Ace the Exam Series®

FIRST EDITION

2022

CCNP ENARSI

IMPLEMENTING CISCO ENTERPRISE ADVANCED ROUTING & SERVICES EXAM: 300-410

EXAM CRAM NOTES

IPSpecialist

CCNP ENARSI: Implementing Cisco Enterprise Advanced Routing and Services

Exam: 300-410

Exam Cram Notes

First Edition

CHAPTER 01: INTRODUCTION

Introduction

By obtaining a Cisco CCNP Enterprise (ENARSI 300-410) certification, you can ensure that you have a firm grasp of Cisco's device design and configuration and common industry protocols. With a large global presence, Cisco has a significant market share of routers and switches.

You can acquire the skills required to install, administer, operate, and troubleshoot a business network by passing the Implementing Cisco Enterprise Advanced Routing and Services (ENARSI) exam. The advanced routing and infrastructure technologies covered in this course go beyond the subjects covered in the Implementing and Operating Cisco Enterprise Network Core Technologies (ENCOR) course. This course covers the Implementing Cisco® Enterprise Advanced Routing and Services (ENARSI) exam and the new CCNP® Enterprise and Cisco Certified Specialist - Enterprise Advanced Infrastructure Implementation certifications.

Your success in passing the CCNP Implementing Cisco Enterprise Advanced Routing and Services (ENARSI) 300-410 test is this book's most important and obvious goal. However, the methods used in this book to assist you in passing the exam are intended to increase your knowledge of carrying out your work. The title would be misleading if this book's main objective were something else.

Networking

Computer networking is known as the process of transmitting and exchanging data between nodes via a common medium in an information system. A private Wide Area Network (WAN) or the internet's Local Area Network (LAN) allows for the connection of devices and endpoints. This function is essential for service providers, enterprises, and customers worldwide to share resources, use or supply services, and communicate.

Networking simplifies everything, from phone conversations to text messages to streaming video to the Internet of Things (IoT).

The design, construction, and use of a network are all parts of networking, as are the management, operation, and maintenance of the network's hardware, software, and protocols. The complexity of a network directly affects the level of expertise needed to run it. For example, the management of skilled network administrators is necessary when a major organization has thousands of nodes and strict security requirements, such as end-to-end encryption. In short, networking technology has changed the world and opened up new possibilities for the overall growth of all the world's areas.

Security

A broad notion, network security includes various tools, systems, and procedures. In a nutshell, it is a collection of guidelines and configurations created to safeguard the privacy, accessibility, and integrity of computer networks and data by utilizing software and hardware technologies. Regardless of size, sector, or architecture, every organization needs network security solutions to safeguard it from the ever-growing environment of online threats.

Network infrastructure is complicated in the current threat landscape, where vulnerabilities are continually sought after and exploited. Users, devices, data, applications, and locations are just a few of the scenarios in which these vulnerabilities may be present. Numerous network security management tools and applications are currently used to address specific risks and exploits. These security precautions are important since even a brief interruption could cause considerable harm to an organization's

reputation and baseline.

Cisco Course

With a focus on networking and communications products and services, Cisco Systems, Inc. is a pioneer in global technology. The company's business switching and routing products, which route data, voice, and video traffic across networks worldwide, are presumably well recognized.

An IT infrastructure specialist can obtain Cisco certificates, which are well-known and useful qualifications. Whether you are preparing for your CCNA, CCNP, CCIE, or CCENT tests, you will need the Cisco certification training courses to succeed.

Cisco's training and certification programs have been updated to consider today's rapidly changing technologies and to help students, engineers, and software developers succeed in the most crucial positions in the sector.

What is ENARSI?

ENARSI is a CCNP Enterprise domain "Specialist" level exam, and the certificate launched on June 9th, 2019. It is the first ENARSI exam version to participate in the CCNP Enterprise certification and award the applicant a Cisco Certified Specialist – Enterprise Advanced Infrastructure Implementation certificate.

Concerning the CCNP Enterprise Specialty, Cisco also unveiled a brand-new domain of expertise and hierarchical structure in addition to ENARSI.

ENARSI can be your first and best option in one of two scenarios.

- If you want to dive deep into routing protocols and services based on enterprise-level networks
- If you already have a working knowledge of the old CCNP RS and want to review related material

Why can one do CCNP-ENARSI?

This book aims to significantly raise your chances of passing the ENARSI 300-410 exam. Even though it can be used for that purpose, this book is not intended to be a general networking subject's book. Although this book can be used to accomplish other goals, its primary mission is to assist you in passing the exam.

In light of this, why would you wish to pass the ENARSI 300-410 exam? Because obtaining the CCNP Enterprise certification, which is no small task, depends on passing this milestone. What would you gain personally from earning the CCNP Enterprise certification? A pay increase, a promotion, or acclaim? Why not improve your resume? Proving that you are serious about learning more and are not content to sit back and take it all in? Pleasing your reseller-employer, who needs more qualified staff to receive a greater Cisco discount? The CCNP Enterprise certification may be something you want for several reasons, including one of those listed above.

Benefits?

This book's core methodology is to assist you in identifying the exam topics you need to review in more detail, and how to properly comprehend, remember, and demonstrate to yourself that you still know the material. This book aims to help you learn the material effectively rather than memorize it. For a routing/switching engineer or expert to be skilled, they must possess the knowledge covered in the ENARSI 300-410 exam, which covers foundational topics in the CCNP certification. If this book did not try to make you learn the subject, it would be a disservice. To that aim, the following strategies are included in the book to assist you in passing the test:

- Assisting you in identifying the test subjects you need more practice with

- Offering justifications and details to close any knowledge gaps
- Providing exercises and scenarios that improve your capacity for memory and deduction of test question answers
- Offering test questions on the companion website that serve as practice exercises for the subjects and the examination procedure

Future of CCNP-ENARSI

Only if you possess superior enterprise networking knowledge and abilities will you be able to contribute to the dynamic technological environment of today and the future. With your CCNP Enterprise certification, you will access a broad range of narrowly focused skills in significant technological areas. This illustrates the value of a Cisco certification to your career progress. The future need for cutting-edge talents will increase due to the organization's growing need for networking solutions. As a result, a Cisco CCNP Enterprise certification guarantees that there are enough professionals on hand to address this growing need. And considering the benefits they stand to gain, no one could pass up this vital chance.

What is the Cost of the Cisco 300-410 CCNP ENARSI?

The 90-minute Cisco 300-410 ENARSI test contains a time constraint and 55–65 questions. Japanese and English are the two languages offered for the exam. On the Pearson VUE website, students can register for the exam. The cost of the Cisco 300-410 ENARSI test is $300, and people can take it online or in a testing facility. The CCNP 300-410 ENARSI test has no prerequisites, although it is assumed that candidates have 3-5 years of experience working in the IT networking field and sufficient knowledge and background.

Demand in 2022

The CCNP ENARSI has significantly changed recently, yet it will still be valuable in 2022. The CCNP ENARSI certification will be useful in 2022 since

it attests to your proficiency in managing and configuring enterprise-level networks. Administrators who work for companies that rely on sizable, campus-wide networks to demonstrate their knowledge and competence will benefit from the CCNP ENARSI in 2022.

With the help of this course:

- You will learn the skills to set up, configure, run, and troubleshoot an enterprise network
- Be eligible for positions at the professional level in advanced routing and services
- Acquire 40 CE points for recertification

Strategies for Exam Preparation

Your method for passing the ENARSI 300-410 exam may be slightly different from that of other readers depending on your skills, knowledge, and experience. For example, suppose you took the CCNP Implementing Cisco Enterprise Advanced Routing and Services (ENARSI) 300-410 course. In that case, you might approach routing differently from someone who learned it through on-the-job training.

This book is intended to assist you in reaching the point where you can pass the exam in the shortest amount of time possible, regardless of your approach or background. For example, if you already fully comprehend IP addressing and subnetting, there is no need for you to practice or study it. However, many people want to review information they already know to ensure they understand a subject. Several book features will give you the assurance you need to believe that you already understand certain stuff and assist you in identifying the subjects you need to learn more about.

Prerequisites

The prerequisites for this course are:

- Basic knowledge of network principles
- Basic understanding of LAN implementation
- General knowledge of network device management
- General knowledge of network device security
- Understanding the basics of network automation

To assist you in completing these requirements, these Cisco courses are advised:

- Implementing and Operating Cisco Enterprise Network Core Technologies (ENCOR) v1.0
- Interconnecting Cisco Networking Devices, Part 1 (ICND1) v3.0
- Interconnecting Cisco Networking Devices, Part 2 (ICND2) v3.0

CHAPTER 02: LAYER 3 TECHNOLOGIES

Introduction

Layer 3 includes the network's switching and routing technologies, which establish logical paths for data transmission between network nodes. Routing and forwarding, internetworking, addressing, packet sequencing, congestion control, and additional error handling are among Layer 3's primary responsibilities. Many protocols are used in Layer 3, which requires troubleshooting if they do not work properly. Hence, in this chapter, we will learn how to troubleshoot issues in routing protocols.

This chapter will focus on troubleshooting EIGRP for IPv4, EIGRPv6 (EIGRP for IPv6), and named EIGRP. As OSPF can route for IPv4 and IPv6 Protocols, this chapter will discuss troubleshooting OSPFv2 and troubleshooting OSPFv3. OSPFv3 is designed for routing IPv6 networks. This chapter will focus on troubleshooting OSPFv3 using classic configurations and the OSPF address family configurations.

Additionally, this chapter will discuss the different problems you could encounter when attempting to set up an IPv4, IPv6, and internal border gateway protocol (iBGP) neighbor adjacency and how to recognize and fix these problems.

Troubleshooting Administrative Distance

Misconfigured route redistribution can result in problems like routing loops and suboptimal routing. Users may have delayed connectivity due to suboptimal routing, and connectivity may be lost due to routing loops. All redistribution concerns come down to two distinct problems:

- Metric-related problems
- Administrative distance-related problems

Troubleshooting Administrative Distance of EIGRP

EIGRP already distinguishes between routes learned from within the autonomous system and routes learned from outside the autonomous system by providing a separate administrative distance:

- External EIGRP: 170;
- Internal EIGRP: 90

Use the **EIGRP configuration** command distance eigrp ad-internal ad-external to change the administrative distance on IOS routers from its default setting. The range of acceptable values for the AD is 1-255; a value of 255 prevents the route from being installed into the RIB.

With the command **distance ad source-ip source-ip-wildcard [acl-number | acl-name],** Cisco IOS (Internetwork Operating System) routers can enable selective AD adjustment for particular internal networks. The optional ACL limits the modification to a specified network prefix, while the source-ip option limits it to routes in the EIGRP database learned from a particular router. Be aware that for external EIGRP routes, EIGRP does not permit selective AD adjustment based on prefixes.

Troubleshooting Administrative Distance of OSPF

The default AD 110 value used by OSPF is the same for routes discovered inside and outside the OSPF routing domain. You can change the default AD on IOS routers with the OSPF configuration command distance ospf external | inter-area | intra-area ad. Each OSPF network type can have a distinct AD set with the command.

IOS routers' **distance ad source-ip source-ip-wildcard [acl-number | acl-name]** command allows for selective AD change for particular networks. With the source-ip option, only routes in the OSPF LSDB (Link-State

Database), which were learned from the LSA's advertising router, can be modified (Link-State Advertisement). The source-ip-wildcard address fields align with the advertising route's RID (Router-ID). Use the optional ACL (Access Control List) to limit access to a particular network prefix.

The example below shows how to change R1 and R2 so that the external OSPF routes' AD is set to 171, higher than the external EIGRP routes' AD of 170.

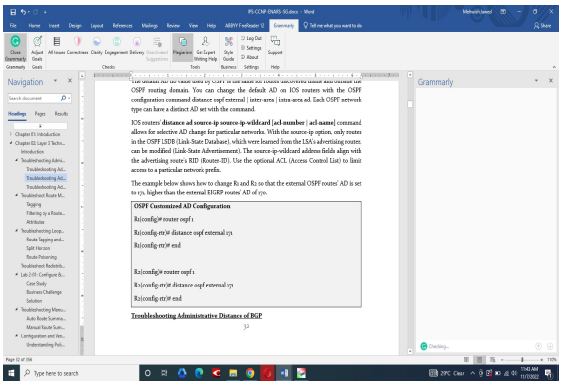

Troubleshooting Administrative Distance of BGP

BGP distinguishes between routes learned locally and from eBGP peers and routes learned from IBGP peers. Use the address family command **distance ad source-ip source-wildcard [acl-number | acl-name]** on IOS routers to change the AD for routes received from a specific neighbor and the BGP configuration command **distance bgp external-ad internal-ad local-routes** to set the AD for each BGP network type.

Troubleshoot Route Map (Tagging, Filtering, Attributes)

The "if-then" programming solution for Cisco devices is route maps. Using a route map, you can check for specific match circumstances and (optionally) specify a value.

Route maps are used in conjunction with other services and features to offer a finer level of control not present with the services or features on their own. For example, all routes are redistributed and handled equally when you switch from one routing protocol to another. On the other hand, you can redistribute each route or a collection of routes differently by including a route map in the procedure. Route maps are also a key component of PBR and are used extensively with BGP for path management.

Consequently, you must be able to troubleshoot the route map when troubleshooting a service or feature that has a route map associated to it to identify whether it is the root of the problem.

Tagging

You do not want the routes transferred from EIGRP to OSPF to be transferred back to the EIGRP autonomous system. This may result in routing problems like loops, which prevent packets from being delivered to their destination properly (in addition to wasting CPU and memory resources on various devices in the network). Route tags are the most effective solution to this problem. The given figure demonstrates how, when a route is redistributed, R1 and R2 can add a tag—basically, any value that can be used to identify the route—to it. Maps of the route are used to achieve this. In this example, R1 adds a tag of 10 while redistributing the 10.1.1.0/24 route into the OSPF domain. R2 adds a tag of 20 when redistributing the 10.1.1.0/24 route into the OSPF domain.

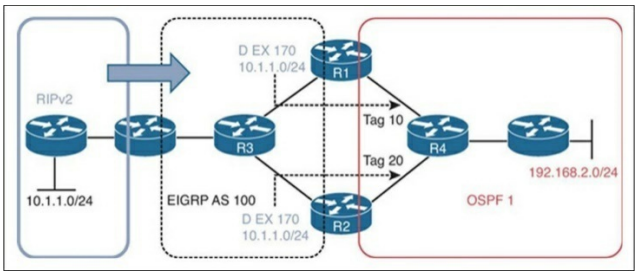

Filtering by a Route Map

Route filtering can be implemented for a variety of reasons. For example, since no site in the branch would ever need to send packets to a host in the transit network, it would be wise to block routes to transit networks from being broadcast to branch routers.

The distribute-list router subcommand in EIGRP is used to set up filters. It can use a prefix list, an ACL, or a route map to decide whether or not routes should be filtered, in which direction, and on which interface.

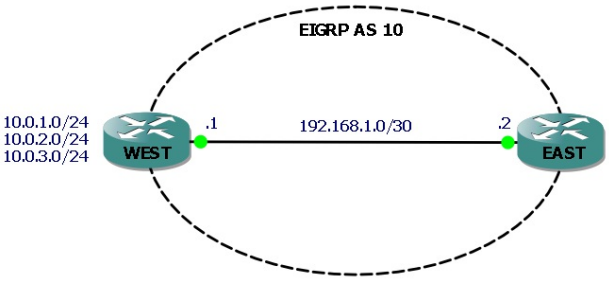

We will filter routes to the three networks connected to router WEST (represented by loopback interfaces) so that they do not show up in the routing table of router EAST. The distribute-list command, a prefix list, an ACL, and a route map will be used to filter each route individually.

Each route is displayed in the routing database of router EAST without any filtering:

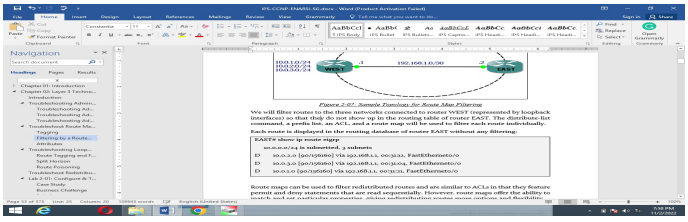

Attributes

Attributes refer to BGP attributes, as BGP has a lot of attributes, and if you have been working with BGP earlier, you surely have been in touch with route maps many times.

```
as-path           Prepend string for a BGP AS-path attribute
automatic-tag     Automatically compute TAG value
clns              OSI summary address
comm-list         set BGP community list (for deletion)
community         BGP community attribute
dampening         Set BGP route flap dampening parameters
default           Set default information
extcommunity      BGP extended community attribute
interface         Output interface
ip                IP specific information
ipv6              IPv6 specific information
level             Where to import route
local-preference  BGP local preference path attribute
metric            Metric value for destination routing protocol
metric-type       Type of metric for destination routing protocol
mpls-label        Set MPLS label for prefix
nlri              BGP NLRI type
origin            BGP origin code
tag               Tag value for destination routing protocol
traffic-index     BGP traffic classification number for accounting
vrf               Define VRF name
weight            BGP weight for routing table
```

Troubleshooting Loop Prevention Mechanisms

With different kinds of networks, the routing loop is a common problem. They are created when a routing algorithm malfunctions, leading to a group of nodes' route to a particular destination and forming a loop. The routing loop will vanish when the new network topology is flooded to all routers within a routing area in link state routing protocols like IS-IS or OSPF. The loop prevention is built into the more recent distance vector protocols, including Babel, DSDV, Eigrp, and BGP. It uses an algorithm that guarantees that routing loops never happen, not briefly. Do not implement a novel loop prevention method in the more ancient routing protocols like RIP. It uses mitigation techniques like route poisoning, split horizons, route filtering, and route tagging.

Route Tagging and Filtering

In the mesh network, the routing loop may easily happen unless a protocol provides an inherent fix. Due to incorrect routing information circulated in a network, the routing loop prevents some packets from being properly routed. Then counting to infinity is a symptom of such routing loops. The use of route tagging to prevent advertising from the routing protocol is recommended when configuring multipoint redistribution.

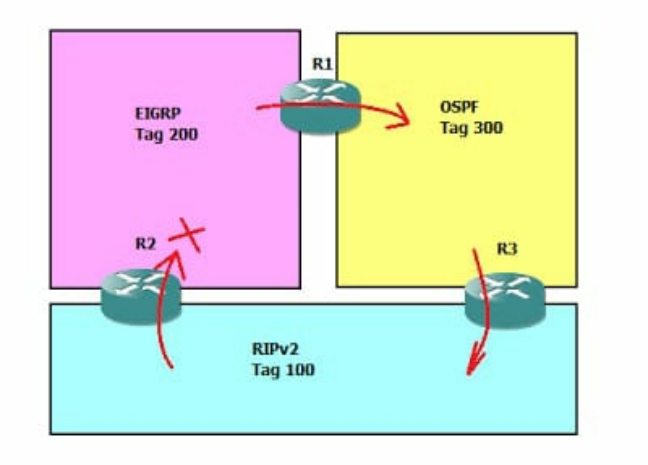

Split Horizon

The routing loops pose a serious threat to distance vector protocols. One of the characteristics of distance vector routing protocols that prevent them is the split horizon. The router is prevented from announcing the path back onto an interface from which it learned it.

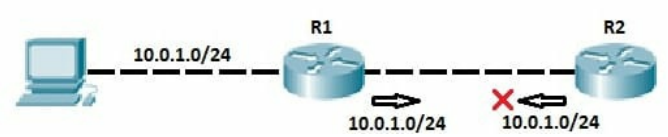

Route Poisoning

Another technique distance vector routing protocols use to avoid routing loops is route poisoning. A router transmits the route's advertisement with an infinite metric when it notices that one of its directly related routes has failed (poisoning the route). When a router receives an update, it is aware that the route has failed and stops using it.

