Exam

300-410

Implementing Cisco Enterprise Advanced Routing and Services (300-410 ENARSI)

Free demo	80 sample questions
Full exam	636 questions with answers and explanations
Contents	Questions with answers and explanations
Generated	23 October 2026

Exam domains

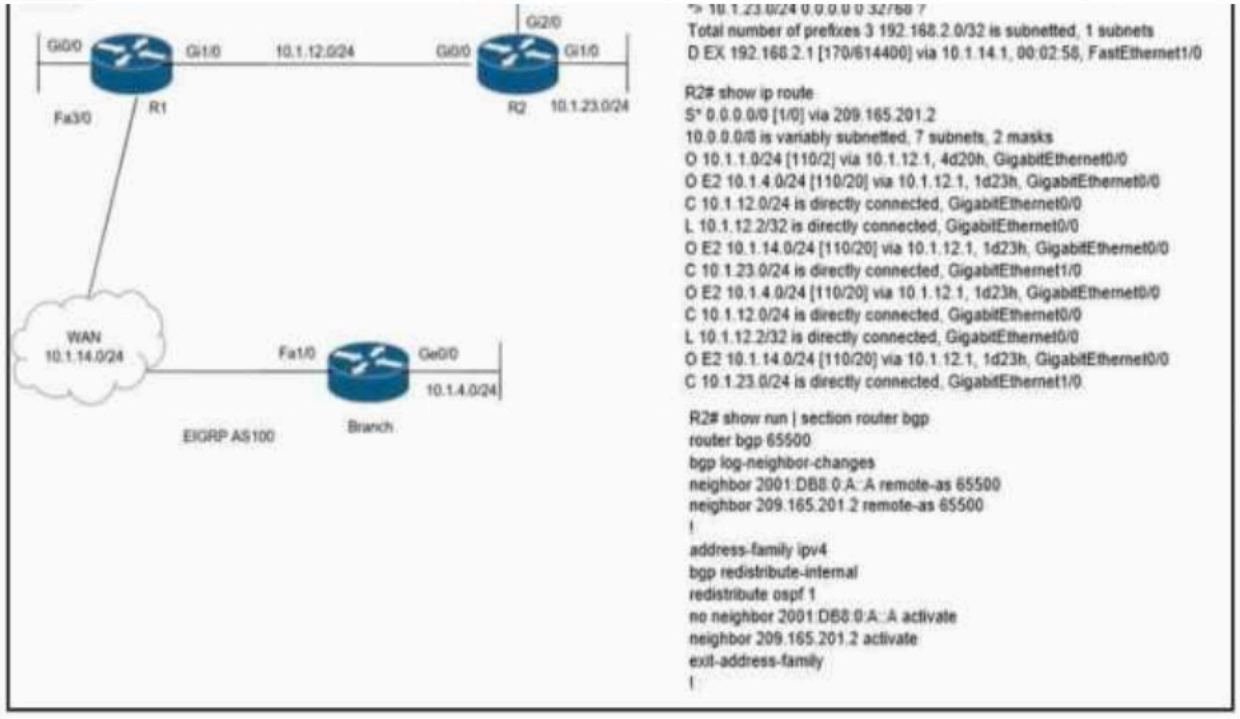
Domain	Exam weight	Questions
D1 Layer 3 Technologies	35%	34
D2 VPN Technologies	20%	15
D3 Infrastructure Security	20%	15
D4 Infrastructure Services	25%	16

This is a free demo from Dumpsgate: 80 of the exam's 636 questions. **Get the full 300-410 exam** with every question, answer and explanation.

Question 1

Single choice · Domain: D1 Layer 3 Technologies › 1.4 Troubleshoot redistribution between any pair of routing protocols or routing sources

Refer to the exhibit.



Refer to the exhibit. Users in BGP autonomous system 65500 cannot access IPv4 resources in the EIGRP IPv4 autonomous system 100.

Which action resolves the issue?

A. Redistribute BGP routes into OSPF.
B. Redistribute OSPF external routes into BGP.
C. Redistribute OSPF internal routes into BGP.
D. Redistribute EIGRP routes into OSPF.

Answer: B. Redistribute OSPF external routes into BGP.

Explanation

The EIGRP IPv4 resources are represented in OSPF as external routes, as shown by the OSPF E2 route in the exhibit. To make those prefixes reachable from BGP AS 65500, redistribute the OSPF external routes into BGP. This advertises the EIGRP resources to the BGP autonomous system so its users can reach them.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.pdf

✗ **A** Redistribute BGP routes into OSPF.
This sends routes in the opposite direction: it would make BGP routes available to OSPF, but would not

advertise the EIGRP resources to BGP AS 65500. The problem is that BGP users need routes to the EIGRP IPv4 resources.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.pdf

✓ **B** Redistribute OSPF external routes into BGP.
The EIGRP resources appear as OSPF external routes in the exhibit. Redistributing those routes into BGP advertises the resources to AS 65500, allowing its users to route to them.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.pdf

✗ **C** Redistribute OSPF internal routes into BGP.
The relevant EIGRP prefixes are not OSPF internal routes; the exhibit identifies an OSPF external route. Redistributing only OSPF internal routes would not advertise those external EIGRP resources to BGP.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.pdf

✗ **D** Redistribute EIGRP routes into OSPF.
The exhibit already shows an OSPF external route for a resource in the EIGRP domain, indicating that the EIGRP route is represented in OSPF. The missing step is advertising that OSPF external route into BGP, not redistributing EIGRP into OSPF again.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5242-bgp-ospf-redis.pdf

Dumpsgate

Question 2

Single choice · Domain: D1 Layer 3 Technologies › 1.9.f Metrics

While examining log messages from a router, an administrator sees this error message: *Mar 13 11:32:42.453: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 10.10.10.1 (Ethernet0/1) is down: K-value mismatch.

Both routers are using default values.

Which action resolves this issue?

A. Configure metric weights between the peers to calculate the valid K-value.
B. Resolve metric differences between the peers to calculate the valid K-value.
C. Upgrade the peer device IOS to support the goodbye message that establishes adjacency.
D. Match authentication parameters to support the goodbye message that establishes adjacency.

Answer: A. Configure metric weights between the peers to calculate the valid K-value.

Explanation

The log explicitly identifies the adjacency failure as an EIGRP K-value mismatch. EIGRP peers must use matching metric weights (K-values) to form an adjacency, so configure the weights consistently between the routers. Since both routers are said to use defaults, investigate and correct the effective metric-weight configuration so both peers agree; differences in calculated metrics are not the same as a K-value mismatch.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_ire-enhanced-igrp.html

Free demo · full exam: 636 questions

https://www.cisco.com/c/en/us/td/docs/routers/ios/config/17-x/ip-routing/b-ip-routing/m_ire-enhanced-igrp.html

✓ **A** Configure metric weights between the peers to calculate the valid K-value.
The message reports a K-value mismatch, which means the EIGRP metric-weight settings used by the peers do not match. Configure matching metric weights on both routers; matching those settings allows them to establish the adjacency. [1]

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_ire-enhanced-igrp.html

✗ **B** Resolve metric differences between the peers to calculate the valid K-value.
Differences in route metrics do not cause the reported K-value mismatch. The issue is that the peers' metric-weight settings must match, not that their calculated route metrics must be made equal. [1]

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_ire-enhanced-igrp.html

✗ **C** Upgrade the peer device IOS to support the goodbye message that establishes adjacency.
The log names a K-value mismatch, not an unsupported goodbye message or IOS compatibility problem. An upgrade would not resolve mismatched EIGRP metric weights. [1]

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html

✗ **D** Match authentication parameters to support the goodbye message that establishes adjacency.
Authentication settings are not the cause identified by this log message; it explicitly reports a K-value mismatch. Matching authentication parameters would not correct inconsistent EIGRP metric weights. [1]

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html

Dumpsgate

Question 3

Single choice · Domain: D1 Layer 3 Technologies › 1.8 Describe Bidirectional Forwarding Detection

Which type of BFD session is created on a broadcast Ethernet interface with OSPF?

A. With the DR and BDR, but not any DROTHER routers
B. With all OSPF routers
C. Only on P2P Ethernet interfaces
D. Only with the DR

Answer: A. With the DR and BDR, but not any DROTHER routers

Explanation

On a broadcast Ethernet network, OSPF establishes BFD sessions with the designated router (DR) and backup designated router (BDR), but not with DROTHER routers. Therefore, option A correctly describes the BFD sessions created for this interface type.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-3s/asr903/17-1-1/b-irb-xe-17-1-asr903.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-17/irb-xe-17-book-cat8000.pdf

✓ **A** With the DR and BDR, but not any DROTHER routers
This is the broadcast-Ethernet behavior described in the question: BFD sessions are created with the DR and BDR, while DROTHER routers are excluded. It is the only option that matches that behavior.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-3s/asr903/17-1-1/b-irb-xe-17-1-asr903.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-17/irb-xe-17-book-cat8000.pdf

✗ **B** With all OSPF routers
BFD sessions are not created with every OSPF router on a broadcast Ethernet interface. DROTHER routers do not receive BFD sessions in this scenario.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-3s/asr903/17-1-1/b-irb-xe-17-1-asr903.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-17/irb-xe-17-book-cat8000.pdf

✗ **C** Only on P2P Ethernet interfaces
The question specifically concerns a broadcast Ethernet interface, and BFD sessions are created on that interface type as well. This option incorrectly limits the behavior to point-to-point Ethernet interfaces.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-3s/asr903/17-1-1/b-irb-xe-17-1-asr903.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-17/irb-xe-17-book-cat8000.pdf

✗ **D** Only with the DR
The BFD sessions are created with both the DR and the BDR, not just the DR. This option leaves out the BDR.

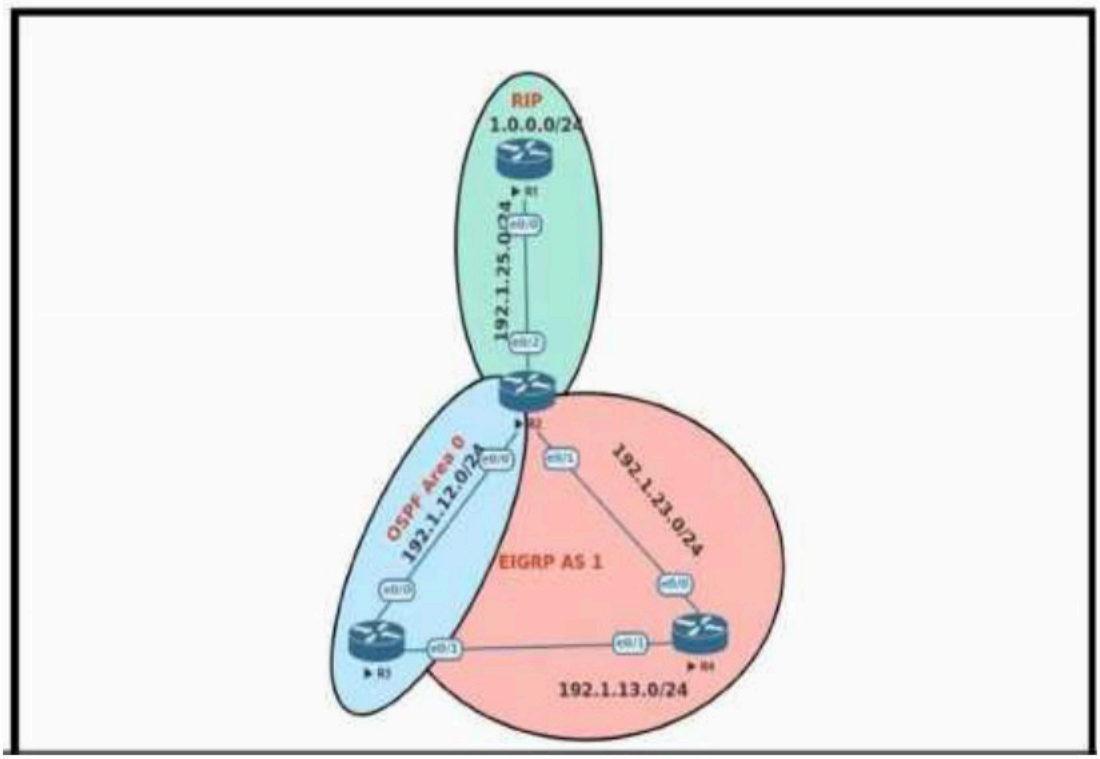
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-3s/asr903/17-1-1/b-irb-xe-17-1-asr903.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-17/irb-xe-17-book-cat8000.pdf

Dumpsgate

Question 4

Single choice · Domain: D1 Layer 3 Technologies › 1.1 Troubleshoot administrative distance (all routing protocols)

Refer to the exhibit.



Refer to the exhibit. R3 is learning the 1.0.0.0/24 route through OSPF instead of EIGRP. Which action causes R3 to choose EIGRP to reach the 1.0.0.0/24 network?

A. Configure OSPF administrative distance to 200
B. Configure OSPF administrative distance to 120
C. Configure EIGRP administrative distance to 110
D. Configure EIGRP administrative distance to 120

Answer: A. Configure OSPF administrative distance to 200

Explanation

R3 selects the route with the lower administrative distance when competing routes to the same network are available. EIGRP's internal administrative distance is 90, while OSPF's is 110, so raising OSPF's distance to 200 makes the EIGRP route preferred. This meets the requirement for R3 to reach 1.0.0.0/24 through EIGRP.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-ip-prot-indep-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/command/ire-cr-book/ire-a1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/m_ospf-a1.html

✓ **A** Configure OSPF administrative distance to 200

Free demo · full exam: 636 questions

An administrative distance of 200 is higher than EIGRP's default internal distance of 90. R3 will therefore prefer the EIGRP route over the OSPF route to 1.0.0.0/24.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-ip-prot-indep-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/command/ire-cr-book/ire-a1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/m_ospf-a1.html

✗ **B** Configure OSPF administrative distance to 120
Although 120 is higher than OSPF's default distance of 110, it is still higher than EIGRP's default internal distance of 90, so this would also favor EIGRP—not explain why this is not the answer?

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-ip-prot-indep-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/command/ire-cr-book/ire-a1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/m_ospf-a1.html

✗ **C** Configure EIGRP administrative distance to 110
Changing EIGRP's distance to 110 would make it tie OSPF's default administrative distance of 110. It would not make EIGRP's distance lower and thus does not ensure EIGRP is selected.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-ip-prot-indep-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/command/ire-cr-book/ire-a1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/m_ospf-a1.html

✗ **D** Configure EIGRP administrative distance to 120
An EIGRP distance of 120 is higher than OSPF's default distance of 110. R3 would continue to prefer the OSPF route, contrary to the requirement.

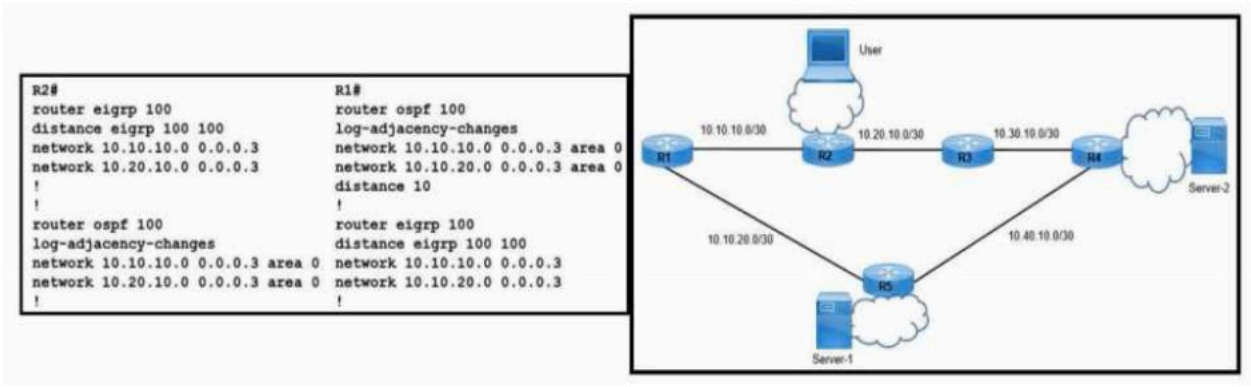
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-ip-prot-indep-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/command/ire-cr-book/ire-a1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/m_ospf-a1.html

Dumpsgate

Question 5

Single choice · Domain: D1 Layer 3 Technologies › 1.1 Troubleshoot administrative distance (all routing protocols)

Refer to the exhibit.



Refer to the exhibit. The user experiences random packet drops when accessing Server-1 but does not have any issues when reaching out to Server-2.

Which action resolves the issue?

A. Remove the distance 10 under the OSPF process on R1.
B. Configure the distance under the EIGRP process to 10 on R2.
C. Configure the distance under the EIGRP process to 10 on R1.
D. Modify the distance 10 under the OSPF process on R1 to 90.

Answer: A. Remove the distance 10 under the OSPF process on R1.

Explanation

R1 has an OSPF administrative distance of 10, making OSPF-learned routes more preferred than EIGRP-learned routes even when EIGRP provides the intended path toward Server-1. Removing the custom OSPF distance restores its default preference, allowing internal EIGRP routes (AD 90) to be preferred over OSPF routes (AD 110). This addresses the route-selection issue affecting Server-1 while leaving the working path to Server-2 unchanged.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/m_ospf-a1.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-ip-prot-indep-0.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_eigrp/command/reference/ire_book/ire_a1.html

✓ **A** Remove the distance 10 under the OSPF process on R1.
This removes the unusually low administrative distance that causes R1 to prefer OSPF routes over EIGRP routes. With the defaults restored, internal EIGRP routes are preferred, resolving the problematic route selection toward Server-1.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/m_ospf-a1.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-ip-prot-indep-0.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_eigrp/command/reference/ire_book/ire_a1.html

✗ **B** Configure the distance under the EIGRP process to 10 on R2.
The problematic custom distance is configured under OSPF on R1, so changing EIGRP's distance on R2 does not correct R1's route selection. It could instead make EIGRP routes unusually preferred on R2.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/m_ospf-a1.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-ip-prot-indep-0.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_eigrp/command/reference/ire_book/ire_a1.html

✗ **C** Configure the distance under the EIGRP process to 10 on R1.
This changes the wrong protocol setting: the exhibit shows the low distance under OSPF on R1. Lowering EIGRP's distance to 10 would make EIGRP routes more preferred, but it would leave the inappropriate OSPF distance in place rather than correcting the configuration causing the issue.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/m_ospf-a1.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-ip-prot-indep-0.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_eigrp/command/reference/ire_book/ire_a1.html

✗ **D** Modify the distance 10 under the OSPF process on R1 to 90.
Changing the OSPF distance to 90 still leaves it far below OSPF's default distance of 110 and does not restore the intended protocol preference. It also places OSPF at the same distance as internal EIGRP, rather than clearly preferring EIGRP.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/m_ospf-a1.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-ip-prot-indep-0.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_eigrp/command/reference/ire_book/ire_a1.html

Dumpsgate

Question 6

Single choice · Domain: D1 Layer 3 Technologies › 1.11.b Neighbor relationship and authentication (next-hop, multihop, 4-byte AS, private AS, route refresh, synchronization, operation, peer group, states and timers)

Refer to the exhibit.

```
R2#show run | begin bgp
--------
router bgp 64512
neighbor 10.1.1.1 remote-as 1
network 10.10.10.0 mask 255.255.255.0 route-map COMM
---------
route-map COMM permit 10
set community 64512:600
---------
R2#show ip bgp 10.10.10.0
BGP routing table entry for 10.10.10.0/24, version 1320 Paths: (1 available, best #1)
Local 0.0.0.0 from 0.0.0.0 (1.1.1.1)
Origin IGP, metric 0, localpref100, weight 32768,
valid, sourced, local, best
Community 64512:600
```

Refer to the exhibit. The eBGP neighbor 10.1.1.1 of router R2 receives and installs the prefix, but the community value is not found in the BGP routing table.

How must this issue be resolved?

A. R1 must have a route map assigned to neighbor R2, which is clearing the community values.
B. R2 is not sending the community value to neighbor R1, configure R2 to send community.
C. R1 must have a route map assigned to neighbor R2, which is not matching the community values.
D. Community values are nontransitive and R2 is an eBGP peer, configure R2 to send community.

Answer: B. R2 is not sending the community value to neighbor R1, configure R2 to send community.

