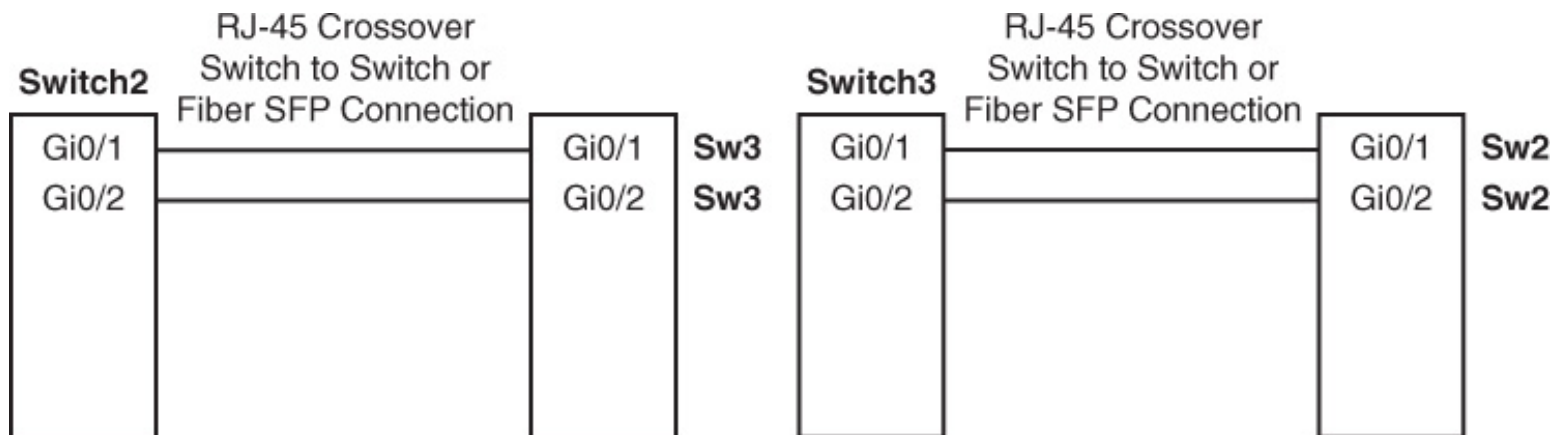


Figure 2-2 Switch-to-Switch Connectivity

Note

For simplicity in this lab, R2 and R3 use loopback interfaces to simulate user-connected interfaces.

IP Address Instructions

For this exercise, you configure your IP addresses as shown in Figure 2-3 or load the initial router configurations supplied.

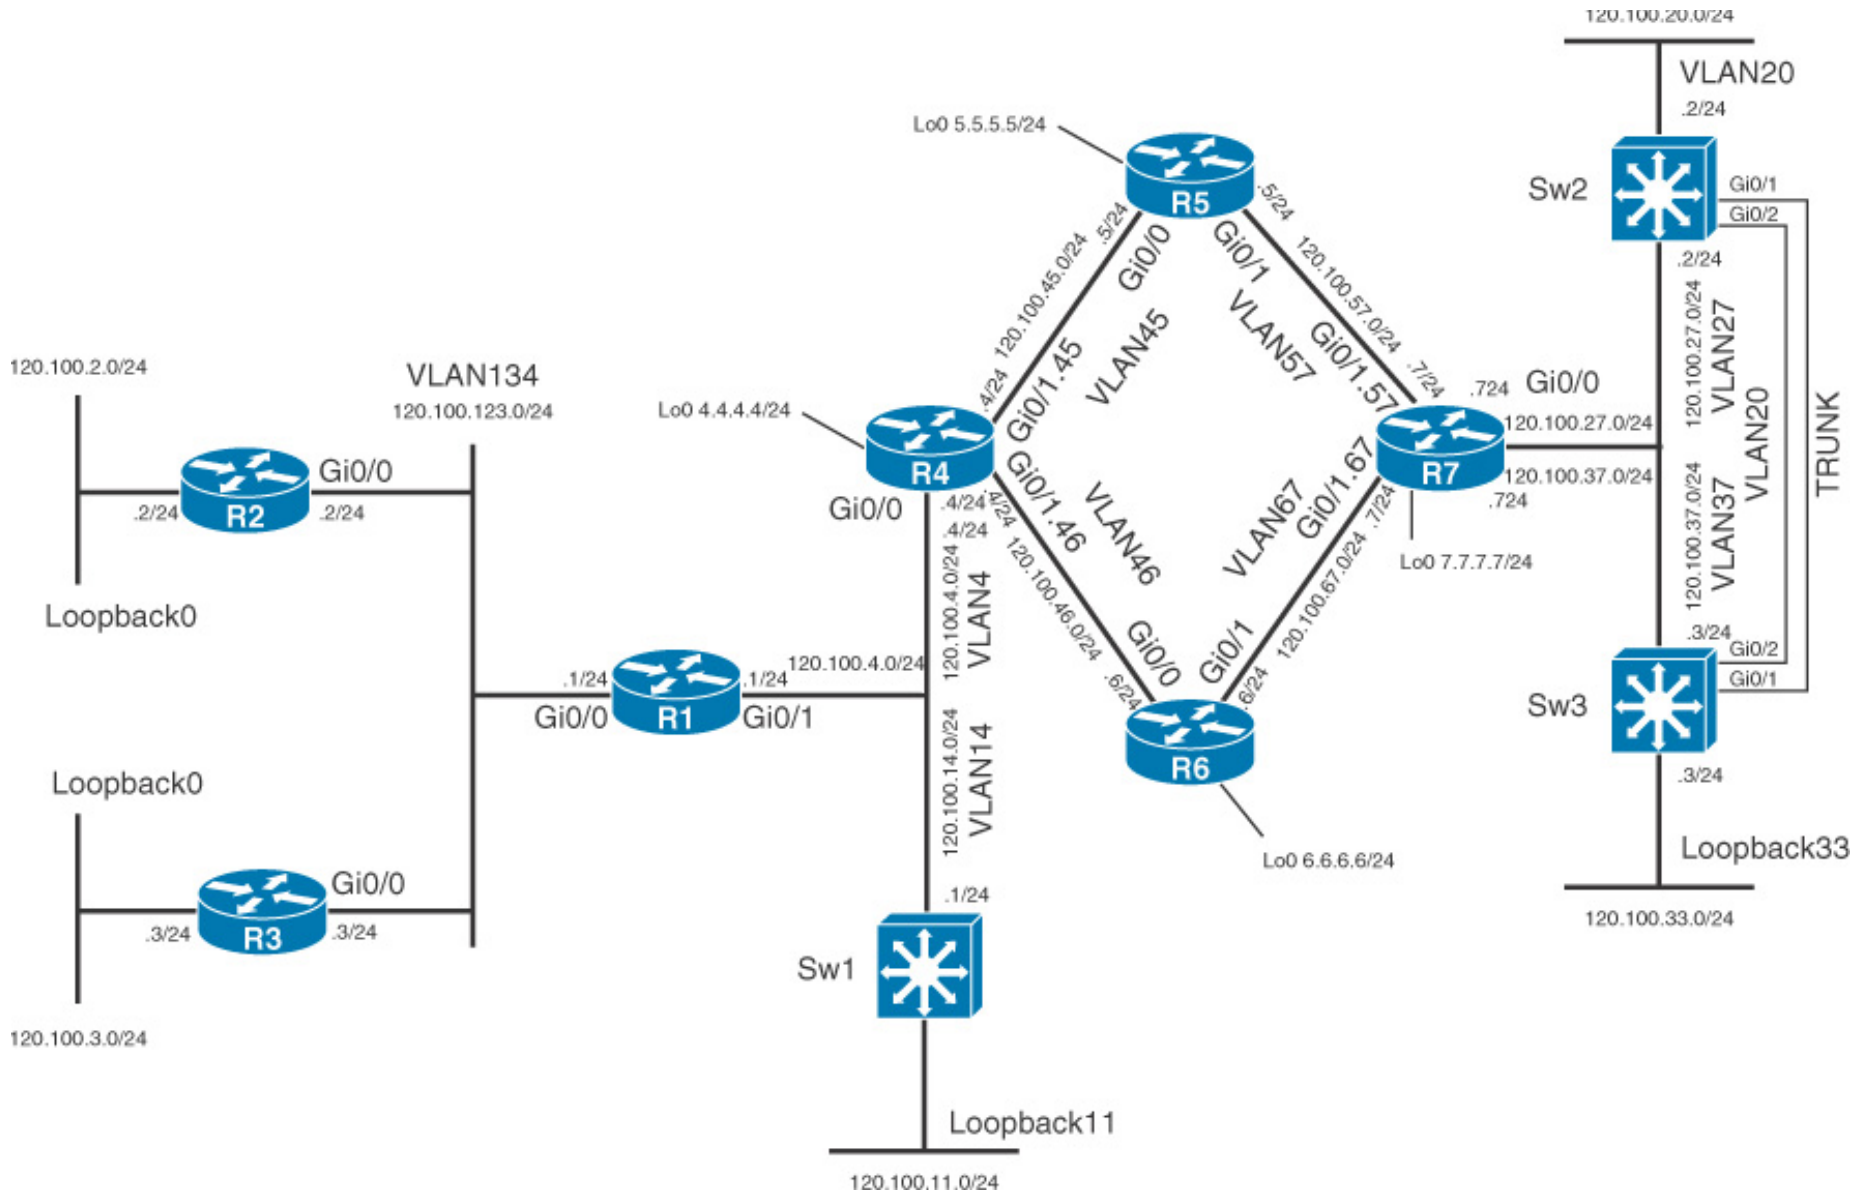


Figure 2-3 IP Addressing Diagram

PRE-LAB TASKS

- Build the lab topology as per [Figure 2-1](#) and [Figure 2-2](#).
- Configure the IP addresses on each router as shown in [Figure 2-3](#), alternatively you can load the initial <https://t.me/learningnets>

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.
- Tackle questions sequentially. You might find that one trouble ticket needs to be resolved before moving on to the next ticket. (This will not be the case, however, on your real lab exam, which will have a far greater number of devices with noninterconnected faults.)
- 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 that you follow these items correctly to maximize your score.
- Do not remove any configured feature to resolve an incident. You must resolve the misconfiguration rather than remove 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.)
- Access the latest documentation from <http://www.cisco.com/cisco/web/psa/configure.html>.

Note

Access only the website at the URL shown here, not the whole [Cisco.com](http://www.cisco.com) website (because if you are permitted to use documentation during your CCIE lab exam, your access will be restricted). Well-prepared candidates should not allow themselves to lose time during the exam searching for information.

TROUBLESHOOTING LAB 2

You will now be answering questions in relation to the network topology shown in [Figure 2-4](#).

