https://www.leads4pass.com/300-410.html

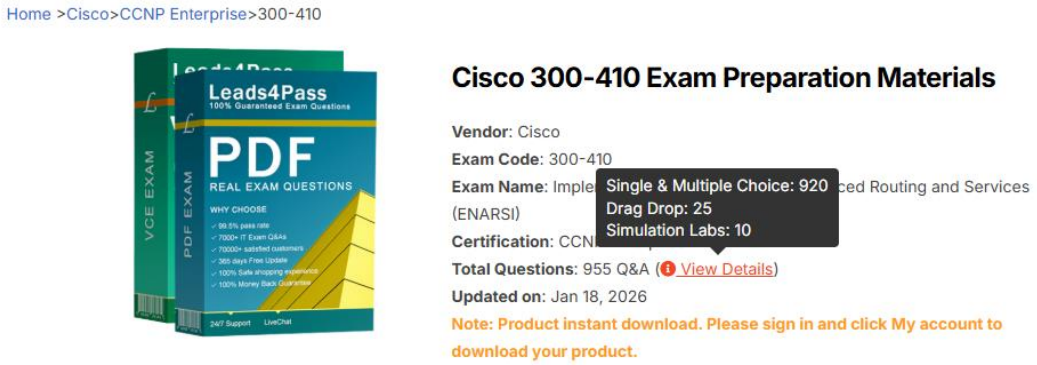

Question 1:
What are two functions of IPv6 Source Guard? (Choose two.)

A. It works independent from IPv6 neighbor discovery.

B. It denies traffic from unknown sources or unallocated addresses.

C. It uses the populated binding table for allowing legitimate traffic.

D. It denies traffic by inspecting neighbor discovery packets for specific patterns.

E. It blocks certain traffic by inspecting DHCP packets for specific sources.

Correct Answer: BC

IPv6 source guard is an interface feature between the populated binding table and data traffic filtering. IPv6 source guard can deny traffic from unknown sources or unallocated addresses.

Question 2:
Refer to the exhibit.

```
R1

ip prefix-list ccnp1 seq 5 permit 10.1.48.0/24 le 24
ip prefix-list ccnp2 seq 5 permit 10.1.80.0/24 le 32
ip prefix-list ccnp3 seq 5 permit 10.1.64.0/24 le 24

route-map ospf-to-eigrp permit 10
   match ip address prefix-list ccnp1
   set tag 30
route-map ospf-to-eigrp permit 20
   match ip address prefix-list ccnp2
   set tag 20
route-map ospf-to-eigrp permit 30
   match ip address prefix-list ccnp3
   set tag 10
```

An engineer wanted to set a tag of 30 to route 10 1.80.65/32 but it failed How is the issue fixed?

A. Modify route-map ospf-to-eigrp permit 30 and match prefix-list ccnp2.

B. Modify route-map ospf-to-eigrp permit 10 and match prefix-list ccnp2.

C. Modify prefix-list ccnp3 to add 10.1.64.0/20 le 24

D. Modify prefix-list ccnp3 to add 10.1.64.0/20 ge 32

Correct Answer: B

Question 3:
Consider the partial output of the show ip route eigrp command:

```
rtrA# show ip route eigrp

Gateway of last resort is not set

15.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D 15.11.78.0/24 [90/5345354] via 10.10.70.41, 01:43:05, S0/0
D 15.200.16.0/24 [90/1723780] via 10.10.78.23, 00:50:37, S0/0
[90/1723780] via 10.10.19.40, 01:04:58, S0/0
[90/1723780] via 10.10.70.41, 01:20:37, S0/0
D 15.90.4.0/16 [90/4869420] via 10.10.19.40, 01:13:17, S0/0

172.161.0.0/16 is variably subnetted, 6 subnets, 3 masks
D 172.161.50.0/24 [90/4531003] via 10.10.70.41, 00:53:10, S0/1
D 172.161.98.1/30 [90/1723695] via 10.10.78.23, 01:27:03, S0/1
D 172.161.11.0/27 [90/1723695] via 10.10.19.45, 00:56:17, S0/1
[90/1723695] via 10.10.19.40, 00:50:58, S0/1
D 172.161.2.74/27 [90/6356189] via 10.10.70.41, 01:36:31, S0/1
D 172.161.4.47/30 [90/88258329] via 10.10.78.23, 01:44:20, S0/1
D 172.161.150.6/24 [90/3285083] via 10.10.70.41, 02:13:55, S0/1
D*EX 0.0.0.0/0 [170/2645987] via 10.10.70.41, 00:05:12, Ethernet0/0
[170/2645987] via 10.10.70.23, 00:05:12, Ethernet0/0
```

Which of the following destination subnets have equally load-balanced routes? (Choose all that apply.)