Explanation

R2 applies route-map COMM to the prefix and sets community 64512:600, as shown in its local BGP table. However, BGP communities are not sent to an eBGP neighbor by default, so R2 must be configured to send the community attribute to neighbor 10.1.1.1. This allows the neighbor to receive the prefix with the community value attached. [2, 10]

https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp3.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/28784-bgp-community.html

✗ **A** R1 must have a route map assigned to neighbor R2, which is clearing the community values.
The exhibit shows R2 has the community locally, while the problem is that its eBGP neighbor does not receive it. There is no evidence that a route map on R1 is stripping the community; configure R2 to advertise the community to its peer. [2]

https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp3.html

✓ **B** R2 is not sending the community value to neighbor R1, configure R2 to send community.
R2's BGP table confirms community 64512:600 is attached to the prefix locally. Community attributes are not propagated to eBGP peers by default, so R2 must be configured to send the community to neighbor 10.1.1.1. [2, 10]

https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp3.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/28784-bgp-community.html

✗ **C** R1 must have a route map assigned to neighbor R2, which is not matching the community values.
A route map that fails to match a community would not explain why the community is absent from the neighbor's received route. The community exists on R2, and the relevant issue is that R2 is not sending it to its eBGP peer. [2]

https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp3.html

✗ **D** Community values are nontransitive and R2 is an eBGP peer, configure R2 to send community.
This option gives the right configuration action, but its stated reason is inaccurate: standard BGP communities are transitive attributes. The issue is that Cisco BGP does not send communities to eBGP neighbors by default, so they must be explicitly enabled. [2, 10]

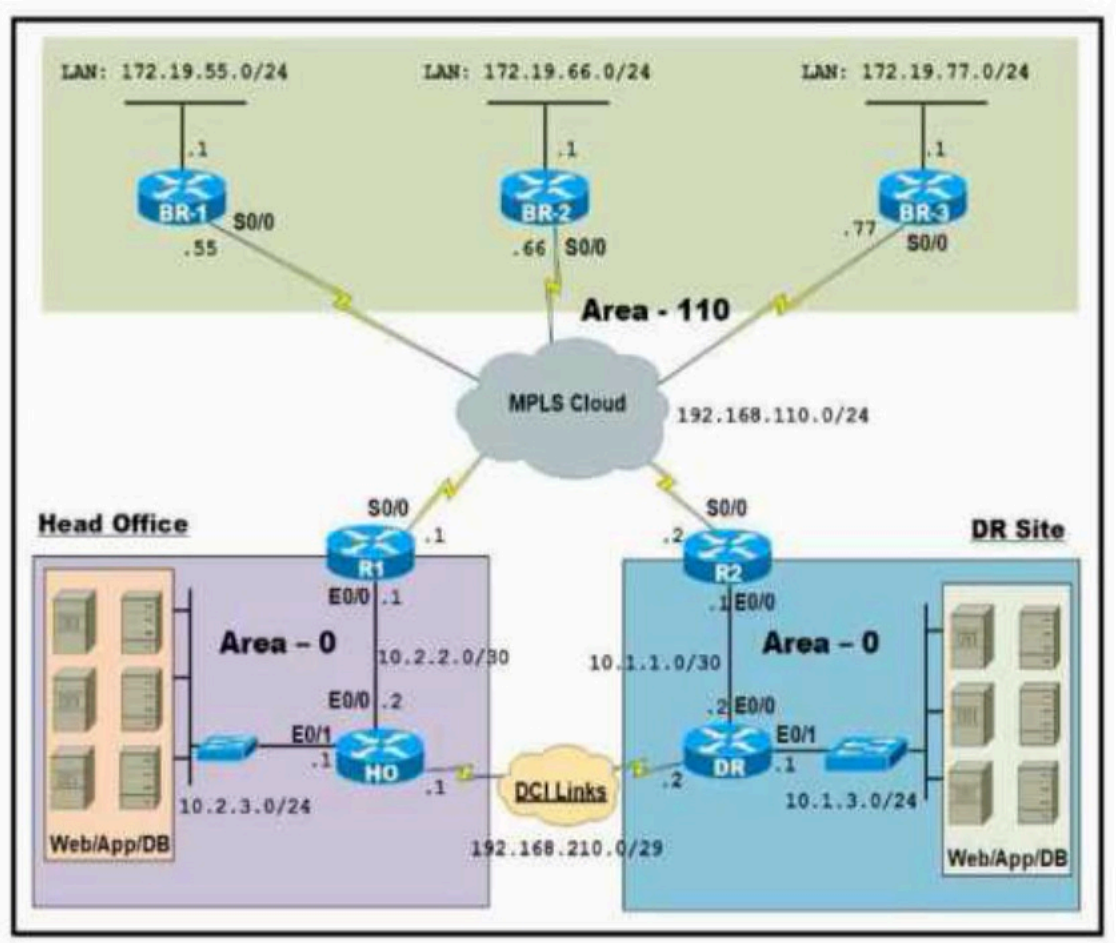
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp3.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/28784-bgp-community.html

Dumpsgate

Question 7

Single choice · Domain: D1 Layer 3 Technologies › 1.10.d Virtual link

Refer to the exhibit.



Refer to the exhibit. During an unannounced link-maintenance window at the ISP, the DCI link went down, which caused a significant service outage.

What action must the network engineer take at the head office to ensure Area-0 connectivity without intervention from the ISP?

A. Create a virtual link in Area-0 between HO E0/0 and DR E0/0.
B. Create a GRE tunnel interface in Area-110 between R1 S0/0 and R2 S0/0.
C. Create a GRE tunnel interface in Area-0 between HO E0/0 and DR E0/D.
D. Create a virtual link in Area-110 between R1 S0/0 and R2 S0/0.

Answer: D. Create a virtual link in Area-110 between R1 S0/0 and R2 S0/0.

Explanation

The DCI failure removes the direct Area 0 connection between the head office and DR site. R1 and R2 are

Free demo · full exam: 636 questions

reachable through the existing ISP network in Area 110, so configuring a virtual link across Area 110 can logically connect their Area 0s without requiring the ISP to restore the DCI link. A virtual link uses a nonbackbone transit area to provide an Area 0 connection between the ABRs [1][7].

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/47866-ospfdb7.html
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/212607-ospf-virtual-links-transit-capability.html

✗ **A** Create a virtual link in Area-0 between HO E0/0 and DR E0/0.
The virtual link must use a transit area that provides reachability between the participating ABRs. Area 0 is the backbone area, not the transit area for this connection, and the HO–DR DCI path is the failed link.

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/47866-ospfdb7.html
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/212607-ospf-virtual-links-transit-capability.html

✗ **B** Create a GRE tunnel interface in Area-110 between R1 S0/0 and R2 S0/0.
This option proposes creating a GRE tunnel, rather than using the existing Area 110 path to establish the required OSPF virtual link. The required action is a virtual link between the ABRs over the transit area.

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/47866-ospfdb7.html
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/212607-ospf-virtual-links-transit-capability.html

✗ **C** Create a GRE tunnel interface in Area-0 between HO E0/0 and DR E0/D.
The proposed tunnel endpoints are at the HO and DR sides of the failed DCI connection, so this does not use the existing Area 110 path between the ABRs. It also calls for a GRE tunnel rather than the virtual link needed to connect the areas through the transit area.

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/47866-ospfdb7.html
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/212607-ospf-virtual-links-transit-capability.html

✓ **D** Create a virtual link in Area-110 between R1 S0/0 and R2 S0/0.
Area 110 is the transit area linking R1 and R2 across the ISP network, and the routers connect that path to their respective Area 0 sites. A virtual link over Area 110 provides the Area 0 connectivity lost when the DCI link goes down, without requiring ISP intervention to restore that link.

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/47866-ospfdb7.html
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/212607-ospf-virtual-links-transit-capability.html

Question 8

Single choice · Domain: D1 Layer 3 Technologies › 1.6 Configure and verify policy-based routing

A network engineer wants to redirect HTTP traffic originating from subnet 192.168.12.1/24 to the next-hop IP address of 10.2.12.1. This configuration was implemented:

```
route-map pbr permit 10
 match ip address 100
 set ip next-hop 10.2.12.1
!
interface GigabitEthernet0/0
 description - To ISP1 link ID 64:f9fc:3421:5f2f:16ff:4f20:7f0f
 ip policy route-map pbr
```

Which command completes the configuration?

A. Router(config)# access-list 100 permit tep any eq 80 192.168.12.0 0.0.0.255
B. Router(config)# access-list 100 permit tcp 192.168.12.0 0.0.0.255 eq any 80
C. Router(config)# access-list 100 permit tcp 192.168.12.0 0.0.0.255 any eq 80
D. Router(config)# access-list 100 permit tcp any 192.168.12.0 0.0.0.255 eq 80

Answer: C. Router(config)# access-list 100 permit tcp 192.168.12.0 0.0.0.255 any eq 80

Explanation

The route-map matches traffic using access list 100, so the ACL must identify HTTP packets from the specified source subnet. Option C uses TCP, matches source 192.168.12.0/24 with wildcard mask 0.0.0.255, and matches destination port 80, which is HTTP. This lets the policy-based routing rule set the next hop for the intended traffic.

https://www.cisco.com/c/en/us/support/docs/ip/ip-routed-protocols/47121-pbr-cmds-ce.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/lyr3-fwd/pbr/policy-based-routing-configuration-guide.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/lyr3-fwd/pbr/policy-based-routing-configuration-guide/policy-based-routing.html

✗ **A** Router(config)# access-list 100 permit tep any eq 80 192.168.12.0 0.0.0.255
This does not use the valid TCP ACL syntax: the protocol is shown as “tep,” and the source/destination fields and port qualifier are not arranged correctly. It therefore does not correctly match HTTP traffic originating from the required subnet.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/security/acls/acls-configuration-guide/access-control-lists.html

✗ **B** Router(config)# access-list 100 permit tcp 192.168.12.0 0.0.0.255 eq any 80
The source subnet is specified, but the port qualifier is malformed: “eq any 80” is not a valid way to identify destination TCP port 80. The ACL must specify the destination address and then its destination port condition.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/security/acls/acls-configuration-guide/access-control-lists.html

✓ **C** Router(config)# access-list 100 permit tcp 192.168.12.0 0.0.0.255 any eq 80
This extended ACL matches TCP traffic sourced from 192.168.12.0/24, using the correct wildcard mask, and destined to any address on port 80. Those matches identify the HTTP traffic that the route-map redirects to next hop 10.2.12.1.

Free demo · full exam: 636 questions

https://www.cisco.com/c/en/us/support/docs/ip/ip-routed-protocols/47121-pbr-cmds-ce.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/lyr3-fwd/pbr/policy-based-routing-configuration-guide.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/lyr3-fwd/pbr/policy-based-routing-configuration-guide/policy-based-routing.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/security/acls/acls-configuration-guide/access-control-lists.html

✗ **D** Router(config)# access-list 100 permit tcp any 192.168.12.0 0.0.0.255 eq 80
This puts the subnet in the destination field and allows any source, so it matches traffic directed to 192.168.12.0/24 rather than HTTP traffic originating from that subnet. It reverses the required direction of the source match.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/security/acls/acls-configuration-guide/access-control-lists.html

Dumpsgate

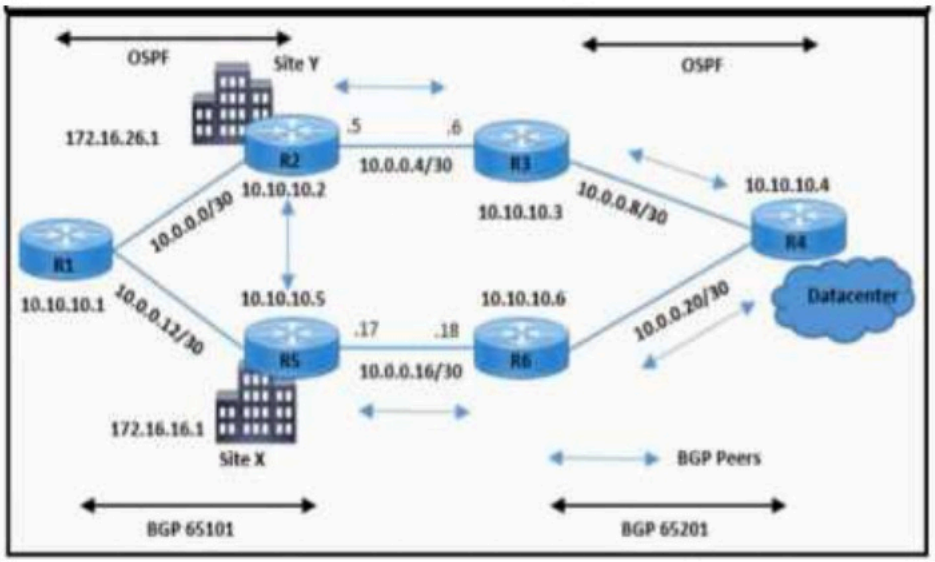
Question 9

Single choice · Domain: D1 Layer 3 Technologies › 1.11.b Neighbor relationship and authentication (next-hop, multihop, 4-byte AS, private AS, route refresh, synchronization, operation, peer group, states and timers)

Refer to the exhibit.

```
*Sep 19 08:29:51.120: BGP: 10.10.10.2 read request no-op
*Sep 19 08:29:51.124: BGP: 10.10.10.2 open failed: Connection refused by
remote host, open active delayed 12988ms (20000ms max, 60% jitter)

R2#show  ip bgp neighbors 10.10.10.5
BGP neighbor is 10.10.10.5,  remote AS 65101, internal link
  BGP version 4, remote router ID 0.0.0.0
  BGP state = Active
  Last read 00:01:18, last write 00:01:18, hold time is 15, keepalive
interval is 3 seconds
  Configured hold time is 15, keepalive interval is 3 seconds
  Minimum holdtime from neighbor is 0 seconds
  Address tracking is enabled, the RIB does have a route to 10.10.10.5
  Connections established 13; dropped 13
  Last reset 00:01:18, due to User reset
  Transport(tcp) path-mtu-discovery is enabled
  No active TCP connection
```



Refer to the exhibit. A customer reports a failure and intermittent disconnection between two office buildings, site X and site Y. The network team determines that both sites are exchanging email application traffic with the data center network. Which configuration resolves the issue between site X and site Y?

A. R5(config-router)#neighbor 10.10.10.2 next-hop-self R5(config)#router bgp 65101
B. R2(config-router)#no timers bgp 3 15 R2(config)#router bgp 65101
C. R2(config-router)#neighbor 10.10.10.5 update-source loopback 0 R2(config)#router bgp 65101
D. R2(config-router)#neighbor 10.10.10.5 next-hop-self R2(config)#router bgp 65101

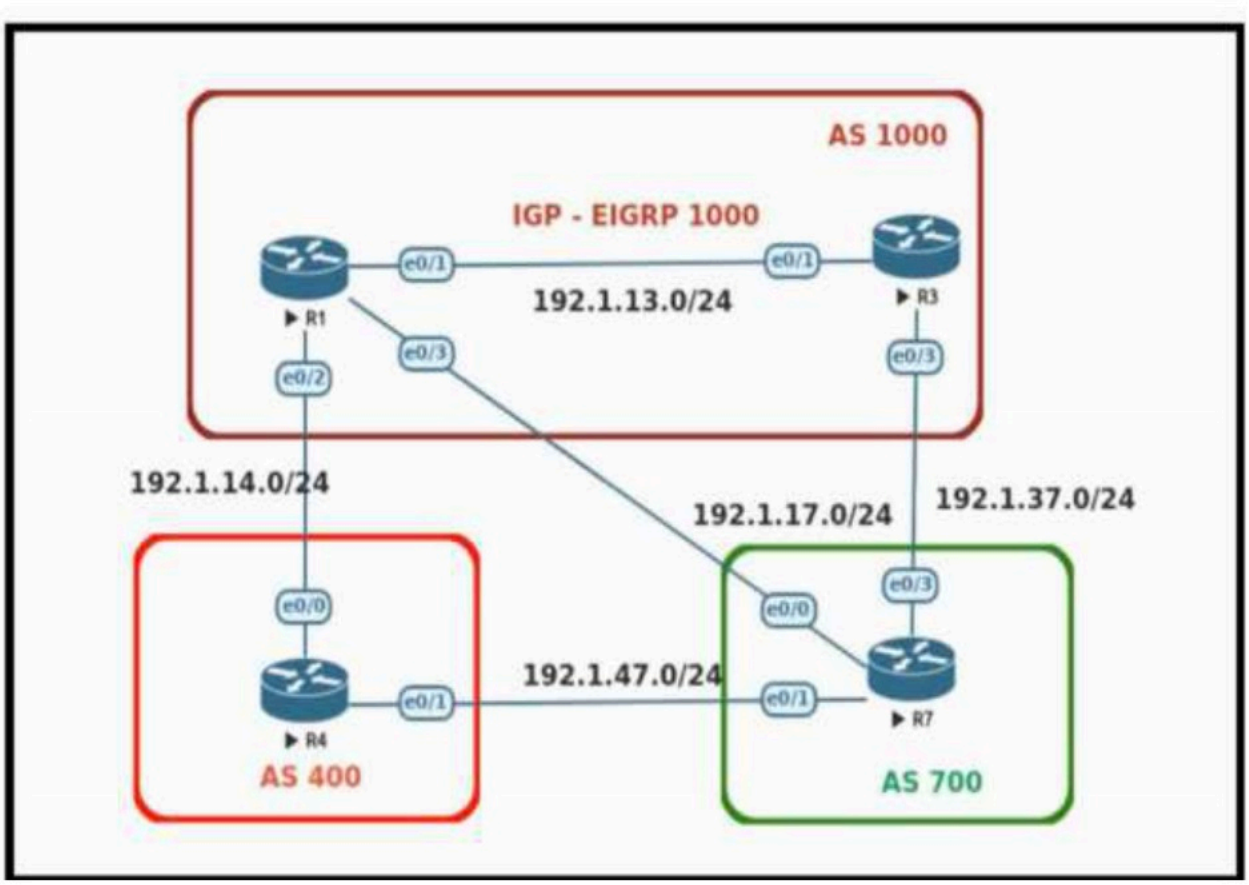
Answer: B. R2(config-router)#no timers bgp 3 15 R2(config)#router bgp 65101

Dumpsgate

Question 10

Single choice · Domain: D1 Layer 3 Technologies › 1.11.c Path preference (attributes and best-path)

Refer to the exhibit.



RPKI validation codes: V valid, I invalid, N not found

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	1.0.0.0	0.0.0.0	0		32768	i
*>	4.0.0.0	192.1.14.4	0		0	400 i
*>	11.11.11.0/24	0.0.0.0	0		32768	i
*>	44.44.44.0/24	192.1.14.4	0		0	400 i
*>	77.77.77.0/24	192.1.17.7	0		0	700 i
*>	77.77.78.0/24	192.1.17.7	0		0	700 ?
*>	77.77.79.0/24	192.1.17.7	0		0	700 ?
*>	77.77.80.0/24	192.1.17.7	0		0	700 ?

Refer to the exhibit. R1 should have BGP routes marked with an "i" (IGP origin code). Which configuration resolves the issue?

A. R1 route-map RC permit 10 match ip address 1 set metric 1
B. R1 route-map RC permit 10 match ip address 1 set origin igp
C. R7 route-map RC permit 10 match ip address 1 set origin igp
D. R7 route-map RC permit 10 match ip address 1 set metric 1

Answer: B. R1 route-map RC permit 10 match ip address 1 set origin igp

Explanation

R1 must change the BGP origin attribute on the routes so they display the IGP origin code, `i`. Option B uses `set origin igp` in a route map on R1, matching the requirement. Setting a metric changes a different BGP attribute, and configuring R7 would not directly change the routes as they appear on R1.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-n1.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp4.html

✗ **A** R1 — route-map RC permit 10; match ip address 1; set metric 1
This configures the route map on R1, but `set metric 1` changes the BGP metric (MED), not the origin attribute. The routes therefore would not be marked with the IGP origin code `i`.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-n1.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp4.html

✓ **B** R1 — route-map RC permit 10; match ip address 1; set origin igp
This sets the origin attribute to IGP for the routes matched by the route map on R1. The BGP origin code for those routes is therefore `i`, as required.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-n1.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp4.html

✗ **C** R7 — route-map RC permit 10; match ip address 1; set origin igp
Although `set origin igp` changes the origin attribute, this option applies the route map on R7 rather than on R1. It does not provide the specified correction on R1.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-n1.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp4.html

✗ **D** R7 — route-map RC permit 10; match ip address 1; set metric 1
This configures R7 and changes the metric, not the BGP origin attribute. It therefore does not make the routes on R1 display the IGP origin code `i`.