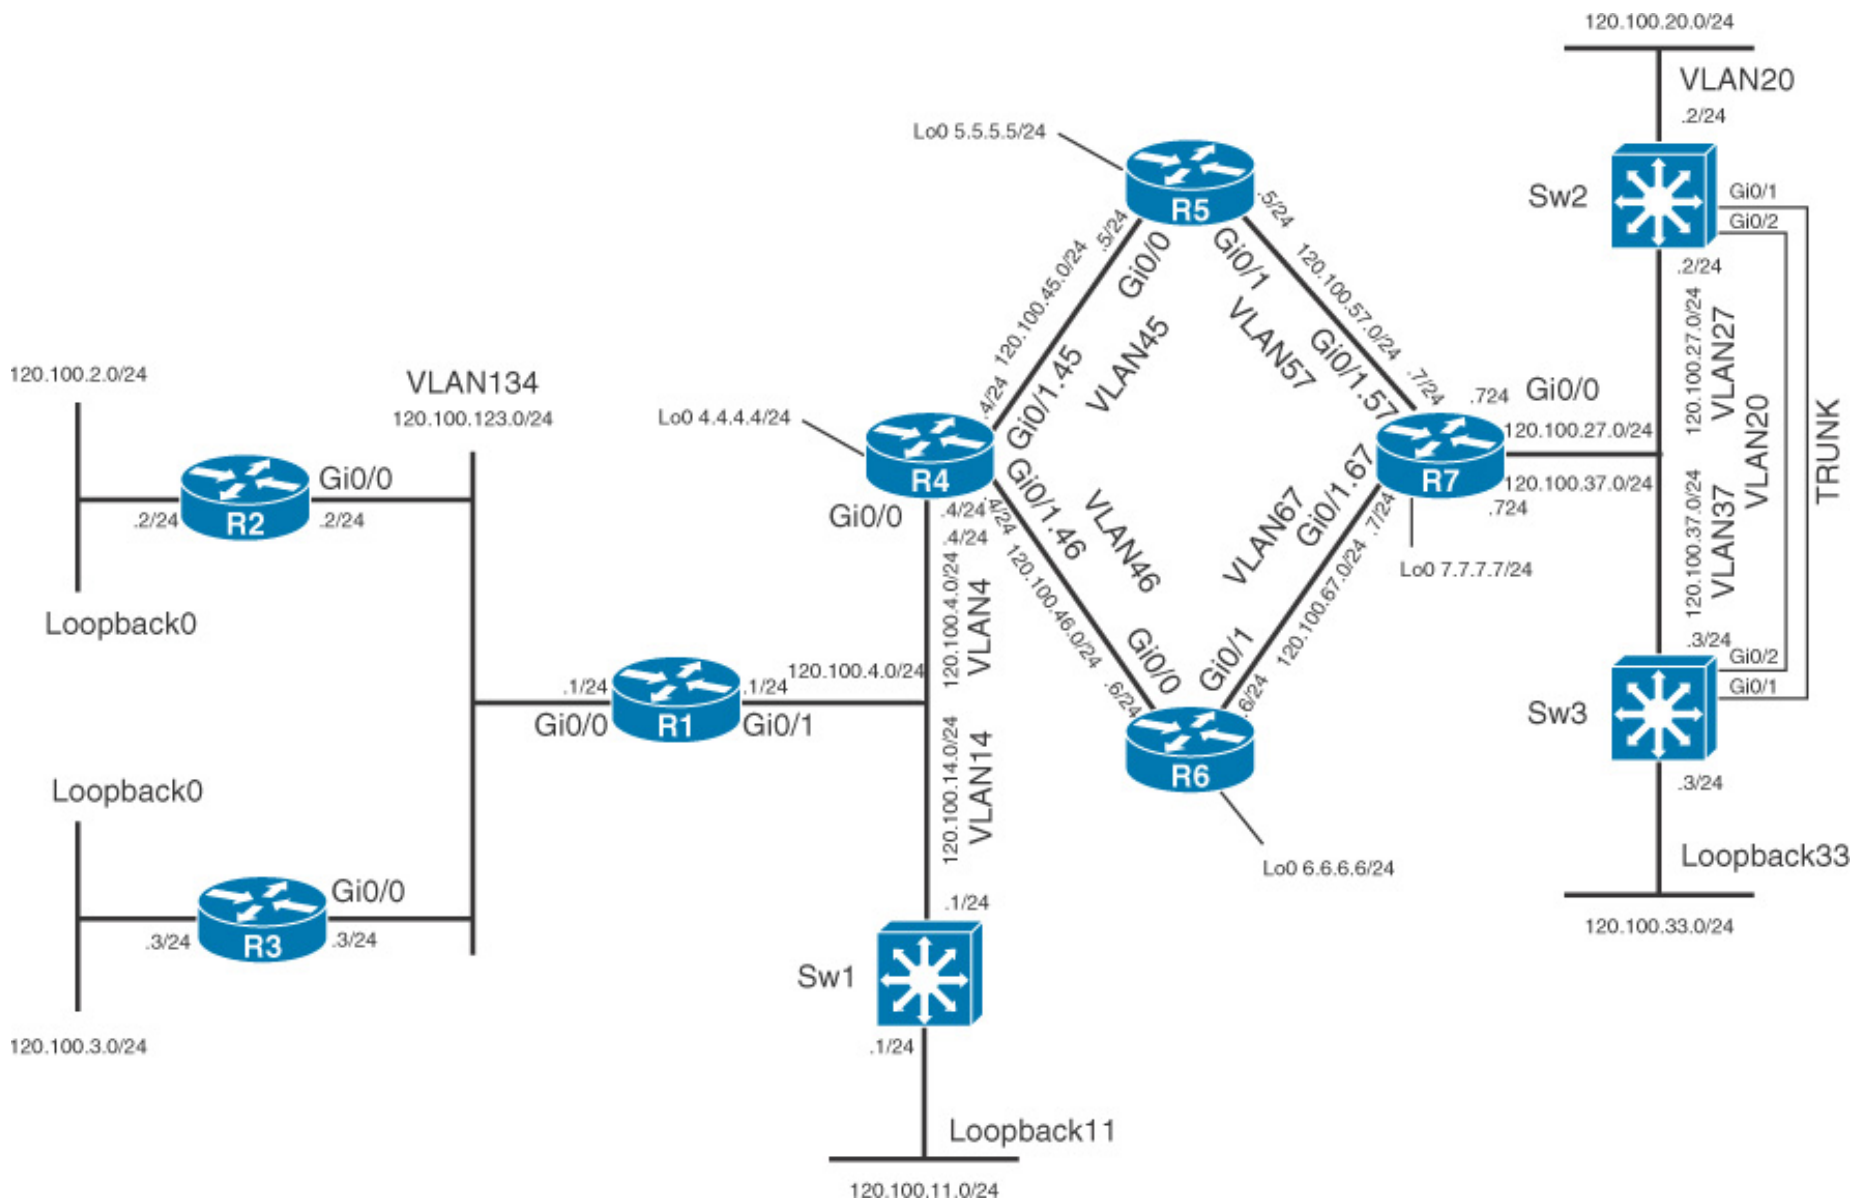


Figure 2-4 Lab Topology Diagram

SECTION 1

Incident 1

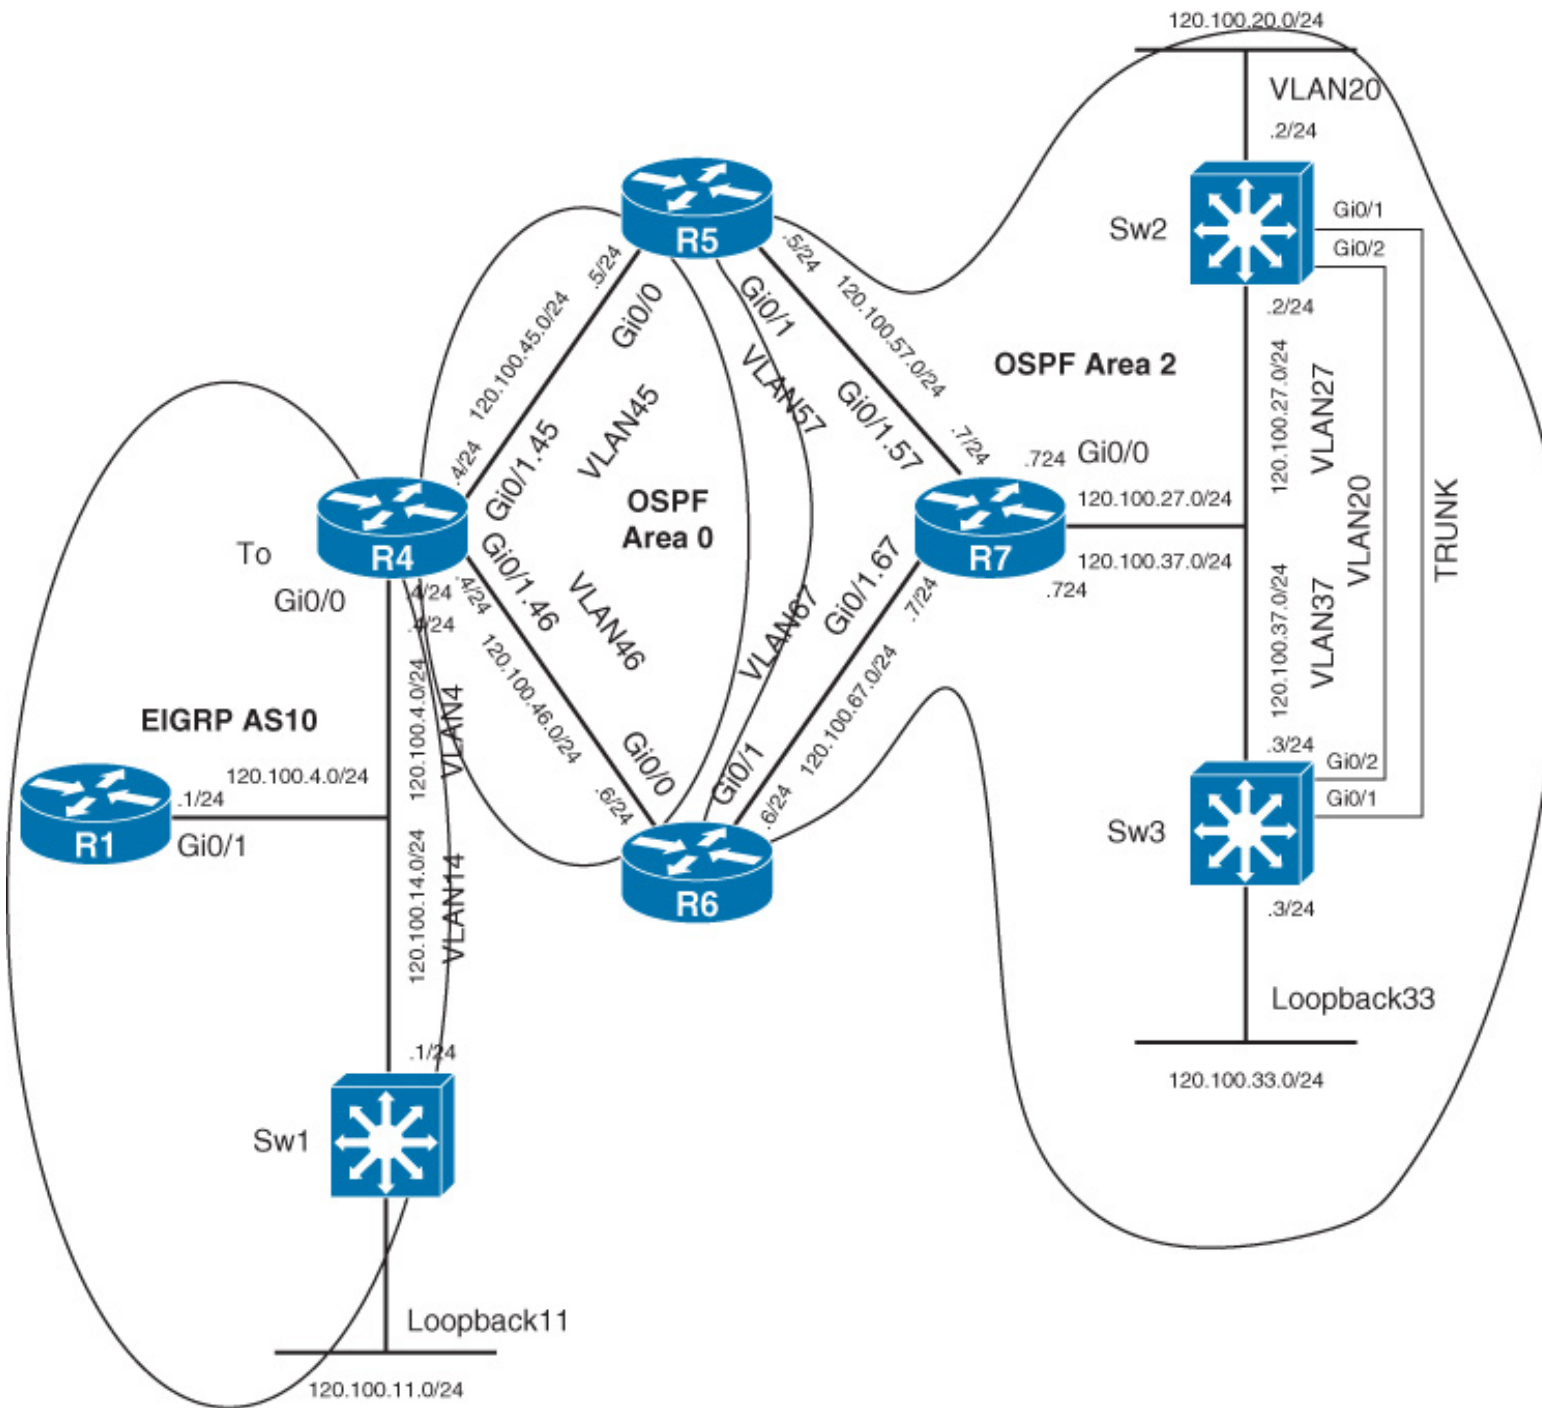


Figure 2-5 Incident 1 Diagram

- Users on Sw1 Lo11 120.100.11.0/24 are not able to ping users on VLAN20 Sw2 120.100.20.0/24.
- Fix the problem so that Sw1 Lo11 120.100.11.1 and Sw1 VLAN14 120.100.14.1 can ping Sw2 VLAN20 120.100.20.2, as follows:

[Click here to view code image](#)

```
Sw1# ping 120.100.20.2 source 120.100.11.1
```

```
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 120.100.11.1
```

```
!!!!
```

```
Sw1# ping 120.100.20.2 source 120.100.14.1
```

```
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 120.100.14.1
```

```
!!!!
```

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

1 point

Incident 2

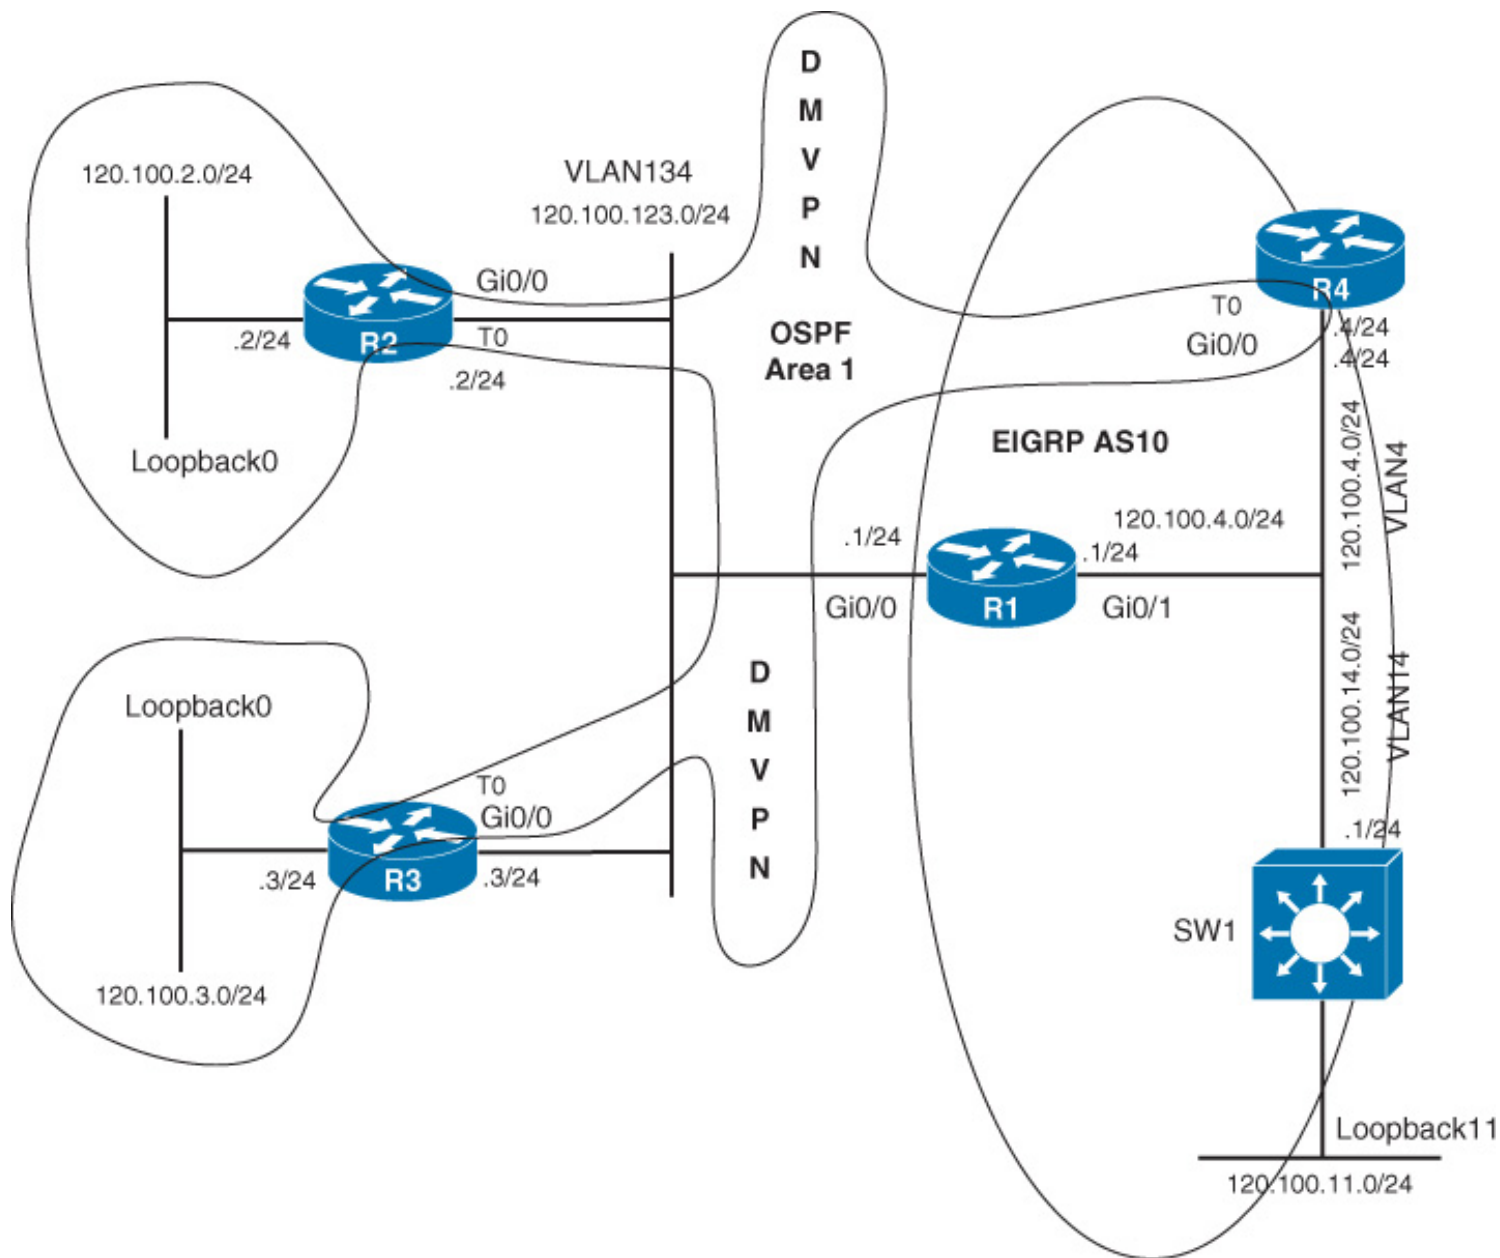


Figure 2-6 Incident 2 Diagram

- DMVPN access between spoke routers R2 and R3 to hub router R4 is down. Fix the access with DMVPN connectivity, as follows:
<https://t.me/learningnets>

[Click here to view code image](#)

```
R2# show dmvpn
```

```
Tunnelo, Type:Spoke, NHRP Peers:1,
```

```
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
```

```
-----
```

```
1 120.100.4.4 120.100.40.4 UP 00:06:40 S
```

```
R3# show dmvpn
```

```
Tunnelo, Type:Spoke, NHRP Peers:1,
```

```
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
```

```
-----
```

```
1 120.100.4.4 120.100.40.4 UP 00:06:40 S
```

- Do not modify or remove the existing default route on R2 and R3 pointing to R1 120.100.123.1.