A. 172.161.4.47/30

B. 172.161.11.0/27

C. 15.200.16.0/24

D. 15.11.78.0/24

E. 0.0.0.0/0

Correct Answer: BCE

The 172.161.11.0/27 and 15.200.16.0/24 networks have equally load-balanced routes. A default route, 0.0.0.0/0, has been configured for load balancing as well. These three subnets are each load balanced on multiple routes. The output

entry for the 172.161.11.0/27 subnet is as follows:

D 172.161.11.0/27 [90/1723695] via 10.10.19.45, 00:56:17, S0/1 [90/1723695] via 10.10.19.40, 00:50:58, S0/1

This subnet can be reached by rtrA through two routes: 10.10.19.45 and 10.10.19.40 next-hop addresses. Both these routes have the same metric (1723695), and so are equally load balanced.

In the output, the 15.200.16.0/24 subnet has three equal-metric routes: 10.10.78.23, 10.10.19.40, and 10.10.70.41. These three routes are equally used to balance the load from rtrA to the 15.200.16.0/24 subnet.

The default route 0.0.0.0/0 is load balanced through two interfaces, as shown in the output:

D*EX 0.0.0.0/0 [170/2645987] via 10.10.70.41, 00:05:12, Ethernet0/0 [170/2645987] via 10.10.70.23, 00:05:12, Ethernet0/0

This load balancing of the default route could be tested by using traceroute to any IP address not represented in the routing table and verifying the path taken.

Subnets 172.161.4.47/30 and 15.11.78.0/24 are not participating in load balancing. In the given output, there is a single route (line) for both of these subnets. The rtrA router uses the route through the next-hop 10.10.78.23 to reach the

172.161.4.47/30 destination subnet. Similarly, rtrA uses the next-hop 10.10.70.41 to transmit packets to the 15.11.78.0/24 subnet.

Objective:

Layer 3 Technologies

Sub-Objective:

Configure and verify EIGRP load balancing

References: Cisco IOS IP Routing: Protocol-Independent Command Reference > show ip route Cisco > Support > Technology Support > IP > IP Routing > Technology Information > Technology Whitepaper > Enhanced Interior Gateway Routing Protocol > Document ID: 16406 > Load Balancing

---

Question 4:
You want to change the Administrative Distance of external EIGRP routes from the default of 170 to 130 instead on router R1 while leaving the default AD value for internal EIGRP routes. Which set of command will accomplish this?

A. R1(config)#router eigrp R1(config-router)#distance 170

B. R1(config)#router eigrp 1 R1(config-router)#distance eigrp 90 130

C. R1(config)#router eigrp 1 R1(config-router)#distance eigrp 130 90

D. R1(config)#router eigrp 1 R1(config-router)#distance 90 130

Correct Answer: B

Question 5:
How does an MPLS Layer 3 VPN function?

A. set of sites use multiprotocol BGP at the customer site for aggregation

B. multiple customer sites interconnect through service provider network to create secure tunnels between customer edge devices

C. set of sites interconnect privately over the Internet for security

D. multiple customer sites interconnect through a service provider network using customer edge to provider edge connectivity

Correct Answer: D

A Multiprotocol Label Switching(MPLS) Layer 3 Virtual Private Network (VPN) consists of a set of sites that are interconnected by means of an MPLS provider core network. At each customer site, one or more customer edge (CE) routers attach to one or more provider edge (PE) routers.

Reference:
https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k-r6-5/lxvpn/configuration/guide/b-l3vpn-cg-asr9000-65x/b-l3vpn-cg-asr9000-65x_chapter_010.pdf

---

Question 6:
Refer to the exhibit.