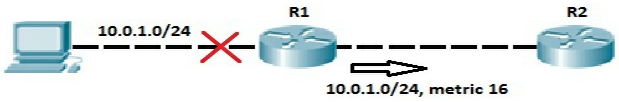

Troubleshoot Redistribution between OSPF and EIGRP

When redistributing routes from one routing source into another, the data

from the original routing source is lost unless the seed metric is injected at the redistribution point. The destination routing source thus loses or is unable to observe the complete network. This is not an issue when there is only one redistribution point between two sources. However, if there are several locations where redistribution occurs between two sources, as illustrated in the given figure, the less-than-ideal path may be used to reach routes.

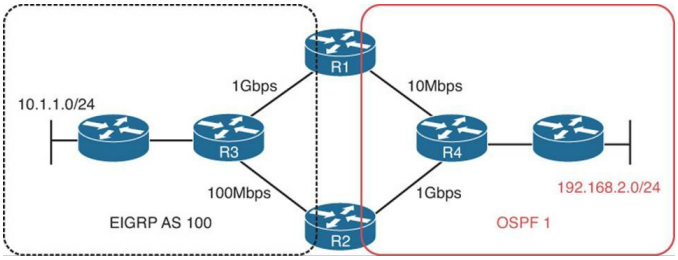

Troubleshooting Manual and Auto Route Summarization with EIGRP

Route summarizing is a technique for combining several networks into one summary address. Because it minimizes routing updates and cuts down on the routes the routers must maintain, it is frequently employed in big networks with numerous subnets. Both automatic and manual summation are available as ways of summarizing routes.

Auto Route Summarization

The default state of the EIGRP auto-summary feature may vary depending on your IOS version (it is more likely to be disabled in modern IOS versions). When auto-summary is enabled, routes are summarized in the routing updates to the classful border.

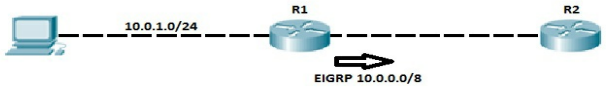

Manual Route Summarization

Manual summarization is possible on any router in a network, which is one advantage EIGRP has over other routing protocols (like OSPF). The number of routes in a routing table can be reduced by using a single route to represent several routes.

Each interface has a unique setting for manual summarization. The command is written as follows:

```
(config-if) ip summary-address eigrp ASN SUMMARY_ADDRESS SUBNET_MASK
```

An illustration will clarify the idea of manual summarization:

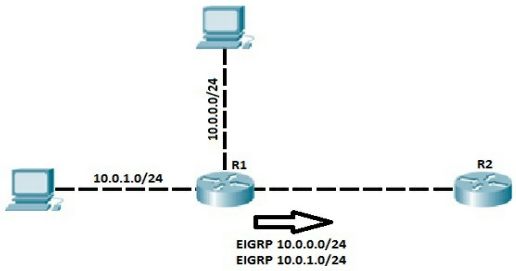

Configuration and Verification of Policy-Based Routing

Understanding Policy-Based Routing

A router's data plane processing logic performs numerous steps to process a packet after it enters the incoming interface. The inbound packet comes inside a data link layer frame; thus, the router must examine the frame's Frame Check Sequence (FCS) and discard the frame if transmission issues occur. The router discards the data-link header and trailer of the incoming frame and only keeps the Layer 3 packet if the FCS check is successful. The router selects the longest-prefix route that matches the destination IP

address, which is equal to checking the packet's destination IP address with the IP routing database.

A router's built-in destination-based forwarding logic is overridden by Policy-Based Routing (PBR). Before the router runs the CEF table lookup, PBR intercepts the packet after de-encapsulation on the incoming interface. Then, PBR decides how to forward the packet based on factors other than comparing the packet's destination address with the CEF table.

A route map, which commonly refers to an IP access control list (ACL), is used by PBR to determine how to forward the packet using matching logic. The next-hop IP address or outgoing interface is defined in the same route map as the forwarding instructions for packets that fit the route map. The figure illustrates the basic idea, with PBR on interface Fa0/0 subverting the standard routing logic and sending packets out to three outgoing interfaces.

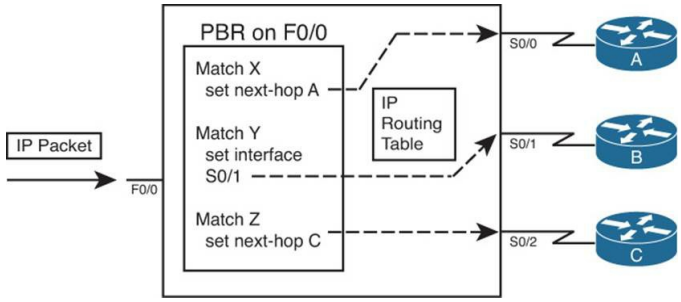

Configuration and Verification of VRF-Lite

Virtual Routing and Forwarding (VRF) allows several virtual routers to be created on a single real router. The isolation of router interfaces, routing tables, and forwarding tables on a VRF-by-VRF basis prevents traffic from one VRF from interfering with that of another VRF. In place of using numerous devices, VRFs boost router capability by segmenting traffic and are a crucial part of the MPLS L3VPN design.

Overview of VRF-Lite

The Global VRF is associated with every router interface, routing table, and forwarding table. So, what you have been referring to as your routing table is the Global VRF's routing table. Create more VRFs, which generate more routing and forwarding tables if you divide your router into many virtual routers.

VRF-Lite Configuration

Consider Figure 2-18 to show a basic VRF-Lite configuration. Isolating the voice, data, and video networks into distinct VRF instances is one of the objectives of the network topology displayed. You will see that the COMMON router's Fa 0/0 interface is broken up into three sub-interfaces (Fa 0/0.2, Fa 0/0.3, and Fa 0/0.4). Following that, an 802.1Q trunk connects the COMMON router to switch SW1. The switch port connected to each router is part of a distinct VLAN (i.e., VOICE VLAN = 2, DATA VLAN = 3, and VIDEO VLAN = 4), which is how Switch SW1 links to the VOICE, DATA, and VIDEO routers.

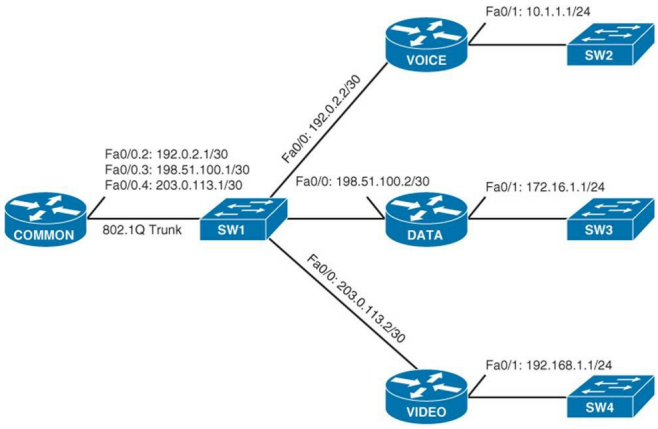

Bidirectional Forwarding Detection

Bidirectional Forwarding Detection, or BFD, is a very quick technology that may identify link problems in milliseconds or even microseconds. Every (routing) protocol has some type of method to identify failed links. EIGRP utilizes hello packets and a hold-down timer, while OSPF uses hello packets

with a dead interval.

Fast convergence times are necessary for networks that use real-time traffic, such as VoIP. Routing protocols like OSPF or EIGRP can immediately choose a different path after losing a neighbor, but it takes some time to figure out what is wrong.

We can adjust timers to converge quickly; for example, OSPF can be set to use a dead interval of just one second. The issue is that none of these protocols were intended for sub-second failover. There is a lot of overhead since the control plane processes hello packets and other things. BFD was created with speed; a few interface modules or line cards can process its packets, minimizing overhead.

Any other routing protocols are not necessary for BFD to function. Once it is operational, protocols like OSPF, EIGRP, BGP, HSRP, MPLS LDP, etc., can be configured to use BFD rather than their systems for link failure detection. BFD will let the protocol know when the link breaks. How to picture this is as follows:

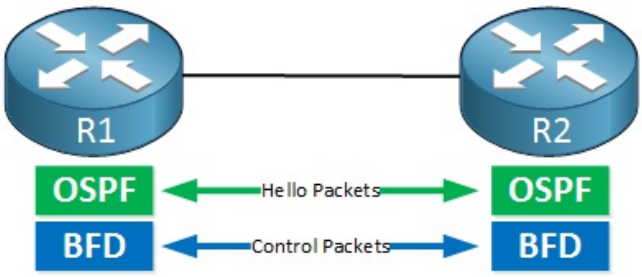

With BFD set up, R1 and R2 will communicate by exchanging control packets. As before, OSPF is still sending its OSPF packets. This will occur once the link fails:

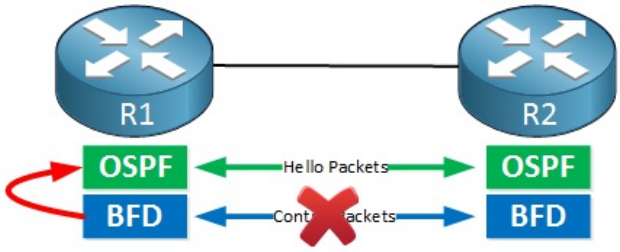

Troubleshooting EIGRP (Classic and Named Mode)

Enhanced Interior Gateway Routing Protocol (EIGRP) from Cisco is a sophisticated distance vector routing protocol. EIGRP establishes neighbor associations like a link-state routing protocol while advertising routes to directly associated neighbors.

Both IPv4 and IPv6 can be routed via EIGRP. This chapter's main goal is to troubleshoot these protocols using classic and EIGRP configurations.

Classic Mode

When using the traditional EIGRP configuration mode, the majority of settings are configured in the EIGRP process, though some options are configured in the interface configuration sub-mode. This can make deployment and troubleshooting more challenging because users have to scroll back and forth between the EIGRP process and specific network interfaces. A few examples of specialized choices include the hello advertisement interval, split-horizon, authentication, and summary route advertisements.

Troubleshooting EIGRP for IPv4 Addresses

EIGRP establishes neighbor connections by sending hello packets to the multicast address 224.0.0.10 from participating interfaces. In router EIGRP configuration mode, the network ip address wildcard mask command is used to enable the EIGRP process on a certain interface. To enable EIGRP, for

example, use the network command network 10.1.1.0 0.0.0.255 on all interfaces with an IP address between 10.1.1.0 and 10.1.1.255. The network command network 10.1.1.65 0.0.0.0 allows the EIGRP process to be enabled only on the interface with that IP address. Although it looks pretty straightforward, which it is, there are several reasons why a neighbor connection might not establish. You must be aware of all of them if you intend to successfully troubleshoot EIGRP-related issues.

Troubleshooting EIGRP for IPv4 Neighbor Adjacencies

To confirm EIGRP neighbors, run the **show ip eigrp neighbors** command. The output of the show ip eigrp neighbors command is displayed in the example below.

The length of time the local router will consider the neighboring router to be a neighbor, the length of time the routers have been neighbors for, the average amount of time the neighbors spend communicating, the number of EIGRP packets in a queue ready to be transmitted to a neighbor, and other information are all included in this message.

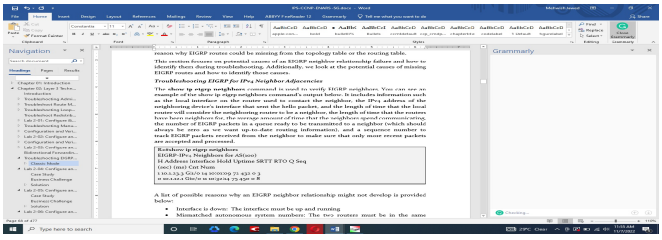

The following is a list of potential causes why an EIGRP neighbor relationship might not materialize:

- **Interface is down:** The interface must be up and running
- **Mismatched autonomous system numbers:** The two routers must be in the same autonomous system
- **Network statement error:** The interface's IP address

must be listed in the network statement if you want to include it in the EIGRP process

- **Mismatched K values:** The K values used by the two routers must be the same
- **Passive interface:** This feature prevents hello packets from being sent and received while allowing the interface's network to be advertised
- **Different subnets:** Hello packets must be sent and received on the same subnet; otherwise, they are ignored
- **Authentication:** For authentication to be successful, the key ID and key string must match, and the key itself must be genuine (if configured)
- **ACL:** An ACL blocking packets from reaching the EIGRP multicast address 224.0.0.10
- **Timers:** Timers do not need to match, but if they are not set up properly, your neighbor adjacencies will be affected

Interface Is Down

The interface must be active if you intend to establish an EIGRP neighbor adjacency. The **show ip interface brief** command allows you to check the status of an interface.

Mismatched K Values (Troubleshooting Metrics)

To produce an adjacency, the K values employed in the metric calculation must agree between neighbors. Show ip protocols can be used to check whether K values match, as illustrated below. Normally, the K values do not need to be changed. If they are altered, make sure they are identical on each router in the autonomous system. The spot-the-difference technique can be used to check whether K values are inconsistent between routers.

Additionally, you will get a warning that looks something like this if you are logging syslog messages with a severity level of 5:

%DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.12.2 (GigabitEthernet1/0)
is down: K-value mismatch

Passive Interface

All organizations must have passive interface functionality. It performs two tasks:

- Reduces the EIGRP-related traffic on a network
- Enhances the security of EIGRP

Different Subnets

The router interfaces must be on the same subnet to establish an EIGRP neighbor adjacency. There are numerous ways to verify this. The **show run interface interface_type interface_number** command is the simplest way to view the interface settings in the running configuration. The configuration of Gig1/0 on R1 and Gig0/0 on R2 is shown below. Do they belong to the same subnet? Yes! Based on the IP address and subnet mask, they would be on the 10.1.12.0/24 subnet. However, if you have syslog configured for a severity level of 6 and they are not in the same subnet, you will see a message that looks something like this:

%DUAL-6-NBRINFO: EIGRP-IPv4 100: Neighbor 10.1.21.2 (GigabitEthernet1/0) is blocked: not on a common subnet (10.1.12.1/24)

ACL

Access control lists have a lot of strength; what they control in your network will depend on how they are implemented. A neighbor relationship will not

form if an interface has an ACL applied and the ACL denies EIGRP packets. Use the **show ip interface interface_type interface_number** command, as shown in the following snippet, to see if an ACL is applied to an interface. As you can see, interface Gig 1/0 has ACL 100 applied inbound. Use the command **show access-list 100** to confirm the ACL 100 entries.

Timer

Even though EIGRP timers are not required to synchronize, the adjacency will flap if the timers are significantly off. Consider, for example, that R1 sends hello packets every five and fifteen seconds, whereas R2 sends them every twenty seconds. Before receiving another hello packet from R2, R1's hold time will have expired, ending the neighbor connection. The neighbor relationship is established five seconds after the hello packet arrives and lasts for 15 seconds.

Authentication

Authentication makes ensuring that your EIGRP routers only connect to authorized routers as neighbors and that they only receive EIGRP packets coming from authorised routers. Therefore, if authentication is used, both routers must concur on the settings for a neighbor relationship to develop. The spot-the-difference method can be used with authentication. An output of the commands **show run interface interface_type interface_number** and **show ip eigrp interface detail *interface type interface* number** is shown below, indicating whether or not EIGRP authentication is enabled on the interface. It is, as stated in the text, that has been highlighted. Remember that the proper interface must be used for the authentication, and the appropriate autonomous system number must be connected. It would not be activated for the correct autonomous system if you input the wrong number for that system.

Troubleshooting Named EIGRP Configurations (Address Family IPv6)

To give you a single location on the local router to carry out all EIGRP for IPv4 and IPv6 configurations, EIGRP named configurations to serve this purpose. An example of an EIGRP configuration is given below under the name TSHOOT EIGRP. In this EIGRP configuration, both an IPv4 and an IPv6 unicast address family are present. Although it is not necessary, they both employ autonomous system 100.

Named EIGRP Verification Commands

The named EIGRP supports all EIGRP show commands that are applicable to both classic EIGRP for IPv4 and classic EIGRP for IPv6. You might be interested in finding out about a fresh set of EIGRP show commands.

The EIGRP for IPv4 address family and the EIGRP for IPv6 address family, along with each autonomous system number, are both displayed by the command show eigrp protocols. The K values, router ID, stub router status, AD, maximum pathways, and variation are also displayed.

Troubleshooting OSPF

The Open Shortest Path First (OSPF) dynamic routing protocol is a link-state routing protocol using Dijkstra's shortest path first (SPF) algorithm. Due to the implementation of a hierarchical design, it is a very scalable routing protocol. OSPF can route both the IPv4 and IPv6 protocols. This chapter focuses on troubleshooting OSPFv2 and OSPFv3 using the older OSPF address family configurations and the more recent ones.

An OSPF neighbor connection must be established before any routes on the same LAN or across a WAN can be exchanged between OSPF routers. You need to understand why a neighborly relationship would not flourish if you

want to solve problems. This part goes in-depth on these factors and gives you the knowledge and abilities needed to identify them and successfully settle neighbor disputes.

Neighboring routers will communicate when neighbor relationships are established by exchanging OSPF LSAs containing route-related data. Routes may go missing for a variety of reasons, so you need to be able to pinpoint them. This chapter covers how OSPF routes could disappear, how to find out why they vanished, and how to resolve problems with routes.

Troubleshooting OSPFv2

By broadcasting hello packets to all participating interfaces, OSPF creates neighbor connections. To enable the OSPF process on an interface and add it to an OSPF area, use the **network ip address wildcard mask area *area-id*** command in router OSPF configuration mode or the **ip ospf process-id area area-id** command in interface configuration mode. For example, the network 10.1.1.0 0.0.0.255 area 0 command enables OSPF and places all interfaces with IP addresses between 10.1.1.0 and 10.1.1.255 in area 0. **Ip ospf 1 area 51 is the interface configuration** command that activates the OSPF process on the interface and assigns it to area 51.

Troubleshooting OSPFv2 Neighbor Adjacencies

The **show ip ospf neighbor** command is used to confirm OSPFv2 neighbors. You can see an example of the output below of the show ip ospf neighbor command. Router ID (RID), priority in the Designated Router/Backup Designated Router (DR/BDR) election process, state (which will be discussed shortly), and whether the neighbor is a DR, BDR, or DROther are all displayed for the neighbor. There are two ways to enable OSPFv2 on an interface, so you must be very careful while debugging neighbor

adjacencies. It also displays the dead time, which shows how long the neighborhood router will remain up if it does not. Otherwise, you can make a mistake and think the OSPF process was turned off when it actually was enabled on an interface. This is a reminder to check both places. You should receive another greeting packet by then (the default is 40 seconds on a LAN). You can also see the local router interface of the neighbor and the IP address from which the hello packet was delivered.

Troubleshooting OSPFv2 Neighbor Adjacencies

The **show ip ospf neighbor** command is used to confirm OSPFv2 neighbors. You can see an example of the output below of the show ip ospf neighbor command. Router ID (RID), priority in the Designated Router/Backup Designated Router (DR/BDR) election process, state (which will be discussed shortly), and whether the neighbor is a DR, BDR, or DROther are all displayed for the neighbor. It also displays the dead time, which represents the amount of time the local router will remain operational in the event that it does not receive another hello packet during that period (default is 40 seconds on a LAN). You can also see the local router interface of the neighbor and the IP address from which the hello packet was delivered.