2 points

Incident 3

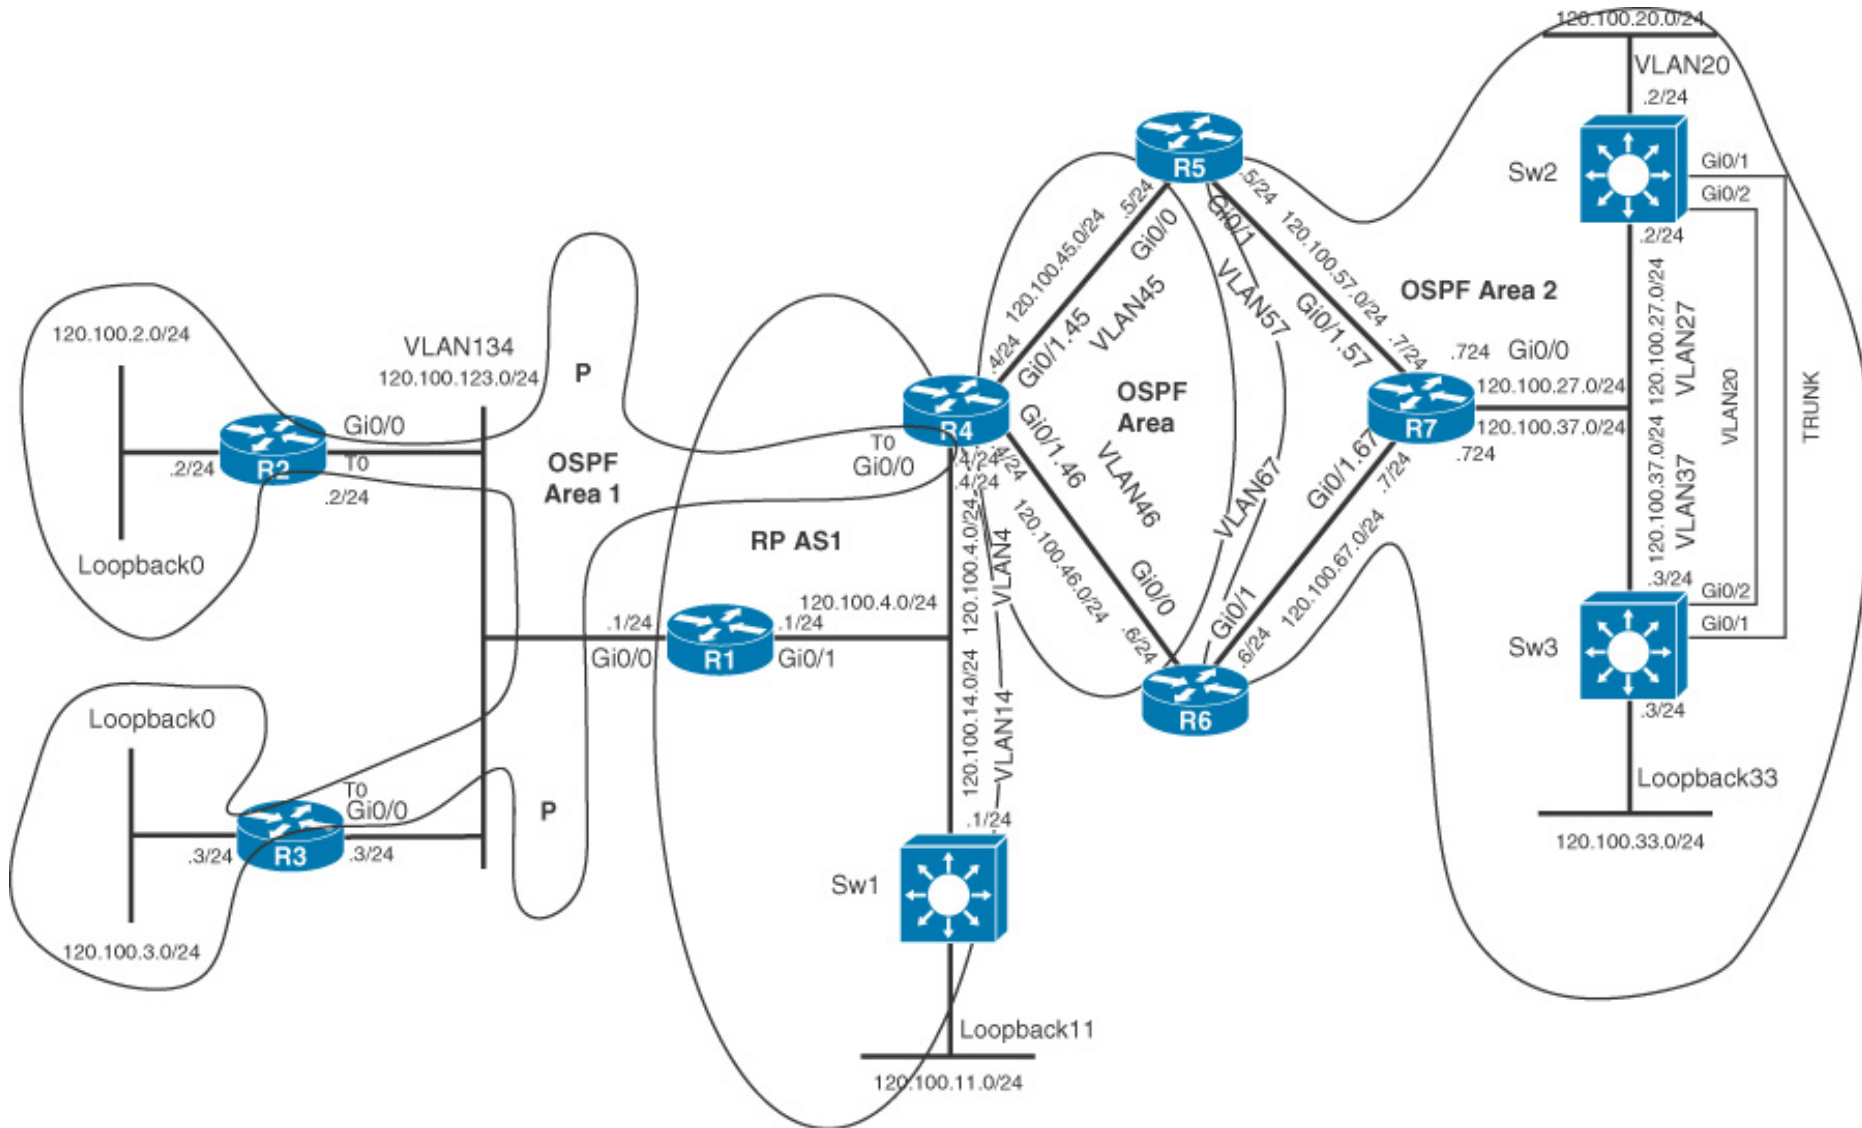


Figure 2-7 Incident 3 Diagram

- Home workers on R2 120.100.2.0/24 and R3 120.100.3.0/24 subnets are unable to reach services located on Switch2 VLAN20 120.100.20.0/24 through their DMVPN connection.
- Fix the problem so that R2 and R3 users have DMVPN access to VLAN20 via the DMVPN hub router R4, as <https://t.me/learningnets>

follows:

[Click here to view code image](#)

```
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 = 1/2/4 ms

```
R2# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	0	FULL/ -	00:00:34	120.100.40.4	Tunnel0

```
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 = 1/2/4 ms

```
R3# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	0	FULL/ -	00:00:34	120.100.40.4	Tunnel0

- Do not modify or remove the existing default route on R2 and R3 pointing to R1 120.100.123.1. Do not add further routing protocols to R2, R3, and R4. OSPF connectivity should be maintained over the DMVPN network. Area 1 should remain a stub area.

3 points

Incident 4

- BGP routes are not being discovered between Sw1 in AS63513 and Sw2 in AS654513.
- Fix the problem so that Sw1 and Sw2 have the following BGP table entries:

[Click here to view code image](#)

```
Sw2# show ip bgp
```

```
BGP table version is 3, local router ID is 120.100.27.2
```

```
Status codes: s suppressed, d damped, h history, * valid, > best,
```

```
i - internal,
```

```
    r RIB-failure, S Stale, m multipath, b backup-path,
```

```
    x best-external, f RT-Filter
```

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

Network	Next Hop	Metric	LocPrf	Weight	Path
r>i120.100.11.0/24	4.4.4.4	0	100	0	63513 i
*> 120.100.20.0/24	0.0.0.0	0			32768 i

Sw1# **show ip bgp**

BGP table version is 9, local router ID is 120.100.12.1

Status codes: s suppressed, d damped, h history, * valid, > best,

i - internal,

r RIB-failure, S Stale, m multipath, b backup-path,

x best-external, f RT-Filter

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

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 120.100.11.0/24	0.0.0.0	0			32768 i
*>i120.100.20.0/24	5.5.5.5	0	100	0	64513 i

2 points

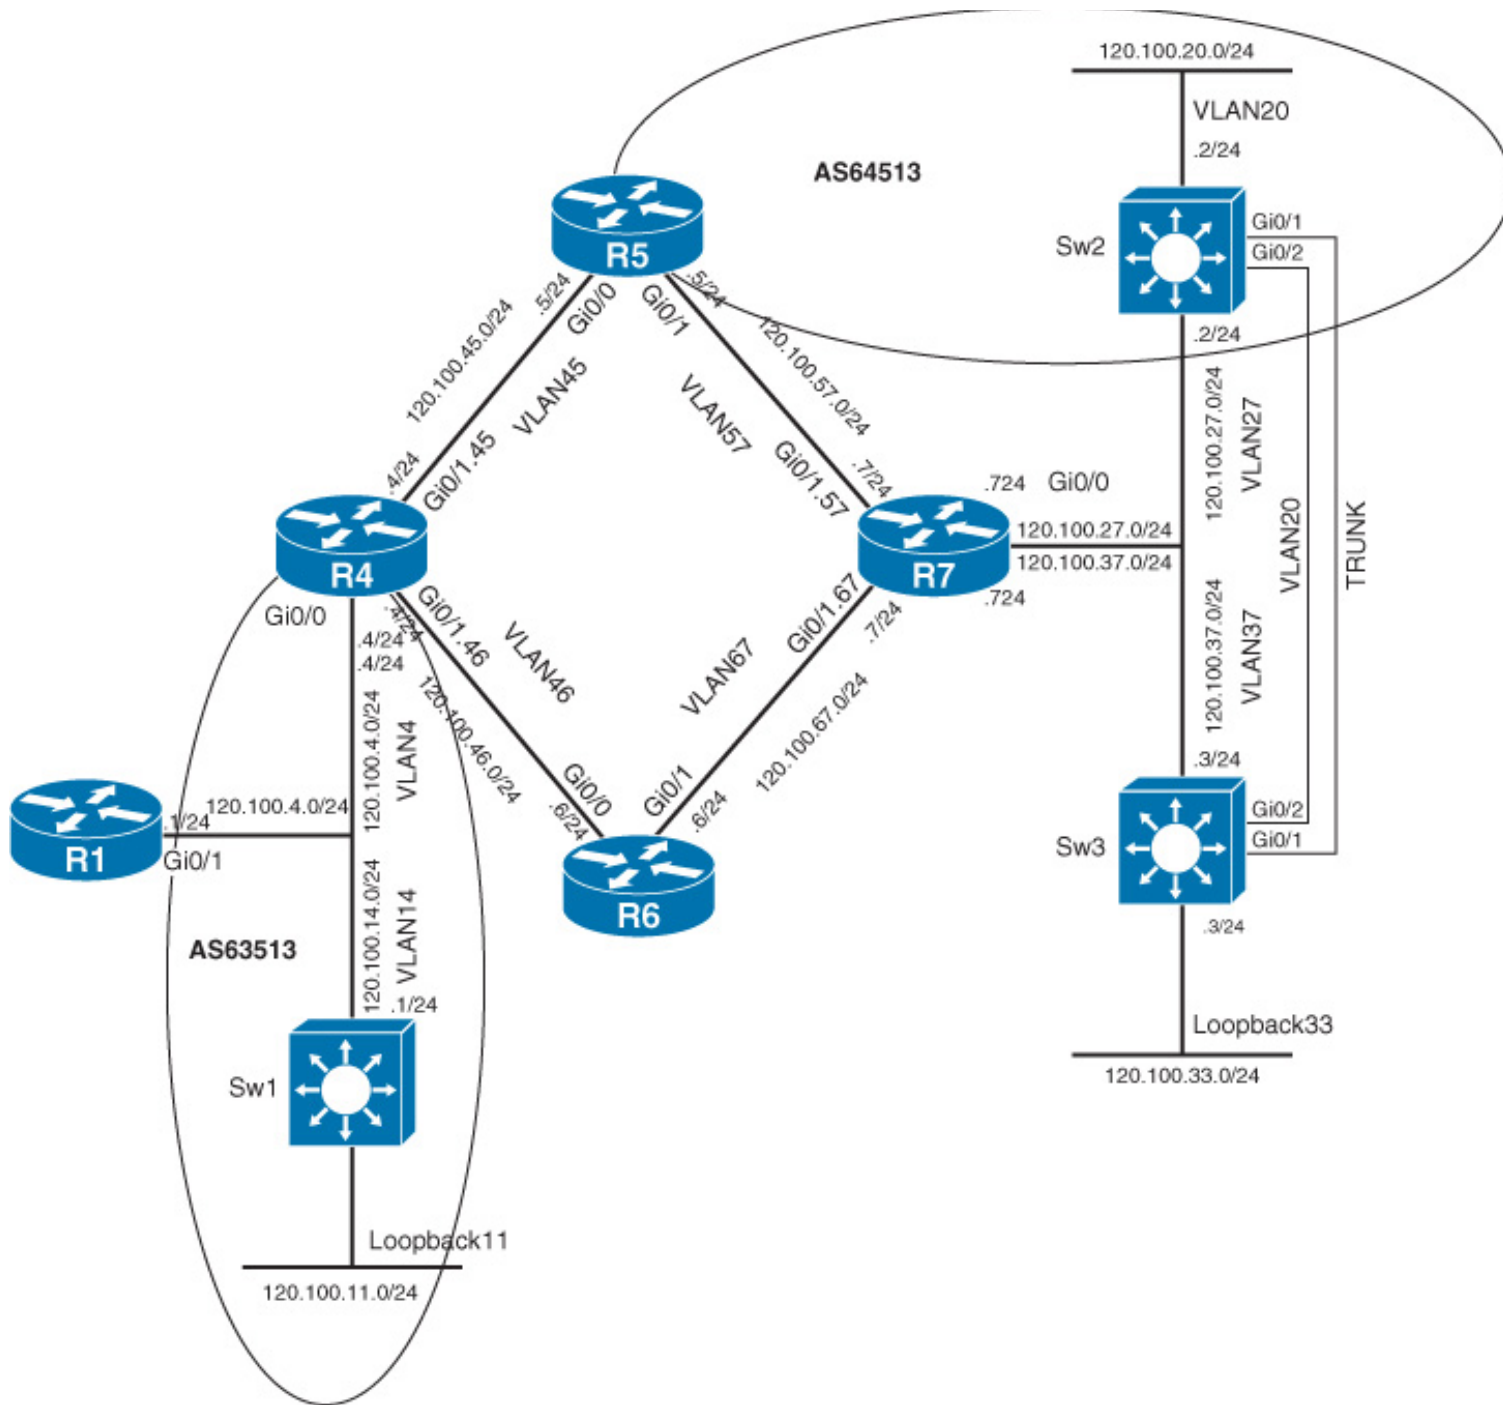


Figure 2-8 Incident 4 Diagram

Incident 5

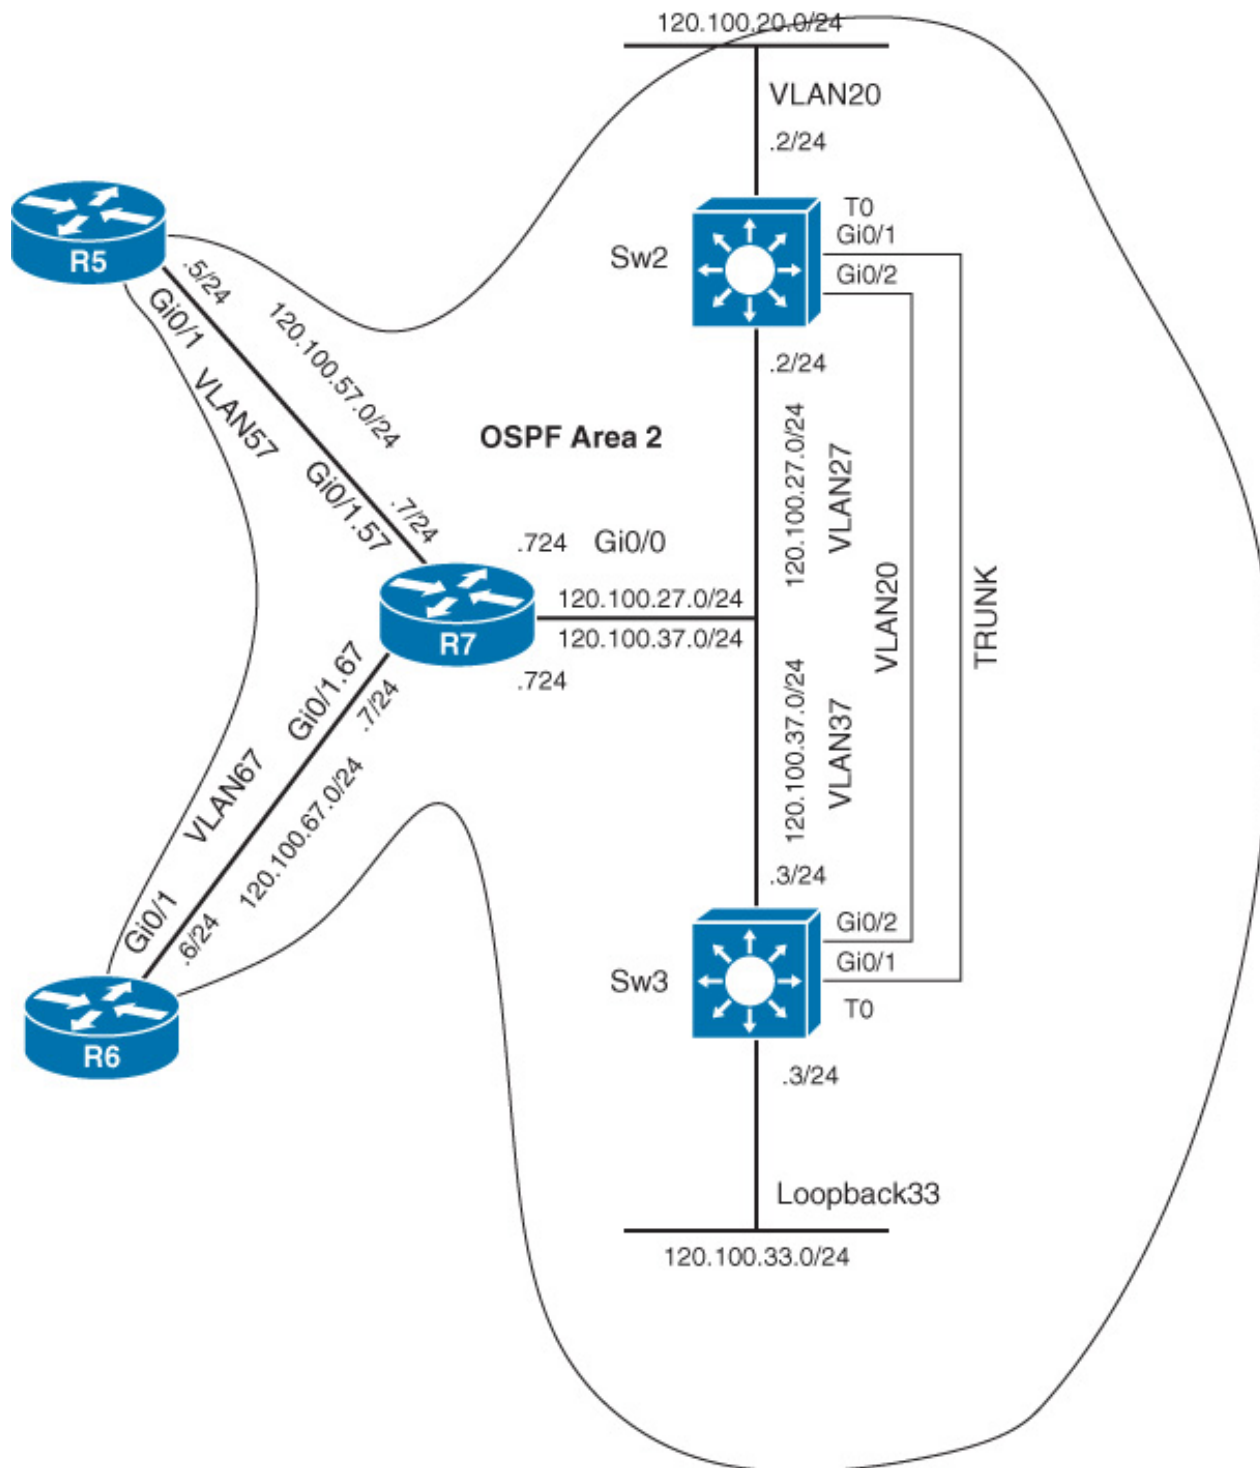


Figure 2-9 *Incident 5 Diagram*

- Sw3 is not receiving routes from its OSPF neighbor R7.
- Fix the problem so that Sw2 has the following routing table entry and connectivity:

[Click here to view code image](#)

```
Sw3# show ip route 120.100.57.0
```

```
Routing entry for 120.100.57.0/24
```

```
Known via "ospf 1", distance 110, metric 2, type intra area
```

```
Last update from 120.100.37.7 on Vlan37, 01:54:32 ago
```

```
Routing Descriptor Blocks:
```

```
* 120.100.37.7, from 7.7.7.7, 01:54:32 ago, via Vlan37
```

```
Route metric is 2, traffic share count is 1
```

```
Sw3# ping 120.100.57.5
```

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

```
Sending 5, 100-byte ICMP Echos to 120.100.57.5, timeout is 2 seconds:
```

```
!!!!
```

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

- You may not modify any parameter under the OSPF process of Sw3.

1 point

Incident 6

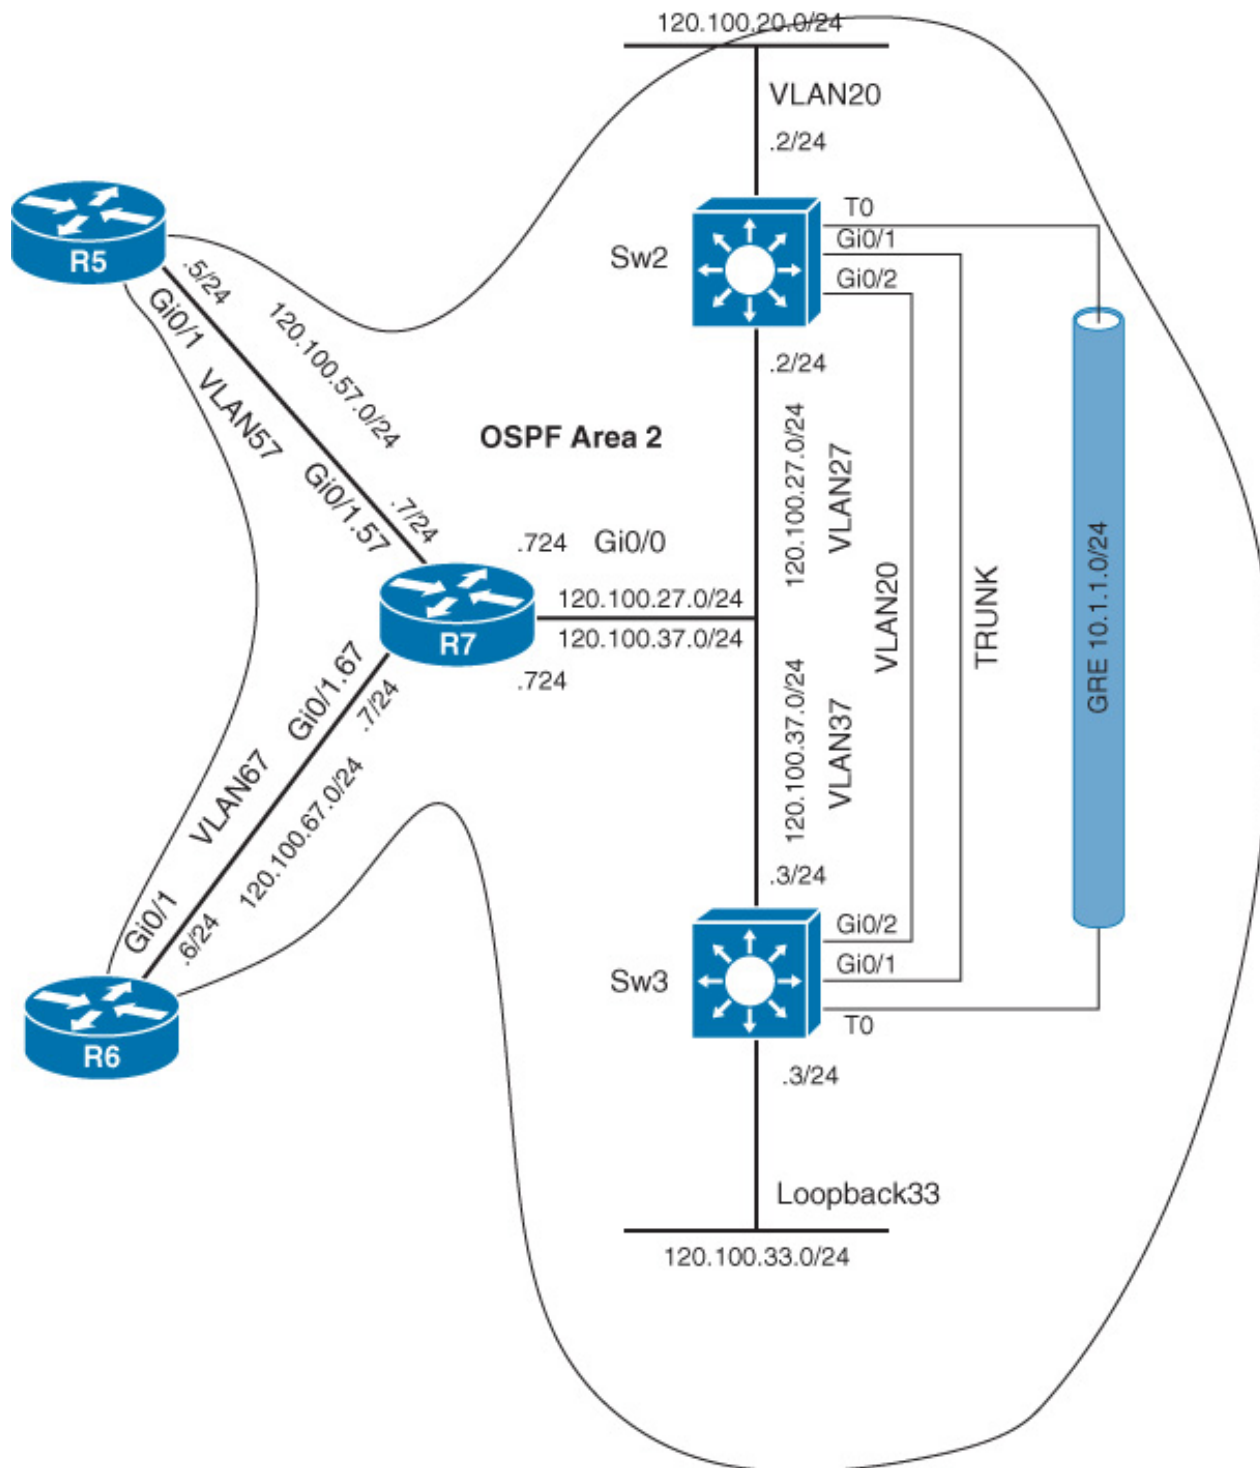


Figure 2-10 *Incident 6 Diagram*

- A GRE tunnel configured between Sw2 and Sw3 is flapping. OSPF is configured to run between the two routers.
- Fix the problem so that Sw2 and Sw3 have the following reliable OSPF adjacency with each other:

[Click here to view code image](#)

```
SW2# sh ip ospf neighbor | include Tunnel0
120.100.33.1  0 FULL/ -   00:00:32  10.1.1.3   Tunnel0
```

```
SW3# sh ip ospf neigh | include Tunnel0
120.100.27.2  0 FULL/ -   00:00:33  10.1.1.2   Tunnel0
```

1 point

Incident 7

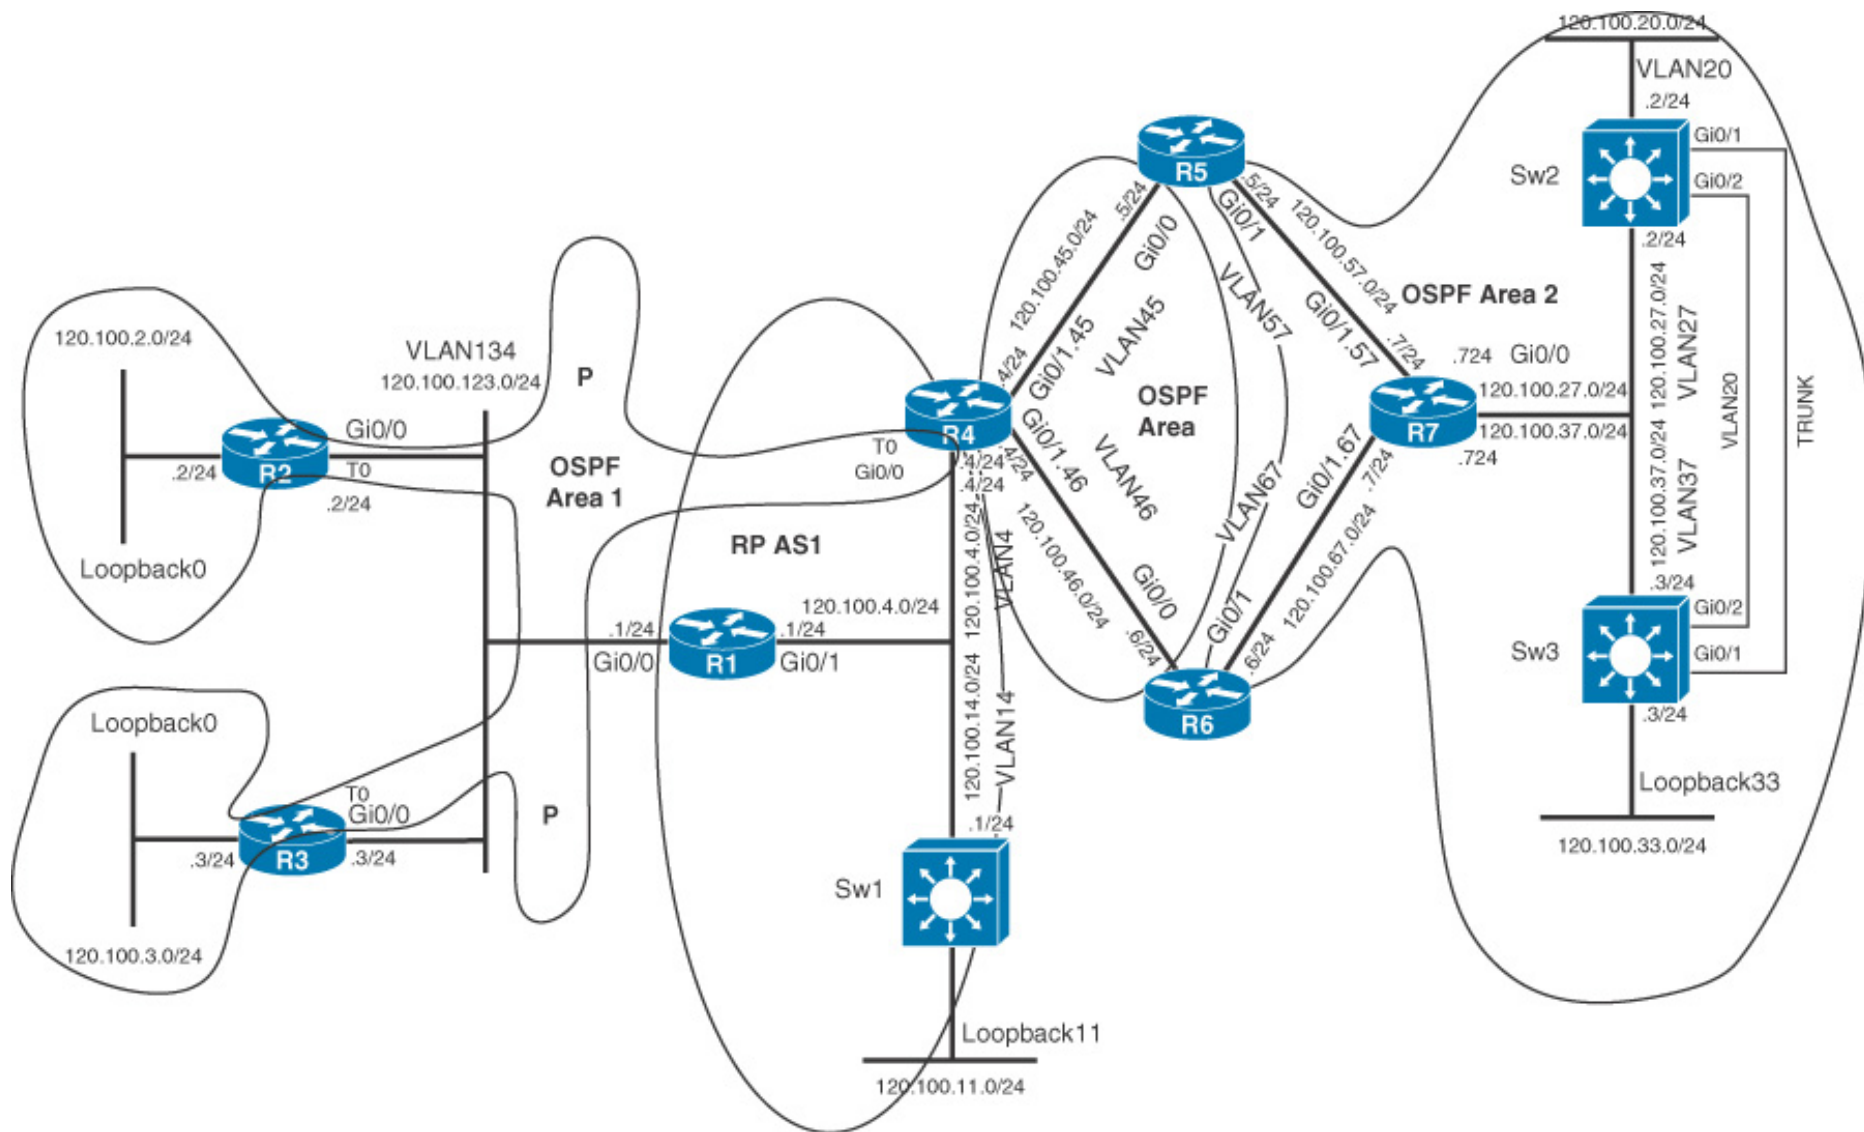


Figure 2-11 Incident 7 Diagram

- Sw3 cannot communicate with networks in the EIGRP domain.
- Fix the problem so that Sw3 has the following OSPF route and connectivity:

Click here to view code image

<https://t.me/learningnets>

```
Sw3# show ip route | include 120.100.11.0
```

```
O E1 120.100.11.0/24 [110/5003] via 120.100.37.7, 00:01:48, Vlan37
```

```
Sw3# ping 120.100.11.1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.11.1, timeout is 2 seconds:

```
!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/17 ms

2 points

Incident 8




Figure 2-12 Incident 8 Diagram

- Sw3 cannot telnet to Sw2 using IPv6 on VLAN20. The network administrator of Sw2 requires that ICMP ping testing should not be allowed to Sw3 but that telnet access is allowed.
- Fix the problem so that Sw3 has the following response to an ICMP ping and a Telnet session to Sw2:
<https://t.me/learningnets>

[Click here to view code image](#)

```
Sw3# ping ipv6 CC1E:CAFE::1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to CC1E:CAFE::1, timeout is 2 seconds:

```
AAAAA
```

Success rate is 0 percent (0/5)