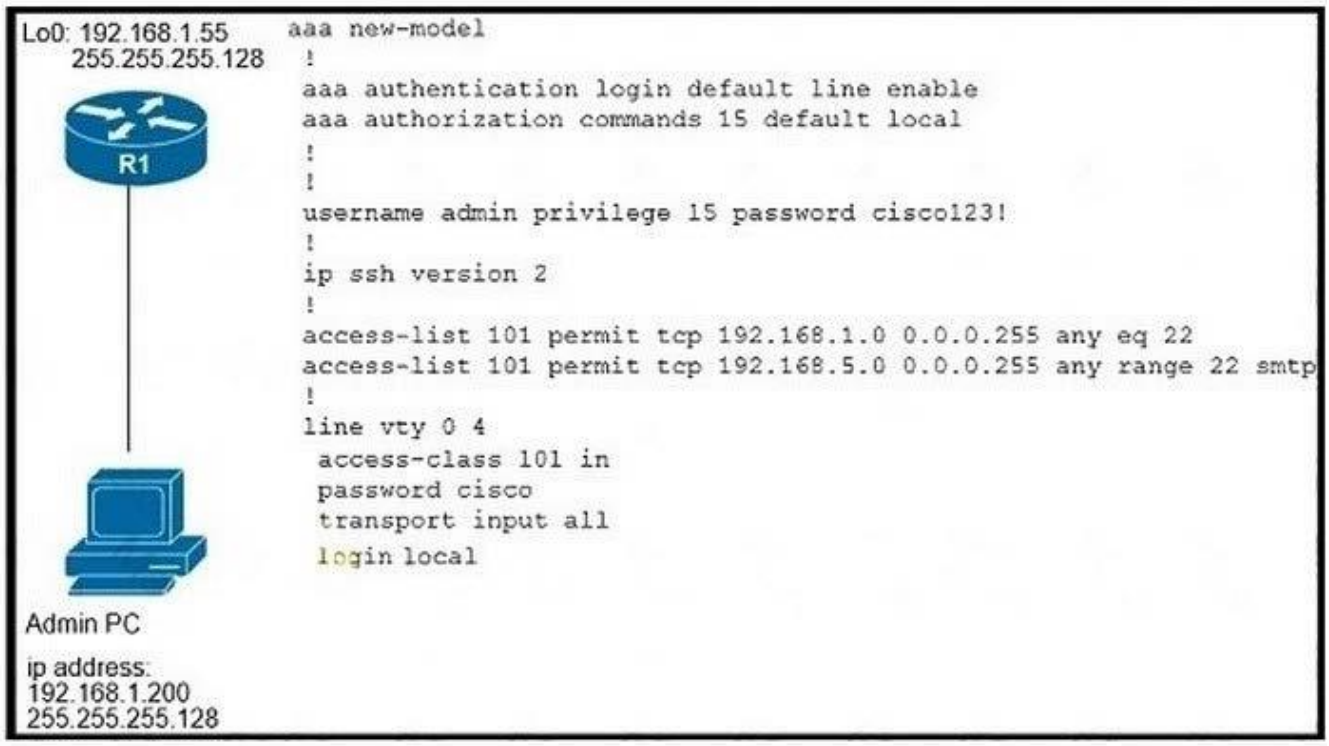

An engineer configured user login based on authentication database on the router, but no one can log into the router. Which configuration resolves the issue?

A. aaa authentication login default enable

B. aaa authorization network default local

C. aaa authentication login default local

D. aaa authorization exec default local

Correct Answer: C

Question 7:
Refer to the exhibit.