Verifying OSPF Neighbors with show ip ospf neighbor

```
R1#show ip ospf neighbor

Neighbor ID Pri State Dead Time Address Interface
10.1.23.2 1 FULL/BDR 00:00:37 10.1.12.2 GigabitEthernet1/0
```

A syslog message resembling the following will be sent when an OSPF neighbor adjacency is successfully formed:

```
%OSPF-5-ADJCHG: Process 1, Nbr 10.1.23.2 on GigabitEthernet1/0 from
LOADING to FULL, Loading Done
```

Here is a listing of reasons why an OSPFv2 neighbor relationship might not form:

- **Interface is down:** The interface has to be up/up
- **Interface not running the OSPF process:** If the interface is not enabled for OSPF, it will not send hello packets or form an adjacency
- **Mismatched timers:** Hello and dead timers have to match between neighbors
- **Mismatched area numbers:** Both ends of a link must be in the same OSPF area
- **Mismatched area type:** Besides a normal OSPF, an area type could be either stub or Not-So-Stubby Area (NSSA). The routers have to agree on the type of area they are in
- **Different subnets:** Neighbors have to be in the same subnet
- **Passive interface:** The passive interface feature suppresses the sending and receiving of hello packets while allowing the interface's network to be advertised
- **Mismatched authentication information:** If one OSPF interface is configured for authentication, the OSPF interface at the other end of the link has to be configured with matching authentication information
- **ACLs:** An ACL denying packets to the OSPF multicast address 224.0.0.5
- **MTU mismatch:** The maximum transmission unit of neighboring interfaces must match
- **Duplicate router IDs:** Router IDs must be unique
- **Mismatched network types**: Two neighbors configured

with a different OSPF network type may not generate an adjacency based on that network type's default values and features.

Mismatched Timers

To create a neighbor adjacency in OSPF, timers must match between neighbors, unlike in EIGRP (Enhanced Interior Gateway Routing Protocol). For broadcast and point-to-point network types, the default hello timer is set to 10 seconds; for non-broadcast and point-to-multipoint network types, it is set to 30 seconds. For broadcast and point-to-point network types, the dead timer defaults to 40 seconds; for non-broadcast and point-to-multipoint network types, it defaults to 120 seconds.

Mismatched Area Numbers

To make OSPF a highly scalable dynamic routing protocol, areas are used. The neighboring interfaces' vicinity is required for OSPF routers to form a neighbor adjacency. You can find out to which area an OSPF interface belongs by using the **show ip ospf interface interface_type interface_number or show ip ospf interface brief** commands.

Troubleshoot Area Type (Mismatched Area Type)

The kind of area that OSPF, by default, specifies is a normal area. You can convert a regular area into a stub area or NSSA area to regulate the types of LSAs transmitted into the area via an Area Border Router (ABR). For adjacencies to form, all of the routers in the area must accept the area type. The hello packet contains a stub area flag to indicate the neighborhood type the neighbor is in.

Different Subnets

The router interfaces must be on the same subnet to establish an OSPF

neighbor adjacency. There are numerous ways to confirm this. **The show run interface interface_type interface_number** command is the simplest way to view the interface configuration in the running configuration.

Passive Interface

All organizations must have the passive interface feature. It performs two tasks: 1) lessens OSPF-related network traffic; 2) increases OSPF security.

Mismatched Authentication Information

Because of authentication, your OSPF routers only establish neighbor relationships with legitimate routers and only accept OSPF packets from legitimate routers. As a result, both routers must concur on the configuration for a neighbor relationship to form if authentication is used. When troubleshooting authentication, you can employ the spot-the-difference technique. OSPF accepts three different kinds of authentication:

- **Null:** Type 0 and denoting a lack of authentication
- **Plain text:** This is type 1 and uses clear text to send credentials
- **MD5:** A type 2 protocol that sends a hash

ACLs

The power of Access Control Lists (ACLs) is immense. What they control in your network will depend on how they are implemented. A neighbor relationship will not form if an ACL is applied to an interface and the ACL blocks OSPF packets.

MTU Mismatch

Each router's interface forming the adjacency must have the same MTU for OSPF routers to become neighbors and reach the full adjacency state.

Otherwise, the routers will detect one another but remain in the exchange/start states. R1 and R2 are both stuck in the exchange state, according to the output of the **show ip ospf neighbor** command.

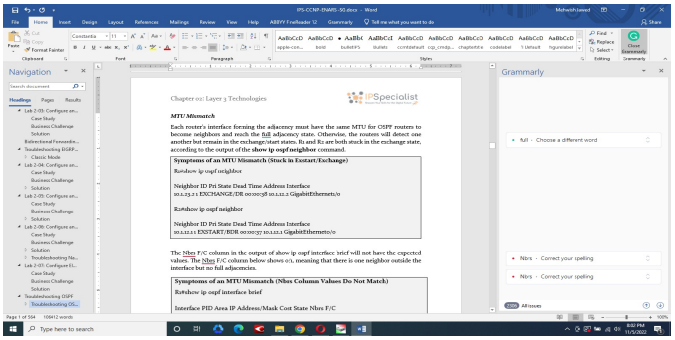

Duplicate Router IDs

RIDs must be distinct for a variety of reasons. One of the reasons is that if two routers have the same RID, a neighbor relationship will not form between them. A message similar to that displayed below will appear in your syslog when a duplicate RID is present:

```
.%OSPF-4-DUP_RTRID_NBR: OSPF detected duplicate router-id 10.1.23.2 from 10.1.12.2 on interface GigabitEthernet1/0
```

Troubleshooting Network Types (Mismatched Network Types)

OSPF supports multiple network types. Different types of networks have various default settings. Therefore, a neighbor relationship will not form if two OSPF routers trying to form a neighbor adjacency are configured with incompatible network types.

Area Types, Network Types, Router Types

Area Types

OSPF operates within a single Autonomous System (AS). However, networks within this single AS can be divided into several areas. By default, Area 0 is created. Area 0 can either function alone or act as the OSPF backbone for a

larger number of areas. Each OSPF area is named using a 32-bit identifier which in most cases is written in the same dotted-decimal notation as an IP4 address. For example, Area 0 is usually written as 0.0.0.0.

The topology of an area is maintained in its own link state database and is hidden from other areas, which reduces the amount of traffic routing required by OSPF. A connecting router then shares the topology in a summarized form between areas.

Backbone Area

The center of an OSPF network is the backbone area (Area 0). It connects to all other areas and is the only route vehicles can use between them. The backbone area serves as the distribution point for all routing between regions. While all other OSPF areas must connect to the backbone area, this connection need not be direct.

Normal OSPF Area

In a normal OSPF area, there are no restrictions; the area can carry all types of routes.

Stub Area

A stub area does not receive routes from other autonomous systems. The stub area is routed through the default route to the backbone area.

Totally Stubby Areas

Total stubby regions at the ABR are off-limits to Type 3 LSAs (interarea), Type 4 LSAs (ASBR summary LSAs), and Type 5 LSAs (external routes). When it receives a Type 3 or Type 5 LSA, an ABR of a fully stubby region constructs a default route for the completely stubby area.

NSSA

The Not So Stubby Area (NSSA) is a type of stub area that can import

external routes, with some limited exceptions.

Network Types

Broadcast Network

Broadcast multi-access is a better term to distinguish broadcast media like Ethernet from Non-Broadcast Multi-Access (NBMA) networks. Broadcast networks are multi-access in the sense that they can connect more than two devices, and broadcasts sent out to one interface can reach all interfaces on that segment.

For Ethernet ports, the OSPF network type is defaulted to broadcast. Because numerous nodes may exist on a segment and LSA flooding must be regulated, this OSPF network type requires a DR. The default value for the hello timer is 10 seconds, as stated by RFC 2328.

An interface is statically configured as an OSPF broadcast network type when the interface parameter command ***ip ospf network broadcast*** is used to override the automated configuration.

Nonbroadcast

In that, they can connect more than two devices, Frame Relay, ATM, and X.25 are termed Non-Broadcast Multi-Access (NBMA), and broadcasts sent out one interface might not always be able to reach all the interfaces attached to the segment. Dynamic virtual circuits may offer connection. However, the architecture might only offer a hub-and-spoke structure rather than a full mesh.

Point-to-Point Network

A Point-to-Point (P2P) network is a network circuit that allows only two devices to communicate. Point-to-point networks do not require Address Resolution Protocol (ARP) because of the nature of the medium, and

broadcast traffic is not a limiting concern.

The OSPF network type is set to point-to-point by default for point-to-point Frame Relay Sub-Interfaces, Generic Routing Encapsulation (GRE) Tunnels, and Serial Interfaces (HDLC or PPP encapsulation). Since there can only be two nodes on this kind of network medium, OSPF does not waste CPU time on DR functionality. The hello timer on OSPF point-to-point network types is set to 10 seconds.

Point-to-Multipoint Networks

For any medium, the OSPF network type point-to-multipoint is not by default enabled. Manual configuration is necessary. This OSPF network type does not have a DR enabled, and the hello timer is set to 30 seconds. A point-to-multipoint OSPF network type is widespread in Frame Relay and Layer 2 VPN (L2VPN) topologies and provides hub-and-spoke communication while utilising the same IP subnet.

Interfaces configured for OSPF point-to-multipoint networks add the interface's IP address as a /32 network to the OSPF LSDB. Even if the next-hop IP address is on the same IP subnet, it is set to the interface's IP address when advertising routes to OSPF peers on that interface. An interface can be manually set to be an OSPF point-to-multipoint network type using the IOS interface parameter command ip ospf network point-to-multipoint. R1, R2, and R3 are all using Frame Relay point-to-multipoint sub-interfaces using the same subnet in the topology example shown in Figure 2-46.

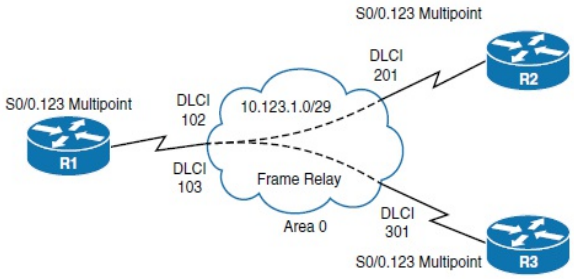

Router Types

Within an OSPF area, routers are divided into the following categories.

Internal Router

Only devices in the same area have OSPF neighbor relationships with the router.

Area Border Router (ABR)

A router that is connected to other OSPF neighbors through numerous OSPF areas. ABRs collect and distribute topological data to the backbone area from their connected areas.

Backbone Router

A router that uses OSPF and has at least one interface connected to the OSPF backbone area is referred to be a backbone router. ABRs are usually categorized as backbone routers since they are always linked to the backbone.

Autonomous System Boundary Router (ASBR)

A router known as an ASBR is one that connects to many routing protocols and communicates routing data between them.

Troubleshooting OSPFv3 for IPv6

With a few minor differences due to IPv6, OSPFv3 is based on OSPFv2, so you will be dealing with similar problems when it comes to troubleshooting. Knowing that you do not need to learn a lot of new information for OSPFv3 should be a relief. However, to troubleshoot any given OSPFv3-related issue, you need to be familiar with the show commands that will display the necessary information.

The Link LSA (Type 8) and the Intra Area Prefix LSA (also known as Type 9) are two new LSA types that can

be seen in the example above and defined below. Both of these LSAs are described below for OSPFv3. Also, note that the Type 3 LSA (Summary LSA) is now referred to as the Inter-Area Prefix LSA in the above example.

LSA 8: The link LSA provides information to neighbors about link-local addresses and the IPv6 addresses associated with the link. Therefore, it is only flooded on the local link and is not reflooded by other OSPF routers.

LSA 9: The intra-area prefix LSA provides information for two different scenarios. First, it provides information about IPv6 address prefixes associated with a transit network by referencing a network LSA. Then, it provides information about IPv6 address prefixes associated with a router by referencing a router LSA. Type 9 LSAs are flooded only within an area.

Troubleshoot OSPFv3 Address Families

With the help of OSPFv3 Address Families (AFs), you can set up a single process to support IPv4 and IPv6. Additionally, both IPv4 and IPv6 are maintained in the same database. On the other hand, Adjacencies are established uniquely for each AF, and settings can be customized for each AF separately.

OSPF Path Preference

To build a loop-free topology of shortest paths, OSPF applies Dijkstra's shortest path first (SPF) algorithm. To determine the shortest path for each network, all routers employ the same methodology.

Path selection uses the following logic to prioritize paths:

1. Intra-area

2. Interarea

3. External Type 1

4. External Type 2

Each component is thoroughly explained in the sections that follow.

Intra-Area Routes

Type 1 LSAs are always preferred over Type 3 LSAs for routes advertised in a given area. If multiple intra-area routes exist, the OSPF Routing Information Base (RIB) installs the path with the lowest total path metric, which is then displayed to the router's global RIB. Both routes are added to the OSPF RIB if there is a tie in the metric. R1 is calculating the route to 10.4.4.0/24 in Figure 2-48. R1 travels to R4 via the slower serial link (R1⟶R3⟶R4) because that is the intra-area path, instead of the faster Ethernet connection (R1⟶R2⟶R4).

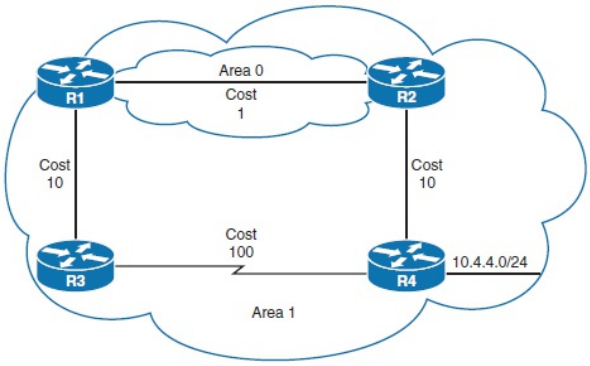

Interarea Routes

The next priority when selecting a path to a network should be the way with the lowest overall path metric to the destination. If the metrics are tied, both routes are added to the OSPF RIB. All interarea pathways must traverse Area 0 in order for a route to be taken into consideration. In the accompanying image, R1 is calculating the route to R6. R1 uses the path R1R3R5R6 since its total path metric is 35 as compared to the path R1R2R4R6 with a metric of 40.

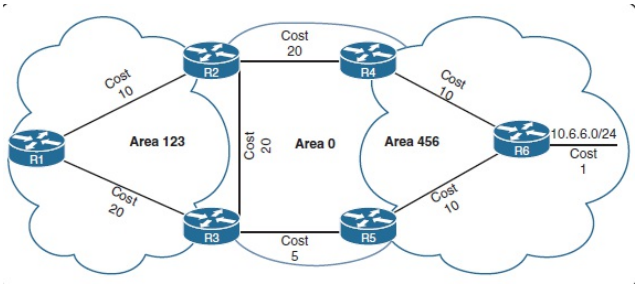

External Route Selection

Type 1 or Type 2 external routes are the two categories. Following are the primary distinctions between Type 1 and Type 2 external OSPF routes:

- Routes of Type 1 are preferred to those of Type 2.
- The redistribution metre plus the overall path metric to the ASBR make up the Type 1 metric. In other words, the measure rises as the LSA spreads away from the initial ASBR.
- Only the redistribution metric is equal to the Type 2 metric. The router nearest to the ASBR and the router 30 hops away from the originating ASBR use the same measure. This is the type of external measure that OSPF employs by default.

Troubleshooting iBGP and eBGP

The Internet protocol's name is Border Gateway Protocol (BGP). It seeks to facilitate routing data transfer among numerous autonomous systems (networks under different administrative control). It is therefore categorized as an Exterior Gateway Protocol (EGP). In contrast to Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), or Routing Information Protocol, it bases its best path decisions on variables such as local preference, the length of the autonomous system path, and even the BGP router ID (RID) (RIP). The best, most reliable, and easiest to maintain protocol is BGP. But there is a cost associated with that.

Organizations mainly use BGP to connect to their Internet service provider (ISP). Static routes are employed if not. ISPs, however, frequently exchange Internet routes with one another using BGP. The ISP-to-ISP BGP connectivity

is not the foundation of the 300-135 TSHOOT exam and is based on connectivity between businesses and ISPs. Therefore, you should concentrate on resolving basic BGP connectivity and route advertising issues.

Address Families (IPv4)

The idea behind several address-family commands is this. When we designate someone as a neighbor under a specific address family, we intend to trade routes from that address family with that neighbor. If a neighbor is not identified under a certain address family, we do not intend to share information from that family with that neighbor.

The neighbors with whom we want to exchange standard IPv4 unicast routes are declared by the address-family ipv4. This may come as a surprise as exchanging IPv4 routes with neighbors only requires defining a neighbor by their address. The BGP implicitly allocates all defined neighbors to an unseen address-family ipv4 section in order to maintain backward compatibility with earlier BGP versions that were not multiprotocol capable. So you do not have to manually add it, as soon as you define a neighbor, it is immediately added to an invisible address-family ipv4 section.

Troubleshooting BGP for IPv6 (Address Families)

The same BGP autonomous system configuration mode is used to configure BGP for IPv4 and IPv6. Multiprotocol BGP, or MP-BGP for short, is the term used for this. Utilizing address families and turning on neighbors for those address families are necessary for implementing BGP for IPv4 and IPv6 on the same router.

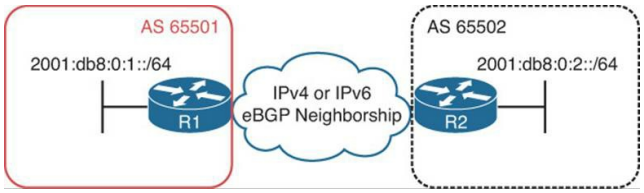

BGP Packets Sourced from Wrong IP Address

A BGP router will have several active IP addresses configured across various interfaces in a redundant topology. There are two BGP autonomous systems shown in the given figure. Observe that the multiple paths allow R2, R3, and R4 to establish a BGP peering with one another using any physical interface. For example, R2 and R4 could establish a peering over a direct connection or a connection through R3.

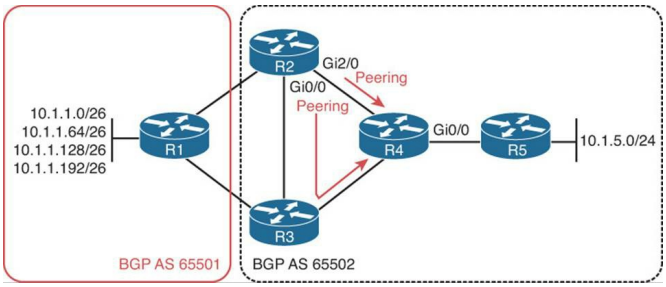

TTL of BGP Packet Expires

Between two routers that are directly connected, an eBGP peering normally takes place. This indicates that the distance between the routers forming the eBGP peering is predicted to be one router hop. The routers in an iBGP peering can be up to 255 router hops apart and still form a peering. The result of **show bgp ipv4 unicast neighbors** reveals that the iBGP neighbor at 5.5.5.5 can be up to 255 hops away from the eBGP neighbor at 10.1.12.1, and that the eBGP neighbor at 10.1.12.1 must be within one router hop of being accessed. If the neighbor cannot be reached within the allotted number of hops, the BGP packet expires and no neighbor relationship is created.

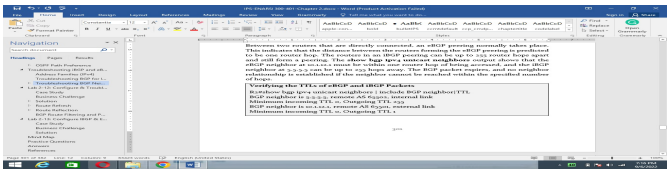

The packet will be discarded if the TTL cannot support the necessary distance to establish a BGP peering. For illustration, let us use the loopback interfaces of R1 and R2 to establish an eBGP peering in the given figure. The loopback interfaces for R1 and R2 are 1.1.1.1 and 2.2.2.2, respectively. Ping has been used to test Layer 3 connectivity, and it is effective and not via a default route.

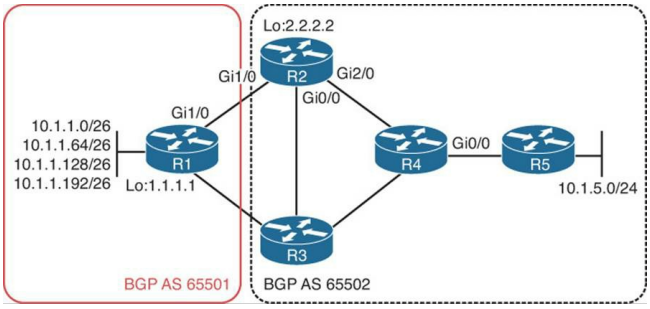

BGP 4-byte ASN

The 2-byte (16-bit) BGP Autonomous System Number (ASN) is an entity. IANA has reserved the digits 0, 23456, 65535, 64512, and 65534 (private ASN) out of the 65536 available numbers for 2-bytes. 39000+ ASNs of the remaining ones have already been used.