```
Sw3# telnet CC1E:CAFE::1
```

```
Trying CC1E:CAFE::1 ... Open
```

- Do not modify the IPv6 address assignment on Sw3.
- The network administrator has chosen to suppress RAs on Sw2, so do not remove this suppression.

2 points

Incident 9

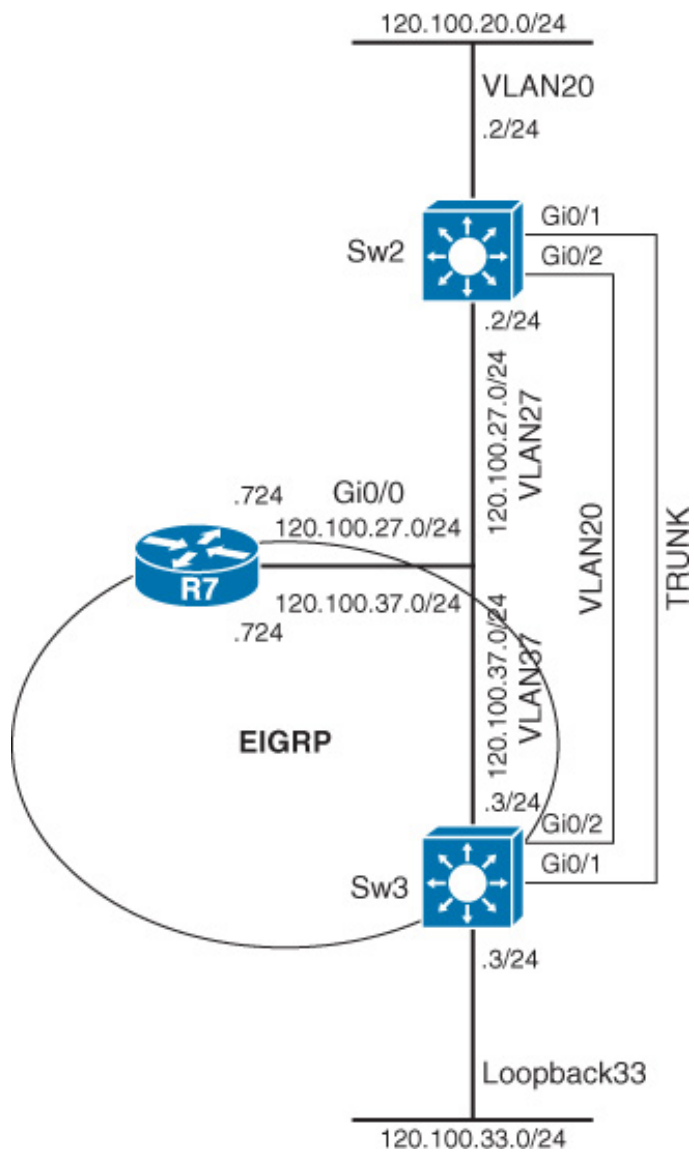


Figure 2-13 Incident 9 Diagram

- Request a colleague to reconfigure a preconfigured EIGRP process of JAKE running on Sw3 VLAN37 that connects to R7 on network 120.100.37.0/24. Have your colleague change the autonomous system number on Sw3 from 11 to a number of his choice.

<https://t.me/learningnets>

■ Use the embedded packet-capture facility on R7 to inspect the EIGRP traffic received from Sw3 on network 120.100.37.0/24 to determine which autonomous system Sw3 has been changed to and modify the autonomous system on R7 to match to form a successful neighbor relationship between R7 and Sw3. (Should a question similar to this arise on the real lab exam, it would, of course, relate to an unknown parameter on a device outside of your administrative control.)

3 points

“ASK THE PROCTOR”

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 will be present to ensure that you do not have problems with the lab environment and to maintain the timing element of the exam.

Section 1

Incident 1

Q . I don't see visibility of the 120.100.14.0/24 network within my OSPF domain. Would this be the default behavior of the redistribution, or is this an injected fault?

A . This would be the default behavior, but you need to fix it. Consider whether this network is learned via <https://t.me/learningnets>

EIGRP or is just a connected interface with EIGRP running on it.

Q . Ah, okay, so if I enabled this network in OSPF and ran a passive interface on it, would this be acceptable?

A . Well, it would work, but there is a simpler method of redistributing this network into EIGRP without actually running OSPF on it. Keep your solution as simple as possible.

Incident 2

Q . From the information provided within the incident, I can see that items such as my DMVPN peering and NHRP addresses are incorrect. After changing them, the DMVPN is still down. Is there something else non-DMVPN-related that would cause an issue with this incident?

A . It is all related to the overall DMVPN solution; remember that you require a crypto element to this.

Incident 3

Q . I can see that there is an MTU issue between OSPF neighbors, yet I still have a connectivity problem after I have fixed this. I believe this is down to having a static o/o route and receiving a default route over the OSPF network. Can I change the OSPF area type so that I do not receive a default route?

A . The incident instructions state that you must leave Area 1 as a stub area. Think about how you can maintain the stub area while still sending specific OSPF routes over it.

Q . Okay, can I change the area to just a normal stub area instead of being totally stubby so that I lose the

default route and receive just type 3 LSAs?

A . Yes.

Incident 4

Q . My iBGP within each autonomous system looks fine. I can see that there is an issue with eBGP between autonomous systems, but am I able to change the peering or can I just enable eBGP multihop between neighbors?

A . Keep your solution as simple as possible with minimal configuration changes to successfully bring up your eBGP peers.

Q . I can see there is a route map that is changing the next-hop address on R4 advertised routes. Can I modify this route map or remove it?

A . If the route-map is solely changing the next-hop and you believe this is the sole cause of the issue then you can just remove the route-map from the neighbor configuration.

Incident 5

Q . Can I change the prefix-list “allow” that OSPF is using inbound on Sw3?

A . Yes.

Incident 6

Q . I don't see any flapping issue reported on the command line of Sw2 and Sw3. Is this intermittent?

A . Make sure that you are running "logging monitor" within your switch configuration in addition to

"terminal monitor" on any Telnet or SSH session to the devices.

Q . I can see an OSPF cost has been assigned. Can I modify or simply remove this cost on Sw2 VLAN27?

A . You can remove it.

Incident 7

Q . I am receiving OSPF routes on Sw3 from R7, so everything looks fine. Is this correct?

A . You will not be receiving external routes via R7, so you will not have full connectivity to the EIGRP domain.

Q . The redistribution issue was fixed in an earlier incident. Is this related to something specific with Sw3?

A . Check your OSPF database. Consider why the external routes would not make it into the routing table from the database.

Q . Okay, I can see this is a duplicate route ID issue between R4 and Sw3. Should I change R4's or R7's router ID?

<https://t.me/learningnets>

A . Either is acceptable.

Incident 8

Q . If Sw3 has been configured to assign its IPv6 prefix automatically from RA advertisements from Sw2 and Sw2 has been configured to block outbound RAs, can I remove the blocking or manually assign an IPv6 prefix to Sw3?

A . No, the incident states that you are not permitted to perform either of these actions. Look for a way to allow requests from Sw3 for a prefix to be accepted by Sw2.

Q . I believe the prefix assignment would be facilitated by use of ICMP, and this is all blocked from entering Sw2 as ping should be denied. So, how can the address request be accepted?

A . Modify the ACL to just allow the prefix request and block all subsequent ICMP packets.

Q . How would I know which ICMP ports are required for prefix assignment?

A . Create a new ACL on Sw2 and log the traffic or run a debug to see the traffic between devices.

Incident 9

Q . I have had a colleague change the EIGRP autonomous system on Sw3 and have run the embedded packet capture successfully on R7. However, there is no way I can decode this capture to tell where the autonomous system information resides. Surely, I am not expected to know each hex code value per packet, am I?

<https://t.me/learningnets>

A . Of course you aren't. Time to think outside of the box. Consider setting up another EIGRP instance

between Sw2 and R7 (for example, on VLAN27 with a known EIGRP autonomous system) and run the same packet capture. Then look for differences in the packet captures.

Q . I have completed a packet capture of a known autonomous system (decimal 10, so hex A) between Sw2

and R7. There are multiple A values within the capture dump, so it is impossible to tell where the autonomous system field may be within the capture. Surely this approach will not work, will it?

A . Change the autonomous system number into something more visible and less likely to be within the

capture by default. An autonomous system number of 255 would be hex FF, for example.

TROUBLESHOOTING LAB 2 DEBRIEF

The debrief aims to achieve the following items:

- Definition of the problem and identification of the symptoms (questioning to the void to identify meaningful symptoms)
- Definition of hypothesis and proof-testing of possible causes
- Design and implementation of a final solution (prior to application of the configuration)
- Verification of 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 SW1: L011 120.100.11.0/24 cannot ping users on VLAN20 Sw2 120.100.20.0/24.

- Fix the problem so that Sw1 Lo11 120.100.11.1 and Sw1 VLAN14 120.100.14.1 can ping Sw2 VLAN20 120.100.20.2, as follows:

[Click here to view code image](#)

```
Sw1# ping 120.100.20.2 source 120.100.11.1
```

```
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 120.100.11.1
```

```
!!!!
```

```
Sw1# ping 120.100.20.2 source 120.100.14.1
```

```
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 120.100.14.1
```

```
!!!!
```

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

1 point

[Figure 2-5](#) shows that Sw1 resides within the EIGRP domain and Sw2 resides within the OSPF domain. Therefore, the best location to begin testing is the redistribution router R4. [Example 2-1](#) shows that initial testing from R4 to each end system is successful. This indicates a redistribution problem is likely to be the cause of any end-to-end issue between networks. A **show ip route 120.100.20.0** on Sw1 proves that OSPF is being redistributed correctly into EIGRP and **show ip route 120.100.11.0** and **show ip route 120.100.14.0** on Sw2 <https://t.me/learningnets>

proves that EIGRP is not being redistributed correctly into OSPF. Inspection of the redistribution configuration on R4 shows that it is running EIGRP AS10 yet is redistributing AS1 into OSPF. After this has been corrected, both network 120.100.11.0/24 and 120.100.14.0/24 are reachable. If you have answered this correctly, as detailed in [Example 2-1](#), you have scored 1 point.

Example 2-1 *Redistribution Verification and Reconfiguration*

[Click here to view code image](#)

```
R4# ping 120.100.11.1
```

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

```
Sending 5, 100-byte ICMP Echos to 120.100.11.1, timeout is 2 seconds:
```

```
!!!!
```

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

```
R4# ping 120.100.14.1
```

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

```
Sending 5, 100-byte ICMP Echos to 120.100.14.1, timeout is 2 seconds:
```

```
!!!!
```

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

```
R4# 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 = 1/2/4 ms
```

```
Sw1# show ip route 120.100.20.0
```

```
Routing entry for 120.100.20.0/24
```

```
Known via "eigrp 10", distance 170, metric 2562816, type external
```

```
Redistributing via eigrp 10
```

```
Last update from 120.100.14.4 on Vlan14, 1w6d ago
```

```
Routing Descriptor Blocks:
```

```
* 120.100.14.4, from 120.100.14.4, 1w6d ago, via Vlan14
```

```
Route metric is 2562816, traffic share count is 1
```

```
Total delay is 110 microseconds, minimum bandwidth is 1000 Kbit
```

```
Reliability 255/255, minimum MTU 1500 bytes
```

```
Loading 1/255, Hops 1
```

```
Sw2# sh ip route 120.100.11.0
```

```
% Subnet not in table
```

```
Sw2# sh ip route 120.100.14.0
```

```
% Subnet not in table
```

```
R4# show run | section eigrp JAKE
```

```
router eigrp JAKE
```

```
!
```

```
address-family ipv4 unicast autonomous-system 10
```

```
topology base
```

```
redistribute ospf 1 metric 1000 10 255 1 1500
```

```
exit-af-topology
```

```
network 120.100.4.0 0.0.0.255
```

```
network 120.100.14.0 0.0.0.255
```

```
exit-address-family
```

```
R4# show run | section redistribute eigrp
```

```
redistribute eigrp 1 metric 5000 metric-type 1 subnets
```

```
R4(config)# router ospf 1
```

```
R4(config-router)# no redistribute eigrp 1 metric 5000 metric-type 1 subnets
```

```
R4(config-router)# redistribute eigrp 10 metric 5000 metric-type 1 subnets
```

```
Sw2# sh ip route 120.100.11.0
```

```
Routing entry for 120.100.11.0/24
```

```
Known via "ospf 1", distance 110, metric 5003, type extern 1
```

```
Last update from 120.100.27.7 on Vlan27, 00:00:08 ago
```

```
Routing Descriptor Blocks:
```

```
* 120.100.27.7, from 4.4.4.4, 00:00:08 ago, via Vlan27
```

```
Route metric is 5003, traffic share count is 1
```

```
Sw1# ping 120.100.20.2 source 120.100.11.1
```

```
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 120.100.11.1

!!!!

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

```
Sw1# ping 120.100.20.2 source 120.100.14.1
```

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 120.100.14.1

!!!!

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

Incident 2

- DMVPN access between spoke routers R2 and R3 to hub router R4 is down. Fix the access with DMVPN connectivity, as follows:

[Click here to view code image](#)

```
R2# show dmvpn
```

Tunnelo, Type:Spoke, NHRP Peers:1,

```
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
```

```
-----
```

```
1 120.100.4.4 120.100.40.4 UP 00:06:40 S
```

```
R3# show dmvpn
```

```
Tunnel0, Type:Spoke, NHRP Peers:1,
```

```
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
```

```
-----  
1 120.100.4.4 120.100.40.4 UP 00:06:40 S
```

- Do not modify or remove the existing default route on R2 and R3 pointing to R1 120.100.123.1.

2 points

Inspection of [Figure 2-6](#) shows that R2 and R3 connect to the network over a DMVPN connection to R4 that belongs to the OSPF domain from where VLAN20 resides. The incident states that R2 and R3 have a default gateway to R1 that cannot be modified. A logical starting point is to select either R2 or R3 and make sure that you have reachability to R4 and then to confirm whether the DMVPN peering is up.

[Example 2-2](#) shows that R2 has ICMP reachability at least to R4 120.100.4.4 (you can tell the peering is to this address on R4 rather than a loopback due to the information provided within the incident) but that the DMVPN connection is down (State - never) and that the peer IP address is set to 120.100.4.1 rather than 120.100.4.4 within the incident. At this point, it is worth comparing the DMVPN configuration between the hub and spokes. [Example 2-2](#) shows that there are differences between the two. First, R2's NHRP map is pointing to 120.100.4.1. This is the IP address of R1 VLAN4 and not of R4, which is 120.100.4.4, so this needs changing. Second, the NHRP multicast IP has also been configured for R1 on 120.100.4.1 rather than 120.100.4.4, so this also needs changing. Third, the NHRP IP address on R2 is pointing to a tunnel address of <https://t.me/learningnets>

120.100.40.14. Inspection of R4's DMVPN tunnel interface shows this should actually be 120.100.40.4. Once these three items have been reconfigured the DMVPN session is correctly pointing to R4 but the VPN is down, at this point it would be worth examining the crypto side of the configuration to see if this is causing an issue. [Example 2-2](#) shows that the crypto session is down; negotiating and a debug reveal errors on key exchange, which would indicate a password type of issue. The **show** command **show crypto isakmp key** on both R2 and R4 reveals that R2 is using a pre-shared key of CC1E and R4 is using a key of CCIE. Changing one or the other to ensure that the same key is used would cure the crypto issue. [Example 2-2](#) shows reconfiguration of the key on R2 and the resulting operational DMVPN connectivity. If you have solved this incident successfully, as detailed in [Example 2-2](#), you have scored 2 points. R3 has identical issues to R2 within this incident.

Example 2-2 *R2 and R4 DMVPN Testing and Configuration*

[Click here to view code image](#)

```
R2# ping 120.100.4.4
```

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

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

```
!!!!
```

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

```
R2# show dmvpn
```

```
Legend: Attrb --> S - Static, D - Dynamic, I - Incompletea
```

```
      N - NATed, L - Local, X - No Socket
```

```
      # Ent --> Number of NHRP entries with same NBMA peer
```

Tunnel0, Type:Spoke, NHRP Peers:1,

```
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
```

```
-----
```

```
1 120.100.4.1 120.100.40.4 NHRP never S
```

R2# **show run int tunnel0**

```
interface Tunnel0
```

```
ip address 120.100.40.2 255.255.255.248
```

```
no ip redirects
```

```
ip mtu 1419
```

```
ip nhrp authentication SECRET
```

```
ip nhrp map 120.100.40.4 120.100.4.1
```

```
ip nhrp map multicast 120.100.4.1
```

```
ip nhrp network-id 10
```

```
ip nhrp nhs 120.100.40.14
```

```
ip ospf 1 area 1
```

```
tunnel source GigabitEthernet0/0
```

```
tunnel mode gre multipoint
```

```
tunnel key 1
```

```
tunnel protection ipsec profile IPSEC
```

```
end
```

```
R4# show run int tunnel0
```

```
interface Tunnel0
```

```
ip address 120.100.40.4 255.255.255.0
```

```
no ip redirects
```

```
ip mtu 1416
```

```
ip nhrp authentication SECRET
```

```
ip nhrp map multicast dynamic
```

```
ip nhrp network-id 10
```

```
ip ospf 1 area 1
```

```
tunnel source GigabitEthernet0/0.4
```

```
tunnel mode gre multipoint
```

```
tunnel key 1
```

```
tunnel protection ipsec profile IPSEC
```

```
end
```

```
R2(config)# int t0
```

```
R2(config-if)# ip nhrp map 120.100.40.4 120.100.4.4
```

```
R2(config-if)# no ip nhrp map multicast 120.100.4.1
```

```
R2(config-if)# ip nhrp map multicast 120.100.4.4
```

```
R2(config-if)# no ip nhrp nhs 120.100.40.14
```

```
R2(config-if)# ip nhrp nhs 120.100.40.4
```

Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete

N - NATed, L - Local, X - No Socket

Ent --> Number of NHRP entries with same NBMA peer

Tunnel0, Type:Spoke, NHRP Peers:1,

```
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb
```

```
-----  
1 120.100.4.4 120.100.40.4 NHRP never S
```

R2# **show crypto session**

Crypto session current status

Interface: Tunnel0

Session status: DOWN-NEGOTIATING

Peer: 120.100.4.4 port 500

IKE SA: local 120.100.123.2/500 remote 120.100.4.4/500 Inactive

IPSEC FLOW: permit 47 host 120.100.123.2 host 120.100.4.4

Active SAs: 0, origin: crypto map

R2# **debug dmvpn all all**

*Jan 12 12:13:23: ISAKMP:(1028): retransmitting phase 1 MM_KEY_EXCH...