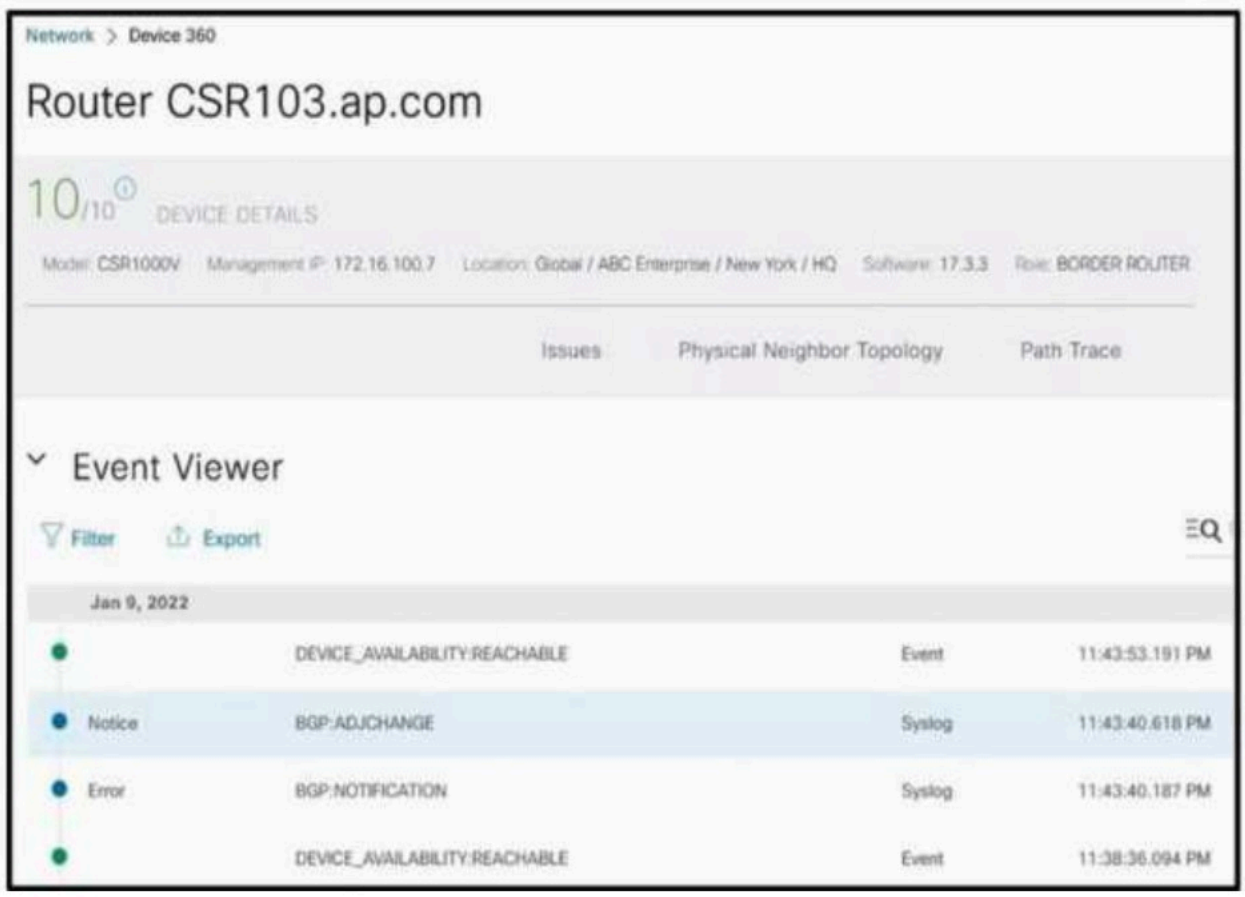
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-n1.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp4.html

Question 11

Single choice · Domain: D1 Layer 3 Technologies › 1.11.b Neighbor relationship and authentication (next-hop, multihop, 4-byte AS, private AS, route refresh, synchronization, operation, peer group, states and timers)

Refer to the exhibit.

Free demo · full exam: 636 questions




```
CSR301#sh ip bgp vpnv4 all summ
BGP router identifier 172.16.3.1, local AS number 65003
BGP table version is 29, main routing table version 29
20 network entries using 5120 bytes of memory

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.103.88.2     4           88       0       0        1    0    0 never    Active
10.103.88.6     4           88       0       0        1    0    0 never    Active
10.103.88.10    4           88       0       0        1    0    0 never    Active

CSR301#sh ip bgp summ
BGP router identifier 172.16.3.1, local AS number 65003
BGP table version is 2, main routing table version 2
1 network entries using 248 bytes of memory

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.103.88.14    4           88       0       0        1    0    0 never    Active
```

Refer to the exhibit. The BGP connection went from Established to Active when a new router was provisioned between CSR301 and its eBGP peers. Which command resolves the IPv4 peering problem?

A. neighbor 10.103.88.14 ebgp-multihop
B. neighbor 10.103.88.14 next-hop-self all
C. neighbor 10.103.88.14 disable-connected-check
D. neighbor 10.103.88.14 allowas-in

Answer: A. neighbor 10.103.88.14 ebgp-multihop

Dumpsgate

Question 12

Single choice · Domain: D1 Layer 3 Technologies › 1.4 Troubleshoot redistribution between any pair of routing protocols or routing sources

Refer to the exhibit.

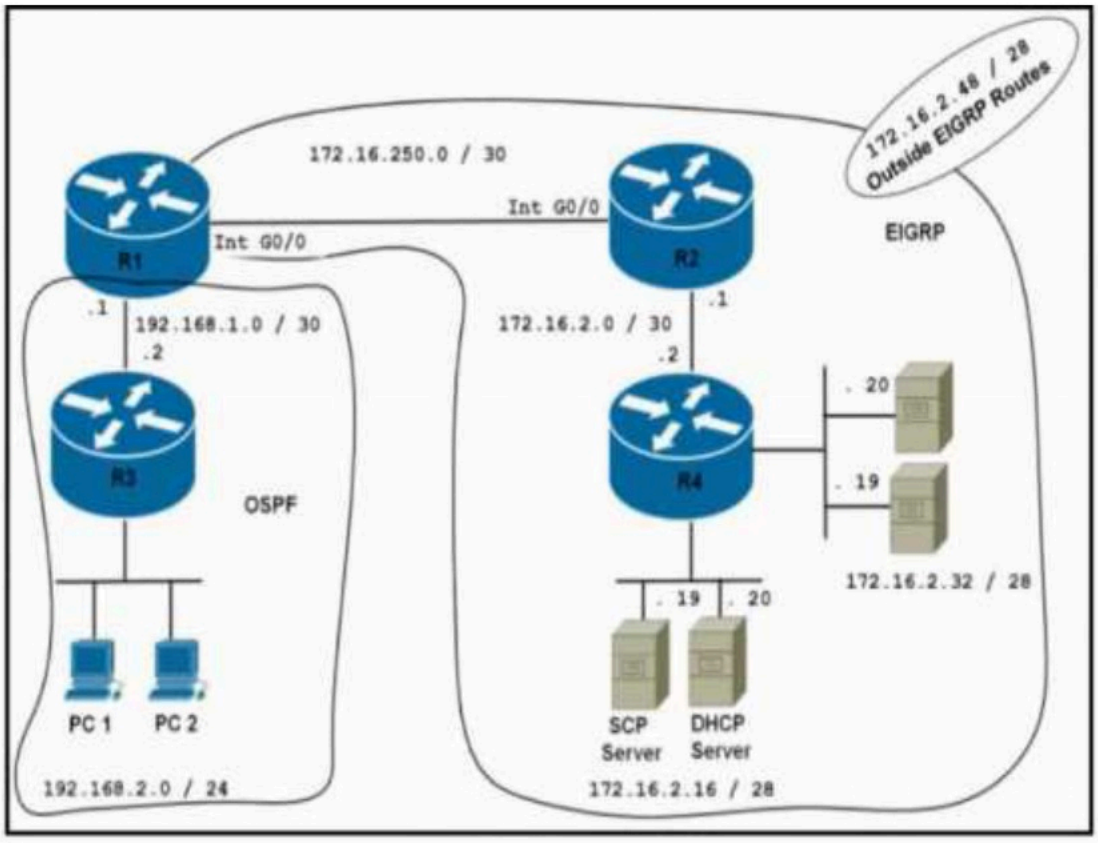
```
R1#show run | begin router eigrp 100
router eigrp 100
 network 172.16.250.0 0.0.0.3
 redistribute ospf 10 metric 1 1 1 1 1
!
router ospf 10
 network 192.168.1.0 0.0.0.3 area 0
!
ip forward-protocol nd
!
!
no ip http server
```

```
R3#traceroute 172.16.2.48
Type escape sequence to abort.
Tracing the route to 172.16.2.48
VRF info: (vrf in name/id, vrf out name/id)
  1  *  *  *
  2  *  *  *
  3  *  *  *
```

```
R3#show ip route
Gateway of last resort is not set

      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.1.0/30 is directly connected, GigabitEthernet0/1
L        192.168.1.2/32 is directly connected, GigabitEthernet0/1
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.2.0/24 is directly connected, Loopback2
L        192.168.2.33/32 is directly connected, Loopback2
      192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.3.0/24 is directly connected, Loopback1
L        192.168.3.17/32 is directly connected, Loopback1
R3#
```

```
R4#show running-config | begin router eigrp
router eigrp 100
 network 172.16.2.0 0.0.0.3
 network 172.16.2.16 0.0.0.15
 network 172.16.2.32 0.0.0.15
 redistribute static metric 100 1 1 1 1 route-map CCNP
!
ip forward-protocol nd
!
no ip http server
no ip http secure-server
ip route 172.16.2.48 255.255.255.240 172.16.2.34
!
route-map CCNP permit 10
 match ip address 10
 set tag 200
!
!
access-list 10 permit 172.16.2.48 0.0.0.15
```



Refer to the exhibit. An engineer must troubleshoot a connectivity issue impacting the redistribution

Free demo · full exam: 636 questions

of the subnet 172.16.2.48/28 into the OSPF domain. Which configuration on router R1 advertises this subnet into OSPF?

A. R1(config-router)#redistribute eigrp 100 subnets route-map CCNP R1(config)#router ospf 10 R1(config-route-map)#match route-type level-2 R1(config)#route-map CCNP permit 10

B. R1(config-router)#redistribute eigrp 100 subnets route-map CCNP R1(config)#router ospf 10 R1(config)#route-map CCNP permit 10 R1(config-route-map)#match tag 200 R1(config)#route-map CCNP deny 10

C. R1(config-router)#redistribute eigrp 100 subnets route-map CCNP R1(config)#router ospf 10 R1(config-route-map)#match route-type internal R1(config)#route-map CCNP permit 10

D. R1(config-router)#redistribute eigrp 100 subnets route-map CCNP R1(config)#router ospf 10 R1(config-route-map)#exit R1(config-route-map)#match tag 200 R1(config)#route-map CCNP permit 10

Answer: B. R1(config-router)#redistribute eigrp 100 subnets route-map CCNP R1(config)#router ospf 10 R1(config)#route-map CCNP permit 10 R1(config-route-map)#match tag 200 R1(config)#route-map CCNP deny 10

Explanation

R1 must redistribute EIGRP 100 into OSPF process 10, while the route map must select the route carrying tag 200. The exhibit shows the EIGRP route map uses `match ip address 10` and sets tag 200, and ACL 10 identifies 172.16.2.48/28. Option B configures the redistribution route map to match that tag and includes a deny sequence afterward, so the tagged route is permitted into OSPF while other unmatched routes are not redistributed by this map.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/49111-route-map-bestp.html
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/ios-xe-17/223145-best-practices-for-route-map.html

✗ **A** redistribute eigrp 100 subnets route-map CCNP; match route-type level-2; route-map CCNP permit 10
The subnet is an EIGRP route, so matching OSPF route type level-2 does not select it for redistribution. The route map must match the tag set on the EIGRP route in the exhibit, which is 200.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/49111-route-map-bestp.html
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

✓ **B** redistribute eigrp 100 subnets route-map CCNP; route-map CCNP permit 10; match tag 200; route-map CCNP deny 10
The exhibit's EIGRP configuration applies route map CCNP and sets tag 200 on the route identified by ACL 10, which permits 172.16.2.48/28. Matching tag 200 in OSPF sequence 10 permits that route to be redistributed; the later deny sequence prevents unmatched routes from passing the route map.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/49111-route-map-bestp.html
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/ios-xe-17/223145-best-practices-for-route-map.html

✗ **C** redistribute eigrp 100 subnets route-map CCNP; match route-type internal; route-map CCNP permit 10
This matches route type internal rather than the tag 200 assigned by the EIGRP route map shown in the exhibit. It therefore does not specifically select the subnet required for redistribution.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/49111-route-map-bestp.html
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

✗ **D** redistribute eigrp 100 subnets route-map CCNP; exit; match tag 200; route-map CCNP permit 10
The `exit` command leaves route-map configuration mode, so the following `match tag 200` is not configured within route-map sequence CCNP permit 10. Without that match in the route map, this configuration does not select the tagged subnet as required.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/49111-route-map-bestp.html
https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/ios-xe-17/223145-best-practices-for-route-map.html

Dumpsgate

Question 13

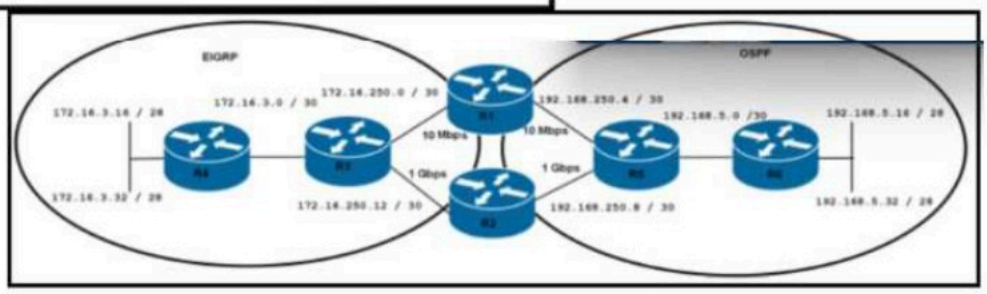
Single choice · Domain: D1 Layer 3 Technologies › 1.4 Troubleshoot redistribution between any pair of routing protocols or routing sources

Refer to the exhibit.

```
router ospf 1
R1#sh run | b eigrp 100 subnets
router eigrp 168.250.4 0.0.0.3 area 0
 network 172.
 redistributeotocol nd
!
router ospf 1
 redistribute eigrp 100 subnets
 network 192.168.250.4 0.0.0.3 area 0
!
ip forward-protocol nd
!
no ip http server
no ip http secure-server
```

```
router ospf 1
R2#sh run | b eigrp 100 metric 100 subnets
router eigrp 168.250.8
 network 172.          0.0.0.3
 redistributcotocol ndetric 1 1 1 1 1
!
router ospf 1
 redistribute eigrp 100 metric 100 subnets
 network 192.168.250.8 0.0.0.3 area 0
!
ip forward-protocol nd
!
no ip http server
no ip http secure-server
```

```
R6#traceroute 172.16.3.17
Type escape sequence to abort.
Tracing the route to 172.16.3.17
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.5.1 45 msec 38 msec 21 msec
  2 192.168.250.6 165 msec 19 msec 28 msec
  3 172.16.250.2 38 msec 40 msec 45 msec
  4 172.16.3.2 30 msec 58 msec *
R6#
```



Refer to the exhibit. An engineer is troubleshooting suboptimal communication from the 192.168.5.32/28 subnet to the 172.16.3.16/28 segment using the slowest links. Which configuration resolves the suboptimal routing issue?

A. R1(config-router)#router elgrp 100 R1(config-router)#redistribute ospf 1 metric 1000000 1111

B. R2{config-router)#router ospf 1 R2{config-roiiter)#default-metric 1 R1(config-router)#router ospf 1 R1(config-router)#default-metric 10

C. R2(config-router)#ro liter eigrp 100 R2(config-router)#redistribute ospf 1 metric 1000000 1111

D. R2(config-router)#ro uter ospf 1 R2(config-router)#default-metric 10 R1(config-router)#router ospf 1 R1(config-router)#default-metric 1

Free demo · full exam: 636 questions

Answer: B. R2{config-router)#router ospf 1 R2{config-roiiter)#default-metric 1 R1(config-router)#router ospf 1 R1(config-router)#default-metric 10

Explanation

The issue is that OSPF redistribution into EIGRP on R2 uses a metric that does not make EIGRP prefer the intended slower links. Option B sets the default metric for OSPF on R2 to 1 and for OSPF on R1 to 10, aligning the redistributed routes' metrics so traffic from 192.168.5.32/28 toward 172.16.3.16/28 takes the intended path. EIGRP requires a metric when routes are redistributed into it, and the metric influences route selection [1].

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

✗ **A** R1: configure EIGRP 100 and redistribute OSPF 1 with metric 1000000 1111.
This applies the redistribution metric on R1, but the exhibit shows the relevant OSPF-to-EIGRP redistribution on R2. Configuring the metric on the wrong router does not correct the route preference for the stated source and destination.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

✓ **B** R2: configure OSPF 1 with default-metric 1; R1: configure OSPF 1 with default-metric 10.
These commands set the OSPF redistribution default metrics on the two routers, making the route costs differ as needed to influence EIGRP's path selection. This addresses the suboptimal route so communication uses the intended slower links; redistributed EIGRP routes need an appropriate metric [1].

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

✗ **C** R2: configure EIGRP 100 and redistribute OSPF 1 with metric 1000000 1111.
This adds an explicit EIGRP redistribution metric on R2, but it does not match the configuration required to distinguish the OSPF redistribution costs between R2 and R1. The scenario's correction is to set the differing OSPF default metrics shown in option B [1].

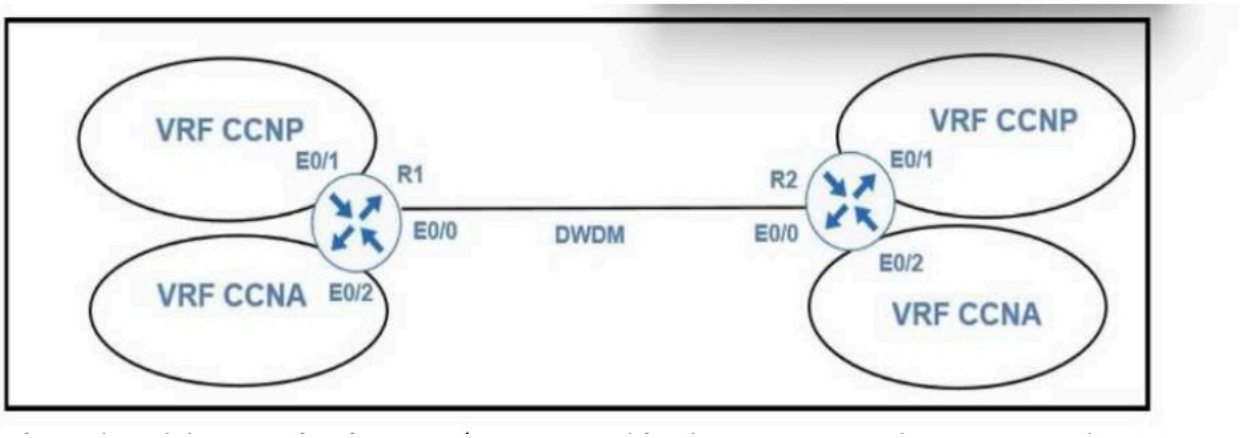
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

✗ **D** R2: configure OSPF 1 with default-metric 10; R1: configure OSPF 1 with default-metric 1.
This reverses the metric assignments in option B, favoring the opposite router for redistributed routes. As a result, it does not resolve the stated suboptimal path over the slowest links.

Question 14

Single choice · Domain: D1 Layer 3 Technologies › 1.7 Configure and verify VRF-Lite

Refer to the exhibit.



Refer to the exhibit. A prefix of 10.1.1.0/31 is assigned for the DWDM circuit between R1 and R2 to connect the two attached VRFs on each router. Which configuration provides the required connectivity on R1?