ASNs in the 4-byte range from 0 to 4294967296. Mappable-ASNs range from 0 to 65535. There are three possible methods to express the 4-byte ASN:

asplain - ASN representation in decimal form is known as asplain. ASN 7747, for example, will be rendered as 7747, and ASN 123456, as 123456.

asdot+ - The function asdot+ divides the number into two 16-bit values, low-order and high-order, and separates them with a dot. The low-order value can be used to represent any 2-byte ASN. ASN 65535, for example, will be

0.65535, 65536, 1.0, 65537, 1.1, and so forth. ASN 4294967296 will be replaced by 65535.65535.

asdot - asplain and asdot+ are combined to form asdot. Any ASN in the 2-byte range is denoted by the symbol asplain, while ASNs above this range are denoted by asdot+. For example, 65536 will be 1.0, and 65535 will be 65535. Cisco employs this technique for implementation.

Route Refresh

The Border Gateway Protocol (BGP) Enhanced Route Refresh feature gives BGP the ability to detect route irregularities and, in the uncommon case that it does, synchronize BGP peers without performing a hard reset.

BGP Enhanced Route Refresh Functionality

BGP peers communicate information about each other's ability to perform the BGP Enhanced Route Refresh functionality during session establishment. The function is turned on by default.

It is not anticipated that the peers will start acting differently toward one another. Only in the most unlikely of circumstances may that occur; in such case, this function aids in identifying it and synchronizing the peers without performing a hard reset.

Assume that two peers have Enhanced Route Refresh capabilities. Then, before and after each peer's advertisement of the Adj-RIB-Out, it will generate a Route-Refresh Start-of-RIB (SOR) message and an EOR message, respectively. When a BGP speaker receives an EOR message from a peer, it deletes any routes that the peer did not re-advertise in the Route Refresh response.

The peers were inconsistent with one another if, in the uncommon case that

the router still contained stale routes after receiving the EOR message or after the EOR period expired. To determine whether the routes are reliable, use this data.

- Timers for BGP Enhanced Route Refresh
- Messages Produced by the BGP Enhanced Route Refresh in Syslog

BGP Enhanced Route Refresh Timers

Normally, these timers do not need to be configured. You could configure one or both timers if you notice constant route flapping to the point where a Route Refresh EOR cannot be created.

When a router is not getting an EOR message, the first timer kicks in. The router must transmit the EOR message according to the second timer.

- Stale path timer – If the router is configured with the bgp refresh stalepath-time command and does not receive a Route-Refresh EOR message following an Adj-RIB-Out, the router deletes the stale routes from the BGP table once the timer has run out. When the router receives a Route-Refresh SOR message, the stale path timer begins.
- Maximum EOR timer: If the router cannot generate a Route-Refresh EOR message due to the configuration of the bgp refresh max-eor-time command, an EOR message is generated when the timer expires.

Both times can be customized. They are both disabled by default (set to 0 seconds).

Syslog Messages Generated by the BGP Enhanced Route Refresh

Here are some illustrations of syslog messages that are produced when a peer deletes stale routes following receipt of the Route-Refresh EOR message or following the expiration of the stale path timer. You can

determine whether the routers were inconsistent with the messages.

After a refresh EOR, net 300:300:3.3.0.0/0 from bgp neighbor IPv4 MDT 10.0.101.1 is stale (rate-limited)

After the refresh stale-path timeout expires, Net 300:300:3.3.0.0/0 from bgp neighbor IPv4 MDT 10.0.101.1 is stale (rate-limited)

The following are examples of messages logged after a Route-Refresh EOR or after the stale path timer expires, indicating the total number of stale paths from the neighbor.

- Following refresh EOR, 3 stale-paths were removed from bgp neighbor IPv4 MDT 10.0.101.
- Following the expiration of the refresh stale-path timeout, 3 stale-paths were eliminated from bgp neighbor IPv4 MDT 10.0.101.1.

Troubleshooting BGP Path Selection

In contrast to OSPF and EIGRP, BGP does not consider a link's bandwidth when choosing a route. Instead, BGP considers several factors when deciding which path is best. You must have a firm grasp of all the attributes when troubleshooting BGP paths to fully understand why BGP made the choice. The decision-making procedure for the BGP best path is covered in this section. Additionally, we look at the quantity of private autonomous systems.

Understanding the Best Path Decision-Making Process

The attributes are listed below in the order that BGP considers them when deciding which path is the best:

1. Prefer the highest weight

2. Prefer highest local preference

3. preference for local router-originating routes

4. Prefer shortest autonomous system path

5. Prefer lowest origin code

6. Prefer lowest MED (metric)

7. Prefer external over the internal path

8. Prefer path through closest IGP neighbor

9. Prefer the oldest route for eBGP paths

10. Prefer the path with the lowest neighbor BGP RID

11. Prefer the path with the lowest neighbor IP address

When choosing the BGP optimal path, keep in mind the above figure and the output of show bgp ipv4 unicast 10.1.1.0 on R5.

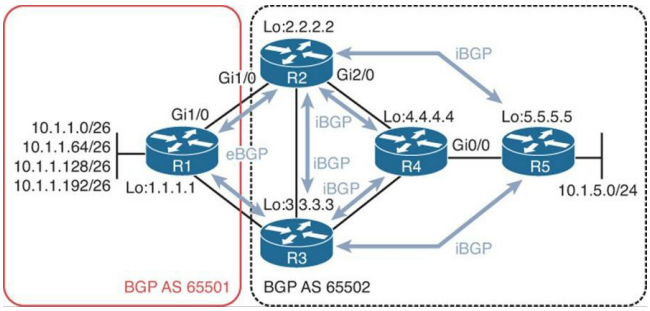

Private Autonomous System Numbers

BGP autonomous system numbers have a private range just like IPv4 addresses do. It is between 64,512 and 65,534 for 2-byte autonomous systems, and between 4,200,000,000 and 4,294,967,294 for 4-byte autonomous systems. The public autonomous system numbers for networks that are multi-homed to several ISPs can still be utilized for single- or dual-

homed networks to the same ISP.

The customer's network may use the private autonomous system numbers, but they must not be included in the AS PATH attribute when the routes are advertised to the Internet (global BGP table) because doing so may cause issues on the Internet because multiple autonomous systems may share the same private autonomous system number.

It is necessary to discontinue sending private autonomous system numbers into the global BGP database. The neighbor ip address remove-private-as command can be used to do this.

Route Reflector

All iBGP peers within an AS must be completely mesh due to the BGP split-horizon requirement (within iBGP). While iBGP neighbors do not include their ASN in the AS PATH when transmitting updates, eBGP neighbors use the AS-PATH for loop avoidance. So, what method of loop prevention does iBGP employ? Divided horizon The iBGP split-horizon rule says:

"Any routes learned from an iBGP neighbor must never be advertised to any other iBGP neighbor."

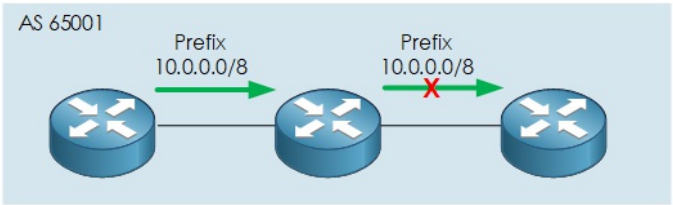

Due to the requirement for numerous distinct BGP sessions, complicates and limits the network's ability to scale and the BGP routers.

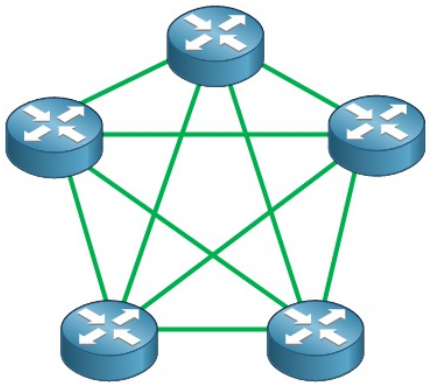

When n = BGP speakers inside the AS, the total number of BGP sessions needed within the AS may be computed using the formula: n*(n-1)/2. For example, 10 BGP routers would need 45 BGP sessions to have fully meshed.

Route Reflection

One method to reduce BGP peering within an AS is route reflection. Instead of each BGP system needing to peer with every other BGP system with the AS, each BGP speaker instead peers with a router reflector. Any routing advertising sent to the route reflector is reflected to all other BGP speakers.

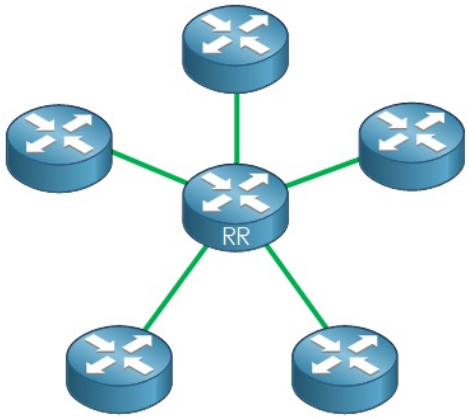

The route reflector reflects routes considered as best only. Additionally, a route reflector is NOT allowed to change any reflected route attribute, including the next-hop attribute.

Peer Types

There are two types of internal peers to a route reflector - Client and Non-

Client. Let us look at the differences,

Peer Type	Full Mesh	Cluster	RR Operation
Client (RR Client)	Not Required	With other clients and RR form a cluster	A route from a client peer reflect to all the non-client peers and client peers
Non-Client (Regular iBGP)	Required	Sits outside of the cluster	A route from a non-client iBGP peer reflect to all the clients

Routing Information Loops

Let's examine the three BGP properties that are employed with route reflectors to avoid routing information loops.

CLUSTER_ID

Multiple RRs can be established to prevent a route reflector from serving as a single point of failover. However, this raises the possibility of a problem with advertising loops between RRs. It is possible to have many RRs in the same cluster by configuring each RR with a 4-byte CLUSTER ID, which allows an RR to ignore routes from other RRs in the same cluster.

CLUSTER_LIST

What happens, though, if our RR is a component of many clusters? In this example, the RR will add its CLUSTER ID to the CLUSTER LIST attribute and use distinct CLUSTER IDs for the pertinent RR neighbors. The RR deletes the update when it receives an update with the router's CLUSTER ID in the CLUSTER LIST.

ORIGINATOR_ID

The ORIGINATOR ID attribute is used by the route reflector to assign the router ID of the advertising router when it reflects a path. The UPDATE is disregarded if any router subsequently receives a UPDATE that includes a path with an ORIGINATOR ID matching its router ID.

BGP Route Filtering and Path Manipulation

A technique for selectively recognizing prefixes published or received from

peers is conditional route selection. Selected routes can be changed or eliminated to alter traffic patterns, make better use of memory, or enhance security. The entire BGP route processing logic is depicted in the given image. Keep in mind that the route policies appear on both the outgoing route advertisement and the inbound route receipt.

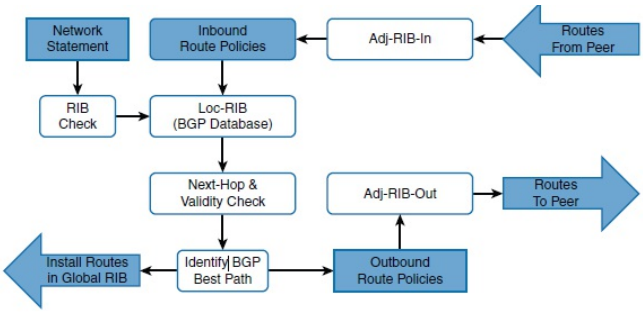

For a specific BGP peer, IOS XE has four options for route filtering, either inbound or outgoing. Each of these approaches can be used independently or in tandem with additional approaches:

- **Distribution list:** Using a standard or enhanced ACL, a distribution list filters network prefixes. Any prefix that is not allowed is connected to an implicit refusal.
- **Prefix list:** A list of prefix-matching specifications that, like an ACL, top-down allow or refuse network prefixes. Any prefix that is not allowed is connected to an implicit refuse.
- **AS_Path ACL/filtering:** A list of regex instructions that allow or disallow a network prefix depending on the current AS Path values. Any prefix that is not allowed is connected to an implicit refusal.
- **Route maps:** Route maps give you a way to conditionally match on different prefix qualities and let you do different things. Simple permit, deny, or alter BGP path properties are examples of actions. Any prefix that is not

allowed is connected to an implicit refusal.

Mind Map

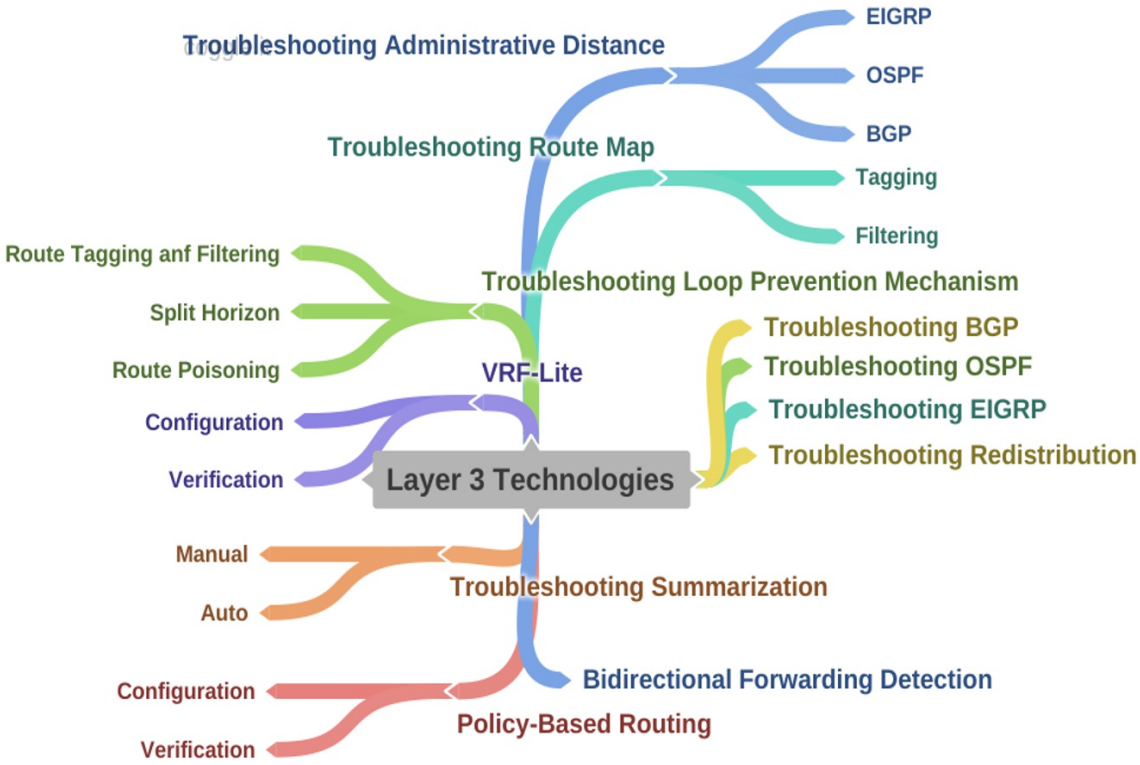
Troubleshooting Administrative Distance
EIGRP
OSPF
BGP
Troubleshooting Route Map
Tagging
Filtering
Route Tagging anf Filtering
Troubleshooting Loop Prevention Mechanism
Split Horizon
Route Poisoning
VRF-Lite
Troubleshooting BGP
Troubleshooting OSPF
Troubleshooting EIGRP
Troubleshooting Redistribution
Configuration
Verification
Layer 3 Technologies
Manual
Troubleshooting Summarization
Auto
Bidirectional Forwarding Detection
Configuration
Policy-Based Routing
Verification

CHAPTER 03: VPN TECHNOLOGIES

Introduction

A Virtual Private Network (VPN) is a technology that connects private sites or networks. A VPN allows you to be virtually present in a remote network. For example, you can connect to your company's network remotely. Your home network and office network are both private networks. These private networks are connected with the help of a VPN.

In this chapter, you will learn about different VPN technologies. Internet Service Providers provide VPN technologies, and organizations buy those services from their ISPs to connect their branches. Organizations mostly buy MPLS services from their ISPs to connect their branches at different locations.

MPLS Operations (LSR, LDP, Label Switching, LSP)

Computer networks use routers to transport packets from one network to another. These routers use routing tables to store the routes and other information, such as source and destination IP addresses. Routers make forwarding decisions to forward the packets with the help of these routing tables, and they look up a route to forward a packet to a certain destination. This IP-based routing is called traditional routing.

Multiprotocol Label Switching (MPLS) is a new technology. MPLS is a packet-forwarding technology that uses labels. Hence the name label switching. When a packet arrives at the router, it applies a label on the packet as an identifier and forwards it to the next router. The next router removes the label applied by the previous router, applies its label, and forwards the packet to the next router. When the last router receives the packet, it removes the label and forwards the IP packet without a label to the destination host or server.

MPLS was originally designed to support various routing protocols. It can run while any other routing protocol, such as OSPF or EIGRP, is already running. MPLS is not much faster than traditional routing, but it is still preferable. One of the major benefits is the less forwarding overhead on core routers because routers are not required to look up the routing table every time due to the use of labels. Internet Service Providers (ISPs) benefit from MPLS and run this technology on their core routers. This makes forwarding somewhat faster than traditional routing.

MPLS has other benefits too. It supports many services, such as multicast routing, unicast routing, Quality of Service (QoS), Traffic Engineering (TE), and Virtual Private Networks. Therefore, it is preferred over traditional routing.

MPLS LIB and LFIB

Every router has a **control plane** and a **data plane**. The Control plane acts like a control unit and handles the processing and computing, and the Data plane handles the actual data packets and forwards them.

The Control plane is also known as **the forwarding plane**. The Control plane has **Routing Information Base (RIB)**, and the data plane has **Forwarding Information Base (FIB)**. RIB and FIB are used in traditional routing, which is IP-based, but in MPLS, something else is used.

A router creates a routing table in the control plane. The routing table has all the routes and other related information. This routing table is part of the Routing Information Base (RIB). Information from the RIB is used in the Forwarding Information Base (FIB), where the packets are forwarded.

In MPLS, labels are used to identify the packets, so a **Label Information Base (LIB)** is created in the control plane. LIB is equivalent to Routing Information

Base (RIB). Information from the LIB is then used in the **Label Forwarding Information Base (LFIB)** to forward the data (packets). Label Forwarding Information Base (LFIB) is equivalent to the Forwarding Information Base (FIB).

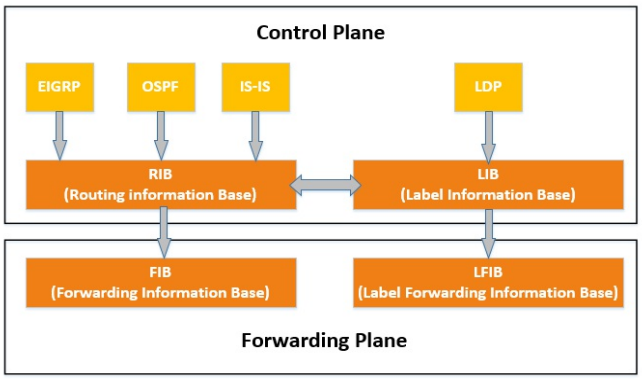

Before enabling MPLS, **Cisco Express Forwarding (CEF)** must be enabled on the routers. In new routers, CEF is enabled by default. It is a technique that makes the switching faster. As discussed above, there is a control plane and a data plane. There is an ARP table in CEF-enabled routers and a routing table in the control plane. Information from the ARP table creates an **Adjacency table** in the data plane, just like a routing table creates a Forwarding Information Base.