*Jan 12 12:13:23: ISAKMP (0:1028): incrementing error counter on sa, attempt 1 of 5:

```
*Jan 12 12:13:23: ISAKMP:(1028): retransmitting phase 1 MM_KEY_EXCH
*Jan 12 12:13:23: ISAKMP:(1028): sending packet to 120.100.4.4 my_port 500 peer_port
500 (I) MM_KEY_EXCH
*Jan 12 12:13:23: ISAKMP:(1028):Sending an IKE IPv4 Packet.
*Jan 12 12:13:23: ISAKMP (0:1028): received packet from 120.100.4.4 dport 500 sport
500 Global (I) MM_KEY_EXCH
*Jan 12 12:13:23: ISAKMP:(1028): phase 1 packet is a duplicate of a previous
```

```
R2# show crypto isakmp key
```

Keyring	Hostname/Address	Preshared Key
default	0.0.0.0 [0.0.0.0	CC1E

```
R4# show crypto isakmp key
```

Keyring	Hostname/Address	Preshared Key
default	0.0.0.0 [0.0.0.0]	CCIE

```
R2# show run | include crypto
```

```
crypto isakmp policy 1
crypto isakmp key CC1E address 0.0.0.0 0.0.0.0
crypto ipsec transform-set DMVPN esp-des esp-md5-hmac
crypto ipsec profile IPSEC
```


[Click here to view code image](#)

```
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 = 1/2/4 ms

```
R2# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	0	FULL/ -	00:00:34	120.100.40.4	Tunnel0

```
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 = 1/2/4 ms

```
R0# show ip ospf neighbor
```

<https://it-ebooks.com/learningnets>

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	0	FULL/ -	00:00:34	120.100.40.4	Tunnel0

■ Do not modify or remove the existing default route on R2 and R3 pointing to R1 120.100.123.1. Do not add further routing protocols to R2, R3, and R4. OSPF connectivity should be maintained over the DMVPN network. Area 1 should remain a stub area.

3 points

The DMVPN connection should have been successfully fixed in the previous incident. If this is still down, you are unlikely to fix this issue. On the real lab exam, the majority of questions are completely isolated. The incident shows that OSPF is required over the DMVPN connection between R2 and R3 as spoke routers to R4 the hub router.

Example 2-3 shows that R2's neighbor relationship to R4 over Tunnel0 is down. Basic OSPF troubleshooting should tell you to compare the two routers' OSPF configuration using the **show ip ospf interface** command. Network types and Hello/Dead intervals are identical, but the interface address mask differs. R4 has a /24 mask, and R2 has /29 mask, which will cause an issue. Debugging the OSPF adjacency also shows an MTU issue between routers with R2 using 1419 and reporting a smaller value for R4.

Example 2-3 shows reconfiguration on R2 for the tunnel interface address mask and MTU to match that of R4 using a /24 and MTU value of 1416, the result of which is full neighbor adjacency between the two routers. However, testing shows that VLAN20 is not reachable from R2. The command **show ip ospf routes** on R2 reveals that there are no routes being learned over the DMVPN tunnel between R2 and R4, even though they are successfully peering with each other. Inspection of the OSPF database on R2 shows that a default route

<https://t.me/learningnets>

0.0.0.0 is being advertised from R4 (4.4.4.4) to R2, which is not showing in the routing table. This, of course, is a behavior of the area type being a stub and the ABR router R4 being configured to send a default route within the area as it has been configured with the **no-summary** parameter. However, the default route present on R2 is a static route pointing to R1 for DMVPN peering (and which cannot be modified as per the incident instructions) with an admin distance of 1, which clearly overrides the OSPF route, which would have an admin distance of 110. Therefore, the incident can be solved only if R4 can send the individual routes (LSA types 1, 2, and 3) rather than just a default route into the area for all nonstub area routes. This is simply achieved by removing the **no-summary** parameter within the area configuration of R4.

Example 2-3 shows inspection of the area type and modification from a totally stubby area on R4 to a standard stubby area to remove the default route and resulting route propagation of VLAN20 to R2, with the resulting successful testing to VLAN20 from R2. If you have solved this incident successfully, as detailed in Example 2-3, you have scored 3 points. R3 has identical issues to R2 within this incident.

Example 2-3 *R2 and R4 OSPF Testing and Configuration*

[Click here to view code image](#)

```
R2# show ip ospf neigh
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	0	DOWN/ -	-	120.100.40.4	Tunnel0

```
R2# show ip ospf int t0
```

```
Tunnel0 is up, line protocol is up
```

Internet Address 120.100.40.2/29, Area 1

Process ID 1, Router ID 120.100.2.2, Network Type POINT_TO_POINT, Cost: 11111

Enabled by interface config, including secondary ip addresses

Transmit Delay is 1 sec, State POINT_TO_POINT

Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

oob-resync timeout 40

Hello due in 00:00:07

Supports Link-local Signaling (LLS)

Cisco NSF helper support enabled

IETF NSF helper support enabled

Index 1/1, flood queue length 0

Next oxo(0)/oxo(0)

Last flood scan length is 0, maximum is 0

Last flood scan time is 0 msec, maximum is 0 msec

Neighbor Count is 1, Adjacent neighbor count is 0

Suppress hello for 0 neighbor(s)

R4# show ip ospf int t0

Tunnel1 is up, line protocol is up

Internet Address 120.100.40.4/24, Area 1

Process ID 1, Router ID 4.4.4.4, Network Type POINT_TO_POINT, Cost: 1000

Topology-MTID	Cost	Disabled	Shutdown	Topology Name
---------------	------	----------	----------	---------------

0	1000	no	no	Base
---	------	----	----	------

<https://t.me/learningnets>

Enabled by interface config, including secondary ip addresses

Transmit Delay is 1 sec, State POINT_TO_POINT

Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

oob-resync timeout 40

Hello due in 00:00:01

Supports Link-local Signaling (LLS)

Cisco NSF helper support enabled

IETF NSF helper support enabled

Index 1/4, flood queue length 0

Next ox0(o)/ox0(o)

Last flood scan length is 1, maximum is 7

Last flood scan time is 0 msec, maximum is 0 msec

Neighbor Count is 1, Adjacent neighbor count is 0

Suppress hello for 0 neighbor(s)

R2# debug ip ospf adjacency

*Jan 12 12:31:11: OSPF: 2 Way Communication to 4.4.4.4 on Tunnel0, state 2WAY

*Jan 12 12:31:11: OSPF: Send DBD to 4.4.4.4 on Tunnel0 seq 0x1418 opt 0x50 flag 0x7
len 32

*Jan 12 12:31:11: OSPF: Rcv DBD from 4.4.4.4 on Tunnel0 seq 0x2D1 opt 0x50 flag 0x7
len 32 mtu 1416 state EXSTART

*Jan 12 12:31:11: OSPF: Nbr 4.4.4.4 has smaller interface MTU

```
R2# show run int t0 | include mtu
```

```
ip mtu 1419
```

```
R2(config)# int t0
```

```
R2(config-if)# ip add 120.100.40.2 255.255.255.0
```

```
R2(config-if)# ip mtu 1416
```

```
R2# sh ip ospf neigh
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	0	FULL/ -	00:00:34	120.100.40.4	Tunnel0

```
R2# 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:

```
U.U.U
```

Success rate is 0 percent (0/5)

```
R2# show ip route ospf
```

```
R2# show ip ospf database
```

Router Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
4.4.4.4	4.4.4.4	120	0x8000002D	0x00D43E	2
120.100.2.2	120.100.2.2	119	0x8000008A	0x00FFA8	3

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	4.4.4.4	301	0x80000025	0x00F019

```
R2# show ip route 0.0.0.0
```

```
Routing entry for 0.0.0.0/0, supernet
```

```
Known via "static", distance 1, metric 0, candidate default path
```

```
Routing Descriptor Blocks:
```

```
* 120.100.123.1
```

```
Route metric is 0, traffic share count is 1
```

```
R4# show ip ospf | begin Area 1
```

```
Area 1
```

```
Number of interfaces in this area is 1
```

```
It is a stub area, no summary LSA in this area
```

Generates stub default route with cost 1

Area has no authentication

SPF algorithm last executed 00:03:57.620 ago

SPF algorithm executed 15 times

Area ranges are

Number of LSA 3. Checksum Sum 0x02C3FF

Number of opaque link LSA 0. Checksum Sum 0x000000

Number of DCbitless LSA 0

Number of indication LSA 0

Number of DoNotAge LSA 0

Flood list length 0

```
R4(config)# router ospf 1
```

```
R4(config-router)# no area 1 stub no-summary
```

```
R4(config-router)# area 1 stub
```

```
R2# show ip route ospf | include 120.100.20.0
```

```
O IA 120.100.20.0 [110/11114] via 120.100.40.4, 00:01:36, Tunnel0
```

```
R2#
```

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

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 = 4/4/8 ms

Incident 4

- BGP routes are not being discovered between Sw1 in AS63513 and Sw2 in AS654513.
- Fix the problem so that Sw1 and Sw2 have the following BGP table entries:

[Click here to view code image](#)

```
SW2# show ip bgp
```

```
BGP table version is 3, local router ID is 120.100.27.2
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i
```

```
- internal,
```

```
    r RIB-failure, S Stale, m multipath, b backup-path, x
```

```
    best-external, f RT-Filter
```

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

```
Network        Next Hop        Metric LocPrf Weight Path
r>i120.100.11.0/24 4.4.4.4        0 100    0 63513 i
*> 120.100.20.0/24 0.0.0.0        0     32768 i
```

```
Sw1# show ip bgp
```

```
BGP table version is 9, local router ID is 120.100.12.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i
```

```
- internal,
```

```
  r RIB-failure, S Stale, m multipath, b backup-path, x
```

```
  best-external, f RT-Filter
```

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

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 120.100.11.0/24	0.0.0.0	0		32768	i
*>i120.100.20.0/24	5.5.5.5	0	100	0	64513 i

2 points

You can establish from the incident details and from issuing the command **show ip bgp** on Sw1 and Sw2 that Sw1 is advertising network 120.100.11.0/24 into BGP and Sw2 is advertising network 120.100.20.0/24. By examining the edge of each autonomous system, you can determine whether the issue of each switch not receiving the other's BGP update is internal to the autonomous system and thus iBGP or external to the autonomous system and thus eBGP related.

Example 2-4 shows the local routes being advertised from each switch and that R4 and R5 are only receiving their iBGP routes. Therefore, the issue appears to be related to eBGP between R4 and R5. The example shows that by issuing the command **show ip bgp summary** on R4 and R5 that the neighbor relationship between the two routers is in the idle state. Inspection of the BGP configuration on each router shows the routers are <https://t.me/learningnets>

peering from loopback interfaces and therefore require eBGP multihop enabled, which is missing from the configuration. Post-eBGP multihop configuration between R4 and R5, it can be seen the neighbor adjacency issue is rectified and that R4 actually learns the BGP network 120.100.20.0/24 as advertised by Sw2. However, R5 does not learn the BGP network 120.100.11.0/24 as advertised by Sw1. Inspection of the advertised routes from R4 shows that the network is, in fact, being advertised by R4 but not received on R5. A debug of the BGP updates from R4 on R5 shows that the route is actually being received yet denied because it contains a next hop that is actually an IP address that resides on R5. Inspection of the BGP configuration on R4 details a route map labeled med is being used toward R5. This route map actually changes the next hop to 120.100.45.5, which belongs to R5. Removal of the route map clears the issue, and BGP routes are exchanged correctly between autonomous systems, leading to the required BGP tables within Sw1 and Sw2. If you have solved this incident successfully, as detailed in [Example 2-4](#), you have scored 2 points.

Example 2-4 *BGP Testing and Reconfiguration*

[Click here to view code image](#)

```
Sw1# show ip bgp
```

```
BGP table version is 4, local router ID is 120.100.12.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,  
r RIB-failure, S Stale, m multipath, b backup-path, x best-external,  
f RT-Filter
```

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

```
Network Next Hop Metric LocPrf Weight Path
```

```
*> 120.100.11.0/24 0.0.0.0 0 32768 i
```

```
Sw2# show ip bgp
```

```
BGP table version is 2, local router ID is 120.100.27.2
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,  
r RIB-failure, S Stale, m multipath, b backup-path, x best-external,  
f RT-Filter
```

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

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 120.100.20.0/24	0.0.0.0	0	32768	i	

```
R4# show ip bgp
```

```
BGP table version is 2, 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, m multipath, b backup-path, x best-external
```

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

Network	Next Hop	Metric	LocPrf	Weight	Path
r>i120.100.11.0/24	120.100.14.1	0	100	0	i

```
R5# show ip bgp
```

```
BGP table version is 7, local router ID is 5.5.5.5
```

```
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
r>i120.100.20.0/24	120.100.27.2	0	100	0	i

R4# **show ip bgp summary**

BGP router identifier 4.4.4.4, local AS number 63513

BGP table version is 2, main routing table version 2

1 network entries using 136 bytes of memory

1 path entries using 52 bytes of memory

1/1 BGP path/bestpath attribute entries using 124 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory

BGP using 312 total bytes of memory

BGP activity 4/3 prefixes, 4/3 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/
									PfxRcd
5.5.5.5	4	64513	0	0	1	0	0	00:31:15	Idle
120.100.14.1	4	63513	39	37	2	0	0	00:31:14	1

R5# **show ip bgp summary**

BGP router identifier 5.5.5.5, local AS number 64513

BGP table version is 7, main routing table version 7

1 network entries using 120 bytes of memory

1 path entries using 52 bytes of memory

2/1 BGP path/bestpath attribute entries using 248 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 1 (at peak 2) using 28 bytes of memory
BGP using 448 total bytes of memory
BGP activity 3/2 prefixes, 3/2 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
4.4.4.4	4	63513	44044	40024	0	0	0	00:31:55	Idle
120.100.27.2	4	64513	50	45	7	0	0	00:41:43	1

R4# **show run | sect bgp**

router bgp 63513

bgp log-neighbor-changes

neighbor 5.5.5.5 remote-as 64513

neighbor 5.5.5.5 update-source Loopback0

neighbor 5.5.5.5 route-map med out

neighbor 120.100.14.1 remote-as 63513

neighbor 120.100.14.1 update-source Loopback0

no auto-summary

R5# **show run | sect bgp**

router bgp 64513

no synchronization

```
bgp log-neighbor-changes
neighbor 4.4.4.4 remote-as 63513
neighbor 4.4.4.4 update-source Loopback0
neighbor 120.100.27.2 remote-as 64513
neighbor 120.100.27.2 update-source Loopback0
no auto-summary
```

```
R4(config)# router bgp 63513
```

```
R4(config-router)# neighbor 5.5.5.5 ebgp-multihop 2
```

```
R5(config)# router bgp 64513
```

```
R5(config-router)# neighbor 4.4.4.4 ebgp-multihop 2
```

```
R4# show ip bgp summary | include 5.5.5.5
```

```
5.5.5.5      4      64513    10    10     3    0    0 00:05:33    1
```

```
R4# show ip bgp 120.100.20.0
```

```
BGP routing table entry for 120.100.20.0/24, version 3
```

```
Paths: (1 available, best #1, table default)
```

```
Advertised to update-groups:
```

```
7
```

```
64513
```

```
5.5.5.5 (metric 2) from 5.5.5.5 (5.5.5.5)
```

```
Origin IGP, localpref 100, valid, external, best
```

```
R5# show ip bgp summary | include 4.4.4.4
4.4.4.4    4 63513 44054 40034    7 0 0 00:05:12    0
```

```
R5# show ip bgp 120.100.11.0
```

```
% Network not in table
```

```
R4# show ip bgp neighbors 5.5.5.5 advertised-routes
```

```
BGP table version is 3, 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, m multipath, b backup-path, x best-external
```

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

Network	Next Hop	Metric	LocPrf	Weight	Path
r>i120.100.11.0/24	120.100.14.1	0	100	0	i