A. interface EthernetO/0 no ip address ! interface Eth erne to. -0.10 encapsulation dotlQ 10 ip vrf forwarding CCNP ip address 10.1.1.0 255.255.255.252 ! interface EthernetO/0.20 encapsulation dotlQ 20 ip vrf forwarding CCNA ip address 10.1.1.0 255.255.255.252

B. interface EthernetO/0 ip address 10.1.1.0 255.255.255.252 ! Interface Loopback 0 ip vrf forwarding CCNP ! Interface Loopback 1 ip vrf forwarding CCNA ! interface TunnelO tunnel source Loopback 0 tunnel mode gre IP Ip vrf forwarding CCNP tunnel destination 10.1.1.1 interface Tunnell tunnel source Loopback 1 tunnel mode gre IP Ip vrf forwarding CCNA tunnel destination 10.1.1.1

C. Interface ElhernetO/0 ip address 10.1.1.0 255.255.255.254 ! Interface Loopback 0 Ip vi f forwarding CCNP ! Interface Loopback 1 ip vrf forwarding CCNA ! interface TunnelO tunnel source Loopback 0 tunnel mode gre IP Ip vrf forwarding CCNP tunnel destination 10.1.1.1 ! interface Tunned tunnel source Loopback 1 tunnel mode gre IP Ip vrf forwarding CCNA tunnel destination 10.1.1.1

D. Interface EthernetO/0 no ip address ! interface EthernetO/0.10 encapsulation dotlQ 10 ip vrf forwarding CCNP ip address 10.1.1.0 255.255.255.254 ! interface EthernetO/0.20 encapsulation dotlQ 20 Ip vrf forwarding CCNA ip address 10.1.1.0 255.255.255.254

Answer: D. Interface EthernetO/0 no ip address ! interface EthernetO/0.10 encapsulation dotlQ 10 ip vrf forwarding CCNP ip address 10.1.1.0 255.255.255.254 ! interface EthernetO/0.20 encapsulation dotlQ 20 Ip vrf forwarding CCNA ip address 10.1.1.0 255.255.255.254

Explanation

D uses 802.1Q subinterfaces to place the DWDM link into the two separate VRFs, CCNP and CCNA. Each subinterface uses the assigned 10.1.1.0/31 prefix, so the configuration matches both the VRF-separation requirement and the specified point-to-point prefix. [6]

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iro-vrf-lite-pe-ce.html

Free demo · full exam: 636 questions

✗ **A** Physical Ethernet0/0 has no address; subinterfaces .10 and .20 are assigned to CCNP and CCNA, but both use 10.1.1.0/30.
The subinterfaces correctly separate the traffic into the two VRFs, but /30 does not match the assigned 10.1.1.0/31 prefix. The mask therefore fails the scenario's addressing requirement.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iro-vrf-lite-pe-ce.html

✗ **B** Ethernet0/0 uses 10.1.1.0/30, with loopbacks and GRE tunnels configured for the two VRFs.
This uses a /30 rather than the assigned /31. It also configures GRE tunnels instead of assigning the DWDM circuit's two VRF connections to tagged subinterfaces.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/lyr3-fwd/gre/gre-configuration-guide/m-gre-tunnel-ip-source-and-destination-vrf-membership.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iro-vrf-lite-pe-ce.html

✗ **C** Ethernet0/0 uses 10.1.1.0/31, with loopbacks and GRE tunnels configured for the two VRFs.
Although the physical interface uses the required /31 mask, it is in the default routing context rather than being separated into the CCNP and CCNA VRFs on subinterfaces. The GRE tunnel configuration shown does not provide the required direct VRF connectivity over the DWDM link.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/lyr3-fwd/gre/gre-configuration-guide/m-gre-tunnel-ip-source-and-destination-vrf-membership.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iro-vrf-lite-pe-ce.html

✓ **D** Physical Ethernet0/0 has no address; subinterfaces .10 and .20 use dot1Q encapsulation, belong to CCNP and CCNA respectively, and use 10.1.1.0/31.
This configuration provides a separate tagged subinterface in each VRF, allowing the two attached VRFs to use the DWDM connection independently. Both subinterfaces use the specified 10.1.1.0/31 prefix, meeting the addressing requirement.

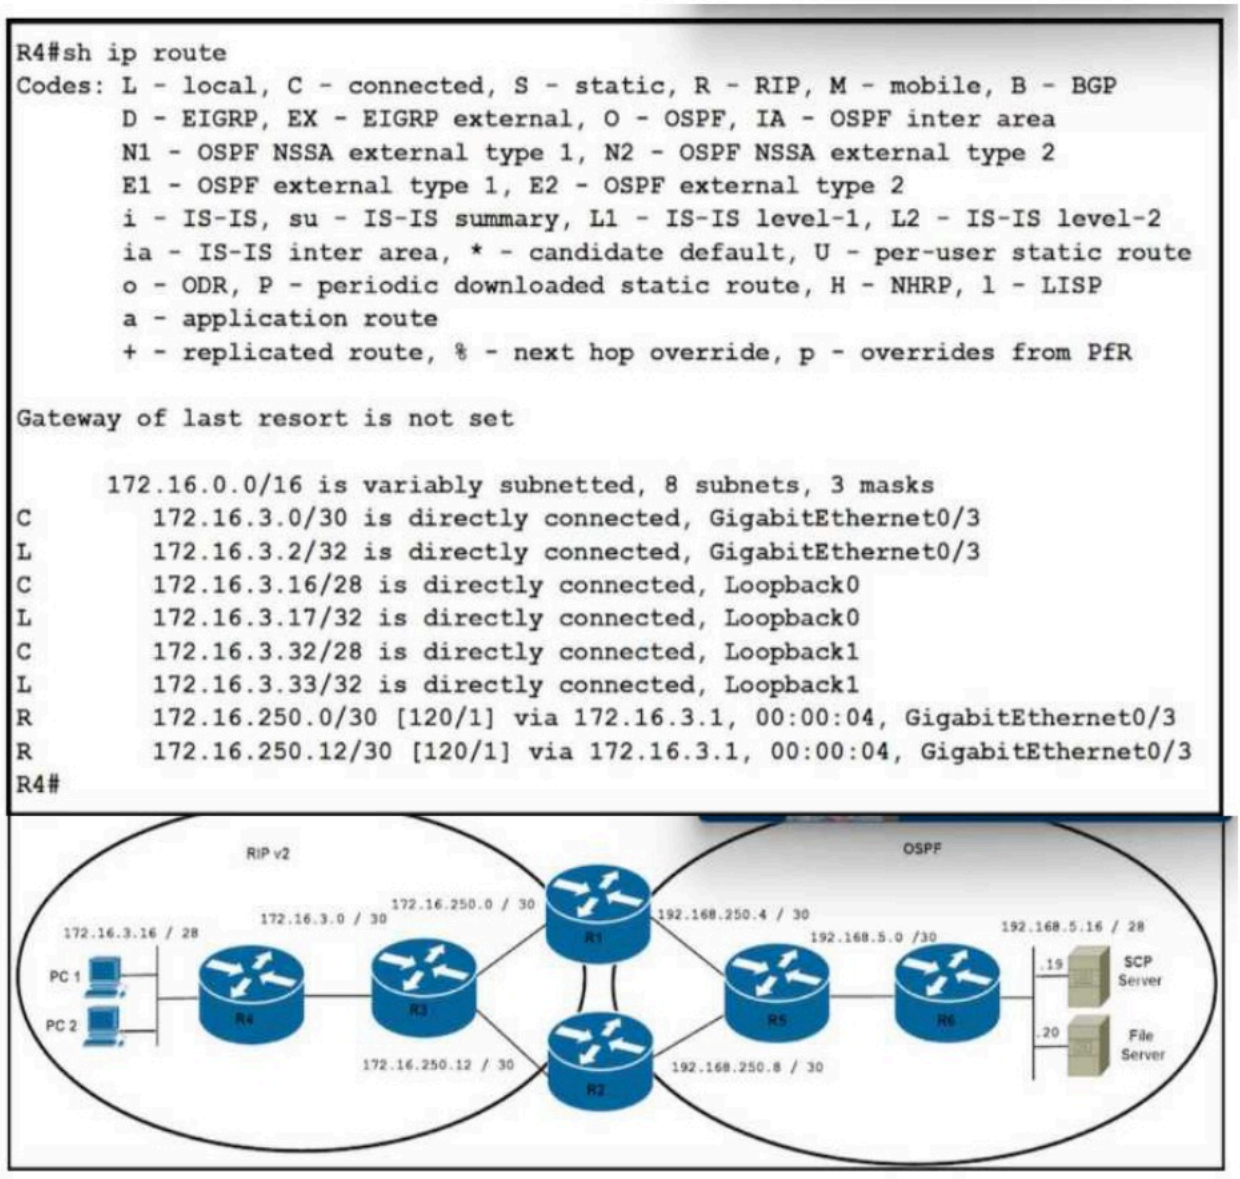
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iro-vrf-lite-pe-ce.html

Dumpsgate

Question 15

Single choice · Domain: D1 Layer 3 Technologies › 1.4 Troubleshoot redistribution between any pair of routing protocols or routing sources

Refer to the exhibit.



```
R4#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 8 subnets, 3 masks
C        172.16.3.0/30 is directly connected, GigabitEthernet0/3
L        172.16.3.2/32 is directly connected, GigabitEthernet0/3
C        172.16.3.16/28 is directly connected, Loopback0
L        172.16.3.17/32 is directly connected, Loopback0
C        172.16.3.32/28 is directly connected, Loopback1
L        172.16.3.33/32 is directly connected, Loopback1
R        172.16.250.0/30 [120/1] via 172.16.3.1, 00:00:04, GigabitEthernet0/3
R        172.16.250.12/30 [120/1] via 172.16.3.1, 00:00:04, GigabitEthernet0/3
R4#
```


Refer to the exhibit. An engineer troubleshoots a connectivity problem that is impacting the communication from the users at segment 172.16.3.16/28 to the server farm at 192.168.5.16/ 28. Which configuration resolves the issue on router R1?

A. router ospf 1 redistribute rip metric 14
B. router rip redistribute ospf 1 metric 16
C. router ospf 1 redistribute rip metric 16
D. router rip redistribute ospf 1 metric 14

Free demo · full exam: 636 questions

Answer: D. router rip redistribute ospf 1 metric 14

Explanation

R1 connects the RIP and OSPF domains, and the server farm is in the OSPF domain. To let the RIP side learn routes to the OSPF networks, R1 must redistribute OSPF routes into RIP. RIP metric 14 is a usable hop count, while 16 represents an unreachable route, so option D resolves the issue.

https://www.cisco.com/c/en/us/td/docs/ios/iproute_rip/command/reference/irr_book/irr_rip.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_rip/command/irr-cr-book.pdf

✗ **A** router ospf 1; redistribute rip metric 14
This redistributes RIP routes into OSPF, not OSPF server-farm routes into RIP. It therefore does not provide the needed OSPF routes to the RIP domain; the selected RIP metric also is not the issue this direction requires.

https://www.cisco.com/c/en/us/td/docs/ios/iproute_rip/command/reference/irr_book/irr_rip.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book.pdf

✗ **B** router rip; redistribute ospf 1 metric 16
This uses the correct redistribution direction, but RIP metric 16 means the route is unreachable. As a result, RIP routers cannot use the redistributed OSPF routes to reach the server farm.

https://www.cisco.com/c/en/us/td/docs/ios/iproute_rip/command/reference/irr_book/irr_rip.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_rip/command/irr-cr-book.pdf

✗ **C** router ospf 1; redistribute rip metric 16
This redistributes routes in the wrong direction for advertising the OSPF server-farm networks to the RIP domain. Additionally, the metric 16 is the RIP unreachable value rather than a usable RIP hop count.

https://www.cisco.com/c/en/us/td/docs/ios/iproute_rip/command/reference/irr_book/irr_rip.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_rip/command/irr-cr-book.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book.pdf

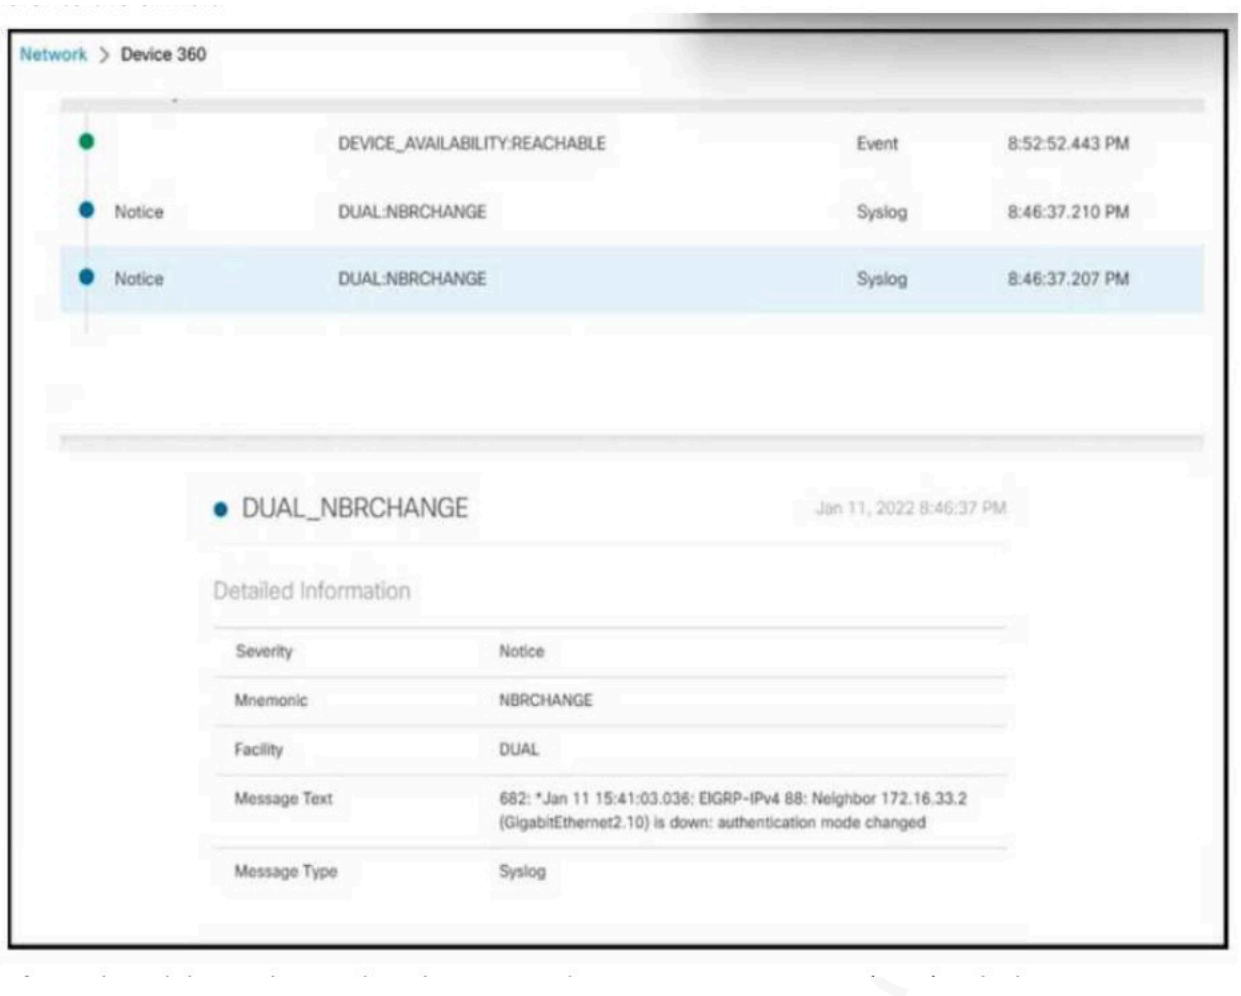
✓ **D** router rip; redistribute ospf 1 metric 14
This redistributes OSPF routes, including the server-farm network, into RIP so the RIP domain can learn routes toward the servers. Metric 14 is below RIP's maximum usable metric of 15; metric 16 would mark the route unreachable.

https://www.cisco.com/c/en/us/td/docs/ios/iproute_rip/command/reference/irr_book/irr_rip.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_rip/command/irr-cr-book.pdf

Question 16

Single choice · Domain: D1 Layer 3 Technologies › 1.9.b Neighbor relationship and authentication

Refer to the exhibit.



Refer to the exhibit. R1 lost its directly connected EIGRP peer 172.16.33.2 (SW1). Which configuration resolves the issue?

A. key chain EIGRP keyl key-string Cisco ! interface GigabitEthernet 2.10 Ip authentication mode elgrp 88 md5 Ip authentication key-chain elgrp 88 EIGRP

B. key chain EIGRP keyl key-string Cisco ! interface GigabitEthernet 2 ip authentication mode elgrp 88 md5 ip authentication key-chain elgrp 88 Cisco

C. key chain EIGRP keyl key-string Cisco ! interface GigabitEthernet 2 ip authentication mode elgrp 88 md5 ip authentication key-chain elgrp 88 EIGRP

D. key chain EIGRP keyl key-string Cisco ! interface GigabitEthernet 2.10 ip authentication mode eigrp 88 md5 ip authentication key-chain eigrp 88 Cisco

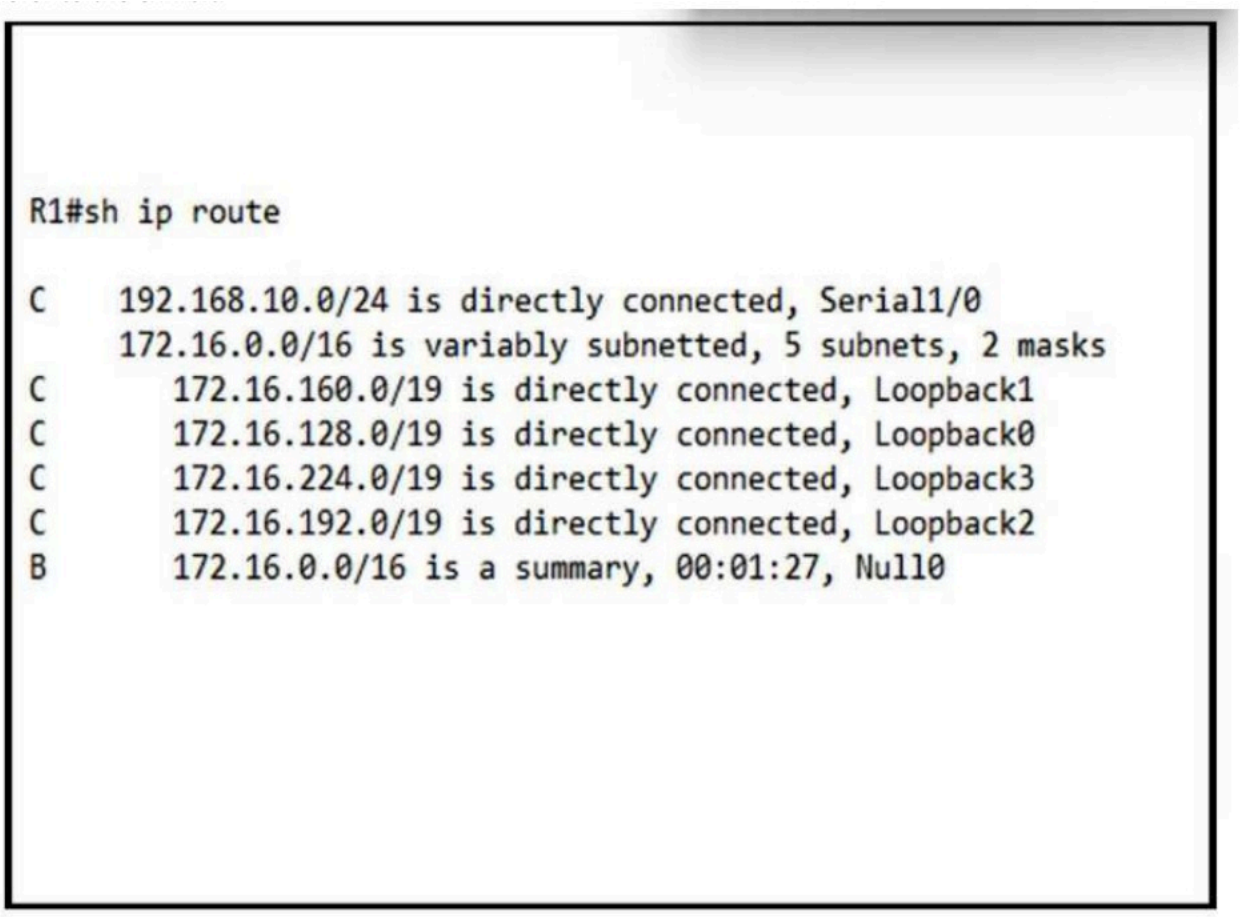
Answer: A. key chain EIGRP keyl key-string Cisco ! interface GigabitEthernet 2.10 Ip authentication mode elgrp 88 md5 Ip authentication key-chain elgrp 88 EIGRP

Dumpsgate

Question 17

Single choice · Domain: D1 Layer 3 Technologies › 1.5 Troubleshoot manual and auto-summarization with any routing protocol

Refer to the exhibit.



Refer to the exhibit. Which configuration advertises more specific routes to R1 without sending a BGP summary route?