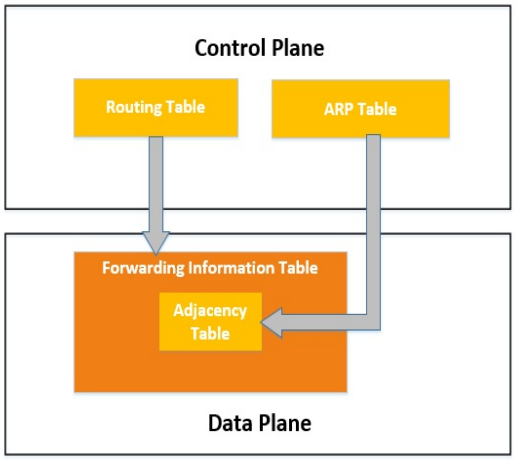

Address Resolution Protocol (ARP) is a layer 2 protocol that helps find the MAC address of a host when only its IP is known. Routing takes place on

layer 3, so the control plane has layer 2 and layer 3 information in a CEF-enabled router. MPLS is configured on a CEF-enabled router and uses its control and data planes. MPLS lies in between layer 2 and layer 3.

Label Switch Router (LSR)

An MPLS architecture consists of core Provider Routers (P) routers. MPLS is mostly used by Internet Service Providers (ISPs). These routers provide services to the customers through the **Provider Edge Routers (PE)**. The routers placed on the customers' side are known as **Customer Edge Routers (CE)**. Provider Routers are connected to the Provider Edge Routers, and Provider Edge Routers are connected to the Customer Edge Routers.

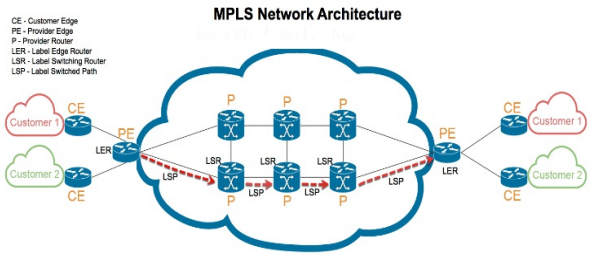

Label-Switched Path (LSP)

Every labeled packet takes a path (route) from a source to a destination within an MPLS domain. This path is known as **Label-Switched Path (LSP)**, a unidirectional one.

In complex networks, a return path may differ from the first source to the destination path, but it is not common. MPLS runs over an underlying routing protocol such as OSPF or EIGRP. These routing protocols ensure symmetric networks and forwarding paths. So, most traffic returns through the path where they come from.

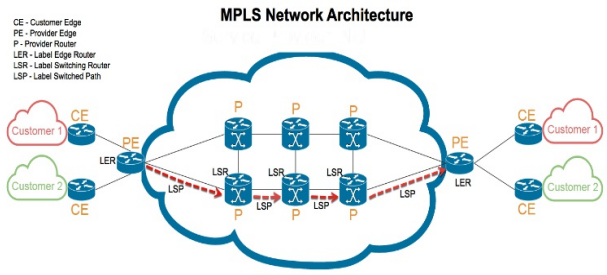

The above figure's red-dashed arrows indicate the Label-Switched Path (LSP). LSP is inside the MPLS domain, where the packet has a label. No label on the packet (unlabeled packet) is not on the LSP and is clearly outside the MPLS domain.

Labels

The label is the most important thing in Multiprotocol Label Switching (MPLS). When a packet is labeled, it is a header placed between the packet's layer 2 frame header and layer 3 IP header. It is called a shim header that is placed between these two headers.

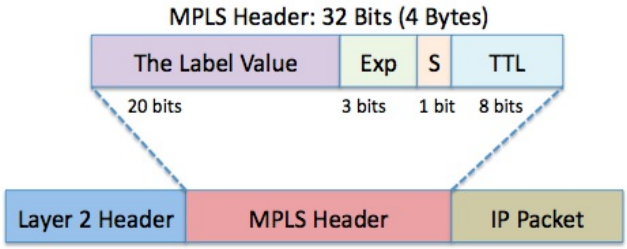

MPLS label header is 32-bit (4-byte) in size and contains four fields. The first 20-bit **Label Value field** has the label value. The next 3-bit **EXP field** is used for Quality-of-Service (QoS). The next 1-bit **S field** tells whether the label is last or not in the stack because there can be multiple labels on a packet (as in VPNs). The last 8-bit **TTL field** has a Time-to-Live value that decreases on every hop till zero.

Label Distribution Protocol (LDP)

MPLS does not do all its work by itself but also needs a protocol. The MPLS

technology uses a common protocol called Label Distribution Protocol (LDP). This protocol helps distribute/share labels between the MPLS-enabled routers. A router uses the label information of its neighbors to populate its LIB.

Label Switching

Label switching is adding or removing the labels to or from the packets. The process in which a router adds the label on the packet is known as **Pushing** the label. The process in which a router removes the label from the packet is known as **Popping** the label. The process in which a label is replaced with another label is known as **Swapping** the label.

Penultimate Hop Popping

Till now, you have learned that the last LSR in the MPLS domain pops the label from the packet because there is no more LSR on the outgoing interface. At last, LSR looks up its Label Forwarding Information Base when a labeled packet is received. If there is no label-out (the label it has to push), it looks up the Forwarding Information Base and forwards the packet without any label to the destination. This method is not efficient.

MPLS Layer 3 VPN

Multiprotocol Label Switching (MPLS) layer 3 Virtual Private Network (VPN) is used to connect private sites over a public network. Most Internet Service Providers (ISPs) use MPLS, and their customers (private sites) can connect using MPLS layer 3 VPN. ISP's MPLS domain is a shared network, and this VPN goes over the shared network of the Internet Service Provider.

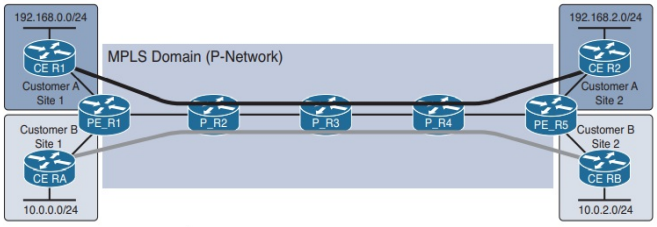

In the above figure, Customer A is connected to its private site over the MPLS domain. Likewise, Customer B is also connected to its private site over the same MPLS shared domain. The MPLS domain is called P-Network, and the customer sites are called **C-Network**.

MPLS Layer 3 VPNs

MPLS Layer 3 VPN architecture consists of core routers known as Intermediate or Provider Routers (P), Edge Routers known as Provider Edge Routers (PE), and routers on the customer sites known as Customer Edge Routers (CE). MPLS domain is sometimes referred to as Provider Network because the service providers mostly use it. The router on the receiving end is known as an Ingress LSR, and the router that forwards the traffic outside the MPLS domain is known as an Egress LSR.

The idea is that Customer A site 1 can exchange its local routing information with customer A site 2, so they can communicate with each other as needed. Due to MPLS Layer 3 VPN, Customer A, and Customer B can have the same IP address spaces. For example, Customer A and Customer B can use the 192.168.0.0 IP address range.

To isolate the customers (because they can be using the same IP space), **Virtual Routing and Forwarding (VRF)** is used. VRF should be enabled on the Provider Edge Router (PE) interface connected to the Customer Edge Router (CE). CE and PE exchange their routing information by using any underlying dynamic routing protocol such as Enhanced Interior Gateway Routing Protocol (EIGRP), Routing Information Protocol (RIP), Open Shortest Path First (OSPF), or Border Gateway Protocol (BGP).

MPLS Layer 3 VPNv4 Address

Customers can use overlapping address spaces, so to differentiate between

the customers' information, **Route Distinguisher (RD)** is used. RD is a unique value that can be added to the IP prefixes so the identical prefixes can be differentiated. RD is generated and used by Provider Edge Routers because the routing information from the Customer Edge Routers first goes to the PE routers. With the help of RD value, MP-BGP can distinguish between the customer routes.

RD has always used whether the customers use the same IP prefixes or not. A Router Distinguisher value is a **64-bit value** that is prepended (added at the beginning of) to the 32-bit IP address, making it a **96-bit unique prefix** called **VPNv4 address**. The MP-IBGP neighboring routers then exchange this VPNv4 address.

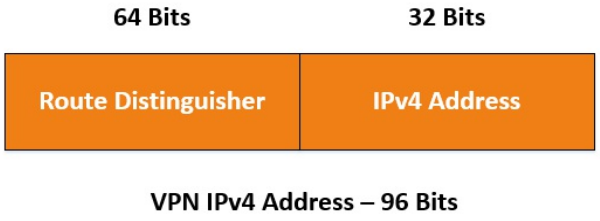

MPLS Layer 3 VPN Label Stack

A label stack is required in the MPLS domain to forward the packets successfully. At least a stack for two labels is required. One label is the **LDP label**, and another is the **VPN label**. Both labels are added by the Provider Edge (PE) routers.

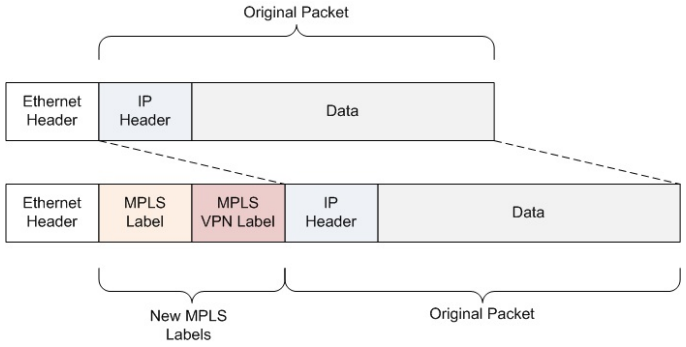

VPN label is for the Egress LSR, so it can identify which traffic belongs to which customer. LDP label is used for the label switching. VPN label is learned from the Provider Edge (PE) router over an MP-IBGP (Multiprotocol Interior Border Gateway Protocol). Label Distribution Protocol (LDP) labels are learned from the Provider Routers in the MPLS domain because they share their label information within the MPLS domain.

DMVPN

MPLS is used to connect the two private sites. If there is a need for a private site to connect to multiple sites, MPLS service will be very costly to buy from the service provider. Dynamic Multipoint VPN is used to connect multiple private sites, allowing a site to connect to other sites without going through the ISP or hub router.

The DMVPN architecture consists of a single hub, the main router, and the spoke routers, the customers or endpoint routers. All the spokes are connected to the hub. When a spoke wants to connect to another spoke in the network, a tunnel is built between the spoke routers. This way, traffic from one spoke does not go through the hub router.

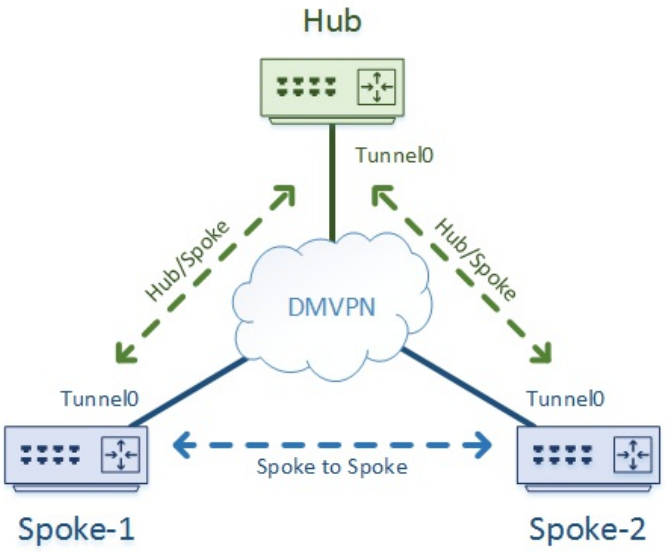

Some of the benefits of using DMVPN are listed below:

- **Zero-Touch Provisioning** – When a new spoke router is added to the network, it does not require additional configuration. Spokes can use a template for their configuration
- **Scalable Deployment** – It provides massively scalable deployment by allowing minimal permanent state and peering on the spoke routers
- **Spoke-to-Spoke Tunnels** – These establish a separate tunnel between the two spoke routers so the traffic does not go through the hub router. Tunnels are created on an on-demand basis and terminated when no longer needed. A spoke router stores the forwarding state of the other spoke router with which it communicates only
- **Multiprotocol Support** - DMVPN supports IPv4, IPv6, and MPLS protocols as the overlay or transport network protocol
- **Multicast Support** – Multicast traffic is allowed to flow on the tunnel interfaces.
- **Standardized Technologies** – DMVPN uses industry-standardized technologies such as GRE, IPSec, and NHRP to build an overlay network

DMVPN was released in three phases. Each phase was built upon the previous phase with additional features. One tunnel interface is required on the router for all three phases, and it should include the endpoints in its network size associated with this tunnel interface. A spoke discovers the IP address of the destination spoke with the help of NHRP.

- Phase 1: Spoke-to-Hub
- Phase 2: Spoke-to-Spoke
- Phase 3: Hierarchical Tree Spoke-to-Spoke

Phase 1: Spoke-to-Hub

The first implementation of DMVPN was Spoke-to-Hub. In this, a tunnel is created only between the spoke and the hub, and traffic from the spoke must go to the hub to reach another spoke. This DMVPN was a zero-touched technology.

Phase 2: Spoke-to-Spoke

DMVPN Phase 2 provides more capabilities than Phase 1. It allows spoke-to-spoke communication by creating a tunnel between them. This tunnel is created on-demand and terminated when no longer required.

It does not support summarization which means next-hop preservation. Consequently, it does not support spoke-to-spoke communication between multiple DMVPN networks.

Phase 3: Hierarchical Tree Spoke-to-Spoke

DMVPN Phase 3 adds more functionality to DMVPN Phase 2 by using NHRP (Next-Hop Resolution Protocol) and a routing table. When a spoke sends a request to the hub to communicate with another spoke, the hub router sends an NHRP redirecting the message to the initiating spoke. This redirect message includes all the information required for the spoke to go to the destination spoke.

GRE/mGRE

Physical networks needed a way to communicate with each other. **Overlay tunneling technologies** came into existence to bridge this communication gap. These technologies are often called overlay because they are virtual (logical) and built over physical (underlay) networks.

Generic Routing Encapsulation (GRE) Tunnels

GRE is a **tunneling protocol** that provides connectivity between physical networks and supports several network-layer protocols. It encapsulates the

data packets from one network and forwards them through the tunnel to the connected network.

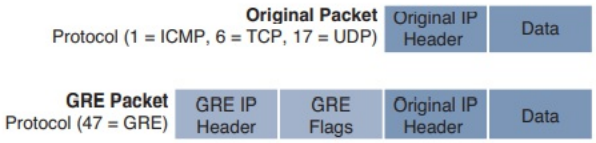

This tunnel is created over an **IP-based network** which is the Internet. There were some non-routable protocols, such as **Internetwork Packets Exchange (IPX)**, a legacy protocol. There was a need to connect the physical networks, and GRE was created to provide transport to this legacy and non-routable protocol over the Internet.

NHRP

Next-Hop Resolution Protocol (NHRP) is an address resolution protocol for the hosts or Non-Broadcast Multi-Access (NBMA) networks such as ATM and Frame Relay. It is defined in RFC 2332. It provides a spoke router with all the necessary information about protocols and NBMA networks to communicate directly with another spoke router.

IPSec

IPSec stands for IP Security. It is an IETF (Internet Engineering Task Force) standard suite of protocols between 2 communication points across the IP network that provide data authentication, integrity, and confidentiality. It also defines the encrypted, decrypted, and authenticated packets. IPsec provides security services that include authentication, data confidentiality, integrity, and anti-replay. It consists of multiple protocols and standards, such as the Internet Security Association and Key Management Protocol (ISAKMP).

Phase 1 Negotiations

The two VPN gateway devices exchange credentials during Phase 1 negotiations, and the devices recognize and negotiate a shared set of Phase

1 settings. The two devices will have a Phase 1 Security Association (SA) once Phase 1 discussions are finished, and this SA is only suitable for a set period. If the two VPN gateways do not finish Phase 2 discussions before the Phase 1 SA expires, they will have to start over with Phase 1.

The gateway endpoints' IKE version determines the Phase 1 negotiation process. IKE authenticates IPSec peers during this phase and negotiates IKE SAs, establishing a secure communications channel for Phase 2 IPSec SA negotiation.

The following steps are included in the first phase of negotiations:

1. The devices agree to utilize the IKE version (IKEv1 or IKEv2). Each device can use IKEv1 or IKEv2, and both devices' IKE versions must be the same.

2. The gadgets pass credentials back and forth. A certificate or a pre-shared key can be used as credentials. Both ends of the gateway must utilize the exact credential mechanism, and the credentials must be identical.

3. The gadgets are able to identify one another. Each device provides a Phase 1 identity, which may be an X500 name, an IP address, a domain name, or other domain-related data. Phase 1 identities for the local and remote devices are supplied in the VPN setup for each device, and the VPN setups must match.

4. For Phase 1 negotiations with IKEv1, VPN gateways choose between Main Mode and Aggressive Mode.

5. The VPN gateways agree upon phase 1 settings.

- Using NAT traversal or not
- Using IKE Keep-Alive or not using IKE Keep-Alive (between Fireboxes only)

- Using Dead Peer Detection or Not (RFC 3706)

NAT Traversal and DPD are always enabled in IKEv2. However, IKE Keep-Alive is not.

6. The VPN gateways agree upon phase 1 Transform settings. IKE negotiations will fail if the settings in the Phase 1 transform on each IPSec device do not match.

The following items can be set in the Phase 1 transform:

- **Authentication** — Authentication is a term that refers to the process of verifying someone's identity (SHA-2, SHA-1, or MD5)
- **Encryption** — The type of encryption technique (DES, 3DES, or AES) and the length of the key
- **SA Life** — SA The time until the Phase 1 Security Association expires is called life.
- **Key Group** — A Diffie-Hellman key group is a key group in cryptography.

Phase 2 Negotiations

Phase 2 negotiations commence once the two IPSec VPN gateways successfully conclude Phase 1 negotiations, and phase 2 discussions are intended to establish the Phase 2 SA. The IPSec Security Association (SA) is a collection of traffic requirements that tells a device what kind of traffic to deliver over a VPN connection and how to encrypt and authenticate it.

The following steps are included in Phase 2 negotiations:

1. The Phase 1 SA is used by VPN gateways to secure Phase 2 discussions. The VPN gateways have agreed on whether or not to deploy Perfect Forward Secrecy (PFS).

The Force Key Expiration setting determines how often VPN encryption keys are refreshed, and the default interval is eight hours. PFS forces the DH computation a second time to prevent SAs from utilizing Phase 1 keys for Phase 2. This means that the keys for Phases 1 and 2 are constantly different, making it harder to crack unless you choose a DH group lower than 14.

We advise you to use PFS in order to protect your data. If you want to use PFS, both VPN gateways must be configured to use the same Diffie-Hellman key groups and have PFS enabled.

2. The VPN gateways have agreed upon a Phase 2 plan. The Phase 2 proposal specifies the algorithm used to authenticate data, the algorithm used to encrypt data, and the frequency with which new Phase 2 encryption keys should be generated.

In a Phase 2 proposal, you can include the following items:

- **Type** — You can choose between Authentication Header (AH) and Encapsulating Security Payload (ESP) for a manual BOVPN. AH and ESP encrypts data and defend against packet tampering and spoofing (replay detection). We recommend you utilize ESP because there are other ways to defend yourself from spoofing. Managed BOVPNs always use ESP, mobile VPNs with IKEv2, mobile VPNs with IPSec, and mobile VPNs with L2TP
- **Authentication** — Authentication ensures that the information received matches the information given precisely. To authenticate IKE messages from each other, VPN gateways can employ SHA-1, SHA-2, or MD5 as the algorithm, and the only secure choice is SHA-2
- **Encryption** — Encryption ensures the privacy of data. You

can choose between DES, 3DES, AES, and AES-GCM; AES and AES-GCM versions are the only safe choices

- **Force Key Expiration** — To ensure that Phase 2 encryption keys are updated regularly, choose a key expiration frequency. By default, the timer is set to 8 hours. The more information an attacker may gather to launch an attack on a Phase 2 encryption key, the longer it will be in use. The Traffic option is not recommended because it creates significant Firebox load, throughput difficulties, packet loss, and frequent, random outages. Most third-party devices are incompatible with the Traffic option

3. Exchange of Phase 2 traffic selections between VPN gateways (tunnel routes). Phase 2 traffic selectors for local and remote VPN gateways can be a host IP address, a network IP address, or an IP address range. One specifies which IP addresses behind the local device are authorized to send traffic over the VPN, and the other specifies which IP addresses behind the remote device are allowed to send traffic over the VPN.