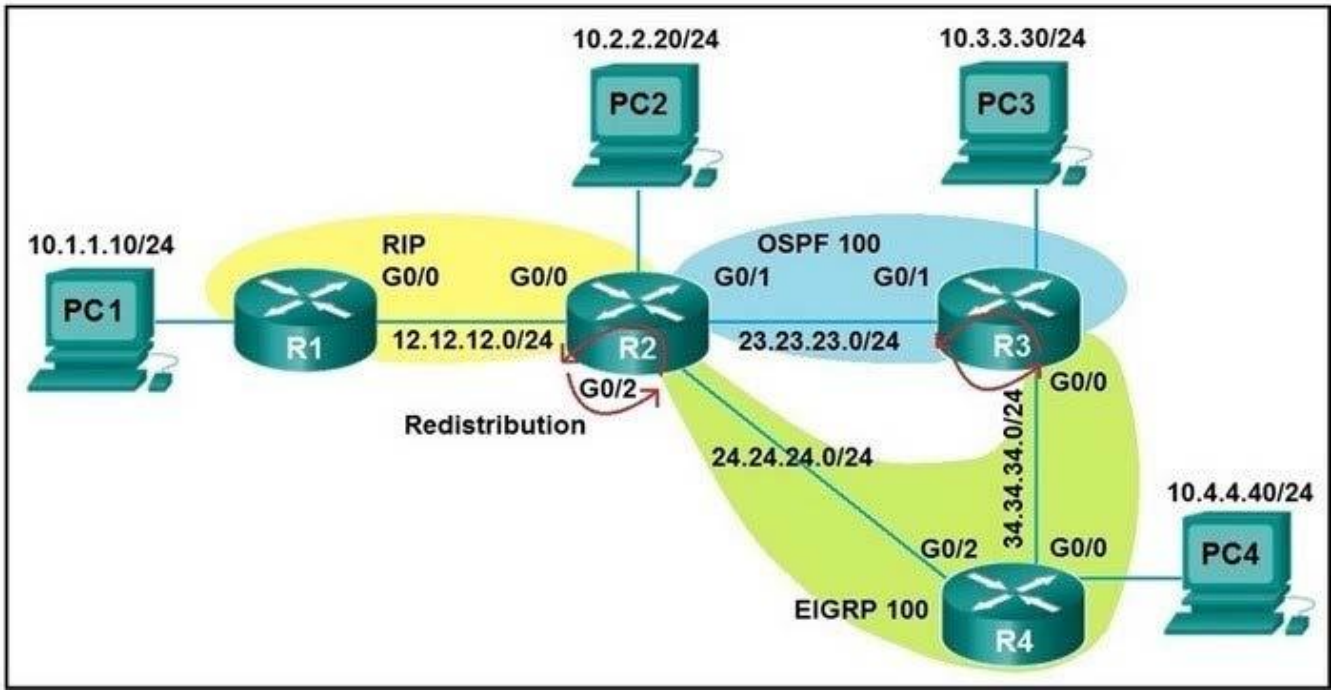

Redistribution is enabled between the routing protocols, and nowPC2 PC3, and PC4 cannot reach PC1. What are the two solutions to fix the problem? (Choose two.)

A. Filter RIP routes back into RIP when redistributing into RIP in R2

B. Filter OSPF routes into RIP from EIGRP when redistributing into RIP in R2

C. Filter all routes except RIP routes when redistributing into EIGRP in R2.

D. Filter RIP AND OSPF routes back into OSPF from EIGRP when redistributing into OSPF in R2

E. Filter all routes except EIGRP routes when redistributing into OSPF in R3.

Correct Answer: AB

Even PC2 cannot reach PC1 so there is something wrong with RIP redistribution in R2. Because RIP has higher Administrative Distance (AD) value than OSPF and EIGRP so it will be looped when doing mutual redistribution.

Question 8:
Refer to the exhibit.

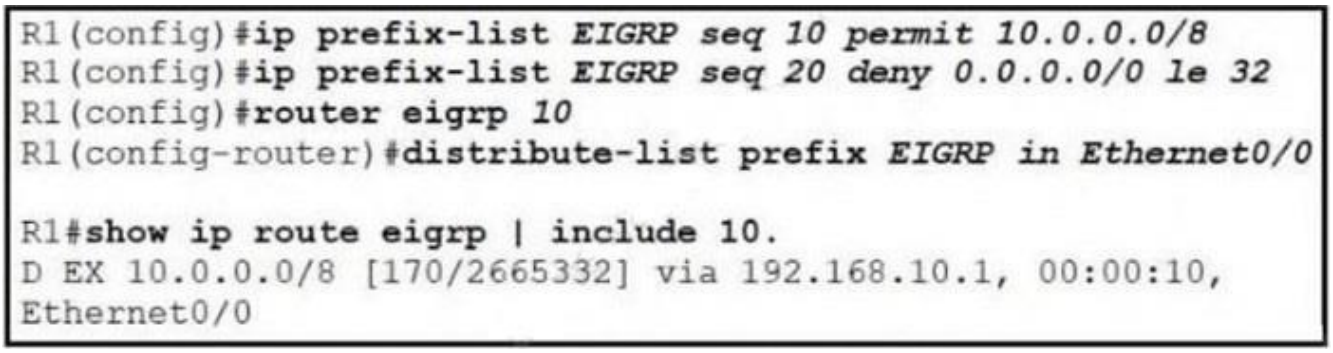

An engineer applies a prefix-list filter that filters most of the network 10 prefixes instead of allowing them. Which action resolves the issue?