A. R1#configure terminal R1 (config)#rouler BGP 100 R1 (conflg-router)#no auto-summary
B. R1#configure terminal R1 (config)#router BGP 100 R1 (config-router)#auto-summary
C. R2#configure terminal R2 (conflg)#router BGP 100 R2 (config-rouler)»no auto-summary
D. R2#configure terminal R2 (config)#router BGP 100 R2 (config-router)#auto-summary

Answer: C. R2#configure terminal R2 (conflg)#router BGP 100 R2 (config-rouler)»no auto-summary

Explanation

The exhibit shows R1 receiving a BGP summary route for 172.16.0.0/16 from R2, while R2 has several connected /19 networks within that range. On R2, `no auto-summary` disables BGP classful auto-

Free demo · full exam: 636 questions

summarization, allowing the more-specific routes to be advertised rather than the summary route. Therefore, option C meets both requirements: configure the setting on R2 and disable auto-summary.

https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp1.html
https://team-development.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5816-bgpfaq-5816.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5816-bgpfaq-5816.html

✗ **A** Configure `no auto-summary` under BGP on R1.
The setting is applied on R1, but R2 is the router advertising the routes to R1. Changing R1's BGP auto-summary behavior does not stop R2 from sending its summary route or cause R2 to advertise its more-specific routes.

https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp1.html
https://team-development.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5816-bgpfaq-5816.html

✗ **B** Configure `auto-summary` under BGP on R1.
This configures the receiving router, R1, rather than the advertising router, R2. Also, enabling auto-summary is contrary to the requirement to avoid sending a BGP summary route.

https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp1.html
https://team-development.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5816-bgpfaq-5816.html

✓ **C** Configure `no auto-summary` under BGP on R2.
R2 is the router advertising routes to R1, so the change belongs on R2. Disabling auto-summary prevents classful summarization and permits the more-specific routes shown in the exhibit to be advertised instead of the 172.16.0.0/16 BGP summary.

https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp1.html
https://team-development.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5816-bgpfaq-5816.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5816-bgpfaq-5816.html

✗ **D** Configure `auto-summary` under BGP on R2.
Although the command is entered on the advertising router, enabling auto-summary would promote summarization rather than prevent a summary route from being sent. It therefore conflicts with the requirement to advertise more-specific routes without the BGP summary route.

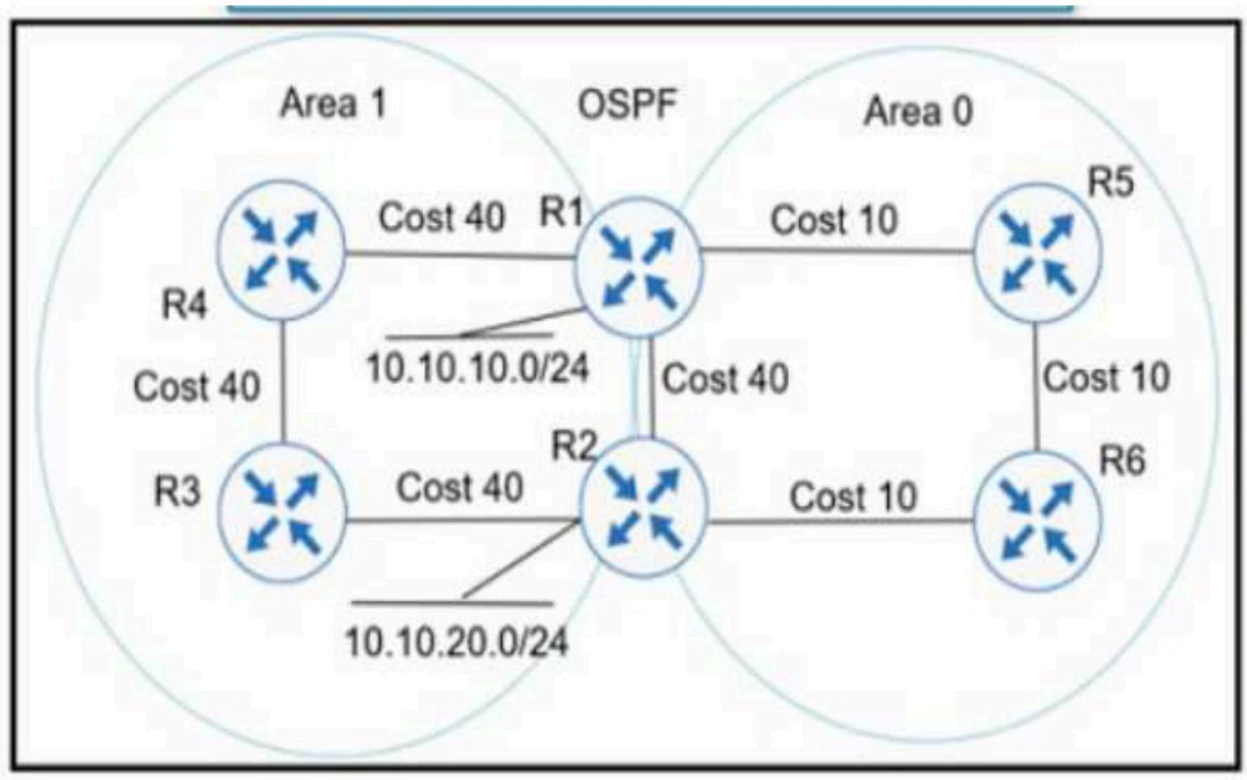
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp1.html
https://team-development.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5816-bgpfaq-5816.html

Dumpsgate

Question 18

Single choice · Domain: D1 Layer 3 Technologies › 1.10.c Network types, area types, and router types

Refer to the exhibit.



Refer to the exhibit Which action ensures that 10 10 10 0/24 reaches 10 10 20 0/24 through the direct link between R1 and R2?

A. Configure R1 and R2 LAN links as nonpassive.
B. Configure R1 and R2 links under area 1
C. Configure OSPF link cost to 1 between R1 and R2
D. Configure OSPF path cost to 3 between R1 and R2

Answer: B. Configure R1 and R2 links under area 1

Explanation

The two networks are in Area 1, while the direct R1–R2 link is shown in Area 0. OSPF neighbors must share an area on a link, and OSPF route calculation is performed within an area; placing the R1–R2 link in Area 1 allows the two Area 1 networks to communicate over that direct link. Changing link costs would not resolve the area mismatch.

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/118879-configure-ospf-00.html
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **A** Configure R1 and R2 LAN links as nonpassive.

Free demo · full exam: 636 questions

The LANs already have OSPF routes, and making LAN interfaces nonpassive would affect neighbor formation on those LANs. It does not correct the area mismatch shown on the direct R1–R2 link.

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/118879-configure-ospf-00.html

✓ **B** Configure R1 and R2 links under area 1
The networks 10.10.10.0/24 and 10.10.20.0/24 are associated with Area 1, but the direct R1–R2 link is shown in Area 0. Configuring that link in Area 1 places both endpoints in the same OSPF area as the networks, enabling the intended direct path.

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/118879-configure-ospf-00.html
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **C** Configure OSPF link cost to 1 between R1 and R2
Lowering the cost cannot make an OSPF adjacency form when the two ends of the link are configured for different areas. The area mismatch must be fixed before the link cost can influence path selection.

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/118879-configure-ospf-00.html
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html?trk=article-ssr-frontend-pulse_little-text-block

✗ **D** Configure OSPF path cost to 3 between R1 and R2
Changing the path cost does not resolve the area mismatch on the R1–R2 link. OSPF must first have the link in a common area before cost can select it as the route between the two networks.

https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/118879-configure-ospf-00.html
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html?trk=article-ssr-frontend-pulse_little-text-block

Dumpsgate

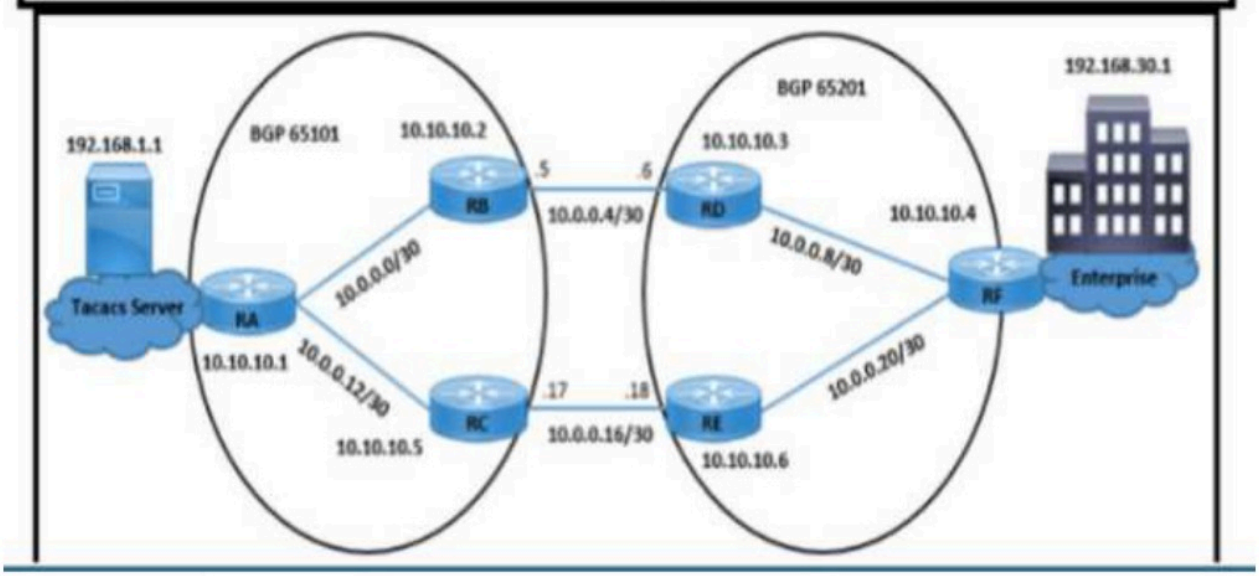
Question 19

Single choice · Domain: D1 Layer 3 Technologies › 1.11.c Path preference (attributes and best-path)

Refer to the exhibit.

```
RF#traceroute 192.168.1.1
  1 10.0.0.9 40 msec 28 msec 24 msec
  2  *  *  *
  3  *  *  *
RE#show  ip prefix-list detail
Prefix-list with the last deletion/insertion: Customer
ip prefix-list Customer:
   count: 2, range entries: 1, sequences: 5 - 10, refcount: 3
   seq 5 deny 192.168.1.1/32 (hit count: 5, refcount: 1)
   seq 10 permit 0.0.0.0/0 le 32 (hit count: 26, refcount: 1)

RC#show  ip prefix-list detail
Prefix-list with the last deletion/insertion: Customer
ip prefix-list Customer:
   count: 1, range entries: 1, sequences: 10 - 10, refcount: 4
   seq 10 permit 0.0.0.0/0 le 32 (hit count: 7, refcount: 1)
```



Refer to the exhibit The enterprise users fail to authenticate with the TACACS server when a direct fiber link fails between RB and RD The NOC team observes Users connected on AS65201 fail to authenticate with TACACS server 192 168 1 1 Users connected on AS65101 successfully authenticate with TACACS server 192 168 1 1 All AS65101 and AS65201 users are configured to authenticate with the TACACS server Which configuration resolves the issue?

A.

```
RC(config)# ip prefix-list Customer seq 5 permit 192.168.30.1/32
```

B.

```
RC(config)#router bgp 65101
RC(config-router)# neighbor 10.0.0.18 prefix-list Customer in
```

C.

```
RF(config)#no ip prefix-list Customer seq 5 deny 192.168.1.1/32
```

D.

```
RF(config)#router bgp 65201
RF(config-router)# neighbor 10.0.0.17 prefix-list Customer out
```

Answer: C

Explanation

Option C removes the deny entry for 192.168.1.1/32 from RF's Customer prefix list. That deny prevents the TACACS server's route from being advertised through the alternate path when the RB–RD link fails, so AS65201 users cannot reach the server. Removing it allows the route to pass the prefix list and supports authentication over the remaining path.

✗ **A** RC(config)# ip prefix-list Customer seq 5 permit 192.168.30.1/32
This adds a permit for 192.168.30.1/32, not for the TACACS server at 192.168.1.1. It does not remove the deny that blocks the TACACS server's route on RF.

✗ **B** RC(config)#router bgp 65101; RC(config-router)# neighbor 10.0.0.18 prefix-list Customer in
This applies the Customer prefix list inbound on RC from neighbor 10.0.0.18. It does not correct the deny entry in RF's prefix list that blocks the TACACS server route from being passed over the alternate path.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-n1.html

✓ **C** RF(config)#no ip prefix-list Customer seq 5 deny 192.168.1.1/32
The displayed RF prefix list explicitly denies 192.168.1.1/32 before permitting other prefixes. Removing that deny allows the TACACS server route to pass the list, so users in AS65201 can reach the server when the direct RB–RD link is down.

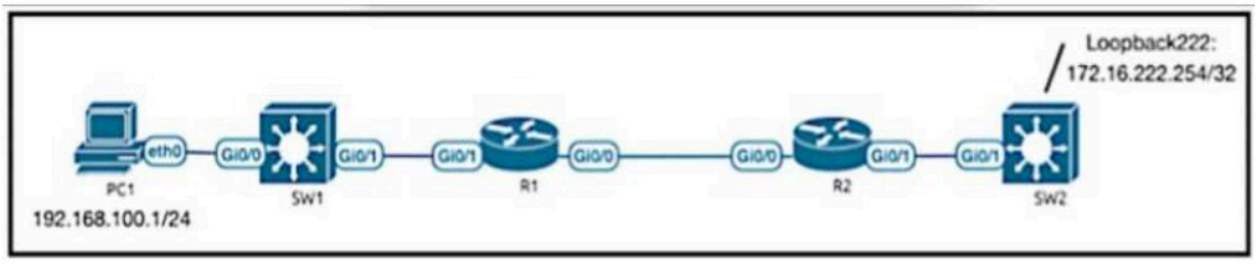
✗ **D** RF(config)#router bgp 65201; RF(config-router)# neighbor 10.0.0.17 prefix-list Customer out
This applies the existing Customer prefix list outbound to a BGP neighbor; it does not remove the entry that denies 192.168.1.1/32. The TACACS server route would therefore remain blocked by the deny in RF's list.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-n1.html

Question 20

Single choice · Domain: D1 Layer 3 Technologies › 1.4 Troubleshoot redistribution between any pair of routing protocols or routing sources

Refer to the exhibit.




```
R1#show run | sec router eigrp
router eigrp VR1
 !
 address-family ipv4 unicast autonomous-system 1
  !
  topology base
  exit-af-topology
  network 172.16.1.1 0.0.0.0
  network 192.168.100.0
  network 192.168.200.0
  network 192.168.255.91 0.0.0.0
 exit-address-family

R2(config)#do show run | sec router eigrp
router eigrp 1
 network 172.16.1.2 0.0.0.0
 network 172.16.222.0 0.0.0.255
 network 192.168.222.254 0.0.0.0
```

Refer to the exhibit R2 can reach Loopback222, but R1 SW1 and PC1 cannot communicate with 172.16.222 254 R1 and R2 configurations are shown here Which EIGRP configuration command resolves the issue?

A. R2(config-router) # redistribute static
B. R1(conftg-router)# network 172.16.222.254 0.0.0.0
C. R1 (config-router)# network 172.16.222.264 255.255.255.255
D. R1(config-router)# redistribute static

Answer: A. R2(config-router) # redistribute static

Explanation

R2's EIGRP configuration includes the loopback network, so R2 can reach Loopback222, but the loopback is a static route and is not being advertised to R1. Enabling static-route redistribution under R2's EIGRP process advertises that route, allowing R1 and the devices behind it to reach 172.16.222.254. Reference 2 covers redistribution into EIGRP.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

✓ **A** R2(config-router)# redistribute static
This enables R2 to redistribute its static route for Loopback222 into EIGRP. R1 can then learn the route and forward traffic from SW1 and PC1 to 172.16.222.254.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

✗ **B** R1(config-router)# network 172.16.222.254 0.0.0.0
This would tell R1 to include a matching local interface in EIGRP; it does not make R2 advertise its static route for the loopback. The destination address is on R2, not R1.

✗ **C** R1(config-router)# network 172.16.222.264 255.255.255.255
The address shown is 172.16.222.264, which is not a valid IPv4 address, and this command would not redistribute R2's static loopback route. The missing route must be advertised by R2.

✗ **D** R1(config-router)# redistribute static
Redistributing static routes on R1 would advertise R1's own static routes, not R2's route to Loopback222. The required change is to redistribute the loopback's static route on R2.

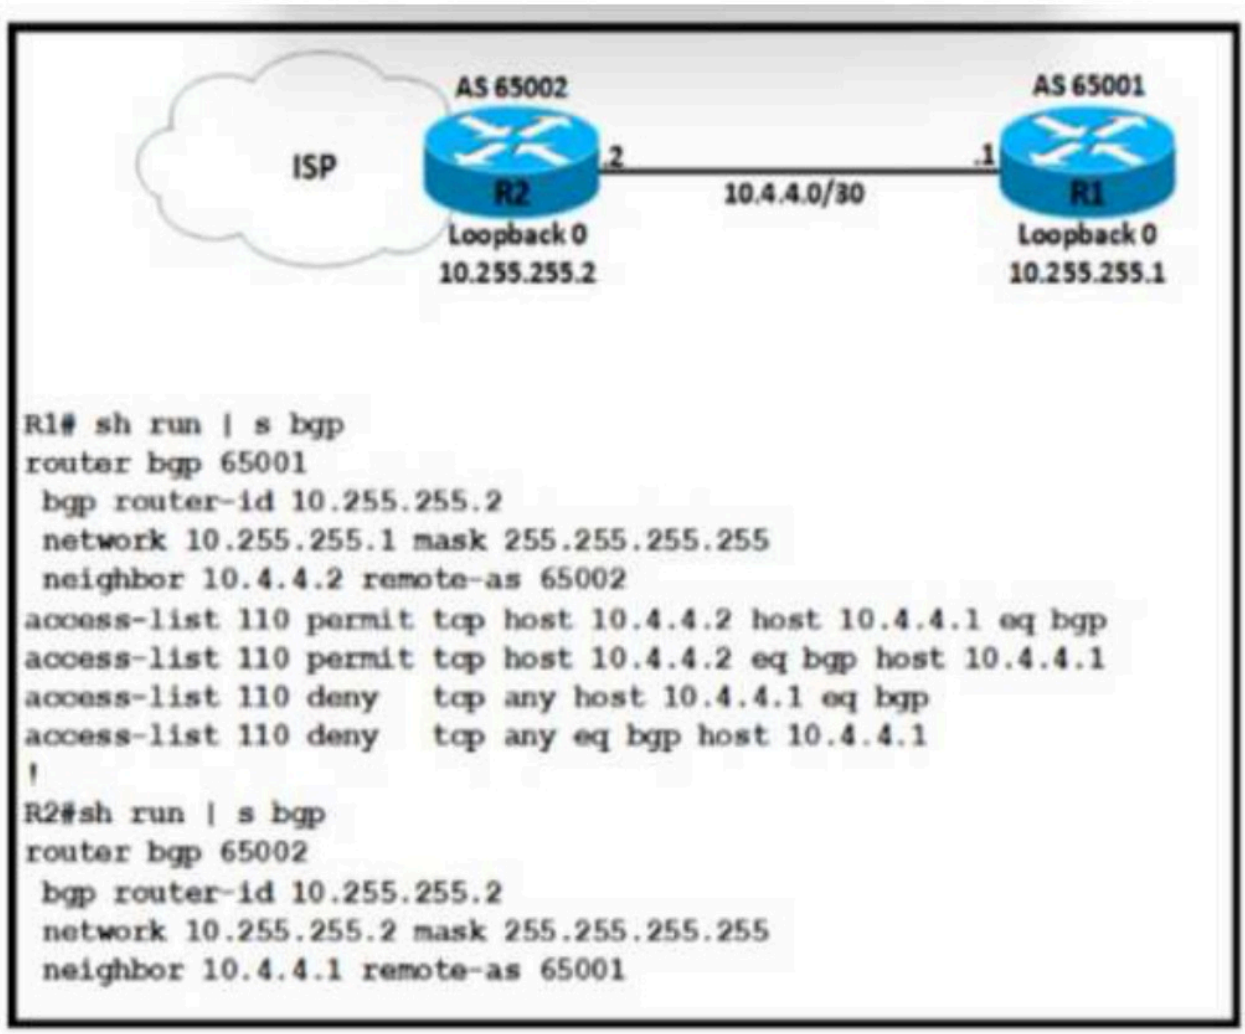
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

Question 21

Single choice · Domain: D1 Layer 3 Technologies › 1.11.b Neighbor relationship and authentication (next-hop, multihop, 4-byte AS, private AS, route refresh, synchronization, operation, peer group, states and timers)

Refer to the exhibit.

Free demo · full exam: 636 questions



*Dec 21 12:08:59.991: BGP: br topo global 10.4.4.2 IPv4 Unicast:base (0x6A8B3998:1) NSF delete stale NSF not active
*Dec 21 12:08:59.995: BGP: br topo global 10.4.4.2 IPv4 Unicast:base (0x44397103:1) NSF no stale paths state is NSF not active
*Dec 21 12:08:59.995: BGP: br topo global 10.4.4.2 IPv4 Unicast:base (0x6A8B3998:1) Resetting ALL counters.
*Dec 21 12:09:09.819: BG-3-NOTIFICATION: sent to neighbor 10.4.4.2 passive 2/3 (BGP identifier wrong) 4 bytes OAFFFF02
*Dec 21 12:09:09.823: BGP-4-MSGDUMP: unsupported or mal-formatted message received from 10.4.4.2:
*Dec 21 12:09:12.443: 8BGP SESSION-5-ADJCHANGE: neighbor 10.4.4.2 IPv4 Unicast topology base removed from session BGP Notification received
*Dec 21 12:09:00.191: BGP: br global 10.4.4.2 Open active delayed 12288ms (35000ms max, 60% jitter)

Refer to the exhibit A network engineer notices that R1 and R2 cannot establish an eBGP peering.
The following messages appear in the log:
Which configuration must the engineer apply to R1 to restore the eBGP peering?