IPsec Rekeying

The rekeying process takes place before SA expiration. IKE SA and IPSec SAs must be **rekeyed** to ensure uninterrupted flow. Rekeying is defined as the generation of a new SA to replace an expired SA well before the expiration date. The process for IKEv2 rekeying with minimal traffic loss is described in RFC 5996. There must be an existing phase 1 SA to rekey Phase 2 (data) tunnels. Idle SA does not have to be rekeyed or just maintained; it should enable idle timeout with crypto IPsec security-association idle-time.

IPsec Tunneling

IPsec is a group of technologies that tunnel data between devices and provide cryptographically secure network communications. Each device in

the VPN has the same IPsec configuration, allowing traffic to flow securely from source to destination between the devices.

IPsec Encapsulation Mode

Tunnel Mode

IPSec protects the entire IP packet, including the IP header and the payload. It uses the entire IP packet to calculate an ESP or AH header and then encapsulates the original IP packet and the ESP or AH header with a new IP header.

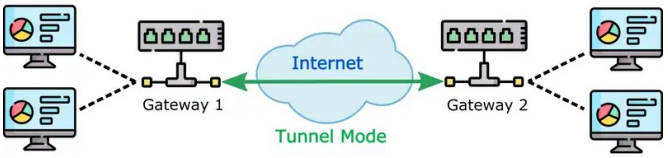

Transport Mode

IPSec only protects the IP payload. It only uses the IP payload to calculate the AH or ESP header and inserts the calculated header between the original IP header and payload. If you look at ESP, an ESP trailer is also encapsulated. The transport mode is usually used to protect communications between hosts or between hosts and gateways.

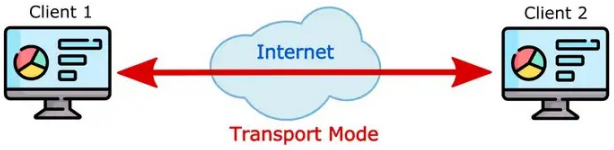

IPsec Security Protocols

IPsec is a collection of methods and protocols for encrypting data sent over open networks like the internet. The IPsec protocols were created by the Internet Engineering Task Force (IETF) in the middle of the 1990s to provide security at the IP layer by authenticating and encrypting IP network packets.

Encapsulating Security Payload and Authentication Header were the first two protocols introduced by IPSecurity for safeguarding IP packets. The former provides anti-replay services and data integrity, while the latter provides data encryption and authentication.

Authentication Header

In RFC 4302, AH is defined. Its services include data integrity and transit security. AH was created to be put into an IP packet to add authentication data and protect the contents from being tampered with.

Header Structure

The AH provides packet authentication and anti-replay services, an important IPSec security mechanism. RFC 2402 defines AH, which uses IP Protocol 51, and the AH can be used in transport or tunnel mode.

Transport mode is typically employed when the client host begins the IPSec communication and protects upper-layer protocols and some IP header fields. The AH is placed after the IP header and before an upper-layer protocol (such as TCP, UDP, or ICMP) or any previously inserted IPSec headers in transport mode.

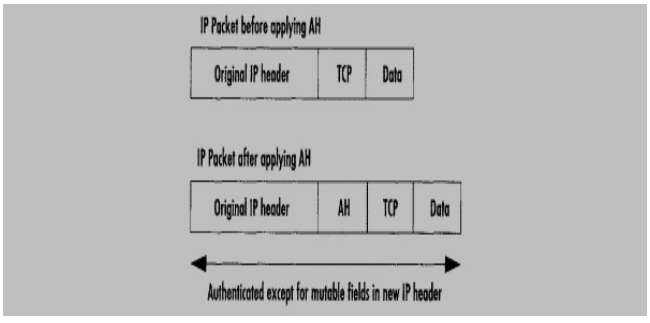

Encapsulating Security Payload

Encapsulating Security Payload provides data confidentiality, integrity, origin authentication, and replay protection. As defined in RFC 4303, ESP encrypts

IP packets to guarantee authentication, integrity, and confidentiality.

Header Structure

The ESP header is usually put after the IP header in an IP network packet. The sequence number, payload data, padding, next header, integrity check, and sequenced numbers are all parts of an ESP header.

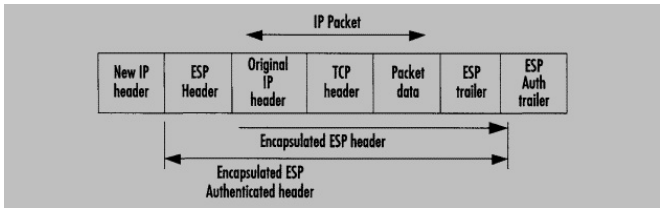

Dynamic Neighbor

Dynamic Multipoint Virtual Private Network is a dynamic protocol that provides a VPN connection to the spoke routers. Since it is a multipoint protocol, it connects multiple client routers. If it were a static protocol, it would be impossible to provide a connection to the hundreds of spoke routers.

It is a dynamic protocol that connects the clients (spoke routers) dynamically with a minimal and simple configuration. When the spoke routers are added to the network, DMVPN lets them in without additional configuration. The spoke routers are **dynamic neighbors** to each other, and these spoke routers use templates for the simplicity of configuration.

Spoke-to-Spoke

It allows spoke-to-spoke communication by creating a tunnel between spoke routers. This tunnel is created on-demand and terminated when no longer required.

In Phase 1, there was an underlying mechanism by which the spoke router was connected to the hub router. The spoke router was configured with the

destination IP address so the encapsulated packets could reach their destination.

Multipoint GRE is used in DMVPN Phase 3, so it relies upon Next-Hop Resolution Protocol (NHRP). **NHRP resolution request messages** and **NHRP redirect messages** are very important to redirect packets from one spoke router to another.

First, packets go to the hub (like in traditional spoke-to-hub packet flow), and then the hub engages the NHRP to find an optimal path to establish a spoke-to-spoke connection.

By the time the spoke router sends the second stream of packets, the hub has found the optimal path for the spoke-to-spoke connection. Now, this path is used as the tunnel between the spoke routers and spoke routers use this tunnel to communicate directly.

Mind Map

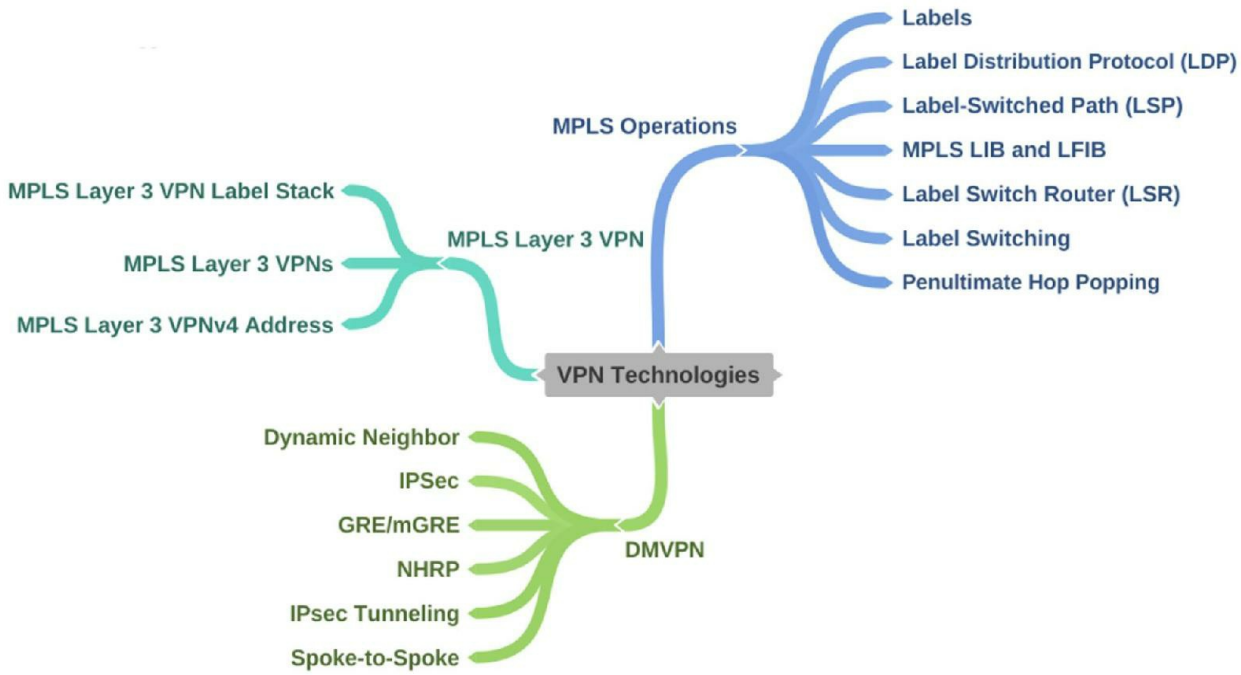
VPN Technologies
MPLS Operations
Labels
Label Distribution Protocol (LDP)
Label-Switched Path (LSP)
MPLS LIB and LFIB
Label Switch Router (LSR)
Label Switching
Penultimate Hop Popping
MPLS Layer 3 VPN
MPLS Layer 3 VPN Label Stack
MPLS Layer 3 VPNs
MPLS Layer 3 VPNv4 Address
DMVPN
Dynamic Neighbor
IPSec
GRE/mGRE
NHRP
IPsec Tunneling
Spoke-to-Spoke

CHAPTER 04: INFRASTRUCTURE SERVICES

Introduction

This chapter focuses on the various reasons for device management and different management tools, such as console/vty access, remote transfer tools like Trivial File Transfer Protocol (TFTP), Hypertext Transfer Protocol (HTTP), Hypertext Transfer Protocol Secure (HTTPS), and Secure Copy Protocol (SCP). It also discusses network management tools like Syslog, SNMP, Cisco IP SLA, Object Tracking, NetFlow, and Flexible NetFlow.

This chapter covers the detection and repair of issues with console and vty access, as well as with remote transfer tools. There are a number of protocols covered, such as Telnet, SSH, TFTP, HTTP, HTTPs, and SCP. The use of and troubleshooting for a number of management tools, such as Syslog, SNMP, Cisco IP SLA, object tracking, NetFlow, and Flexible NetFlow, are also covered in this chapter. Additionally, Cisco DNA Center Assurance is looked at.

Troubleshoot Device Management

A Cisco IOS router can be accessed in several ways for management purposes. You use the console line when using an access server or direct physical access to the device. Additionally, vty lines offer remote communication via Telnet or Secure Shell (SSH), allowing for the management of devices from a distance. Regardless of the technique, you employ for administration, at some point, you will probably find yourself having to fix a device's connection issues to fix another problem that has been brought to your attention. As a result, you might need to resolve one problem before moving on to the next. You will occasionally need to transfer configuration data or IOS images while debugging a problem. Because of this, you must be able to resolve problems with your remote transfers using protocols like TFTP, HTTP(S), and SCP.

This section covers the possible causes of management access to a Cisco IOS router failing, how to identify the issue, and how to resolve it. What to look out for when troubleshooting remote transfers is also covered in this section.

Console Access Troubleshooting

The console port is the default out-of-the-box method for connecting to Cisco routers and switches. When resolving console access issues, keep an eye out for the following:

- Has the terminal program's correct COM port been chosen?
- Are the parameters for the terminal program established properly?
- Is the console authenticated using a line password?
- Does the console require a local username and a password for access?
- Is the console authenticated using a AAA (authentication, authorization, and accounting) server?
- Is the console port being connected with the proper cable and drivers?

vty Access Troubleshooting

Most devices are managed remotely using vty lines, which support remote access protocols like Telnet and SSH. Telnet is not advised since all communication between the management station and the router or switch takes place in plaintext. A malicious user can see all the data sent back and forth if they can capture the packets. The packets will be encrypted if you utilize SSH, ensuring they cannot be read even if captured.

a) Telnet

When troubleshooting Telnet access to a device, keep the following in mind:

- Is it possible to access the remote router's or switch's IP

address?

- Has the line's appropriate transport protocol been defined?
- Is the line set up to request credentials from the user?
- Has a password been provided?
- Does the router/switch have an access control list (ACL) indicating which management stations can access it based on IP address?
- Does port 23 have an ACL blocking it in the path between the client and the device?

b) SSH

The same problems stated with Telnet may occur with Secure Shell (SSH) and a few more. When resolving additional problems with SSH access to a device, take into account the following:

- Is the right SSH version mentioned?
- Was the login command entered correctly?
- Was the key size specified correctly? Is port 22 being blocked by an ACL along the path from the client to the device?

Remote Transfer Troubleshooting

Although IOS and other data are preloaded on Cisco devices, you will likely want to upgrade the IOS image or any other files kept on the device at some point. You can achieve this using a variety of protocols, and TFTP, SCP, HTTP, and HTTPS are some of these protocols.

a) TFTP

TFTP is an unsecured file transfer protocol when utilizing a TFTP server to send files to and from a Cisco device. UDP port 69 is used by TFTP, and it is categorized as an unreliable protocol as a result. Use a TCP-based protocol if you need dependable transport from source to destination. When troubleshooting TFTP problems, take into account the following:

- Make sure the TFTP server has sufficient storage before copying it
- Make sure that the storage location on the Cisco device has enough space before copying data from a TFTP server. Use the show flash command to check the amount of free space present and compare it to the size of the file you wish to copy. A partial copy of a file larger than the available space will be made, and you will receive the error message "buffer overflow - xxxx /xxxx" as a result. The first four x's indicate how many bytes were read from the source file, while the following four x's indicate how many bytes are available on the destination
- Make sure that the Cisco device can connect to the TFTP server
- Look for access lists that might prevent TFTP traffic from the source to the destination
- To specify that a management interface will be used for sourcing TFTP traffic, use the **ip tftp source-interface interface_type interface_number** command if you use a management interface for TFTP traffic

b) SCP

Another choice for copying files from a storage site to a Cisco device is Secure Copy Protocol (SCP). It utilizes Secure Shell (SSH) to offer a protected and verified way to transfer files. Additionally, AAA must be set for the router to determine whether the user has permission to copy.

When troubleshooting SCP problems, keep the following in mind:

- Verify that the device's SSH, authentication, and authorization configurations are valid
- Verify the availability of an RSA (Rivest-Shamir-Adleman) key that can be utilized for encryption
- Verify that AAA is configured correctly and is operational
- Make sure the Cisco equipment has SCP enabled. Use the ip scp server enable command to enable it if it is not already
- Ensure that the copy command is being executed properly
- Confirm that the proper copying credentials, including the username and password, are being used. Verify the server's credentials if you

are utilizing an external authentication server
- Use the debug ip scp command to get additional help resolving SCP problems

a) *HTTP and HTTPS*

You can use either the insecure HTTP protocol, which uses TCP port 80, or its more secure variant HTTPS, which uses TCP port 443, to copy Cisco IOS image files, core files, configuration files, log files, scripts, and more to or from a remote web server. Keep the following in mind when troubleshooting HTTP(S) access for a device:

- Ensure that your Cisco equipment supports the HTTP client. The command display ip http client all can be used to accomplish this. The client is supported if the command is successful
- Verify the web server connection on your router: Utilize the ping command on the Cisco device to ping the web server's IP address or URL
- Verify that the copy command has the correct web server's IP address or URL. You must give a source and destination for the copy command in that order. As an example, the web server, which is indicated with the IP address 10.0.3.8, would be the source for copying data from a web server to flash:

```
copy http://10.0.3.8/cisco_ios_files/c3900-universalk9-mz.SPA.156-3.M6a.bin flash:c3900-universalk9-mz.SPA.156-3.M6a.bin
```

- The web server would be the destination when copying to a web server from flash, as demonstrated here with IP address 10.0.3.8:

```
Copy flash:c3900-universalk9-mz.SPA.156-3.M6a.bin http://10.0.3.8/cisco_ios_files/c3900-universalk9-mz.SPA.156-3.M6a.bin
```

- Check that the copy command specifies the right filename.
- Ensure that the copy command has the proper username and password. User1 is the username in this case, and mypassword is the password:

```
copy http://user1:mypasswrod@10.0.3.8/cisco_ios_files/c3900-universalk9-mz.SPA.156-3.M6a.bin flash:c3900-universalk9-mz.SPA.156-3.M6a.bin
```

- The **ip http client username username** command and the **ip http client password password** command can both be used to specify the authentication credentials. Remember that the login and password used with the **copy** command take precedence over those used with these other commands.
- Verify that the copy command specifies the right port. You can set your web server to use whatever port you like; HTTP uses port 80 by default, and HTTPS uses port 443. The port number being supplied in this example is 8080:

```
copy http://user1:mypasswrod@10.0.3.8:8080/cisco_ios_files/c3900-universalk9-mz.SPA.156-3.M6a.bin flash: c3900-universalk9-mz.SPA.156-3.M6a.bin
```

- Verify that the Cisco device is using the correct IP address to source packets to the web server. If not, the packets may be dropped by an ACL along the way. Using the **ip http client source-interface interface-id** command, you can configure the source IP address
- Ensure that you choose the appropriate protocol, either HTTP or HTTPS. Your URL should start with http if you connect to an HTTP server, and your URL should start with HTTPS if you connect to an HTTPS server
- Use the **debug ip http client all** command for more assistance in troubleshooting HTTP and HTTPS copy problems

Troubleshoot SNMP

You must be able to ping the server from the agent regardless of whether you are using SNMPv2c or SNMPv3. The Simple Network Management Protocol (SNMP) Network Management Server (NMS) cannot access the

data in the Management Information Base (MIB) on the agent if Layer 3 connectivity is not there. Additionally, SNMP uses UDP port 162 for traps and informs and UDP port 161 for general messages. Therefore, SNMP communication between the NMS and the agent is impossible if an ACL blocks certain ports.

SNMPv2c

When troubleshooting SNMPv2c, keep the following in mind:

- Make sure that community strings correspond: The read community string or the read/write community string between the NMS and the agent must match for the NMS to read from or write to the agent. Ensure that servers classified by ACLs are correct: The ACL must precisely identify the server addresses if you are using it to specify which NMS (based on IP address) is permitted to retrieve objects from the MIB.
- Ensure the proper notification configuration: If your agent is set up to transmit traps or informs, you should make sure that:
  - Make sure traps are turned on
  - Verify the host (NMS) IP address is entered correctly
  - Make sure the right SNMP version is mentioned
  - Make sure the appropriate community string is provided

SNMPv3

SNMPv3 provides significant security enhancements over SNMPv2c. It provides enhanced encryption and authentication. When troubleshooting SNMPv3, keep the following in mind:

- **Nesting of users, views, and groups:** Using SNMPv3, you may construct users with authentication and encryption settings nested into groups that specify the servers permitted to read from or write to the objects in the MIB on the agent. SNMPv3 will not operate as intended if the users, views, and groups are not nested.
- **Wrong security level specified:** The three security levels that

SNMPv3 provides are noAuthNoPriv, authNoPriv, and authPriv. The security settings selected for the group, users, and trap sending must coincide with those employed on the server.