A. Modify the ip prefix-list EIGRP seq 20 permit 10.0.0.0/8 ge 9 command.

B. Modify the ip prefix-list EIGRP seq 10 permit 10.0.0.0/8 le 9 command.

C. Modify the ip prefix-list EIGRP seq 20 permit 0.0.0.0/0 le 32 command.

D. Modify the ip prefix-list EIGRP seq 10 permit 10.0.0.0/8 le 32 command.

Correct Answer: D

Question 9:
Refer to the exhibit.

```
Router#sh ip route ospf
<output omitted>
Gateway is last resort is not set

      10.0.0.0/24 is subnetted, 1 subnets
   o  E2     10.0.0.0 [110/20] via 192.168.12.2, 00:00:10, Ethernet0/0
   o       192.168.3.0/24 [110/20] via 192.168.12.2, 00:00:50, Ethernet0/0
Router#

Router#show ip bgp
<output omitted>
      Network          Next Hop      Metric     LocPrf     Weight     Path
>*    192.168.1.1/32   0.0.0.0            0                32768      ?
>*    192.168.3.0      192.168.12.2      20                32768      ?
>*    192.168.12.0     0.0.0.0            0                32768      ?
Router#show running-config | section router bgp
router bgp 65000
 bgp log-neighbor-changes
 redistribute ospf 1
Router#
```

An engineer is trying to redistribute OSPF to BGP, but not all of the routes are redistributed. What is the reason for this issue?

A. By default, only internal routes and external type 1 routes are redistributed into BGP

B. Only classful networks are redistributed from OSPF to BGP

C. BGP convergence is slow, so the route will eventually be present in the BGP table

D. By default, only internal OSPF routes are redistributed into BGP

Correct Answer: D

If you configure the redistribution of OSPF into BGP without keywords, only OSPF intra-area and inter-area routes are redistributed into BGP, by default.

You can redistribute both internal and external (type-1 and type-2) OSPF routes via this command:

Router(config-router)#redistribute ospf 1 match internal external 1 external 2

Question 10:

An engineer configured a leak-map command to summarize EIGRP routes and advertise specifically loopback 0 with an IP of 10.1.1.1.255.255.255.252 along with the summary route. After finishing configuration, the customer complained not receiving summary route with specific loopback address.

Which two configurations will fix it? (Choose two.)

router eigrp 1 ! route-map Leak-Route deny 10 ! interface Serial 0/0 ip summary-address eigrp 1 10.0.0.0 255.0.0.0 leak-map Leak-Route

A. Configure access-list 1 permit 10.1.1.0.0.0.0.3.

B. Configure access-list 1 permit 10.1.1.1.0.0.0.252.

C. Configure access-list 1 and match under route-map Leak-Route.

D. Configure route-map Leak-Route permit 10 and match access-list 1.

E. Configure route-map Leak-Route permit 20.

Correct Answer: AD

When you configure an EIGRP summary route, all networks that fall within the range of your summary are suppressed and no longer advertised on the interface. Only the summary route is advertised. But if we want to advertise a network that has been suppressed along with the summary route then we can use leak-map feature. The below commands will fix the configuration in this question: R1(config)#access-list 1 permit 10.1.1.0 0.0.0.3 R1(config)#route-map Leak-Route permit 10 // this command will also remove the "route_map Leak-Route deny 10" command. R1(config-route-map)#match ip address 1

Question 11:
Refer to the exhibits.

```
R1#show ip route 5.5.5.0
Routing entry for 5.5.5.0/24
  Known via "eigrp 1", distance 90, metric 158720, type internal
  Redistributing via eigrp 1
  Last update from 192.168.13.3 on Ethernet1/0, 00:00:40 ago
  Routing Descriptor Blocks:
  * 192.168.13.3, from 192.168.13.3, 00:00:40 ago, via Ehernet1/0
      Route metric is 412160, traffic share count is 23
      Total delay is 6100 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
    192.168.12.2, from 192.168.12.2, 00:00:40 ago, via FastEthernet0/0
      Route metric is 158720, traffic share count is 60
      Total delay is 5200 microseconds, minimum bandwidth is 100000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 2
```