A.

```
router bgp 65001
  bgp router-id 10.255.255.2
  neighbor 10.4.4.2 remote-as 65002
access-list 110 permit tcp host 10.4.4.2 host 10.4.4.1 eq 179
access-list 110 permit tcp host 10.4.4.2 eq 179 host 10.4.4.1
access-list 110 deny tcp any host 10.4.4.1 eq 179
access-list 110 deny tcp any eq 179 host 10.4.4.1
```

```
router bgp 65001
  bgp router-id 10.255.255.2
  neighbor 10.4.4.2 remote-as 65002
access-list 110 permit udp host 10.4.4.2 host 10.4.4.1 eq 179
access-list 110 permit udp host 10.4.4.2 eq 179 host 10.4.4.1
access-list 110 deny udp any host 10.4.4.1 eq 179
access-list 110 deny udp any eq 179 host 10.4.4.1
```

C.

```
router bgp 65001
  bgp router-id 10.255.255.1
  neighbor 10.4.4.2 remote-as 65002
access-list 110 permit tcp host 10.4.4.2 host 10.4.4.1 eq 179
access-list 110 permit tcp host 10.4.4.2 eq 179 host 10.4.4.1
access-list 110 deny   tcp any host 10.4.4.1 eq 179
access-list 110 deny   tcp any eq 179 host 10.4.4.1
```

D.

```
router bgp 65001
  bgp router-id 10.255.255.1
  neighbor 10.4.4.2 remote-as 65002
access-list 110 permit udp host 10.4.4.2 host 10.4.4.1 eq 179
access-list 110 permit udp host 10.4.4.2 eq 179 host 10.4.4.1
access-list 110 deny udp any host 10.4.4.1 eq 179
access-list 110 deny udp any eq 179 host 10.4.4.1
```

Answer: C

Explanation

The log reports a BGP notification that the BGP identifier is wrong. R1 is configured with router ID 10.255.255.2, which is also R2's router ID; R1 should use its own identifier, 10.255.255.1. Option C makes that correction and uses TCP port 179, the transport used by BGP, so it restores the intended eBGP configuration.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/218027-troubleshoot-border-gateway-protocol-bas.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-16-10/irg-xe-16-10-book/configuring-bgp-neighbor-session-options.pdf

✗ **A** Configure R1 with router ID 10.255.255.2 and TCP ACL entries for port 179.
This leaves R1 using 10.255.255.2, the same router ID configured on R2, so it does not resolve the BGP identifier error shown in the log. The TCP ACL entries do not fix that underlying problem.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/218027-troubleshoot-border-gateway-protocol-bas.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-16-10/irg-xe-16-10-book/configuring-bgp-neighbor-session-options.pdf

✗ **B** Configure R1 with router ID 10.255.255.2 and UDP ACL entries for port 179.

The router ID remains the same as R2's, so the identifier error persists. In addition, BGP uses TCP rather than UDP, so UDP ACL entries do not permit the BGP session.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/218027-troubleshoot-border-gateway-protocol-bas.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-16-10/irg-xe-16-10-book/configuring-bgp-neighbor-session-options.pdf

✓ **C** Configure R1 with router ID 10.255.255.1 and TCP ACL entries for port 179.
This sets R1's router ID to 10.255.255.1, matching the identifier shown for R1 in the exhibit and making it distinct from R2's 10.255.255.2. Its TCP ACL entries permit the BGP traffic between the peers on port 179.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/218027-troubleshoot-border-gateway-protocol-bas.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-16-10/irg-xe-16-10-book/configuring-bgp-neighbor-session-options.pdf

✗ **D** Configure R1 with router ID 10.255.255.1 and UDP ACL entries for port 179.
The router ID is corrected, but the ACL entries permit UDP, not the TCP traffic BGP uses. Therefore this configuration would still fail to permit the BGP session.

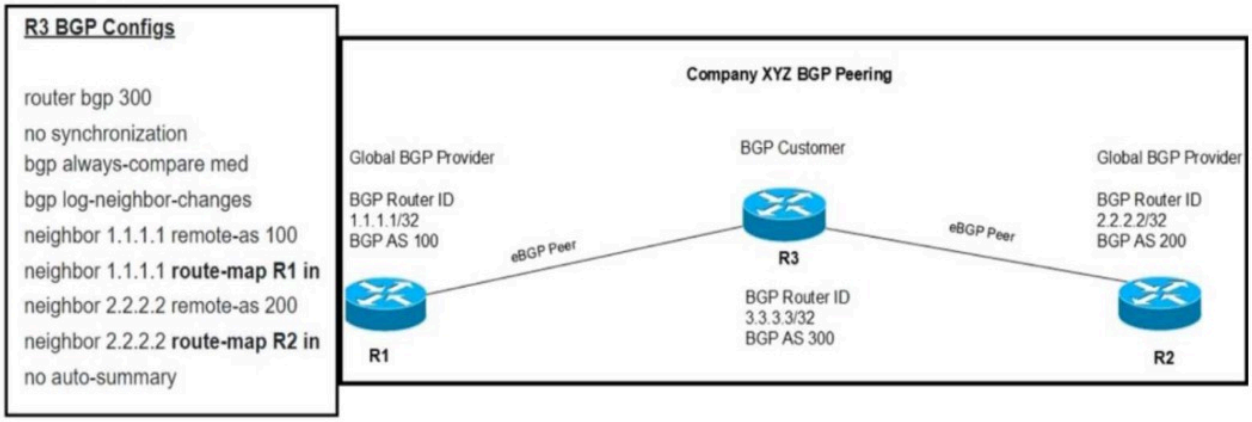
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-16-10/irg-xe-16-10-book/configuring-bgp-neighbor-session-options.pdf

Question 22

Single choice · Domain: D1 Layer 3 Technologies › 1.11.c Path preference (attributes and best-path)

Refer to the exhibit.

Dumpsgate



- route-map R1 permit 10
  set metric 200
  !
  route-map R2 permit 10
  set metric 100

- route-map R1 permit 10
  set metric 100
  !
  route-map R2 permit 10
  set metric 200

- route-map R1 permit 10
  set AS-Path Prepend 200
  !
  route-map R2 permit 10
  set AS-Path Prepend 100

- route-map R1 permit 10
  set AS-Path Prepend 100
  !
  route-map R2 permit 10
  set AS-Path Prepend 200

Refer to the exhibit. R3 is dual-homed to two service providers for traffic redundancy. R3 prefers its outbound traffic via R2. Which set of configurations achieves this goal?

A. Option A
B. Option B
C. Option C
D. Option D

Answer: A. Option A

Explanation

R3 needs to prefer the route learned from R2 for its outbound traffic. Option A sets a higher MED on routes received from R1 (200) than on routes received from R2 (100); because the configuration enables `bgp always-compare-med`, R3 compares these MED values and prefers the lower-MED route via R2. This makes Option A achieve the stated preference.

https://docs.manage.security.cisco.com/cdfmc/c_bgp_path_selection.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html

✓ **A** route-map R1 permit 10; set metric 200. route-map R2 permit 10; set metric 100.
The route maps set MED to 200 for routes received from R1 and 100 for routes received from R2. With `bgp always-compare-med` configured, the lower MED is preferred, so R3 selects R2 for outbound traffic.

https://docs.manage.security.cisco.com/cdfmc/c_bgp_path_selection.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html

✗ **B** route-map R1 permit 10; set metric 100. route-map R2 permit 10; set metric 200.
This assigns the lower MED to routes from R1 and the higher MED to routes from R2. Since lower MED is preferred, R3 would prefer the path via R1 rather than R2.

https://docs.manage.security.cisco.com/cdfmc/c_bgp_path_selection.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html

✗ **C** route-map R1 permit 10; set AS-Path Prepend 200. route-map R2 permit 10; set AS-Path Prepend 100.
AS-path prepending affects the AS-path attribute, but these prepend values are applied to routes received from the providers and do not provide the intended preference for the R2-learned route as directly as setting MED. The configured `bgp always-compare-med` specifically enables comparison of the differing MED values used by Option A.

https://docs.manage.security.cisco.com/cdfmc/c_bgp_path_selection.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html

✗ **D** route-map R1 permit 10; set AS-Path Prepend 100. route-map R2 permit 10; set AS-Path Prepend 200.
This applies the larger prepend value to routes received from R2, but AS-path prepending is primarily used to influence how other autonomous systems choose routes toward the advertising network. It does not use the configured MED comparison to make R3 prefer its outbound route through R2.

https://docs.manage.security.cisco.com/cdfmc/c_bgp_path_selection.html
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html

Question 23

Single choice · Domain: D1 Layer 3 Technologies › 1.9.b Neighbor relationship and authentication

Refer to the exhibit.

```
R1#ping ipv6 ff02::a
Output Interface: fastethernet1/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FF02::A, timeout is 2 seconds:
Packet sent with a source address of FE80::C801:FFF:FE94:1C%FastEthernet1/0

Reply to request 0 received from FE80::C004:22FF:FE78:1, 40 ms
Request 1 timed out
Request 2 timed out
Request 3 timed out
Reply to request 4 received from FE80::C004:22FF:FE78:1, 56 ms
Success rate is 40 percent (2/5), round-trip min/avg/max = 40/48/56 ms
```

Refer to the exhibit An engineer investigates an IPv6 EIGRP neighbor adjacency issue that sees the neighbors flapping and issued a ping from R1 to its directly connected neighbor. The link between the switches is stable at Layer 2. and other connected devices are also functioning. Which action resolves the issue?

A. The switch between the two neighbors is not IPv6 compatible and must be replaced with an IPv6 compatible device.

B. Data is not reliably transmitted between the dynamically assigned link-local addresses of the routers and must be manually assigned.

C. Multicast packets are not reliably transmitted over the link, and the switch must be replaced.

D. The switch between the two neighbors is not configured for IPv6 multicast and must be configured for IPv6 multicast.

Answer: D. The switch between the two neighbors is not configured for IPv6 multicast and must be configured for IPv6 multicast.

Explanation

IPv6 EIGRP uses link-local addresses and multicast to discover and maintain neighbor adjacencies. The ping to the directly connected neighbor succeeds only intermittently, while the Layer 2 link and other devices are reported as functioning, which points to IPv6 multicast not being forwarded reliably through the switch. Configure the switch to handle IPv6 multicast so the EIGRP neighbor traffic can pass reliably. [1][3]

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/lyr2-fwd/eigrp/eigrp-configuration-guide/routing-eigrp.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/multicast/multicast-configuration-guide/mld-

snooping.html

✗ **A** The switch between the two neighbors is not IPv6 compatible and must be replaced with an IPv6 compatible device.
The scenario says the Layer 2 link is stable and other connected devices are functioning, so it does not establish that the switch is incompatible with IPv6. The issue indicated by the intermittent responses is multicast handling, which calls for configuration rather than replacing the switch.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/multicast/multicast-configuration-guide/mld-snooping.html

✗ **B** Data is not reliably transmitted between the dynamically assigned link-local addresses of the routers and must be manually assigned.
IPv6 EIGRP neighbors use link-local addresses, and the exhibit shows replies coming from the directly connected neighbor's link-local address. There is no evidence that the addresses themselves are the cause or that manual assignment would resolve the multicast-dependent adjacency problem.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/lyr2-fwd/eigrp/eigrp-configuration-guide/routing-eigrp.pdf
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/ipaddresses-and-services/configuration-guide/ip-addresses-and-services-config-cisco8000/network-stack-ipv4-and-ipv6-addressing/ipv6-addresses.html

✗ **C** Multicast packets are not reliably transmitted over the link, and the switch must be replaced.
Unreliable multicast forwarding fits the symptoms, but replacing the switch is not required by the scenario. The appropriate remedy is to configure the existing switch to handle IPv6 multicast.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/multicast/multicast-configuration-guide/mld-snooping.html

✓ **D** The switch between the two neighbors is not configured for IPv6 multicast and must be configured for IPv6 multicast.
IPv6 EIGRP relies on multicast for neighbor communication, and the exhibit shows intermittent ping replies despite a stable Layer 2 link. Configuring IPv6 multicast handling on the switch addresses the forwarding issue without replacing otherwise functioning equipment.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/lyr2-fwd/eigrp/eigrp-configuration-guide/routing-eigrp.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/multicast/multicast-configuration-guide/mld-snooping.html

Question 24

Single choice · Domain: D1 Layer 3 Technologies › 1.8 Describe Bidirectional Forwarding Detection

Which IP precedence value does BFD use to prioritize traffic within an infrastructure device?

A. 4
B. 5
C. 6
D. 7

Answer: C. 6

Explanation

BFD uses IP precedence value 6 so its control packets receive priority within infrastructure devices. Therefore, option C matches the requested value.

Free demo · full exam: 636 questions

https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/25xx/routing/configuration/guide/b-routing-cg-asr9000-25xx/implementing-bfd.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/routing/configuration-guide/routing-config-cisco8000/bfd-wrapper/information-about-bfd.html

✗ **A** 4
Precedence 4 is not the value BFD uses for prioritizing its traffic within an infrastructure device.

✗ **B** 5
BFD traffic uses precedence 6, not precedence 5, for this prioritization.

✓ **C** 6
BFD uses IP precedence 6 to mark its traffic for priority treatment within an infrastructure device.

https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/25xx/routing/configuration/guide/b-routing-cg-asr9000-25xx/implementing-bfd.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/routing/configuration-guide/routing-config-cisco8000/bfd-wrapper/information-about-bfd.html

✗ **D** 7
Precedence 7 is not the value specified for BFD traffic; the value used is 6.

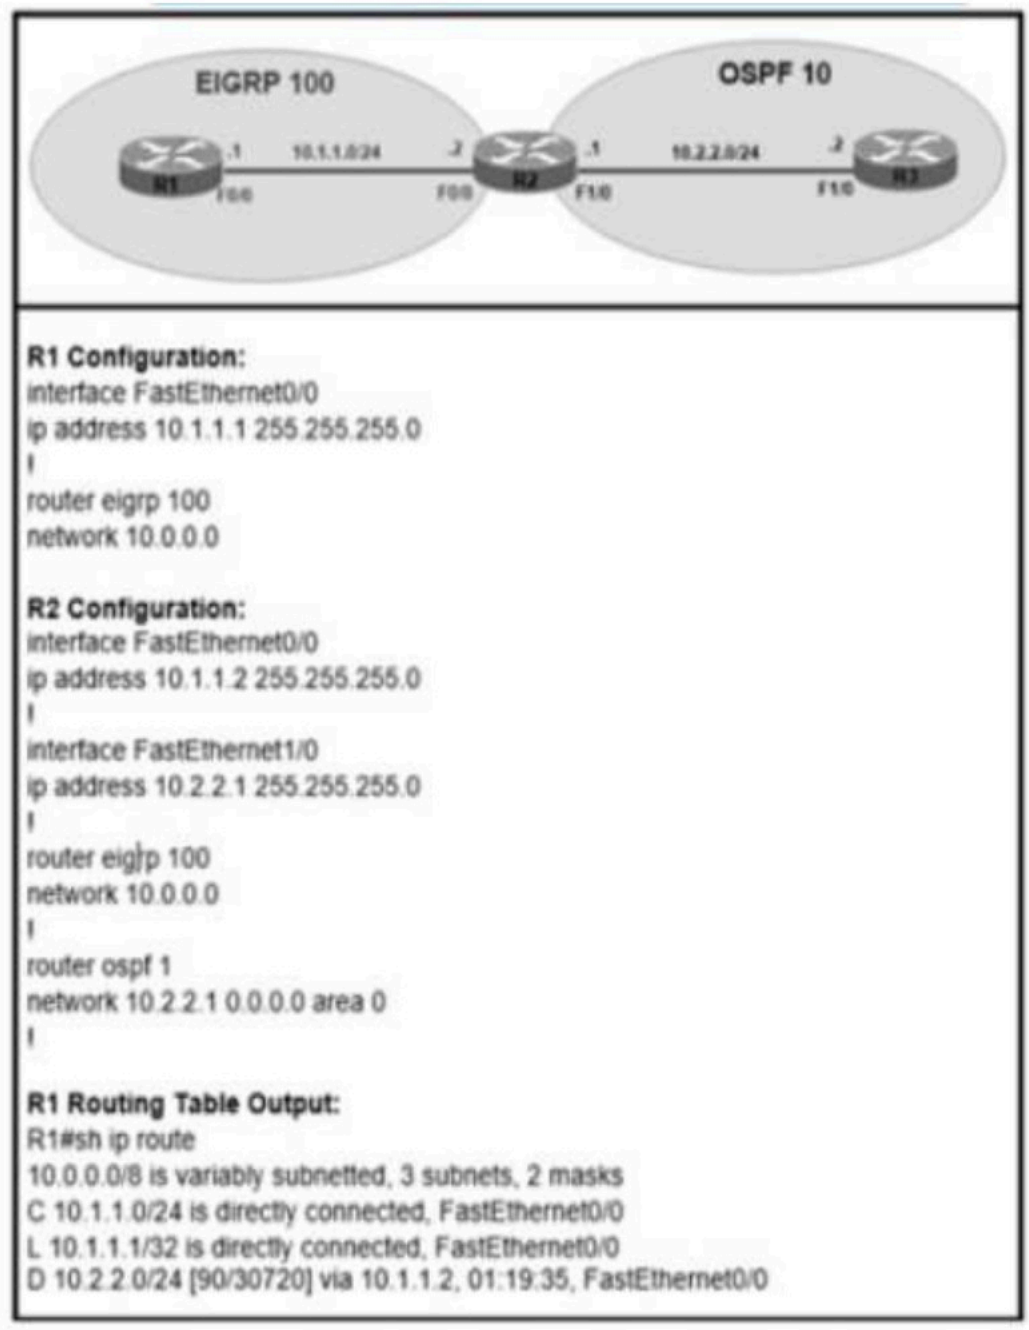
Question 25

Single choice · Domain: D1 Layer 3 Technologies › 1.9 Troubleshoot EIGRP (classic and named mode)

Refer to the exhibit.

Dumpsgate

Free demo · full exam: 636 questions



Refer to the exhibit The R2 OSPF route 10 2 2 0/24 shows in the R1 EIGRP routing table without route redistribution performed between OSPF and EIGRP routing protocols Which configuration is required on router R2 to resolve the issue?

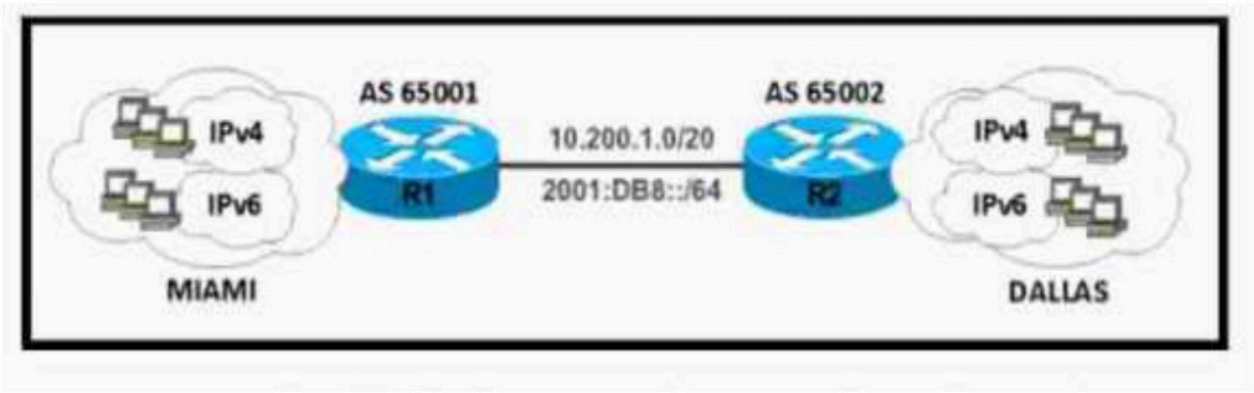
A. passive-interface FastEthernet 0/0 command in OSPF1.
B. Add the no auto-summary command in EIGRP 100.
C. Replace the network 10.0.0.0 command with FastEthernet0/0 network in EIGRP 100.
D. Add the passive-interface FastEthernet 1/0 command in EIGRP 100

Answer: C. Replace the network 10.0.0.0 command with FastEthernet0/0 network in EIGRP 100.

Question 26

Single choice · Domain: D1 Layer 3 Technologies › 1.11.a Address families (IPv4, IPv6)

Refer to the exhibit.