- **Incorrect encryption algorithm, hashing algorithm, or passwords defined:** The hashing algorithm and password used for authentication must match; otherwise, the authentication will fail.
- **Incorrect OIDs specified in the view:** The views list the MIB objects that the NMS can access. SNMPv3 will not give the intended results if the incorrect objects are defined.
- **Notification configuration:** If your agent is set up to send traps or informs, ensure that traps are enabled, the host (NMS) IP address is correct, the SNMP version is valid, the security level is correct, and that you specified traps or informs (default is traps).
- **Index shuffling:** Use the **snmp-server ifindex persist command**, which appears as **snmp ifmib ifindex persist** in the running configuration, to stop index shuffling and ensure index persistence during reboots or minor software upgrades

Troubleshoot Network Problems using Logging

Success depends on the integrity and safety of your network and the clients it serves. You can remain on top of any difficulties that arise if you can monitor your network using various tools. However, you must troubleshoot the tools that aid in troubleshooting when the primary tools malfunction or fail to deliver the desired outcomes.

The implementation of logging is necessary for network administrators to gain insight into their network's activity. Content flow, configuration changes, and the installation of new software, to name a few, can all be recorded in these logs and reports. Logging assists in identifying odd network traffic, network device failures, or just keeping track of the traffic types that go through the network. A router can implement logging locally, but this approach is not scalable. Additionally, all the logs on a router will be lost if reloaded. Implementing logging to an external location is crucial as a

result.

Local Logging

Cisco IOS can store Syslog messages in the internal buffer for local logging. These messages can be seen by using the show logging command. With this command, you can enable internal logging and control the buffer size:

```
Router (config) # logging buffered <size>
```

Syslog

Syslog is the industry standard for logging computer messages. It makes it possible to divide the software that creates messages from the systems that store them from the software that analyzes and reports them. It can be utilized for system administration, message debugging, and security auditing. Numerous devices, including routers and receivers across several platforms, support Syslog. The log data from many different types of systems are integrated into the main repository via Syslog. The Syslog packet is in the HEADER MSG format.

Troubleshoot IPv4 and IPv6 DHCP

DHCP for IPv4

The IPv4 address information for a network host is typically assigned using the Dynamic Host Configuration Protocol (DHCP). A DHCP client can get an IP address, subnet mask, default gateway, DNS server, and other kinds of IP addressing information from a DHCP server specifically due to DHCP. The DHCP server may be local to the network, located further in the subnet, or even the same hardware as the default gateway.

Since using DHCP is the most typical method of deploying IPv4 addresses, you must be well-versed in the DHCP procedure and be able to spot DHCP-related problems. This section describes DHCP's functionality and focuses on

how to troubleshoot DHCP-related problems.

DHCPv4 Operations

Your router most likely receives its IP address from your service provider using DHCP if your home has a cable modem, Digital Subscriber Line (DSL), or fiber connection. In addition, the router serves as a DHCP server for the devices in your house. When a PC boots up on a corporate network, it receives its IP address configuration data from a corporate DHCP server. The figure depicts the message exchange that occurs as a DHCP client receives IP address information from a DHCP server (the Discover, Offer, Request, Acknowledgment [DORA] process).

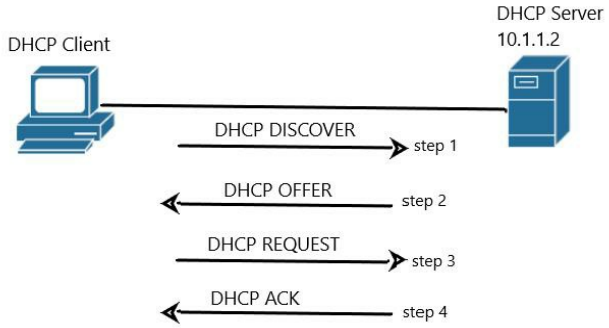

Router Acting as a DHCP Client

A router may perform the function of a DHCP client in addition to a DHCP relay agent. In particular, a DHCP server may provide an IP address to a router's interface. The following figure depicts a router as a DHCP client, receiving an IP address from a DHCP server for its Fast Ethernet 0/1 interface.

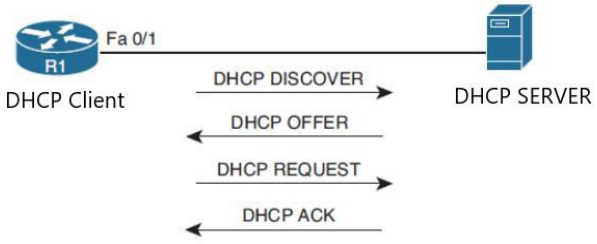

Router Acting as a DHCP Server

A multilayer switch or router can also operate as a DHCP server. A router serving as a DHCP server is shown in the following figure.

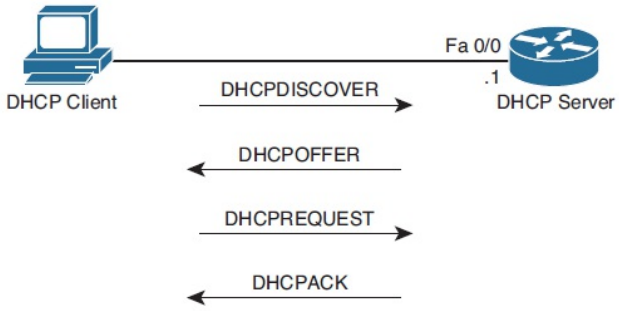

Common DHCP Troubleshooting Issues

Consider the following potential problems while troubleshooting what you think might be a DHCP issue:

A router that does not forward broadcasts: A router does not automatically forward broadcasts, such as DHCPDISCOVER broadcast signals. Therefore, a router must be explicitly configured to operate as a DHCP relay agent if the DHCP client and server are on different subnets.

DHCP pool out of IP addresses: A DHCP pool has a limited amount of addresses. New DHCP requests are denied when a pool is empty.

Misconfiguration: A DHCP server's configuration may be flawed. Inappropriate exclusion of addresses that are statically assigned to routers or DNS servers, for example, or an incorrect range of network addresses distributed by a certain pool.

Duplicate IP addresses: A client may receive an IP address from a DHCP server that is already statically assigned to a different host on the network. The DHCP client and the host statically configured for the IP address may have connectivity problems due to these multiple IP addresses.

Redundant services are not interacting: For redundancy, some DHCP servers coexist with other DHCP servers. These DHCP servers must interact with one another for redundancy to work. The DHCP servers give their clients overlapping IP addresses if the interserver communication breaks down.

The "pull" functionality of DHCP: A DHCP client asks a DHCP server for an IP address when it needs one. After the client receives an IP address, the DHCP server cannot start a change in the client's IP address. In other words, while the DHCP server can transmit information changes to the DHCP client, the DHCP client must pull information from the DHCP server.

Interface not configured with IP address in DHCP pool: To function as a DHCP server, a router or multilayer switch must have an interface set with an IP address that is a component of the pool or subnet it is dispersing IP addresses for. Only clients that can be reached through that interface receive the addresses in the pool from the router. This guarantees that the clients and the router interface are on the same subnet. Notably, this is not the case if a relay agent is DHCP message-forwarding between the client and the router acting as the DHCP server. In that case, the DHCP server does not need to have an IP address on an interface that is a member of the address pool it is managing.

DHCP Troubleshooting Commands

Sample output from the show ip dhcp conflict command is seen in the following snippet:

```
R1# show ip dhcp conflict
IP address Detection method Detection time
172.16.1.3 Ping Oct 15 2018 8:56 PM
```

The output shows that the router used ping to find a duplicate 172.16.1.3 IP address on the network. After fixing the network's duplicate address issue, you clear the information displayed by giving the **clear ip dhcp conflict *** command.

DHCP for IPv6

It is not scalable to manually allocate IP addresses (IPv4 or IPv6). DHCP offers a dynamic addressing solution for IPv4. You can choose between Stateless Address Auto-Configuration (SLAAC), stateful DHCPv6, or stateful DHCPv6 when using IPv6. This section examines potential problems with each and how to fix them.

SLAAC

Without the aid of a DHCPv6 server, SLAAC enables a device to set its own IPv6 address, prefix, and default gateway.

Stateful DHCPv6

Except for their IPv6 address, prefix, and default gateway, a device can only learn these things using SLAAC. The devices in a modern network could also require data from servers supporting the Network Time Protocol (NTP), domain names, DNS, and Trivial File Transfer Protocol (TFTP). Use a DHCPv6 server to distribute IPv6 addressing information along with all optional information.

Stateless DHCPv6

SLAAC and DHCPv6 are combined to create stateless DHCPv6. In this scenario, clients automatically calculate the IPv6 address, prefix, and default gateway using a router's RA. A flag contained in the RA instructs the client to

request additional non-addressing information from a DHCPv6 server, such as the DNS server or TFTP server address.

DHCPv6 Operation

Like IPv4, DHCPv6 uses a four-step negotiation process. But DHCPv6 uses the subsequent messages:

Step 1: SOLICIT: A client sends this message to locate DHCPv6 servers using the multicast address FF02::1:2, which is the address used by all DHCPv6 servers.

Step 2: ADVERTISE: In response to SOLICIT messages, servers send out unicast ADVERTISE signals that provide clients with addressing details.

Step 3: REQUEST: After verifying the addresses and other parameters, the client sends this message to the server.

Step 4: REPLY: The server completes the procedure with this message.

A complete list of DHCPv6 message types you can come across while troubleshooting a DHCPv6 issue is provided for your reference.

SOLICIT: A client sends this message as a petition to find a DHCPv6 server.

ADVERTISE: In response to a SOLICIT, a DHCPv6 server transmits this message to announce its availability.

REQUEST: This message is being delivered from a client to a specific DHCPv6 server as a request for IP configuration information.

CONFIRM: A client sends this message to a server to confirm that the given address is still suitable.

RENEW: To extend the validity of the addresses assigned, a client sends this

message to the server that assigned the address.

REBIND: When a RENEW request is not answered, a client can extend the lifetime of an issued address by sending a REBIND message to a server.

REPLY: A server delivers this message to the client with the assigned address and configuration parameters in response to a SOLICIT, REQUEST, RENEW, or REBIND message received from a client.

RELEASE: A client sends a server this message to inform it that the designated address is no longer required.

DECLINE: This message is sent by a client to a server to inform them that the allocated address is already in use**.**

RECONFIGURE: When a server has new or updated information, it will deliver this message to a client.

INFORMATIONREQUEST: A client will use this message to ask a server for more configuration data when no IP address is required.

RELAY-FORW: Relay agents utilize this message, RELAY-FORW, to forward communications to the DHCP server.

RELAY-REPL: A DHCP server returns this message to the relay agent.

DHCPv6 Relay Agents

The DHCP server has been a part of the same local network in every DHCPv6 example up to this point. The DHCP server is typically placed in a different network, which poses a problem in most networks. Notice a link-local scope multicast address in the SOLICIT message's multicast address. Beginning with FF02. Because of this, the multicast stays within the local network, and the client cannot connect to the DHCPv6 server.

Troubleshoot Network Performance Issues using IP SLA

You can test network availability and evaluate network performance with the aid of Cisco IOS IP SLA by carefully crafting reliable, consistent probes (simulated traffic). How you set up the probe greatly impacts how much information you can obtain. Measurable variables include packet loss, one-way delay, response times, jitter, network resource availability, application performance, server response times, and even audio quality.

The components of IP SLA are an IP SLA source (which transmits the probes) and an IP SLA responder (which replies to the probes). Both are not always required, though. Only the IP SLA source is needed constantly. The IP SLA responder is required when collecting extremely precise statistics for services that are not provided by any specific destination device. The responder can provide the source with precise measurements in response while accounting for its probe processing time.

Enterprise customers can verify service levels, outsourced service level agreements, and network performance for new or existing IP services and applications using IP SLAs. Service provider customers can measure and provide service level agreements using IP SLAs. IP SLAs give extremely accurate, precise service level assurance measurements using unique service level assurance metrics and methods.

Statistics like delay, packet loss, jitter, packet sequence, connectivity, path, server response time, and download time can be tracked inside the Cisco device and saved in both CLI and SNMP MIBs, depending on the individual IP SLAs operation. A URL web address, a VPN routing/forwarding instance (VRF), a source and destination IP address, User Datagram Protocol (UDP)/TCP port numbers, a Type of Service (ToS) byte (including

Differentiated Services Code Point [DSCP] and IP Prefix bits), and other configurable IP and application layer options are all included in the packets.

IP SLA operations gather the following performance metrics:

- Delay (both round-trip and one-way)
- Packet loss (directional)
- Path (per hop)
- Server or website download time
- Jitter (directional)
- Packet sequencing (packet ordering)
- Connectivity (directional)
- Voice quality scores

Object Tracking Troubleshooting

You may control what happens as a result of object tracking dynamically, whether the outcome of the tracking object is up or down. For example, if a static route is added to an object while the object is up, the route is added to the routing database; if the object is down, the route is not put to the routing table. A First-Hop Redundancy Protocol (FHRP) enables you to change the priority based on the item's state. For example, if the tracking object's status is down, the FHRP priority is reduced.

Using Object Tracking, you may monitor IP routes, IP SLA instances, interfaces, and collections of objects. For example, you can follow an IP SLA instance, if it uses ICMP echoes. In the event of an echo failure, the IP SLA instances also fail, and the tracking object becomes inoperative. If the tracking object is linked to FHRP, the priority is reduced; if it is tied to a static route, the static route is removed from the routing table.

Troubleshoot NetFlow

You can gain a great deal of understanding of your network traffic patterns with Cisco IOS NetFlow. Many businesses provide NetFlow collectors,

software programs that can take the raw NetFlow data that is kept in the local device's cache and turn it into useful graphs, charts, and tables illuminating traffic patterns.

NetFlow can be used to discern between various traffic flows. A flow is a collection of packets having common header data, including source and destination IP addresses, protocol and port numbers, and details from the Type of Service (ToS) field. The packets also enter the device through the same interface. NetFlow can record the quantity of packets and bytes seen in each flow. This information is stored in a flow cache that is located in the router's memory.

The NetFlow functionality can be used independently on a single router, and the CLI can be used to see flow cache data. Because you can watch flows being produced as packets reach a router, such a standalone configuration may be helpful for troubleshooting. Instead of using a standalone NetFlow implementation, you can export the data from a router's flow cache to a NetFlow collector, software running on a PC or server in the network. The NetFlow collector's analysis software can generate reports with traffic statistics once it has collected flow data over time.

Consider the following when troubleshooting NetFlow:

- Traffic direction
- Interface
- Export destination
- Export source
- Version

NetFlow Version5

The most used NetFlow version5, or traditional NetFlow, supports Autonomous System (AS) reporting and a few extra fields. When a flow

enters an interface (i.e., when it is inbound), it is calculated, and outbound traffic is reported using inbound flows from the other interfaces. As a result, it is generally recommended that NetFlow v5 be enabled on all of the device's interfaces; otherwise, outbound consumption on some interfaces might be underestimated. Most NetFlow collection and network traffic reporting packages can easily decode the packet format because it is always fixed and consistent.

NetFlow Version9

The Flexible NetFlow technology is known as NetFlow version 9, which is the format for NetFlow flow records. The NetFlow Version 9 format is unique in that it is based on templates. A flexible flow export with user-defined key and non-key fields is provided through templates. It can keep track of various IP packet metadata that is not feasible with traditional NetFlow. This format offers the adaptability required to accommodate additional fields and record kinds. Custom fields like MPLS labels, IPv6 traffic, NBAR protocols, Multicast IP traffic, VLAN ID, and real-time media flow performance are all supported by flexible NetFlow.

Flexible NetFlow

Flexible NetFlow elevates NetFlow by enabling you to adjust the traffic analysis parameters to meet your unique needs. This indicates that there are more settings to check while troubleshooting. You must be able to confirm the flow records, flow monitors, flow exporters, and interface configurations when troubleshooting Flexible NetFlow.

Troubleshoot Network Problems using Cisco DNA Center Assurance

A network's overall health is crucial to its continued success. Thousands of devices, such as routers, switches, wireless LAN controllers, and wireless

access points, may be a part of your network. You cannot keep up with the needs of today's networks by troubleshooting traditionally. It is crucial to have the capacity to assess the overall condition of the network and see any potential problems that need to be fixed.

With Cisco DNA Center, the command and control center for Cisco DNA, you can quickly configure and provision your devices in minutes. Using Artificial Intelligence (AI) and Machine Learning (ML), you can monitor, diagnose, optimize your network, and enhance your operational processes by integrating third-party systems.

One element of Cisco DNA Center is Cisco DNA Center Assurance. Due to proactive monitoring provided by Cisco DNA Center Assurance, you can anticipate issues more immediately and gain information from clients, network devices, network applications, and network services. You will be able to guarantee that implemented policies and configuration changes result in the required business outcomes, give users the experience they want, and spend less time troubleshooting and more time innovating.

Overall Health Page

The Overall Health page is your first helpful troubleshooting resource. At the top of the Cisco DNA Center website, select Assurance to go to this page. This page summarizes the state of the environment's networks and users, and it can be displayed using data from the past three hours, a day, or a week's worth of data. By selecting the Hide/Show button, you can see both hierarchical sites/building maps and health maps. You can obtain an overall network/client score depending on the condition of the network devices and the wired and wireless clients in the Network Devices area and the Client area. The concerns that need to be resolved are listed in the Top 10 Issues

section at the bottom of the page; by clicking on them, you may learn more about the issues, their effects, and possible solutions.

Network Health Page

The Network Health page is another helpful troubleshooting tool you may use in Cisco DNA Center Assurance by choosing Health and then Network. The Network Health page can determine how well the network and the connected devices are doing. You can see the proportion of healthy devices in your entire network depending on categories, including access, core, distribution, router, and wireless. Over the last 24 hours, information can be displayed in 5-minute increments.

Client Health Page

You can see the Client Health page by choosing Health and then Client in Cisco DNA Center Assurance. An overview of the network's and its client devices' operational status is provided on these pages. Any problems with the devices are mentioned, and potential fixes are suggested.

Device 360 and Client 360

The next two essential diagnostic tools are Device 360 and Client 360. By clicking on them, any clients or devices can be chosen to access these features. These features enable you to go down into the device or client and present data about the topology, throughput, and latency from multiple times and applications, giving you a full insight of the performance of the specific device or client over a specified time period. With these tools, you have immediate access to granular troubleshooting.

Path Trace

Path Trace is the next DNA Center Assurance troubleshooting tool that will simplify your life. This functionality is the ping and traceroute you have always wanted. With Path Trace, you can visually view the route that

programs and services running on a client will travel across every network device to get where they are going (a server, for example). You can use this application to quickly complete several troubleshooting activities that would take you five to ten minutes to complete at the command line.

Network Time Travel

Network Time Travel is another fantastic Cisco DNA Center Assurance troubleshooting tool. Instead of attempting to duplicate a network issue, you can use this program to travel back in time and see what caused it. There are numerous timeline views available throughout DNA Center. In Device 360 and Client 360, you may get timeline views for the overall network health, the health of network devices, the health of client devices, and the health of individual devices and clients.

Global Issues Page

You can access all open, resolved, and ignored issues on the Global Issues page.

All Issues Option

Finally, choose the All Issues option from the Issues drop-down, as shown in Figure 4-20, if you want to see every issue that Cisco DNA Center Assurance can monitor in your environment. There are the following categories:

Onboarding: Used to pinpoint problems with wired and wireless client onboarding.

Connectivity: used to identify issues with tunnels, routing protocols (OSPF, Open Shortest Path First, BGP), and other aspects of network connectivity.

Connected: Used to pinpoint client-related problems.

Device: Used to identify problems with the device, including CPU, RAM, fans,

and other components.

Availability: Used to pinpoint any problems with the availability of access points, wireless LAN controllers, and other devices.

Utilization: Used to identify problems with radios, wireless LAN controllers, access points, and other devices.

Application: Used to pinpoint difficulties with the application experience.

Sensor test: Used to pinpoint any global sensor problems.