```
Total number of prefixes 1
```

```
R5# debug ip bgp 4.4.4.4 updates
```

```
R4# clear ip bgp *
```

```
R5#
```

```
*Feb 9 21:33:18: %BGP-5-ADJCHANGE: neighbor 4.4.4.4 Down Peer closed the session
```

```
*Feb 9 21:33:19: %BGP-5-ADJCHANGE: neighbor 4.4.4.4 Up
```

```
*Feb 9 21:33:19: BGP(0): 4.4.4.4 send UPDATE (format) 120.100.20.0/24, next  
5.5.5.5, metric 0, path Local
```

```
R5#
```

```
*Feb 9 21:33:20: BGP: 4.4.4.4 Next hop is our own address 120.100.45.5
```

```
*Feb 9 21:33:20: BGP(0): 4.4.4.4 rcv UPDATE w/ attr: nexthop 120.100.45.5,  
origin i, originator 0.0.0.0, path 63513, community , extended community
```

```
*Feb 9 21:33:20: BGP(0): 4.4.4.4 rcv UPDATE about 120.100.11.0/24 -- DENIED due  
to: NEXTHOP is our own address;
```

```
R4# show run | sect bgp
```

```
router bgp 63513
```

```
bgp log-neighbor-changes
```

```
neighbor 5.5.5.5 remote-as 64513
```

```
neighbor 5.5.5.5 update-source Loopback0
```

```
neighbor 5.5.5.5 route-map med out
```

```
neighbor 120.100.14.1 remote-as 63513
```

```
neighbor 120.100.14.1 update-source Loopback0
```

```
no auto-summary
```

```
R4# show run | section route-map med
```

```
neighbor 5.5.5.5 route-map med out
```

```
route-map med permit 10
```

```
set ip next-hop 120.100.45.5
```

```
R4(config)# router bgp 63513
```

```
R4(config-router)# no neighbor 5.5.5.5 route-map med out
```

```
R4(config-router)# do clear ip bgp *
```

```
R5# show ip bgp
```

BGP table version is 8, local router ID is 5.5.5.5

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
*> 120.100.11.0/24	4.4.4.4		0	63513	i
r>i120.100.20.0/24	120.100.27.2		0	100	o i

```
Sw2# show ip bgp
```

BGP table version is 3, local router ID is 120.100.27.2

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, x best-external,
f RT-Filter

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

```
r>i120.100.11.0/24 4.4.4.4          0 100  0 63513 i
*> 120.100.20.0/24 0.0.0.0        0    32768 i
```

```
Sw1# show ip bgp
```

```
BGP table version is 9, local router ID is 120.100.12.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, x best-external,
               f RT-Filter
```

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

```
   Network      Next Hop      Metric LocPrf Weight Path
*> 120.100.11.0/24 0.0.0.0        0    32768 i
*>i120.100.20.0/24 5.5.5.5        0 100  0 64513 i
```

Incident 5

- Sw3 is not receiving routes from its OSPF neighbor R7.
- Fix the problem so that Sw2 has the following routing table entry and connectivity:

[Click here to view code image](#)

```
Sw3# show ip route 120.100.57.0
```

```
Routing entry for 120.100.57.0/24
```

```
Known via "ospf 1", distance 110, metric 2, type intra area
```

<https://t.me/learningnets>

Last update from 120.100.37.7 on Vlan37, 01:54:32 ago

Routing Descriptor Blocks:

* 120.100.37.7, from 7.7.7.7, 01:54:32 ago, via Vlan37

Route metric is 2, traffic share count is 1

```
Sw3# ping 120.100.57.5
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.57.5, timeout is 2 seconds:

!!!!

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

-
- You may not modify any parameter under the OSPF process of Sw3.

1 point

Begin your troubleshooting by confirming that Sw3 has an OSPF neighbor relationship with R7. Example 2-5 shows that there is a valid neighbor relationship and that the route to network 120.100.57.0/24 is not present in the routing table. The route could be missing due to not being advertised from R7, and it is worth checking the OSPF configuration on R7. However, the incident instructions state that you cannot modify any parameter under the OSPF process of Sw3, so this is likely to be where the problem is located.

Inspection of the OSPF configuration of Sw3 shows that there is a distribute list in operation with a prefix list of “ALLOW” for route updates and a prefix list of “allow” for the gateway (any neighbor). You may suspect the route update prefix list, but it simply allows any update (0.0.0.0/0 le 32). The gateway prefix list, however, is incorrect, and it limits the OSPF process to updates from 120.100.37.3 (Sw3 itself), so this prefix list simply

<https://t.me/learningnets>

needs changing to that of R7 (120.100.37.7). If you have solved this incident successfully, as detailed in [Example 2-5](#), you have scored 1 point.

Example 2-5 *Sw3 OSPF Testing and Configuration*

[Click here to view code image](#)

```
Sw3# show ip ospf neighbor | include Vlan37
7.7.7.7      1 FULL/DR      00:00:34 120.100.37.7 Vlan37
```

```
Sw3# sh ip route ospf | include 120.100.57.0
! The neighbor relationship is up to R7 but the route 120.100.57.0/24 is not
present in the routing table so either not sent or not being received correctl
```

```
Sw3# show run | section ospf 1
ip ospf 1 area 2
router ospf 1
network 120.100.33.0 0.0.0.255 area 2
network 120.100.37.0 0.0.0.255 area 2
distribute-list prefix ALLOW gateway allow in
```

```
! The configuration reveals there is an inbound distribute list with prefix-li
"ALLOW" and a gateway prefix-list of "allow".
```

```
Sw3# show run | include prefix-list
ip prefix-list ALLOW seq 5 permit 0.0.0.0/0 le 32
ip prefix-list allow seq 5 permit 120.100.37.7/32
ip prefix-list allow seq 10 deny 0.0.0.0/0 le 32
```

! Prefix-list "ALLOW" will let all routes into Sw3 but prefix-list "allow" will only allow host 120.100.37.3 to peer with Sw3, this is in fact Sw3's own IP address

```
Sw3(config)# no ip prefix-list allow seq 5 permit 120.100.37.3/32
Sw3(config)# ip prefix-list allow seq 5 permit 120.100.37.7/32
Sw3(config)# ^Z
```

```
Sw3# show ip route ospf | include 120.100.57.0
O    120.100.57.0/24 [110/2] via 120.100.37.7, 02:05:49, Vlan37
```

```
Sw3# ping 120.100.57.5
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.57.5, timeout is 2 seconds:

!!!!

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

Incident 6

- A GRE tunnel configured between SW2 and SW3 is flapping. OSPF is configured to run between the two routers.
- Fix the problem so that SW2 and SW3 have the following reliable OSPF adjacency with each other:
<https://www.cisco.com/learning>

[Click here to view code image](#)

```
Sw2# sh ip ospf neighbor | include Tunnel0
120.100.33.1  0 FULL/ -   00:00:32  10.1.1.3   Tunnel0
```

```
Sw3# sh ip ospf neigh | include Tunnel0
120.100.27.2  0 FULL/ -   00:00:33  10.1.1.2   Tunnel0
```

1 point

[Example 2-6](#) shows that both switches are indeed suffering from a flapping tunnel interface. Switch2 shows that the issue is due to recursive routing. (Remember that some switches will need the command **logging monitor** to display logging output.) This is when the tunnel endpoints are discovered over the tunnel itself, leading to a routing loop until the tunnel is disabled.

[Example 2-6](#) shows that when debugging the routing table the tunnel endpoints are indeed being learned via the tunnel itself rather than via R7. Inspection of the configurations on the VLAN27 interface on Sw2 reveals that a higher OSPF cost of 10000 has been manually defined and hence the tunnel interface appears to be a better option for OSPF path selection because the tunnel cost defaults to 1000. After the OSPF cost has been removed from VLAN27 on Sw2, the neighbors no longer flap because the recursive routing issue has been completely removed. If you have solved this incident successfully, as detailed in [Example 2-6](#), you have scored 1 point.

Example 2-6 *Sw2 and Sw3 Tunnel Testing and Configuration*

[Click here to view code image](#)

```
Sw3# %OSPF-5-ADJCHG: Process 1, Nbr 120.100.27.2 on Tunnel0 from FULL to DOWN,  
Neighbor Down: Dead timer expired
```

```
Sw2#
```

```
Sw2# %OSPF-5-ADJCHG: Process 1, Nbr 120.100.33.1 on Tunnel0 from LOADING to FULL,  
Loading Done
```

```
Sw2# %OSPF-5-ADJCHG: Process 1, Nbr 120.100.33.1 on Tunnel0 from FULL to DOWN,  
Neighbor Down: Dead timer expired
```

```
Sw2#
```

```
%ADJ-5-PARENT: Midchain parent maintenance for IP midchain out of Tunnel0 - looped  
chain attempting to stack
```

```
%TUN-5-RECURDOWN: Tunnel0 temporarily disabled due to recursive routing
```

```
%OSPF-5-ADJCHG: Process 1, Nbr 120.100.33.1 on Tunnel0 from FULL to DOWN, Neighbor  
Down: Interface down or detached
```

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed state to down
```

```
Sw2# debug ip routing
```

```
Sw2# logging monitor
```

```
Sw2# : RT: closer admin distance for 120.100.123.0, flushing 1 routes
```

```
RT: add 120.100.123.0/24 via 10.1.1.3, ospf metric [110/6003]
```

```
1w5d: is_up: Tunnel0 state: 4 sub state: 1 line: 0
1w5d: RT: interface Tunnel0 removed from routing table
1w5d: RT: del 10.1.1.0 via 0.0.0.0, connected metric [0/0]
1w5d: RT: delete subnet route to 10.1.1.0/24
1w5d: RT: del 10.1.1.2 via 0.0.0.0, connected metric [0/0]
1w5d: RT: delete subnet route to 10.1.1.2/32
Sw2#
```

```
Sw2# show run int vlan27
interface Vlan27
ip address 120.100.27.2 255.255.255.0
ip ospf cost 10000
end
```

```
Sw3# show run int vlan37
interface Vlan37
ip address 120.100.37.3 255.255.255.0
end
```

```
Sw3(config)# int vlan27
Sw3(config-if)# no ip ospf cost 10000
```

```
Sw2# show ip ospf neigh | include Tunnel0
```

```
120.100.33.1 0 FULL/ - 00:00:32 10.1.1.3 Tunnel0
```

```
Sw2#
```

```
Sw3# show ip ospf neigh | include Tunnel0
```

```
120.100.27.2 0 FULL/ - 00:00:38 10.1.1.2 Tunnel0
```

Incident 7

- Sw3 cannot communicate with networks in the EIGRP domain.
- Fix the problem so that Sw3 has the following OSPF route and connectivity:

[Click here to view code image](#)

```
Sw3# show ip route | include 120.100.11.0
```

```
O E1 120.100.11.0/24 [110/5003] via 120.100.37.7, 00:01:48, Vlan37
```

```
Sw3# ping 120.100.11.1
```

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

```
Sending 5, 100-byte ICMP Echos to 120.100.11.1, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/17 ms
```

2 points

Usually, the first thing to do in this scenario is to see which routes (if any) are being received by Sw3 from its
<https://t.me/learningnets>

OSPF neighbor within Area 2 (R7). However, the previous question, if answered correctly, would have shown that routes are being learned from R7.

Example 2-7 shows that OSPF routes are being received from R7 and that routes from Area 0 and Area 2 are present but no external routes from the EIGRP domain are present. This could mean that there is a neighbor connectivity, redistribution, or OSPF configuration issue on the ASBR router connecting to the EIGRP domain (R4). It is therefore prudent to check the configuration and connectivity onward within the EIGRP domain on R4.

Example 2-7 shows that R4 is receiving EIGRP routes correctly. Again, if you successfully answered an earlier incident, you could be confident that redistribution was working correctly. Example 2-7 shows that ABR R7 is receiving the redistributed routes correctly. Because no apparent connectivity issue exists between devices, it is time to examine the OSPF database to look for clues.

Example 2-7 details the output from the database for routes originating from EIGRP on Sw3. As shown, the routes show the same router ID on the originating and receiving router. Duplicate router IDs cause routes not to be installed into the routing table because it appears to the routing process that there must be a loop and hence the routes are simply dropped. Example 2-7 shows reconfiguration of the router ID on Sw3 to an IP address of one of its Loopback1 interface to solve the problem. The same result could have been achieved by manually assigning a router ID on R4. If you have solved this incident successfully, as detailed in Example 2-7, you have scored 2 points.

Example 2-7 *Redistribution Testing and Configuration*

[Click here to view code image](#)

```
Sw3# show ip route ospf
```

```
4.0.0.0/24 is subnetted, 1 subnets
```

```
O IA 4.4.4.0 [110/4] via 120.100.37.7, 00:16:29, Vlan37
```

```
5.0.0.0/32 is subnetted, 1 subnets
```

```
O IA 5.5.5.5 [110/3] via 120.100.37.7, 00:16:29, Vlan37
```

```
6.0.0.0/32 is subnetted, 1 subnets
```

```
O IA 6.6.6.6 [110/3] via 120.100.37.7, 00:16:29, Vlan37
```

```
120.0.0.0/8 is variably subnetted, 14 subnets, 2 masks
```

```
O IA 120.100.2.0/24 [110/1004] via 120.100.37.7, 00:16:29, Vlan37
```

```
O 120.100.27.0/24 [110/2] via 120.100.37.7, 00:16:29, Vlan37
```

```
O IA 120.100.40.0/24 [110/1003] via 120.100.37.7, 00:16:29, Vlan37
```

```
O IA 120.100.45.0/24 [110/3] via 120.100.37.7, 00:16:29, Vlan37
```

```
O IA 120.100.46.0/24 [110/3] via 120.100.37.7, 00:16:29, Vlan37
```

```
O 120.100.57.0/24 [110/2] via 120.100.37.7, 00:16:29, Vlan37
```

```
O 120.100.67.0/24 [110/2] via 120.100.37.7, 00:16:29, Vlan37
```

! Area 0 & Area 2 routes are being received from R7 on Sw3 but no external routes (EIGRP).

```
R4# show ip route eigrp
```

```
120.0.0.0/8 is variably subnetted, 20 subnets, 2 masks
```

```
D 120.100.11.0/24
```

[90/130816] via 120.100.14.1, 2dooh, GigabitEthernet0/0.14

D 120.100.123.0/24

[90/28416] via 120.100.4.1, 2dooh, GigabitEthernet0/0.4

! ASBR R4 is learning EIGRP routes correctly.

R7# **show ip route | include O E1**

O E1 120.100.11.0/24

O E1 120.100.123.0/24

! EIGRP routes redistributed into OSPF by ASBR R4 are being learned on ABR R7.

Sw3# **show ip ospf data | include 120.100.11.0**

120.100.11.0 120.100.99.1 3603 0x800001CC 0x001BA3 0

Sw3# **show ip ospf | include ID**

Routing Process "ospf 1" with ID 120.100.99.1

! The external routes are present in the OSPF database but it should be noted that the router-id of the advertising router is 120.100.99.1 which is the same router-id as Sw3.

```
Sw3(config)# router ospf 1
```

```
Sw3(config-router)# router-id 120.100.33.3
```

Reload or use "clear ip ospf process" command, for this to take effect

```
Sw3(config-router)# ^Z
```

```
Sw3# clear ip ospf process
```

```
Reset ALL OSPF processes? [no]: y
```

! Re-assigning the ospf router-id on Sw3 to an IP address of it's Loopback interfaces will cure the duplicate router-id issue.

```
Sw3# show ip route | include 120.100.11.0
```

```
O E1 120.100.11.0/24 [110/5003] via 120.100.37.7, 00:01:48, Vlan37
```

```
Sw3# ping 120.100.11.1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 120.100.11.1, timeout is 2 seconds:

```
!!!!
```

Success rate is 100 percent (5/5), round-trip min/avg/max = 1/7/9 ms

Incident 8

- Sw3 cannot telnet to Sw2 using IPv6 on VLAN20. The network administrator of Sw2 requires that ICMP ping testing should not be allowed to Sw3 but that Telnet access is allowed.

- Fix the problem so that Sw3 has the following response to an ICMP ping and a Telnet session to Sw2:

<https://www.cisco.com/learning>

[Click here to view code image](#)

```
Sw3# ping ipv6 CC1E:CAFE::1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to CC1E:CAFE::1, timeout is 2 seconds:

```
AAAAA
```

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

```
Sw3# telnet CC1E:CAFE::1
```

```
Trying CC1E:CAFE::1 ... Open
```

- Do not modify the IPv6 address assignment on Sw3.
- The network administrator has chosen to suppress RAs on Sw2, so do not remove this suppression.

2 points

You can clearly tell that this issue will be related to an access list because specific protocols are permitted and others denied within the incident description. The incident also tells you that you are not to modify the IPv6 address on Sw3, so it would be a good starting point to check these areas.

Example 2-8 shows the configuration of VLAN20 on both switches and the related IPv6 access list applied to the VLAN on Sw2. From the configuration and incident information, you can tell that Sw3 has been configured to receive its IPv6 prefix via autoconfiguration and that Sw2 has IPv6 neighbor discovery suppressed by the command **ipv6 nd ra suppress**, which will stop router advertisements.