Refer to the exhibit. A network engineer configured routers R1 and R2 with MP-BGP. The engineer noticed that the routers cannot exchange any IPv6 routes, however, the IPv4 neighbor relationship is working fine. Which configuration must the engineer apply to router R2 to exchange IPv6 routes?

A. Ipv6cef I Interface Loopback100Ipv6 address 2001: DB8:128::2/128 Interface GigabitEthernet1/0Ipv6 address 2001:DB8:1::2/64 I router bgp 65002no bgp default ipv4-unlcastneighbor 2001: DB8:1::1 emote-as 65001Iaddress-family ipv6network 2001:DB8:128::2/128neighbor 2001:DB8:1::1 activate

B. Ipv6 unicast-routingipv6 cefInterface Loopback100ipv6 address 2001: DB8:12C:2 129Iinterface GigabitEthernet1/0ipv6 address 2001: DB8:1::2 64description AS65001 ID B466:A83D:3D7::1!router bgp 65002no bgp default Ipv4-unicastneighbor 2001: DB8:1::1 remote-as 65001!address-family Ipv4neighbor 2001:DB8:1::1 activate

C. Ipv6 unicast-routingipvG cefIInterface Loopback100Ipv6 address 2001:DB8:128::2/128 I interface GigabitEthernet1/0ipv6 address 2001:DB8:1::2/64!router bgp 65002 no bgp default ipv4-unicast neighbor 2001: DB8:1::1 remote-as 65001 !address-family ipv6 network 2001:DB8:128::2I128

D. Ipv6 unicast-routing Ipv6 cef Iinterface Loopback100 ipv6 address 2001: DB8:128::2/128 !interface GigabitEthernet1/0 ipv6 address 2001:DB8:1::2/64 irouter bgp 65002 no bgp default ipv4-unicast neighbor 2001: DB8:1::1 remote-as 65001 !address-family ipv6 network 2001:DB8:128::2/128 neighbor 2001:DB8:1::1 activate

Answer: D. Ipv6 unicast-routing Ipv6 cef Iinterface Loopback100 ipv6 address 2001: DB8:128::2/128 !interface GigabitEthernet1/0 ipv6 address 2001:DB8:1::2/64 irouter bgp 65002 no bgp default ipv4-unicast neighbor 2001: DB8:1::1 remote-as 65001 !address-family ipv6 network 2001:DB8:128::2/128 neighbor 2001:DB8:1::1 activate

Explanation

R2 needs to enable IPv6 routing and IPv6 CEF, then configure the BGP IPv6 address family to advertise its

Free demo · full exam: 636 questions

IPv6 loopback route and activate the existing IPv6 neighbor. Option D includes all of these required pieces while retaining the working neighbor relationship to R1 at 2001:DB8:1::1. The other configurations either omit neighbor activation, configure the wrong address family, or contain an invalid or incomplete IPv6 route advertisement.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_irg-bgp4.html
https://www.cisco.com/c/en/us/td/docs/switches/connectedgrid/cg-switch-sw-master/software/configuration/guide/unicast/CGS_1000_Unicast/uni_ipv6.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3e/irg-iproute-bgp-xe-3e-book/ip6-mbgp-ext.html

✗ **A** Enables IPv6 CEF and configures IPv6 addresses, but includes an IPv6 network statement and activates the neighbor.
This configuration does not include `ipv6 unicast-routing`, which is required to enable IPv6 forwarding. It also has a malformed BGP neighbor command (`emote-as`) and an apparently misspelled `ipv4-unicast` command, so it would not correctly establish the intended IPv6 BGP configuration.

https://www.cisco.com/c/en/us/td/docs/switches/connectedgrid/cg-switch-sw-master/software/configuration/guide/unicast/CGS_1000_Unicast/uni_ipv6.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3e/irg-iproute-bgp-xe-3e-book/ip6-mbgp-ext.html

✗ **B** Enables IPv6 routing and CEF, but configures the IPv4 BGP address family.
The BGP address-family section is `ipv4`, so it does not enable exchange of IPv6 routes. Its loopback address is also malformed or has the wrong prefix format compared with the IPv6 /128 route that should be advertised.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_irg-bgp4.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3e/irg-iproute-bgp-xe-3e-book/ip6-mbgp-ext.html

✗ **C** Enables IPv6 routing and CEF and configures the IPv6 address family, but only includes a network statement.
Although it specifies the IPv6 address family and the loopback prefix, it omits `neighbor 2001:DB8:1::1 activate` under that address family. The IPv6 neighbor therefore is not enabled for route exchange.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_irg-bgp4.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3e/irg-iproute-bgp-xe-3e-book/ip6-mbgp-ext.html

✓ **D** Enables IPv6 routing and CEF, advertises the loopback /128 in the IPv6 BGP address family, and activates the IPv6 neighbor.
This configuration enables IPv6 routing and CEF, then enters BGP's IPv6 address family and advertises `2001:DB8:128::2/128`. It also activates the R1 IPv6 neighbor, allowing R2 to exchange IPv6 routes while the existing IPv4 neighbor relationship remains unaffected.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_irg-bgp4.html
https://www.cisco.com/c/en/us/td/docs/switches/connectedgrid/cg-switch-sw-master/software/configuration/guide/unicast/CGS_1000_Unicast/uni_ipv6.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3e/irg-iproute-bgp-xe-3e-book/ip6-mbgp-ext.html

Question 27

Single choice · Domain: D1 Layer 3 Technologies › 1.11.d Route reflector (excluding multiple route reflectors, confederations, and dynamic peering)

Refer to the exhibit.

```
R1#show bgp ipv6 unicast 2001:db8::1/128
BGP routing table entry for 2001:db8::1/128, version 3
Paths: (1 available, best #1, table Global-IPv6-Table)
  Not advertised to any peer
  Local
    2001:db8:33:33::33 (metric 128) from 2001:db8:11:11::11 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 3.3.3.3, Cluster list: 1.1.1.1
```

Refer to the exhibit. An engineer examines the BGP update for the IPv6 prefix 2001:db8::1/128. which should have been summarized into a /64 prefix. Which sequence of actions achieves the summarization?

A. R1 is a route reflector client of a RR with a router ID of 1.1.1.1. and the originator of the prefix has a router ID of 3.3.3.3. Both routers belong to different ASs. The prefix is not advertised to any peer and must be advertised using the network statement on R3.

B. R1 is a route reflector with a router ID of 3.3.3.3. and the originator of the prefix is a route reflector client, which has a router ID of 3.3.3.3. Both routers belong to the same AS Configure an aggregate address on the router with ID 1.1.1.1 for the prefix

C. R1 is a route reflector with a router ID of 1.111. and the originator of the prefix is a route reflector client, which has a router ID of 3.3.3.3. Both routers belong to the same AS Configure an aggregate address on the router with ID 1.1.1.1 for the prefix

D. R1 is a route reflector client of a RR with a router ID of 1.1.1.1. and the originator of the prefix has a router ID of 3.3.3.3. Both routers belong to the same AS. Configure an aggregate address on the router with ID 3 3.3.3 for the prefix.

Answer: D. R1 is a route reflector client of a RR with a router ID of 1.1.1.1. and the originator of the prefix has a router ID of 3.3.3.3. Both routers belong to the same AS. Configure an aggregate address on the router with ID 3 3.3.3 for the prefix.

Explanation

The BGP entry identifies the prefix as an internal route reflected from originator ID 3.3.3.3, with cluster ID 1.1.1.1, and says it is not advertised to any peer. Option D matches that relationship: R1 is a route-reflector client, both routers are in the same AS, and the aggregate is configured on the originator router, 3.3.3.3. This is the required sequence for summarizing the prefix into a /64 aggregate.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5441-aggregation.html

Free demo · full exam: 636 questions

https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/bgp/bgp-config-cisco8000/r-wrapper-bgp-routing-optimisation-and-convergence-techniques/c-bgp-route-reflectors.html

✗ **A** R1 is a route reflector client of a RR with a router ID of 1.1.1.1, and the originator of the prefix has a router ID of 3.3.3.3. Both routers belong to different ASs. The prefix is not advertised to any peer and must be advertised using the network statement on R3.
The route-reflector IDs are consistent with the exhibit, but the routers must be in the same AS for this reflected internal route scenario. The proposed network statement does not perform the required aggregation, and this option incorrectly places the routers in different ASs.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5441-aggregation.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/bgp/bgp-config-cisco8000/r-wrapper-bgp-routing-optimisation-and-convergence-techniques/c-bgp-route-reflectors.html

✗ **B** R1 is a route reflector with a router ID of 3.3.3.3, and the originator of the prefix is a route reflector client with a router ID of 3.3.3.3. Both routers belong to the same AS. Configure an aggregate address on the router with ID 1.1.1.1 for the prefix.
The exhibit's originator ID is 3.3.3.3 and its cluster list contains 1.1.1.1, so 1.1.1.1 is not the originator identified in the route entry. This option also assigns the same router ID to R1 and the originator, rather than describing the client and reflector relationship shown.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5441-aggregation.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/bgp/bgp-config-cisco8000/r-wrapper-bgp-routing-optimisation-and-convergence-techniques/c-bgp-route-reflectors.html

✗ **C** R1 is a route reflector with a router ID of 1.111, and the originator of the prefix is a route-reflector client with a router ID of 3.3.3.3. Both routers belong to the same AS. Configure an aggregate address on the router with ID 1.1.1.1 for the prefix.
The exhibit lists 1.1.1.1 as the cluster ID and 3.3.3.3 as the originator ID; it does not identify R1 as the reflector with router ID 1.111. More importantly, this choice configures the aggregate on 1.1.1.1 instead of the originator router 3.3.3.3 specified by the correct sequence.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5441-aggregation.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/bgp/bgp-config-cisco8000/r-wrapper-bgp-routing-optimisation-and-convergence-techniques/c-bgp-route-reflectors.html

✓ **D** R1 is a route-reflector client of a route reflector with router ID 1.1.1.1, and the originator of the prefix has router ID 3.3.3.3. Both routers belong to the same AS. Configure an aggregate address on the router with ID 3.3.3.3 for the prefix.
This matches the BGP attributes in the exhibit: originator ID 3.3.3.3 and cluster list 1.1.1.1, with the route marked internal. It keeps the routers in the same AS and places the aggregate configuration on the originator router, which produces the requested /64 summary.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5441-aggregation.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/bgp/bgp-config-cisco8000/r-wrapper-bgp-routing-optimisation-and-convergence-techniques/c-bgp-route-reflectors.html

Question 28

Single choice · Domain: D1 Layer 3 Technologies › 1.9.e Load balancing (equal and unequal cost)

Refer to the exhibit.

Free demo · full exam: 636 questions

```
router eigrp 1
 variance 2

R1#show ip eigrp topology 172.16.100.5 255.255.255.255

IP-EIGRP (AS 1): Topology entry for 172.16.100.5/32

  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 409600

  Routing Descriptor Blocks:

  10.4.1.5 (Ethernet1/0), from 10.4.1.5, Send flag is 0x0

      Composite metric is (409600/128256), Route is Internal

      Vector metric:

        Minimum bandwidth is 10000 Kbit

        Total delay is 6000 microseconds

        Reliability is 255/255

        Load is 1/255

        Minimum MTU is 1500

        Hop count is 1

  10.3.1.6 (Serial2/0), from 10.3.1.6, Send flag is 0x0

      Composite metric is (435200/409600), Route is Internal

      Vector metric:

        Minimum bandwidth is 10000 Kbit

        Total delay is 7000 microseconds

        Reliability is 255/255
```

```
        Load is 1/255

        Minimum MTU is 1500

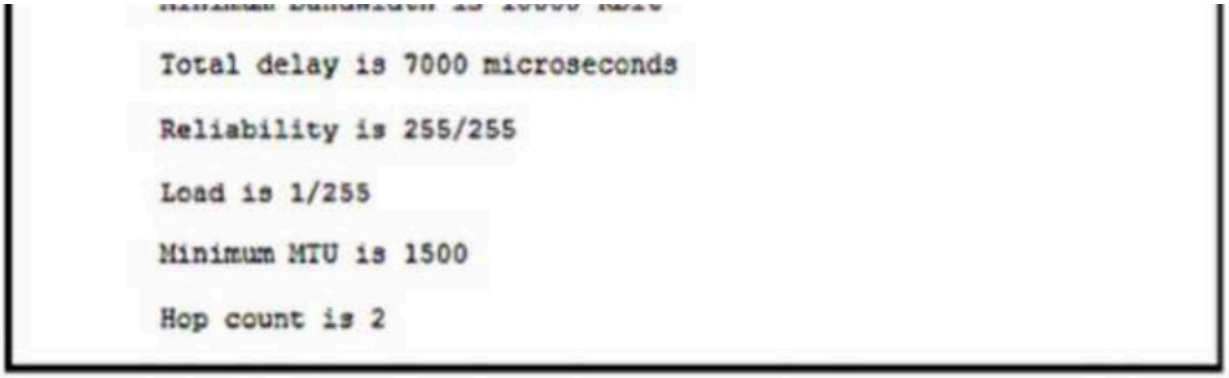
        Hop count is 1

  10.3.1.6 (Serial2/0), from 10.3.1.6, Send flag is 0x0

      Composite metric is (435200/409600), Route is Internal

      Vector metric:

        Minimum bandwidth is 10000 Kbit
```



Refer to the exhibit. A network engineer troubleshooting a packet drop problem for the host 172.16.100.5 notices that only one link is used and installed on the routing table, which saturates the bandwidth. Which action must the engineer take to resolve the high bandwidth utilization problem and share the traffic toward this host between the two available links?

A. Set the eigrp variance equal to 4 to install a second route with a metric not larger than 4 times of the best metric.
B. Change the EIGRP delay metric to meet the feasibility condition.
C. Set the eigrp variance equal to 3 to install a second route with a metric not larger than 3 times of the best metric.
D. Disable the eigrp split horizon loop protection mechanism.

Answer: B. Change the EIGRP delay metric to meet the feasibility condition.

Explanation

The exhibit shows two EIGRP paths, but the second path has a reported distance that does not satisfy the feasibility condition, so it is not eligible as a feasible successor for unequal-cost load balancing.
Increasing the variance alone cannot make an ineligible path usable. The engineer must adjust the EIGRP delay metric so the alternate route meets the feasibility condition, allowing EIGRP to install and use it alongside the best route. [2]

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.pdf

✗ **A** Set the eigrp variance equal to 4 to install a second route with a metric not larger than 4 times of the best metric.
Variance controls how much larger an alternate route's metric may be than the best route's metric, but it does not override the feasibility condition. The alternate shown in the exhibit is not eligible under that condition, so setting variance to 4 alone will not install it. [2]

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.pdf

✓ **B** Change the EIGRP delay metric to meet the feasibility condition.
The exhibit shows that the alternate path is not currently eligible as a feasible successor, preventing it from being installed for unequal-cost load sharing. Adjusting the delay metric can change the route metric and reported distance so that the alternate satisfies the feasibility condition; then EIGRP can use it with the existing path. [2]

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.pdf

✗ **C** Set the eigrp variance equal to 3 to install a second route with a metric not larger than 3 times of the best metric.
Like variance 4, variance 3 only sets the permitted metric ratio for routes that are otherwise eligible for installation. It cannot make this alternate satisfy the feasibility condition, so changing variance alone does not solve the problem. [2]

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.pdf

✗ **D** Disable the eigrp split horizon loop protection mechanism.
The issue shown is that the alternate route does not meet EIGRP's feasibility condition, not that split horizon is preventing traffic from using an eligible route. Disabling loop-prevention behavior is not the required remedy; the delay metric must be adjusted so the alternate qualifies. [2]

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.pdf

Dumpsgate

Question 29

Single choice · Domain: D1 Layer 3 Technologies › 1.11.c Path preference (attributes and best-path)

Refer to the exhibit.

```
Router# show ip interface brief
Interface              IP-Address      OK? Method Status    Protocol
Ethernet0/0            90.0.1.1        YES NVRAM  up        up
Ethernet1/0            118.8.10.1      YES NVRAM  up        up
Loopback6              10.1.7.6        YES NVRAM  up        up
Loopback7              10.1.7.7        YES NVRAM  up        up

Router# show ip bgp
BGP table version is 10, local router ID is 10.1.7.7
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.0.0.0         0.0.0.0                  0         32768 ?
*> 10.1.7.7/32      0.0.0.0                  0         32768 i
*> 90.0.0.0         0.0.0.0                  0         32768 ?
r> 90.0.1.0/24      100.0.1.254              0             0 6 ?
*> 118.0.0.0        0.0.0.0                  0         32768 ?
r>i118.8.10.0/24    118.8.10.254             0    100      0 ?
```

Which action adds the 10 1 7.6-32 route to the BGP table?

A. Add a static roue for the 10 1 7 &32 network
B. Add the network 10.1.7.6 mask 255.255.255.255 backdoor command
C. Add summaty-address10.1.7.6.255.255.2550
D. Add the network 10.1.7.6 mask 255.255.255.255 command

Answer: D. Add the network 10.1.7.6 mask 255.255.255.255 command

Explanation

The BGP table already contains the loopback host route 10.1.7.7/32, but not the 10.1.7.6/32 host route shown by the interface output. The BGP `network` command originates a prefix only when that exact prefix exists in the routing table. Adding `network 10.1.7.6 mask 255.255.255.255` matches the connected Loopback6 address and adds the /32 route to BGP.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/19345-bgp-noad.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/iproute_bgp-xe-3se-3850-cr-book.pdf

✗ **A** Add a static route for the 10.1.7.6/32 network
A static route is not needed: the exhibit shows Loopback6 configured with 10.1.7.6 and up/up, so the host route is already present in the routing table. The required action is to originate that exact route with the BGP network command.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/19345-bgp-noad.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/iproute_bgp-xe-3se-3850-cr-book.pdf

✗ **B** Add the network 10.1.7.6 mask 255.255.255.255 backdoor command
The `backdoor` keyword treats a BGP network as an external route for route-selection purposes; it does not replace the need to originate the prefix with the ordinary network statement. The question asks to add the route to the BGP table, so the plain network command is appropriate.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/19345-bgp-noad.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/iproute_bgp-xe-3se-3850-cr-book.pdf

✗ **C** Add summary-address 10.1.7.6 255.255.255.0
This proposes a /24 summary, not the specific 10.1.7.6/32 loopback route requested. The BGP table must be given the exact host prefix, using a /32 mask.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/5441-aggregation.html

✓ **D** Add the network 10.1.7.6 mask 255.255.255.255 command
The interface output shows Loopback6 has address 10.1.7.6, and the BGP network statement can originate that exact route when it is present in the routing table. The /32 mask specifies the required host route.

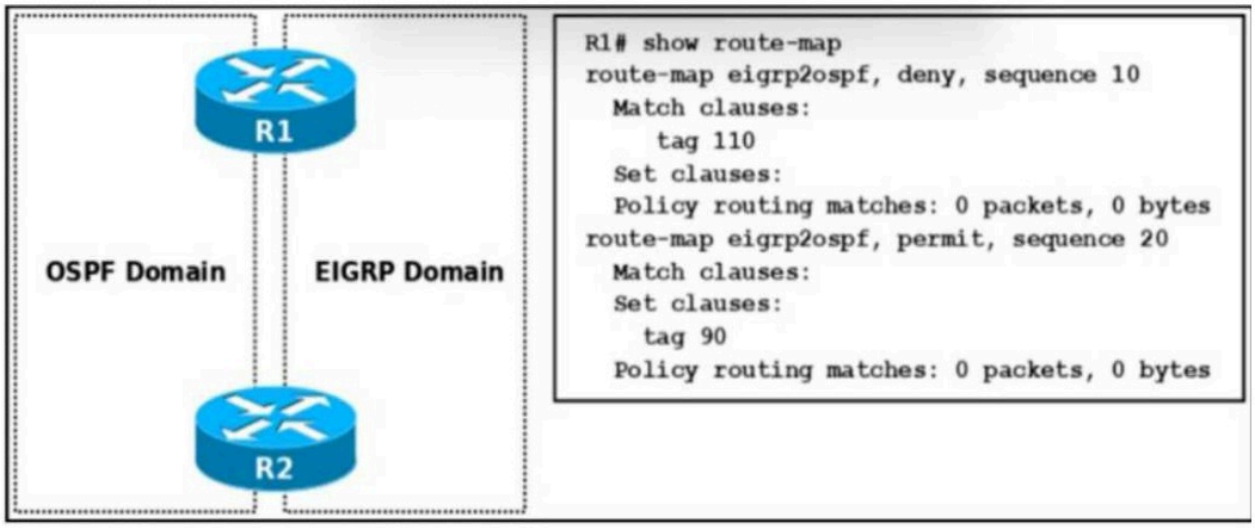
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/19345-bgp-noad.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/iproute_bgp-xe-3se-3850-cr-book.pdf

Dumpsgate

Question 30

Multiple choice · Domain: D1 Layer 3 Technologies › 1.3 Troubleshoot loop prevention mechanisms (filtering, tagging, split horizon, route poisoning)

Refer to the exhibit.