Mind Map

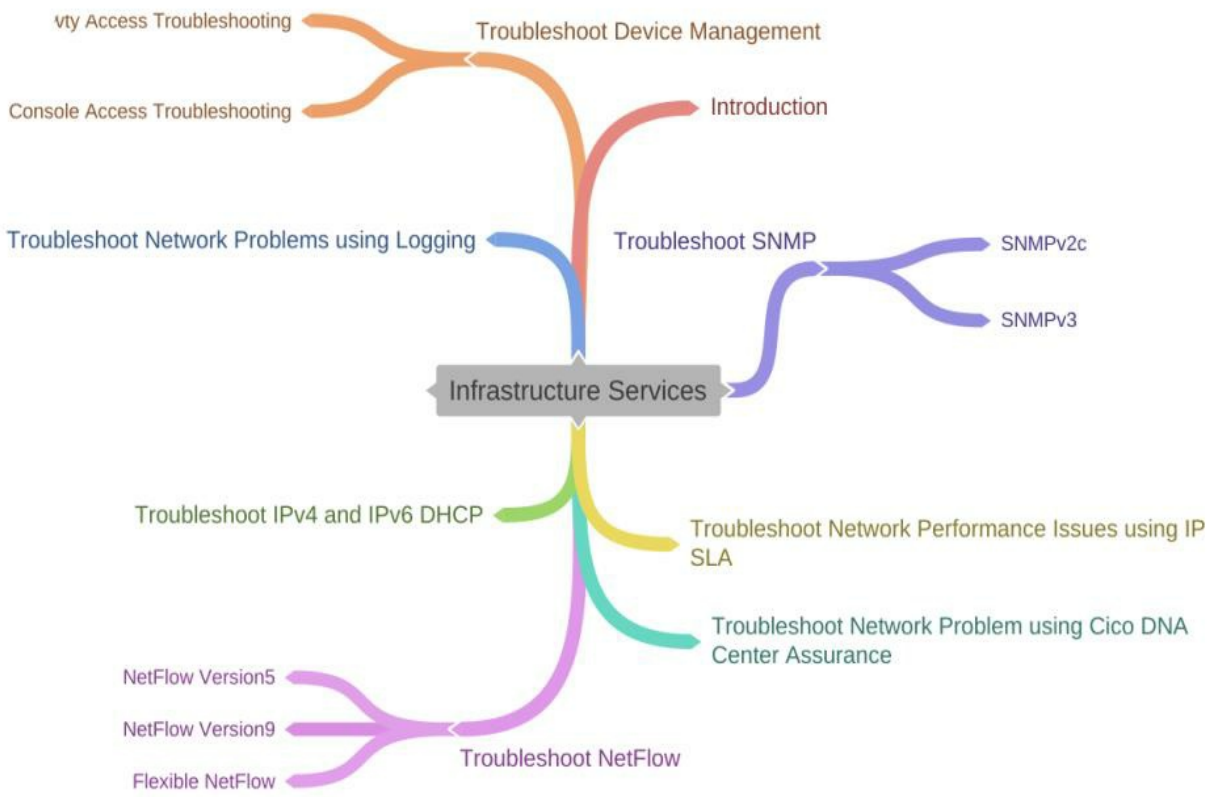
vty Access Troubleshooting
Console Access Troubleshooting
Troubleshoot Device Management
Introduction
Troubleshoot Network Problems using Logging
Troubleshoot SNMP
SNMPv2c
SNMPv3
Infrastructure Services
Troubleshoot IPv4 and IPv6 DHCP
Troubleshoot Network Performance Issues using IP SLA
Troubleshoot Network Problem using Cico DNA Center Assurance
NetFlow Version5
NetFlow Version9
Flexible NetFlow
Troubleshoot NetFlow

CHAPTER 05: INFRASTRUCTURE SECURITY

Introduction

This chapter focuses on the Infrastructure Security of the system or network. It describes the local database, a RADIUS server, and a TACACS+ server as tools for locating and resolving AAA-related problems. In this chapter, we will also be discussing how to troubleshoot router security features, such as how to use Cisco Integrated Services Routers (ISRS) and Cisco Aggregated Services Routers (ASRS) to protect against threats like malware, interruptions, and defiance of authority attacks while maintaining high levels of system and body office performance. Therefore, it is essential to be aware of this crucial characteristic.

You will learn about Control Plane Policing troubleshooting, which looks at CoPP and the factors you should consider when troubleshooting CoPP-related problems. At last, there is the IPv6 First Hop Security section, which goes over security measures for the IPv6 First Hop, including source guard, ND inspection/snooping, RA guard, and DHCP guard.

Troubleshoot Device Security using IOS AAA

The AAA framework offers authentication, authorization, and accounting when securing the management plane. Authentication is represented by the first A in AAA. Authentication aims to identify and confirm the user using something the person knows, has, is, or is something they are. The authority is represented by the second A in AAA. The authenticated user is allowed to do what authorization is all about deciding and managing. Accounting is represented by the last A in AAA. Accounting is gathering data for billing, auditing, and reporting.

The router can be set up using the local AAA server capabilities so that the user authentication and authorization attribute now available on the AAA servers are also available locally on the router. An attribute, such as a

subscriber profile or user database, might supplement the current architecture. A local AAA server provides access to the entire dictionary of attributes supported by Cisco IOS. For several functions, the local database could be the fallback option.

Key Points when Troubleshooting Cisco IOS AAA Authentication

When troubleshooting Cisco IOS AAA authentication, take into account the following:

- AAA needs to be enabled
- The local username and password database or a AAA server like RADIUS or TACACS+ are both used by AAA
- The authentication techniques are described in a method list
- The method list service is wrong
- The lines are subjected to AAA method listings
- The AAA server must be reachable from your router
- The proper preshared key must be entered when configuring the router
- It is necessary to configure the proper accounting and authentication ports
- The AAA server has to be configured with usernames and passwords
- The correct AAA server IP addresses must be in the AAA server group
- The user can log in but cannot run any commands
- Client IP address set up on the AAA server

Troubleshoot Router Security Features

Most network infrastructure comprises switches and routers, both of which are open to assault. Although we often hear about Distributed Denial of Service (DDOS) or mass Denial of Service (DOS) assaults, the network itself

poses as a great risk since data cannot move without it. Although network infrastructure is essential, we also need to defend the networking hardware against intrusion.

Access Control List

Packet filtering is done using Access Control Lists (ACLs) to manage which packets go where in a network. To limit network traffic, control user and device access, and stop traffic from leaving a network, packet filtering helps to offer security. IP access lists permit dynamic, temporary user access over a firewall while lowering the likelihood of spoofing and denial-of-service attacks.

IPv4 Access Control Lists (standard, extended, time-based)

Switches from the Cisco MDS 9000 Family may direct IP form 4 (Ipv4) traffic between Ethernet and Fiber Channel interfaces. To do this, the oddity of the IP static directing course activity between VSANS and every VSAN must be connected to a different IPv4 sub-network. The following services are provided to system administration frameworks (NMS) by each switch in the Cisco MDS 9000 Family:

- IP transmitting on the front board of the administrator modules' out-of-band Ethernet interface (mgmt0)
- Utilizing IP over Fiber Channel (IPFC) capabilities, IP is sent through the in-band Fiber Channel interface. Using exemplification techniques, IPFC specifies how IP edges could be delivered via Fiber Channel. IP casings are typed into Fiber Channel outlines for NMS data to pass via the Fiber Channel system without using an overlay Ethernet system
- IP Routing: (static routing and default directing)- Static routing can be used to set up a default course if your configuration does not call for an outer switch

IPV4 ACLs:

IP addresses or upper-layer protocols may be subject to a series of sequential permit and deny conditions called IP access lists. Routers employ access control lists to identify and manage traffic.

The first three octets of the provided IP address, in this case, 192.168.1, must match according to the mask address 0.0.0.255 to allow IP traffic. The router can ignore the last octet of the filtered IP address if it contains the value 255.

Because every access list has an implicit deny any at the end, if you add an access list to an interface without at least one permit statement, the interface will be essentially shut off.

VTYs (Router command prompts) are accessible from the 192.168.1.0/24 netblock but not from any other location due to the following statements:

```
RTA(config)#access-list 1 permit 192.168.1.0 0.0.0.255
RTA(config)#line vty 0 4
RTA(config-line)#access-class 1 in
```

Standard ACLs:

Standard ACLs regulate traffic by comparing IP packet source addresses to the addresses specified in the ACL.

This is how a standard ACL's command syntax looks.

```
access-list access-list-number {permit|deny}
{host|source source-wildcard|any}
```

Extended ACLs:

Extended ACLs regulate traffic by comparing the IP packets' source and destination addresses to those specified in the ACL. The access-list-number in any software release might range from 100 to 199. Extended ACLs start using different numbers in Cisco IOS Software Release 12.0.1. (2000 to 2699). Expanded IP ACLs is the name given to these extra numbers.

The configuration for an extended ACL example is provided below. Keep in mind that the www is a TCP protocol.

```
access-list 100 deny tcp host 10.0.0.2 host 10.0.1.2 eq www
access-list 100 permit ip any any
interface fastEthernet 0/0
ip access-group 100 in
```

The access list is applied to the interface fe 0/0 by the command "ip access-group 100 in," as you can see.

IP Named ACLs:

Names will be used in place of numbers to identify the standard and extended ACLs.

For IPs named ACLs, the command syntax looks like:

```
ip access-list {extended|standard} name
```

Time-Based ACLs:

With the ability to control access based on time and date, time-based ACLs offer the granular execution of security policies. They obtain the time from the router's system clock; therefore, the Network Time Protocol (NTP) configuration is necessary to guarantee accurate time. In particular, the Cisco router needs to be set up to synchronize with the NTP server, which

gives it accurate time, ensuring that the time-based ACLs you defined take effect as scheduled.

Time-based ACLs can be configured using the time-range command, which is used to provide the time window (which can be recurring or a single occurrence that occurs just once) during which the ACL statement is valid.

Router(config)# time-range time_range_name

After entering this command, you enter ACL sub-configuration mode, where you can choose a one-time-only (absolute) or recurring (periodic) kind of time range.

Absolute: You can specify a beginning time, an ending time, or both to specify a single time period for which the time range is valid.

Periodic: The word periodic designates a recurring time frame for which the time range is appropriate.

Router(config-time-range)# absolute [start_time start_date] [end_time end_date]
Router(config-time-range)# periodic day_of_the_week hh:mm to [day_of_the_week] hh:mm

The time range parameter must be included in the ACL statement to activate the time ranges you have created:

Router(config)# access-list {100-199} {permit | deny} protocol source-addr [source-mask] [operator operand] destination-addr [destination-mask] [operator operand] [established] [log | log-input] [established] [time-range name_of_time_range]

The access list statement will be processed only when the router's time falls within the defined window

IPv6 Traffic Filter

Traffic filtering is a technique used to improve network security by filtering network data according to various parameters. The functionality of an IPv6 standard access control list is identical to an IPv4 standard access control list. These access lists identify the traffic that is banned and the traffic that is routed at the router interface. They also enable filtering based on the source and destination IP addresses for outbound and inbound traffic to a specific interface. Every access list ends with an implicit deny statement. The command ipv6 access list with the permit and deny keywords in the global configuration mode defines the IPV6 ACL and specifies its permit and deny conditions.

Unicast Reverse Path Forwarding (uRPF)

A network security component, uRPF aids in limiting or even completely eliminating spoofed IP packets. An ingress packet's source IP address is examined to see if it is valid in order to do this. The packet will be forwarded if it is valid. The packet will be thrown out if it is invalid. For uRPF to function, CEF (Cisco Express Forwarding) must be activated on the IOS device.

Three different operating modes for uRPF are strict, loose, and Virtual Routing and Forwarding (VRF). You can identify if a packet is valid or invalid depending on the mode you select:

- **Strict:** In strict mode, the router examines the packet's source IP address and records the ingress interface. It next scans the routing table to find the interface (other than a default route) that would be used to connect to the packet's source IP address. The packet is valid and forwarded if the interface is the exact same interface on which it was received and is not the default route. The packet is dropped if the interface is a different interface

- **Loose:** In loose mode, the router just looks at the packet's source IP address. The next step is checking the routing table to see whether an interface (other than a default route) can connect to the packet's reported source IP address. The packet is forwarded if there is one and it is not the default route. Otherwise, the packet gets thrown away
- **VRF:** The sole difference between VRF and loose mode is that only interfaces that are part of the same VRF as the interface on which the packet was received are examined.

Troubleshooting Control Plane Policing (CoPP)

CoPP changes depending on the platform and IOS version. The following actions must be taken when configuring CoPP:

- To identify the traffic, create ACLs
- To define a traffic class, create class maps
- To specify a service policy, create policy maps
- To the control plane, apply the service policy

You should keep an eye out for problems with the application of the service policy, the class maps, the policy maps, and the ACLs when troubleshooting.

1. *Creating ACLs to Identify the Traffic*

With CoPP, ACLs are utilized to identify traffic. The traffic becomes the target of the policy action after it has been matched. As the basis or critical component of CoPP, defining the ACLs is the most crucial phase in the CoPP process. The traffic will not match if the ACL is incorrectly formed, so the policies will not be correctly implemented.

2. *Create class maps to define a traffic class*

Three elements make up a traffic class, defined by class maps. The packets that make up the class are first identified by a name, followed by one or more match commands, and finally, evaluation guidelines for the match command.

3. *Create policy maps to define a service policy*

CoPP employs policy maps to link the traffic class (as defined by the class map) with one or more policies in order to establish a service policy. The trinity consists of a name, a traffic class, and a policy.

When debugging policy maps, keep the following things in mind:

- Order of operations
- Class Map
- Policy
- Default Class
- Case

4. *Apply the service policy to the control plane*

The proper interface must be specified before attaching the service policy.

You should take the following into account while troubleshooting how the service policy is being applied:

- **The correct interface:** The control plane interface is the sole interface to which CoPP can be applied, making it simple to troubleshoot: Either it is applied or not. The command indicated in the example below shows policy-map control-plane [input | output], which can be used to confirm this
- **Direction:** CoPP can be used on packets coming into or out of the control plane interface. Therefore, it is necessary to provide the proper direction. You provide input for incoming packets and output for outgoing packets. Additionally, the show policy-map control-plane output can be used to confirm direction. Not all versions allow for output CoPP; for those that do, you must make sure that the ACLs and class maps are being used to classify the proper traffic. For example, output CoPP is generally used for replies sent due to a previously received packet in BGP, OSPF (Open Shortest Path First), and EIGRP. It would be error and informative reply messages for ICMP. You would deal with replies or traps for Telnet, SSH (Secure

Shell), HTTP (Hypertext Transfer Protocol), or SNMP (Simple Network Management Protocol). The expected outcome will not be obtained if the ACL and class map are not properly set up for the replies

- **Case:** The names of policy maps are case-sensitive. Make sure the names properly match before adding a policy map to the control plane interface

Describe IPv6 First Hop Security Features

Different features of IPv6 FHS (First Hop Security) protect IPv6 on L2 links.

You could assume that the first "hop" refers to the first router, but that is not the case. These functions pertain to switches, specifically the switch positioned between your endpoints and the first router.

Networks are safeguarded by the Cisco IPv6 First Hop Security (FHS) solution, which reduces these threats and configuration problems. It solves vulnerabilities in IPv6 link operations and scalability problems in large Layer2 domains. You get a powerful defense against attack methods and easily accessible tools that exploit weaknesses.

Router Advertisement (RA) Guard

A function called RA Guard examines RAs and can block RAs from coming from unauthorized devices. Routers broadcast their presence on links via RAs. Some RAs might turn out to be undesirable or "rogue"; you would not want them on the network. These undesirable RA messages can be rejected or blocked using RA Guard. The only way to enable RA Guard on an individual interface-by-interface basis is to apply the policy to the interface using the **ipv6 nd raguard attach-policy [policy-name [vlan add | except | none | remove |all vlan [vlan1, vlan2, vlan3...]]** command. RA Guard requires a policy to be defined in RA Guard policy configuration mode.

DHCPv6 Guard

Dynamic Host Configuration Protocol (DHCP) snooping for IPv4 is a feature quite similar to DHCPv6 Guard. It is designed to ensure that malicious DHCPv6 servers cannot assign addresses to clients, reroute client traffic, or starve the DHCPv6 server and launch a Denial-of-Service attack. When it comes to IPv6, DHCPv6 Guard can prevent replies and advertisement messages sent by unapproved DHCPv6 servers and relay agents. On an interface-by-interface basis, DHCPv6 Guard is enabled by applying the policy to the interface with the **ipv6 dhcp guard attach-policy [policy-name [vlan add | except | none | remove | all vlan [vlan1, vlan2, vlan3...]]** command. A policy must be set up in DHCP Guard configuration mode for DHCPv6 Guard to function.

Binding Table

The IPv6 neighbors connected to a device are listed in the binding table database. It includes details like the prefix binding, IPv4 or IPv6 address, and link-layer address. The data in this table is used by other IPv6 First-Hop Security features to prevent spying and redirect attacks.

IPv6 Neighbor Discovery Inspection/IPv6 Snooping

The binding table for stateless autoconfiguration addresses is learned and populated via the IPv6 neighbor discovery inspection/snooping capability. As it examines ND (Neighbor Discovery) messages, it adds any valid bindings to the binding table and discards any messages that do not have valid bindings. A message with a confirmed IPv6-to-MAC mapping is considered a valid ND message.

Source Guard

A Layer 2 snooping interface feature called IPv6 Source Guard is used to confirm the source of IPv6 traffic. IPv6 Source Guard can block traffic from unidentified sources when it arrives on an interface. A source must be listed

in the binding table for traffic to come from that source, and the source is identified and added to the binding table either by ND inspection or IPv6 address gleaning.

Mind Map

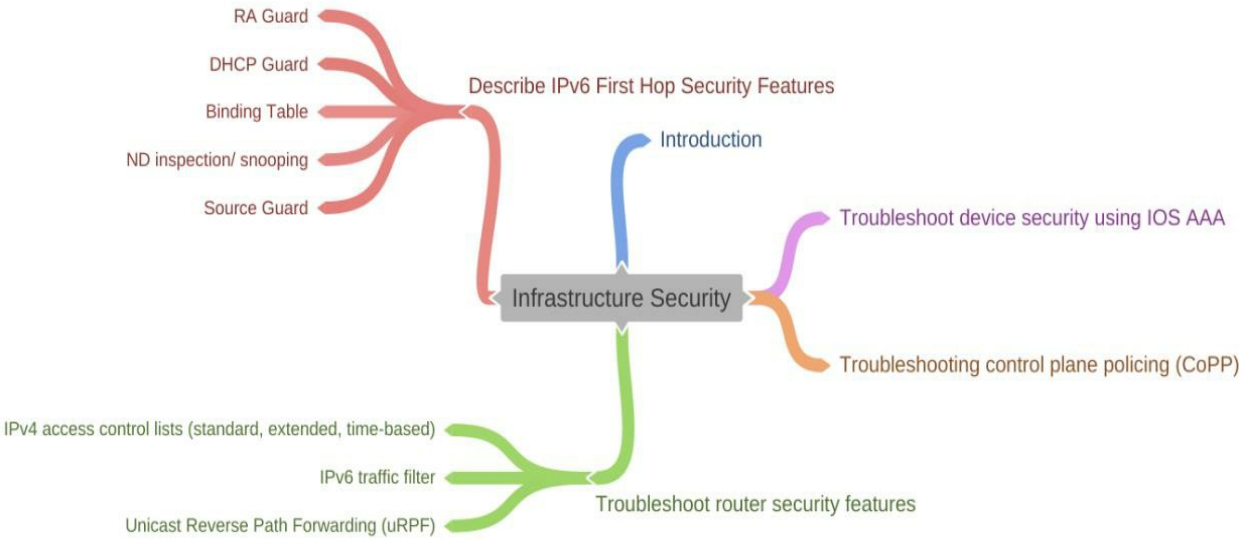
RA Guard
DHCP Guard
Binding Table
ND inspection/ snooping
Source Guard
Describe IPv6 First Hop Security Features
Introduction
Troubleshoot device security using IOS AAA
Infrastructure Security
Troubleshooting control plane policing (CoPP)
IPv4 access control lists (standard, extended, time-based)
IPv6 traffic filter
Unicast Reverse Path Forwarding (uRPF)
Troubleshoot router security features