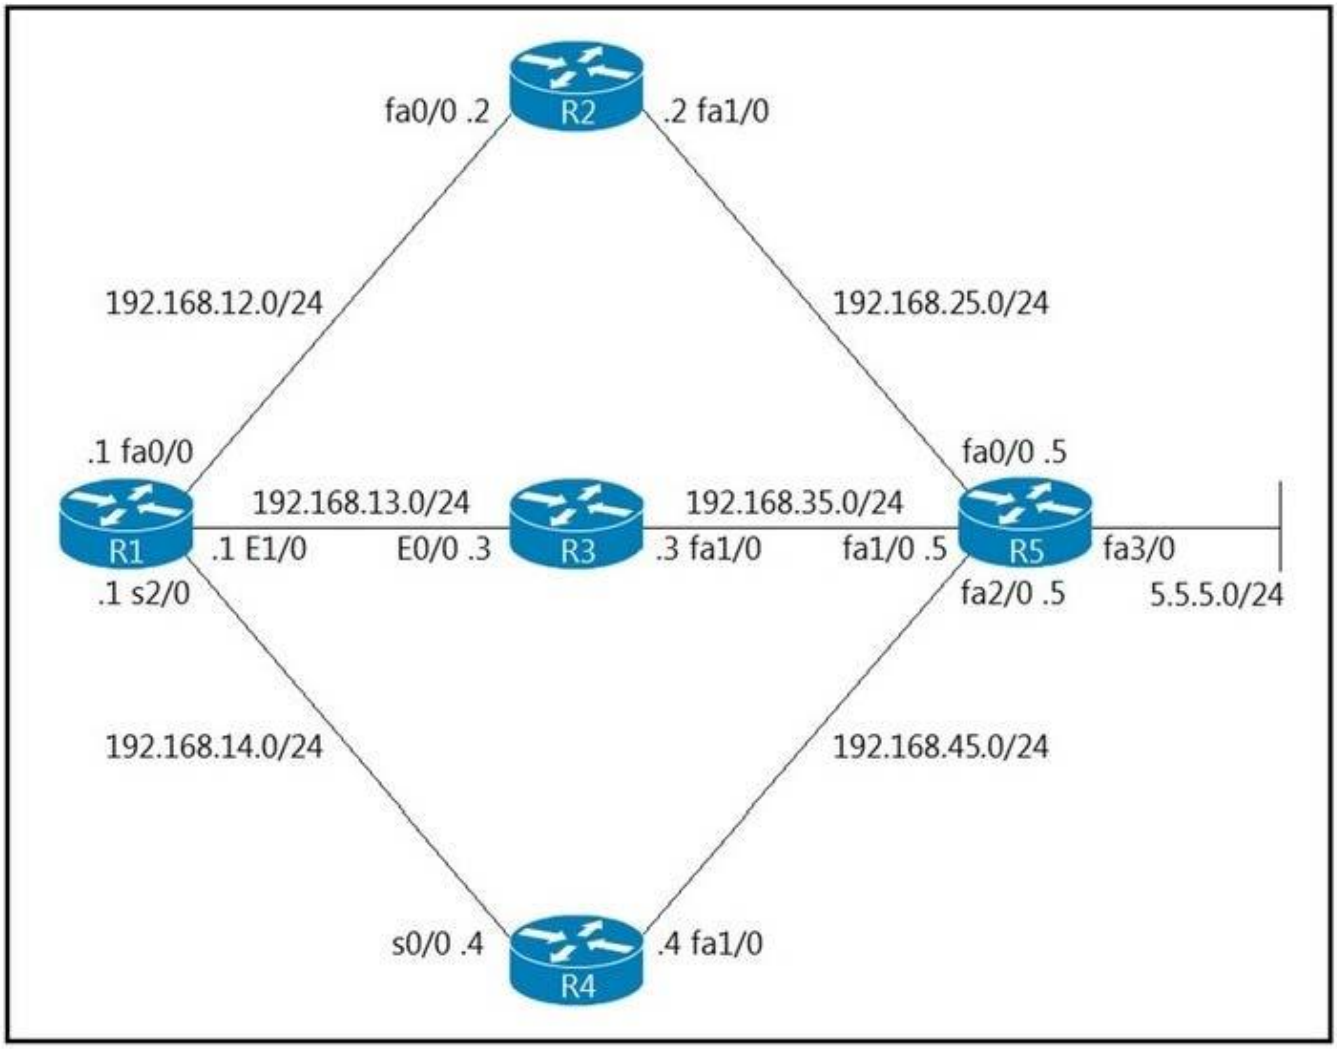

An engineer investigates a routing issue on R1 and finds that traffic destined to 5.5.5.0/24 does not take all of the paths.

An engineer investigates a routing issue on R1 and finds that traffic destined to 5.5.5.0/24 does not take all of the paths.

Which action resolves the issue?