<https://t.me/learningnets>

As detailed in [Example 2-8](#), Sw3 does not have a global IPv6 address on VLAN20, so the access list on Sw2 requires modification to allow IPv6 address assignment via ICMP solicitation, which would not be forthcoming due to the RA suppression in place. Because the access list is denying ICMP, you must modify it to allow the required ICMP ports for address assignment. If you did not know off the top of your head that ports 133, 135, and 136 are required for router solicitation, neighbor solicitation, and neighbor advertisement, you can simply create a new IPv6 ACL on Sw2 to deny all IPv6 traffic and log the hits.

[Example 2-8](#) shows the access list TEST applied to the VLAN20 interface of Sw2 and the resulting logged traffic of ICMP ports 133, 135, and 136 from Sw3. Once the original ACL is modified to permit these ICMP ports, [Example 2-8](#) shows an IPv6 address is now correctly assigned to Sw3 VLAN20 interface and the testing is successful, with ICMP ping blocked and Telnet permitted. If you have solved this incident successfully, as detailed in [Example 2-8](#), you have scored 2 points.

Example 2-8 *IPv6 Address Testing and Configuration*

[Click here to view code image](#)

```
Sw2# sh run int vlan 20
interface Vlan20
 ip address 120.100.20.2 255.255.255.0
 ipv6 address CC1E:CAFE::1/64
 ipv6 enable
 ipv6 nd ra suppress
 ipv6 traffic-filter ALLOW-TELNET in
```

```
Sw2# show ipv6 access-list
```

```
IPv6 access list ALLOW-TELNET
```

```
deny icmp any any sequence 10
```

```
permit ipv6 any any sequence 20
```

```
Sw3# sh run int vlan 20
```

```
interface Vlan20
```

```
ip address 120.100.20.3 255.255.255.0
```

```
ipv6 address autoconfig
```

```
ipv6 enable
```

```
end
```

```
Sw3# show ipv6 int brief vlan 20
```

```
Vlan20 [up/up]
```

```
FE80::215:C6FF:FEF2:ABE0
```

```
Sw2(config)# ipv6 access-list TEST
```

```
Sw2(config-ipv6-acl)# deny ipv6 any any log
```

```
Sw2(config-ipv6-acl)# int vlan 20
```

```
Sw2(config-if)# ipv6 traffic-filter TEST in
```

```
FE80::215:C6FF:FEF2:ABE0 -> FF02::1 (136/o), 1 packet
3do2h: %IPV6_ACL-6-ACCESSLOGDP: list TEST/10 denied icmpv6
FE80::215:C6FF:FEF2:ABE0 -> FF02::2 (133/o), 1 packet
3do2h: %IPV6_ACL-6-ACCESSLOGDP: list TEST/10 denied icmpv6
FE80::215:C6FF:FEF2:ABE0 -> FE80::21B:2AFF:FEFE:49C1 (135/o), 1 packet
```

```
Sw2(config)# ipv6 access-list ALLOW-TELNET
Sw2(config-ipv6-acl)# sequence 5 permit icmp any any 133
Sw2(config-ipv6-acl)# sequence 6 permit icmp any any 135
Sw2(config-ipv6-acl)# sequence 7 permit icmp any any 136
Sw2(config-ipv6-acl)# exit
Sw2(config)# int vlan 20
Sw2(config-if)# ipv6 traffic-filter ALLOW-TELNET in
```

```
Sw3# show ipv6 int brief vlan 20
```

```
Vlan20    [up/up]
FE80::215:C6FF:FEF2:ABE0
CC1E:CAFE::215:C6FF:FEF2:ABE0
```

```
Sw3# ping ipv6 CC1E:CAFE::1
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to CC1E:CAFE::1, timeout is 2 seconds:

Success rate is 0 percent (0/5)

```
Sw3# telnet CC1E:CAFE::1
```

```
Trying CC1E:CAFE::1 ... Open
```

Incident 9

- Request a colleague to reconfigure a preconfigured EIGRP process of JAKE running on Sw3 VLAN37 that connects to R7 on network 120.100.37.0/24. Have your colleague change the autonomous system number on Sw3 from 11 to a number of his choice.
- Use the embedded packet-capture facility on R7 to inspect the EIGRP traffic received from Sw3 on network 120.100.37.0/24 to determine which autonomous system Sw3 has been changed to and modify the autonomous system on R7 to match to form a successful neighbor relationship between R7 and Sw3. (Should a question similar to this arise on the real lab exam, it would, of course, relate to an unknown parameter on a device outside of your administrative control.)

3 points

Usually, this kind of question relates to a device on the backbone outside of your administrative control, so here you need a colleague or an understanding partner (good luck with this) to modify the EIGRP process running on Sw3 and replace AS11 with an autonomous system number of his choice.

Example 2-9 shows the initial verification of neighbor relationship between Sw3 and R7 within AS11. The example then details how an ACL is created to detail the EIGRP traffic specifically received from Sw3 within the embedded packet capture that is required to be run on R7. ACL 100 is applied to the capture buffer created of EIGRPAS, which is associated with the network connecting to Sw3 on interface Go/0.37. The capture point <https://t.me/learningnets>

SWITCH3 is then associated with the capture buffer EIGRPAS and run for a short period of time before being stopped and output to the command line with the command **show monitor capture buffer EIGRPAS dump**.

At this point, the information is fairly meaningless because the intention is usually for the packet capture to be sent to a program for decoding, but you can be certain that the autonomous system field is lurking within the capture, and it is just a case of locating it. As with the majority of problem determination issues, looking for patterns will usually provide you with sufficient clues to rectify a problem, so a good idea here is to compare the packet capture with a known autonomous system packet capture for comparison.

Example 2-9 shows a new EIGRP process being applied to Sw2 and enabled on network 120.100.27.0/24 toward R7 with an autonomous system of 255. The packet capture is enabled once more with new capture buffers and a capture point, but applied to interface G0/0.27 on R7, which connects to network 120.100.27.0/24. Once viewed, it is simply a case of looking for a field that represents 255 (or FF in hex, as the capture is clearly output into hex). Inspection of the packet capture received from Sw2 clearly shows the FF value preceded by a string of 0s. Comparison to the original packet capture from Sw3 shows a hex value of 0F (decimal 15 within the same location). Example 2-9 shows that by changing the autonomous system number on R7 to 15, a successful neighbor relationship is created between R7 and Sw3. If you had not thought to run a comparison between a known autonomous system number, it is highly unlikely that you would have successfully solved this issue. If you did, though, congratulations; you have scored 3 well-earned points.

Example 2-9 EIGRP Autonomous System Testing and Configuration

[Click here to view code image](#)

https://www.cisco.com/itd/learning/jsp/showingipets/eigrp_neighbors

EIGRP-IPv4 VR(JAKE) Address-Family Neighbors for AS(11)

H	Address	Interface	Hold	Uptime	SRTT	RTO	Q	Seq
		(sec)	(ms)	Cnt	Num			
0	120.100.37.7	Vl37	12	00:00:11	7	100	0	3

! The neighbor relationship is up and working between Sw3 & R7 prior to the AS change on Sw3.

```
R7(config)# access-list 100 permit eigrp host 120.100.37.3 any
```

```
R7(config)# exit
```

```
R7# monitor capture buffer EIGRPAS
```

```
R7# monitor capture buffer EIGRPAS filter access-list 100
```

```
R7# monitor capture point ip cef SWITCH3 gigabitEthernet 0/0.37 both
```

```
R7# monitor capture point associate SWITCH3 EIGRPAS
```

```
R7# monitor capture point start SWITCH3
```

```
R7# monitor capture point stop SWITCH3
```

```
R7# show monitor capture buffer EIGRPAS dump
```

```
19:47:38.450 GMT Feb 15 2014 : IPv4 LES CEF : Gio/0.37 None
```

```
7060D430: 01005E00 000A001B 2AF05F42 8100C025 ..^....*p_B..@%
```

```
7060D440: 080045C0 003C0000 00000158 3B397864 ..E@.<....X;9xd
```

```
7060D450: 2503E000 000A0205 F1C30000 00000000 %.`.....qC.....
```

```
7060D460: 00000000 00000000 000F0001 000C0100 .....
```

7060D470: 01000000 00

19:47:42.898 GMT Feb 15 2014 : IPv4 LES CEF : Gio/o.37 None

7060D430: 01005E00 000A001B 2AF05F42 8100C025 ..^.....*p_B..@%

7060D440: 080045C0 003C0000 00000158 3B397864 ..E@.<.....X;9xd

7060D450: 2503E000 000A0205 F1C30000 00000000 %.`.....qC.....

7060D460: 00000000 00000000 000F0001 000C0100

7060D470: 01000000 00

!Configure a further instance of EIGRP with an AS of 255 on Sw2 to compare the packet capture as received at R7 on network 120.100.27.0/24.

```
Sw2(config)# router eigrp JAKE
```

```
Sw2(config-router)# address-family ipv4 autonomous-system 255
```

```
Sw2(config-router-af)# network 120.100.27.0 0.0.0.255
```

```
R7(config)# access-list 101 permit eigrp host 120.100.27.2 any
```

```
R7(config)# exit
```

```
R7# monitor capture buffer EIGRPAS filter access-list 101
```

```
R7# monitor capture point ip cef SWITCH2 gigabitEthernet 0/0.27 both
```

```
R7# monitor capture point associate SWITCH2 EIGRPAS
```

```
R7# monitor capture point start SWITCH2
```

```
R7# monitor capture point stop SWITCH2
```

```
R7# show monitor capture buffer EIGRPAS dump
```

```
01:56:40.948 GMT Feb 16 2014 : IPv4 LES CEF : Gio/o.27 None
```

```
6917D670:          01005E00      ..^.
```

```
6917D680: 000A001B 2AB59BC3 8100C01B 080045C0 ....*5.C..@...E@
```

```
6917D690: 003C0000 00000158 453A7864 1B02E000 .<....XE:xd..`.
```

```
6917D6A0: 000A0205 FoD30000 00000000 00000000 ....pS.....
```

```
6917D6B0: 00000000 00FF0001 000C0100 01000000 .....
```

```
6917D6C0: 00      .
```

```
01:56:45.688 GMT Feb 16 2014 : IPv4 LES CEF : Gio/o.27 None
```

```
6917D670:          01005E00      ..^.
```

```
6917D680: 000A001B 2AB59BC3 8100C01B 080045C0 ....*5.C..@...E@
```

```
6917D690: 003C0000 00000158 453A7864 1B02E000 .<....XE:xd..`.
```

```
6917D6A0: 000A0205 FoD30000 00000000 00000000 ....pS.....
```

```
6917D6B0: 00000000 00FF0001 000C0100 01000000 .....
```

```
6917D6C0: 00
```

```
Sw2(config)# no router eigrp JAKE
```

dump in the same location to view the AS received from Sw3 of HEX 0F which is mal 15.

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

```
R7(config-router)# no address-family ipv4 autonomous-system 11
```

```
R7(config-router)# address-family ipv4 autonomous-system 15
```

```
R7(config-router-af)# network 120.100.37.0 0.0.0.255
```

```
*Feb 15 19:49:51: %DUAL-5-NBRCHANGE: EIGRP-IPv4 15: Neighbor 120.100.37.3 (GigabitEthernet0/0.
```

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 14 points from a potential 17, well done. If you accomplished this within the time frame of 2 hours or less, you are well on your way to being prepared for tickets that you are likely to face during the Troubleshooting section of the real exam.

The questions might have seemed a little vague or, in fact, misleading, but they closely resemble real-world scenarios in which you are given only limited information and have to work out what is really happening yourself to get to the root cause. You might, of course, not have the constraints listed in the real world, but this just ensures that you provide the correct solution to the issues.

Summary

ARE YOU READY?

This well-known Cisco slogan is closely identified with the Internet revolution. By the end of these practice exams, you should have a good idea of whether you are ready. Did you feel confident working through the questions, or was it a complete shock to the system? Are you more used to being spoon fed solitary scenarios than actually having to analyze questions and piece together parts of a complex network jigsaw?

Life is full of challenges. During your education and career, the CCIE certification is as tough as it gets. The exam is designed to test your technical skills, your understanding and analysis of complex topologies, and your capacity to build and troubleshoot a network with IP routing protocols and features. You need to achieve a minimum score of 80 percent to pass.

FURTHER READING

The following Cisco Press titles are on topics appearing on the CCIE exam blueprint. These books are not required study resources, but you can use them to build knowledge in certain areas.

<http://www.ciscopress.com/ccie>

CCIE Routing and Switching v5.0 Official Cert Guide, Volume 1

CCIE Routing and Switching v5.0 Official Cert Guide, Volume 2

CCIE Routing and Switching v5.0 Official Cert Guide Library

CCIE Routing and Switching v5.0 Configuration Practice Labs

Routing TCP/IP, Volume I

<https://t.me/learningnets>

Routing TCP/IP, Volume II (CCIE Professional Development series)

Troubleshooting IP Routing Protocols

Inside Cisco IOS Software Architecture (CCIE Professional Development Series)

Cisco LAN Switching (CCIE Professional Development series)

Cisco OSPF Command and Configuration Handbook (paperback)

Cisco BGP-4 Command and Configuration Handbook

Cisco Field Manual: Router Configuration

Cisco Field Manual: Catalyst Switch Configuration

Developing IP Multicast Networks, Volume I

Internet Routing Architectures

MPLS and VPN Architectures

MPLS and VPN Architectures, Volume II

Cisco Catalyst QoS: Quality of Service in Campus Networks (paperback)

End-to-End QoS Network Design: Quality of Service in LANs, WANs, and VPNs

Deploying IPv6 Networks

Network Security Technologies and Solutions (CCIE Professional Development series)

HELP AND ADVICE

■ Look at http://www.cisco.com/web/learning/certifications/expert/ccie_rs/index.html for the latest information about the CCIE lab exam.

■ Keep your schedule flexible during your rack time. Include time for breaks and relaxation; you will often find <https://timelearning.net>

that 5 minutes away from the keyboard can help you consider possible solutions. Most importantly, do not forget the people you care for; make time for them, too.

- Build your study plan based on a balance between theory and practice. You need to understand the concepts through the theory, and then consolidate this during your rack time.
- Begin with simple topics in isolation, and then work up to complex lab scenarios. Spend as much time repeating your configurations as possible to improve your speed and ability to perform basic configurations with your eyes shut. This practice will save you time for when you need it during the exam.
- Explore the Cisco master indexes documentation at [Cisco.com](https://www.cisco.com).
- Start to plan for your exam at least 6 months before the lab date.
- If you find these practice labs have highlighted weak areas, do not be afraid to postpone your lab date.

HOW CAN I SCHEDULE MY CCIE LAB EXAM?

Go to http://www.cisco.com/web/learning/le3/ccie/rs/lab_exam.html to find all the information you need for how to schedule your exam, including locations, start times, and more. You must have a CCO user ID, your CCIE written exam date, and score to be able to view your profile and schedule your exam.

THE DAY BEFORE

If you are traveling to take your exam, try to arrive the day before to familiarize yourself with the area. Take a tour to the lab location so that you will not be late on the day; the last thing you need is to arrive flustered. The day before is a day to relax, not a day to stay up late studying. Most importantly, save the beer until after the exam; pass or fail, you will feel like one or two, for sure. The CCIE exam might be the reason why Stella Artois is so popular in Brussels!

THE DAY OF THE EXAM

On the day of the exam, plan to arrive for registration at least 15 minutes before the exam begins. The proctor will walk you to the lab and give you a briefing before the exam starts, telling you about the lab and the general guidelines for the day.

The proctor will not discuss solutions or possible solutions for a given question with you. The proctor will be available to help you understand the wording or meaning of the questions and to make sure that the equipment or virtual topology is working correctly so that your exam runs smoothly. Ask the proctor for any assistance or verification if you encounter any issues. To get the “bigger picture,” read the entire exam before you start, ensuring that you fully understand each question and its requirements. Begin by performing easier tasks, leaving the most difficult for later. Take some small breaks during the morning and the afternoon to refresh yourself and relieve the stress.

PASS OR FAIL, WHAT NEXT?

If you pass, you certainly have something to celebrate; you have just joined an elite club that will, without doubt, enhance your career. You have achieved the highest level of certification in the networking world! You should now aim to continue this quest for knowledge that sets you apart from your peers, but first take a break before starting your next CCIE track.

If you failed, don't worry, and don't take it personally; most people fail on their first attempt. Consider it experience, and get back on the keyboard as soon as you can to determine what went wrong. You will more than likely succeed the next time and will ultimately become a better engineer for your extra rack time.

I hope these practice exams and tips prove helpful and guide you to success on your exam.

Where are the Companion Content Files?

Register this digital version of **Cisco CCIE Routing and Switching v5.0 Troubleshooting Practice Labs, Second Edition** to access important downloads.

Register this eBook to unlock the companion configuration files for each practice lab. Follow these steps:

1. Go to www.ciscopress.com/register and log in or create a new account.
2. Enter the ISBN: **9780133786330**
3. Answer the challenge question as proof of purchase.
4. Click on the “Access Bonus Content” link in the Registered Products section of your account page, to be taken to the page where your downloadable content is available.

The Professional and Personal Technology Brands of Pearson



Cisco Press

IBM
Press

InformIT

PEARSON IT Certification



QUE

SAMS

vmware PRESS