A. Increase the variance value in EIGRP.

B. Decrease the variance value in EIGRP.

C. Remove the adjacency of R3 from EIGRP.

D. Stop advertising 192.168.13.0/24 in EIGRP.

Correct Answer: A

Reference:
https://community.cisco.com/t5/networking-documents/troubleshooting-eigrp-variance-command/ta-p/3129662#:~:text=EIGRP%20provides%20a%20mechanism%20to,means%20equal%2Dcost%20load%20balancing

---

Question 12:
Refer to the exhibit. Which command must be configured to make VRF CCNP work?

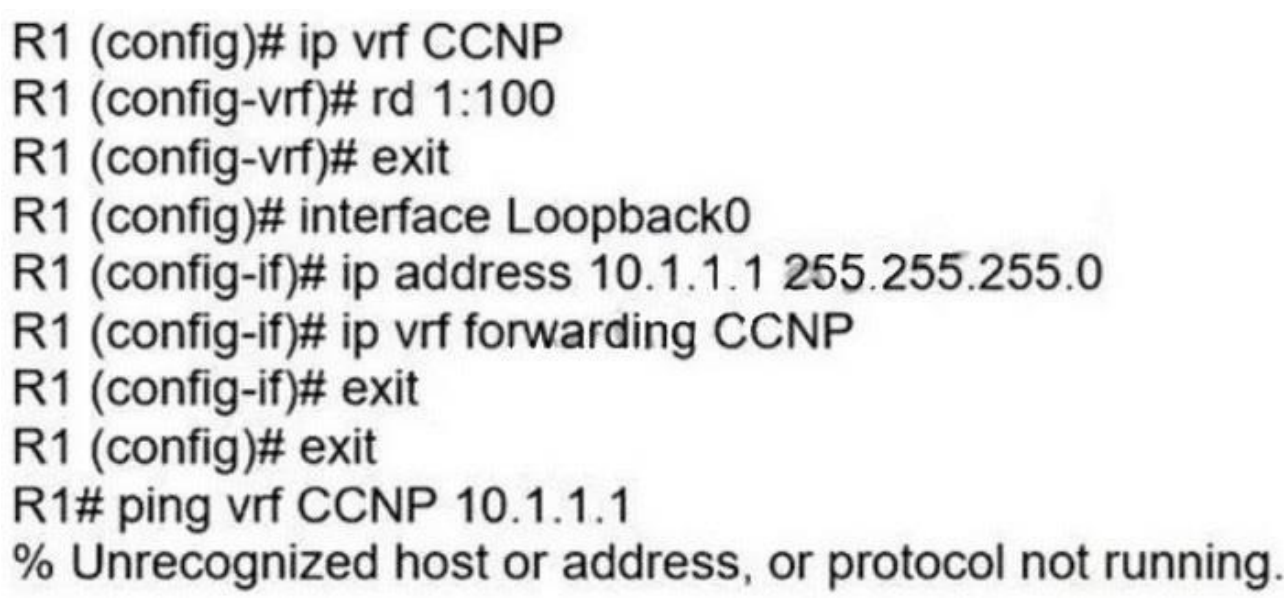

A. interface Loopback0 vrf forwarding CCNP

B. interface Loopback0 ip address 10.1.1.1 255.255.255.0

C. interface Loopback0 ip address 10.1.1.1 255.255.255.0 vrf forwarding CCNP

D. interface Loopback0 ip address 10.1.1.1 255.255.255.0 ip vrf forwarding CCNP

Correct Answer: B

Question 13:
Which of the following statements are TRUE about manually configured IPV4-to-IP6 tunnels and GRE tunnels? (Choose two.)

A. Manually configured tunnels use the tunnel mode ipv6ip command, while GRE tunnels use the tunnel mode gre ip command.

B. Manually configured tunnels support IPv6 IGPs, while GRE tunnels do not.

C. Manually configured tunnels block IPv6 multicasts, while GRE forwards them.

D. Manually configured tunnels do not support multiple passenger protocols, while GRE tunnels support them.

Correct Answer: AD

The following statements are TRUE about manually configured tunnels and GRE tunnels:

Manually configured tunnels use the tunnel mode ipv6ip command, while GRE tunnels use the tunnel mode gre ip command.

Manually configured tunnels do not support multiple passenger protocols, while GRE tunnels support them.

Manually configured tunnels and Generic Routing Encapsulation (GRE) tunnels are static point-to-point tunneling methods. Both of these tunneling methods provide a permanent link between two IPv6 networks that are separated by an IPv4

backbone. For each link between two IPv6 networks, a separate tunnel needs to be created.

Manually configured tunnels use a particular passenger protocol and do not support multiple passenger protocols at the same time. However, GRE tunnels can simultaneously use various passenger protocols.

It is incorrect to state that manually configured tunnels support IPv6 IGPs, while GRE tunnels do not. GRE tunnels also support IPv6 IGPs, such as OSPF, RIP, and IS-IS.

It is incorrect to state that manually configured tunnels block IPv6 multicasts, while GRE forwards them.

Manually configured tunnels also forward IPv6 multicasts.

Objective:

Network Principles

Sub-Objective:

Recognize proposed changes to the network

References:

Cisco IOS IPv6 Configuration Guide, Release 12.4 > Implementing Tunneling for IPv6 > Configuration Examples for Implementing Tunneling for IPv6 > Example: Configuring Manual IPv6 Tunnels

---

Question 14:
Refer to the exhibit.

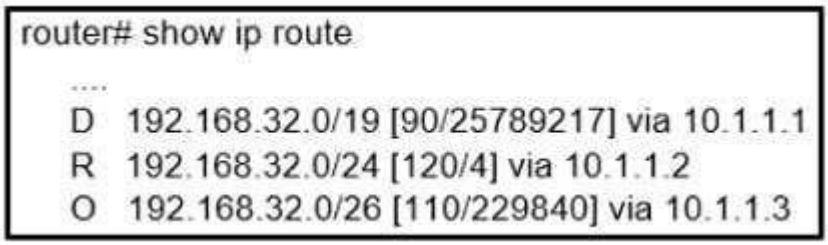

An engineer is trying to get 192.168.32.100 forwarded through 10.1.1.1, but it was forwarded through 10.1.1.2.

What action forwards the packets through 10.1.1.1?

A. Configure EIGRP to receive 192.168.32.0 route with lower admin distance.

B. Configure EIGRP to receive 192.168.32.0 route with longer prefix than /19.

C. Configure EIGRP to receive 192.168.32.0 route with lower metric.

D. Configure EIGRP to receive 192.168.32.0 route with equal or longer prefix than /24.

Correct Answer: D

Question 15:
You are configuring NAT64 to allow communication between a host running IPv6 and a server running IPv4. The router R1 sits between the host and the server. The router\'s Fa0/2/7 interface is connected to the IPv6 host, and the Fa0/2/6 interface is connected to the IPv4 server.

The IPv6 host has an IPv6 address of 2001::a00:1/128 and the IPv4 server is at 10.0.0.1. Below is the relevant configuration on R1:

```
interface FastEthernet0/2/6
ip address 10.0.0.2 255.255.255.0
     nat64 enable
!
interface FastEthernet0/2/7
     no ip address
     ipv6 address 2001::A00:B/128
          nat64 enable

nat64 prefix stateful 3001::/96
nat64 v6v4 static 2001::A00:A 10.0.0.10
```

When the IPv4 server responds to the IPv6 host, what IPv6 address will be in the source address in the packet?

A. 2001::a001

B. 2001::A00:B

C. 3001::a00:1

D. 2001::A00:A

Correct Answer: C

NAT64 is a solution when IPv6 hosts need to communicate with IPv4-only servers. When the translation occur on the router the IPv4 address 10.0.0.1 will converted to hex as a00:1 and will be attached to the end of the stateful prefix of

3001::/96 that was configured on the router interface connected to the IPv4 server. The result will be 3001::a00:1.

The address will not be 2001::a001. The prefix that will be attached to the hex version of 10.0.0.1 will not be that of the interface fa0/2/7 but will be the prefix that was configured on that interface for nat64 translation which is 3301::/96. The

address will not be 2001::a00:b. That is the IPv6 address on the interface connected to the IPv6 host, but that address is not used for IPv4 to IPv6 communication. A translated address will be generated by converting the IPv4 address of the

IPv4 host to hex and attaching it to the IPv6 prefix configured on the interface connected to the IPv4 server.

The address will not be 2001::A00:A. That is the IPv6 address of the IPv6 host. That was statically mapped to 10.0.0.10 in the configuration and as such will be the IPv4 address used by the IPv6 host on the IPv4 side of the router.

Objective:

Infrastructure Services

Sub-Objective:

Describe IPv6 NAT

References:

Stateful Network Address Translation 64 (PDF)