- ☐ route-map eigrp2ospf permit 10
  set tag 90
  route-map eigrp2ospf deny 20
  match tag 110

- ☐ route-map ospf2eigrp deny 10
  match tag 90
  route-map ospf2eigrp permit 20
  set tag 110

- ☐ router ospf 1
  redistribute eigrp 1 route-map ospf2eigrp

- ☐ router eigrp 1
  redistribute ospf 1 metric 100000 1 255 1 1500 route-map eigrp2ospf

- ☐ router eigrp 1
  redistribute ospf 1 metric 100000 1 255 1 1500 route-map ospf2eigrp

Refer to the exhibit. Which two configurations are required on R2 to prevent a routing loop caused by the redistribution from OSPF back into EIGRP? (Choose two.)

A. Option A
B. Option B
C. Option C
D. Option D
E. Option E

Answer

- B. **Option B**
- E. **Option E**

Explanation

R1 tags routes redistributed from EIGRP into OSPF with tag 90. On R2, the OSPF-to-EIGRP route map must reject routes carrying tag 90 to prevent them from being redistributed back into EIGRP, while tagging other permitted routes with 110. Option B defines that filtering and tagging policy, and option E applies it to OSPF-to-EIGRP redistribution. [1][6]

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
https://community.cisco.com/kxiwq67737/attachments/kxiwq67737/4461-docs-network-infrastructure/7191/1/Routes%20Redistribution%20Filtering%20Suboptimal%20Routing%20and%20Routing%20Loop.pdf

✗ **A** route-map eigrp2ospf permit 10; set tag 90; route-map eigrp2ospf deny 20; match tag 110
This defines a map named for the EIGRP-to-OSPF direction, not the OSPF-to-EIGRP redistribution on R2. It also permits and tags routes before the later deny clause, so it does not implement the required rejection of routes already tagged 90.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
https://community.cisco.com/kxiwq67737/attachments/kxiwq67737/4461-docs-network-infrastructure/7191/1/Routes%20Redistribution%20Filtering%20Suboptimal%20Routing%20and%20Routing%20Loop.pdf

✓ **B** route-map ospf2eigrp deny 10; match tag 90; route-map ospf2eigrp permit 20; set tag 110
This route map denies routes tagged 90, which identifies routes previously redistributed from EIGRP into OSPF. It permits other routes and marks them with tag 110, so they can be recognized if encountered in a later redistribution.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
https://community.cisco.com/kxiwq67737/attachments/kxiwq67737/4461-docs-network-infrastructure/7191/1/Routes%20Redistribution%20Filtering%20Suboptimal%20Routing%20and%20Routing%20Loop.pdf

✗ **C** router ospf 1; redistribute eigrp 1 route-map ospf2eigrp
This applies redistribution in the wrong direction: it redistributes EIGRP routes into OSPF. The loop-prevention policy must be applied when R2 redistributes OSPF routes back into EIGRP.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

✗ **D** router eigrp 1; redistribute ospf 1 metric 100000 1 255 1 1500 route-map eigrp2ospf
Although this configures OSPF-to-EIGRP redistribution, it applies the eigrp2ospf route map rather than the ospf2eigrp policy that denies tag 90. Therefore, it does not provide the required filter for routes returning from OSPF.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
https://community.cisco.com/kxiwq67737/attachments/kxiwq67737/4461-docs-network-infrastructure/7191/1/Routes%20Redistribution%20Filtering%20Suboptimal%20Routing%20and%20Routing%20Loop.pdf

✓ **E** router eigrp 1; redistribute ospf 1 metric 100000 1 255 1 1500 route-map ospf2eigrp
This applies the ospf2eigrp route map to the required OSPF-to-EIGRP redistribution on R2. Together with option B, it prevents routes carrying tag 90 from being redistributed back into EIGRP.

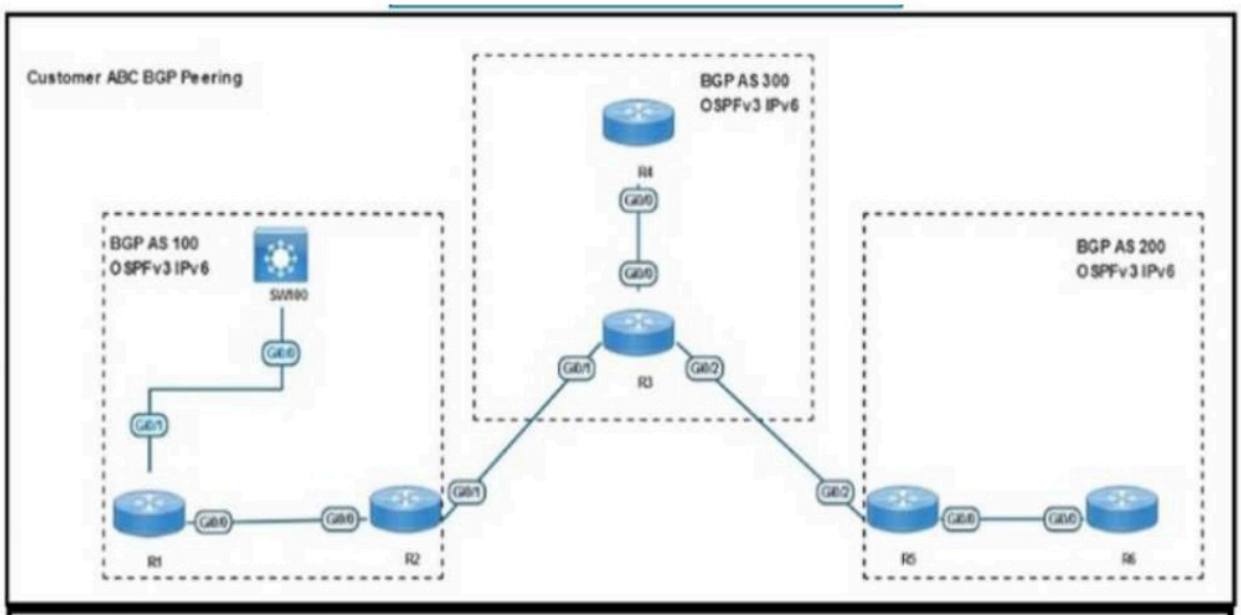
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
https://community.cisco.com/kxiwq67737/attachments/kxiwq67737/4461-docs-network-infrastructure/7191/1/Routes%20Redistribution%20Filtering%20Suboptimal%20Routing%20and%20Routing%20Loop.pdf

Question 31

Single choice · Domain: D1 Layer 3 Technologies › 1.11.c Path preference (attributes and best-path)

Refer to the exhibit.

Dumpsgate




```
R2#sh ip bgp ipv6 uni
BGP table version is 45, local router ID is 2.2.22.22
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

```
              t secondar
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>   2001::5/128      2001::5                  0             0 300 200 i
 *>   2001::4/128      2001::4                  0             0 300 i
 *>   2002::2/128      ::                                 0    32768 i

R2#sh run | section bgp
router bgp 100
   address-family ipv6
neighbor 2001::4 route-map Filter in

ip as-path access-list 1 permit _300_[0-9]

route-map Filter permit 10
match as-path 1
```

Refer to the exhibit R2 has been receiving routes from R4 that originated outside AS300 A network engineer configured an AS-Path ACL to avoid adding these routes to the R2 BGP table but the routes are still present in the R2 routing table Which action resolves the issue?

Free demo · full exam: 636 questions

A. Replace as-path access-list 1 with the ip as-path access-list 1 permit A300$ command
B. Replace as-path access-list 1 with the ip as-path access-list 1 permit ..300." command
C. Replace as-path access-list 1 with the ip as-path access-list 1 permit A300_ command.
D. Replace as-path access-list 1 with the ip as-path access-list 1 permit A300." command

Answer: A. Replace as-path access-list 1 with the ip as-path access-list 1 permit A300$ command

Explanation

The configured AS-path ACL, `permit _300_[0-9]`, matches paths containing AS 300 followed by a digit, rather than identifying routes whose AS path ends at AS 300. In this scenario, R2 must permit paths that end with AS 300 so the route-map can match and filter the routes learned from R4 that originated outside AS 300. The `$` anchor in option A matches the end of the AS path, so replacing the ACL with `permit ^300$`? No—the shown option is `A300$`; this expresses the intended end-of-path match for AS 300 in the exhibit's notation.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13754-26.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp2.html
https://community.cisco.com/t5/networking-knowledge-base/how-to-configure-bgp-route-filtering-using-regular-expression/ta-p/3132699

✓ **A** Replace as-path access-list 1 with the ip as-path access-list 1 permit A300$ command
The `$` anchor matches AS 300 at the end of the AS path, which is the needed pattern for identifying paths that terminate at AS 300 in this scenario. The existing expression `_300_[0-9]` instead looks for AS 300 followed by another AS number, so it does not properly match the intended paths for the route-map filter.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13754-26.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp2.html
https://community.cisco.com/t5/networking-knowledge-base/how-to-configure-bgp-route-filtering-using-regular-expression/ta-p/3132699

✗ **B** Replace as-path access-list 1 with the ip as-path access-list 1 permit ..300." command
This pattern does not anchor AS 300 at the end of the AS path, so it does not precisely select paths ending in AS 300. It therefore would not correct the mismatch between the intended AS-path filtering and the configured ACL.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13754-26.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp2.html
https://community.cisco.com/t5/networking-knowledge-base/how-to-configure-bgp-route-filtering-using-regular-expression/ta-p/3132699

✗ **C** Replace as-path access-list 1 with the ip as-path access-list 1 permit A300_ command.
The underscore is a separator/boundary symbol in BGP regular expressions, not an end-of-path anchor. This expression does not ensure that AS 300 is the final AS in the path, which is required for the intended filter.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13754-26.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp2.html
https://community.cisco.com/t5/networking-knowledge-base/how-to-configure-bgp-route-filtering-using-regular-expression/ta-p/3132699

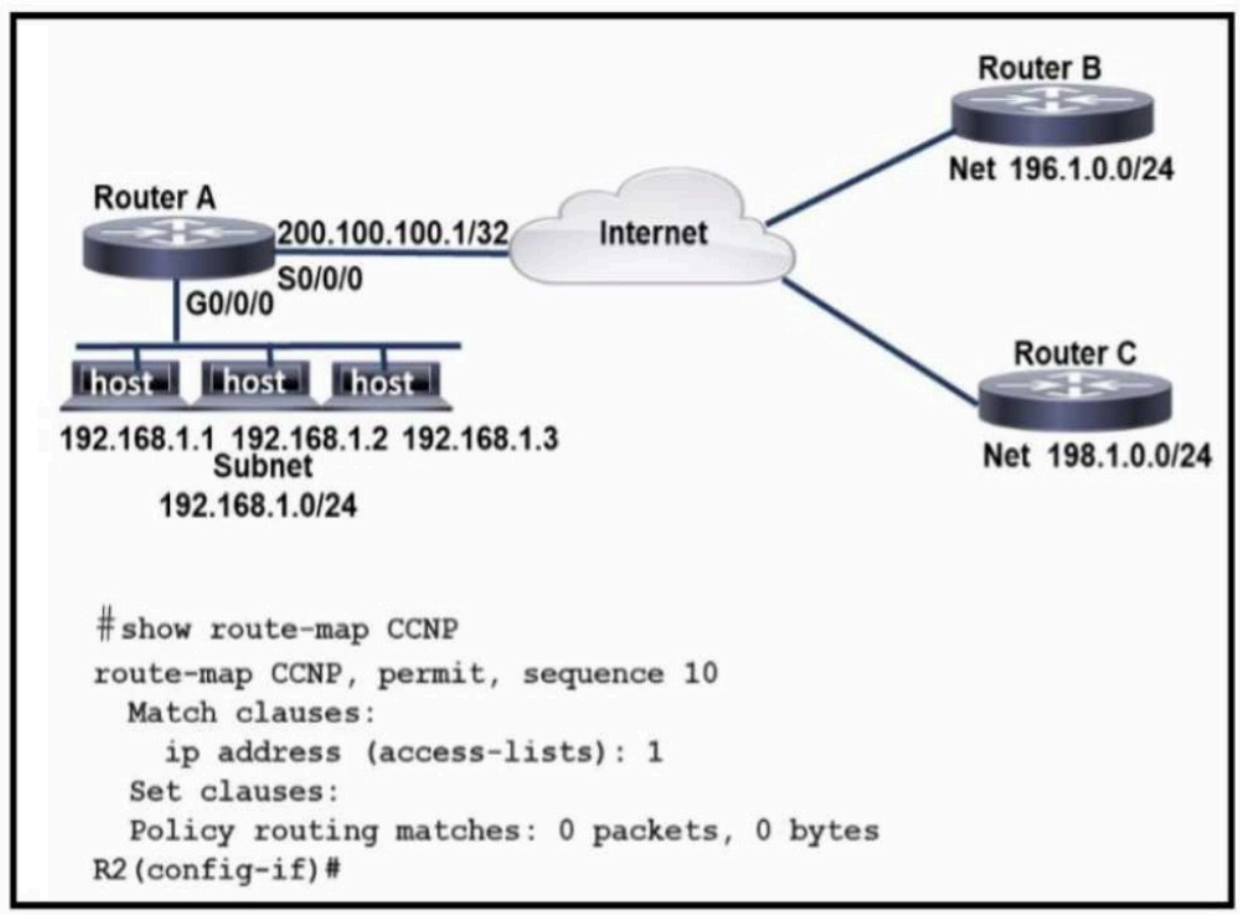
✗ **D** Replace as-path access-list 1 with the ip as-path access-list 1 permit A300." command
The expression does not use the `$` end-of-path anchor and therefore does not identify paths that end specifically with AS 300. The required correction is the end-anchored pattern in option A.

https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13754-26.html
https://www.cisco.com/c/en/us/td/docs/ios/iproute_bgp/command/reference/irg_book/irg_bgp2.html
https://community.cisco.com/t5/networking-knowledge-base/how-to-configure-bgp-route-filtering-using-regular-expression/ta-p/3132699

Question 32

Single choice · Domain: D1 Layer 3 Technologies › 1.6 Configure and verify policy-based routing

Refer to the exhibit.



Refer to the exhibit. An engineer configures router A to mark all inside to outside traffic from network 192 168 1 0, except from host 192 168 1 1. with critical IP precedence. The policy did not work as expected Which configuration resolves the issue?

A. Option

```
RouterA(config)#access-list 1 deny host 192.168.1.1
RouterA(config)#route-map CCNP permit 10
RouterA(config)#match ip address 1
RouterA(config)#set ip precedence critical
RouterA(config)#route-map CCNP permit 20
RouterA(config)# interface g0/0/0
RouterA(config-if)#ip address 192.168.1.4 255.255.255.0
RouterA(config-if)#ip policy route-map CCNP
```

B. Option

Free demo · full exam: 636 questions

```
RouterA(config)#access-list 1 deny host 192.168.1.1
RouterA(config)#access-list 1 permit any any
RouterA(config)#route-map CCNP deny 10
RouterA(config)#match ip address 1
RouterA(config)#set ip precedence critical
RouterA(config)#route-map CCNP permit 20
RouterA(config)# interface g0/0/0
RouterA(config-if)#ip address 192.168.1.4 255.255.255.0
RouterA(config-if)#ip policy route-map CCNP
```

C. Option

```
RouterA(config)#access-list 1 deny host 192.168.1.1
RouterA(config)#access-list 1 permit any any
RouterA(config)#route-map CCNP permit 10
RouterA(config)#match ip address 1
RouterA(config)#set ip precedence critical
RouterA(config)#route-map CCNP permit 20
RouterA(config)#set ip precedence critical
RouterA(config)# interface g0/0/0
RouterA(config-if)#ip address 192.168.1.4 255.255.255.0
RouterA(config-if)#ip policy route-map CCNP
```

D. Option

```
RouterA(config)#access-list 1 deny host 192.168.1.1
RouterA(config)#access-list 1 permit any any
RouterA(config)#route-map CCNP permit 10
RouterA(config)#match ip address 1
RouterA(config)#set ip precedence critical
RouterA(config)# interface g0/0/0
RouterA(config-if)#ip address 192.168.1.4 255.255.255.0
RouterA(config-if)#ip policy route-map CCNP
```

Answer: D. Option

Explanation

Option D fixes the policy by permitting traffic that matches the ACL and setting critical IP precedence in that route-map entry. The ACL denies host 192.168.1.1, then explicitly permits other addresses, so those other packets can match the route map; the denied host does not match and therefore is not marked. Applying the route map to interface G0/0/0 lets it process traffic entering from the inside network.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-pbr.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/15-s/iri-15-s-book.pdf

✗ **A** ACL denies host 192.168.1.1 without permitting other addresses; route map has a permit entry matching that ACL and setting critical precedence, followed by an empty permit entry.

The ACL has an implicit deny for all addresses not explicitly permitted, so the other inside hosts do not match the route-map entry and are not marked. The later empty permit entry does not set critical precedence for them.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-pbr.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/15-s/iri-15-s-book.pdf

✗ **B** ACL denies host 192.168.1.1 and permits any; route map sequence 10 is deny and matches the ACL, followed by an empty permit sequence 20.
The ACL's permit-any entry allows the other hosts to match the ACL, but the matching route-map entry is a deny entry. Those packets are therefore not handled by a permit entry that sets critical precedence, so this does not mark the required traffic.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-pbr.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/15-s/iri-15-s-book.pdf

✗ **C** ACL denies host 192.168.1.1 and permits any; route map sequence 10 matches the ACL and sets critical precedence, while sequence 20 also sets critical precedence without a match condition.
The first sequence can mark the permitted traffic, but the second unconditional permit sequence also sets critical precedence for packets that did not match sequence 10—including the excluded host. Thus, host 192.168.1.1 is not reliably excluded from marking.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-pbr.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/15-s/iri-15-s-book.pdf

✓ **D** ACL denies host 192.168.1.1 and permits any; a permit route-map entry matches that ACL and sets critical precedence, and the route map is applied to interface G0/0/0.
The explicit permit-any ACL line makes the other hosts eligible to match the permit route-map entry, which sets their IP precedence to critical. The excluded host is denied by the ACL and does not match that entry; applying the policy on G0/0/0 processes the inside-to-outside traffic as it enters the router.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-routing/b-ip-routing/m_iri-pbr.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pi/configuration/15-s/iri-15-s-book.pdf

Question 33

Single choice · Domain: D1 Layer 3 Technologies › 1.5 Troubleshoot manual and auto-summarization with any routing protocol

Refer to the exhibit.

Free demo · full exam: 636 questions



```
R1#
interface Loopback0
 ip address 100.1.1.1 255.255.255.255
!
interface Loopback10
 ip address 10.100.1.10 255.255.255.255
!
interface Loopback20
 ip address 10.100.1.20 255.255.255.255
!
interface Loopback30
 ip address 10.100.1.30 255.255.255.255
!
interface Loopback40
 ip address 10.100.1.40 255.255.255.255
!
interface Loopback50
 ip address 10.100.1.50 255.255.255.255
!
interface FastEthernet0/0
 ip address 10.1.1.1 255.255.255.252
 ip authentication mode eigrp 100 md5
 ip authentication key-chain eigrp 100 EIGRPKEY
 ip summary-address eigrp 100 10.100.1.0 255.255.255.0 5 leak-map LOOPBACK50
!
router eigrp 100
 network 10.1.1.0 0.0.0.3
 network 10.1.1.12 0.0.0.3
 neighbor 10.1.1.2 FastEthernet0/0
 neighbor 10.1.1.14 FastEthernet1/0
 default metric 1000000 10 255 1 1500
 no auto-summary
!
ip prefix-list LOOPBACK50 seq 5 permit 10.100.1.50/32
!
route-map LOOPBACK50 permit 10
 match ip address prefix-list LOOPBACK50
```

Refer to the exhibit. An engineer configured summarization for the R1 loopback addresses and failed.

Which action permits the successful advertisement of the R1 loopback addresses?

Free demo · full exam: 636 questions

A. Configure EIGRP 100 with a network statement for loopback 0 and configure and redistribute the static route for loopback 50.

B. Configure EIGRP 100 with a network statement for loopback 0 and redistribute the connected route for loopback 50.

C. Configure EIGRP 100 with a network statement for loopback 0 and remove the leak map for loopback 50.

D. Configure 100.1.1.1 permit statement in the prefix list LOPPBACK50 and redistribute the connected route for loopback 50.

Answer: B. Configure EIGRP 100 with a network statement for loopback 0 and redistribute the connected route for loopback 50.

Explanation

Loopback0 is not included in EIGRP by the shown network statements, so it needs to be added to the EIGRP process. Loopback50 is also not advertised by EIGRP merely because it is connected: the summary's leak map selects its /32 prefix, but that route must be injected into EIGRP first. Redistributing the connected route for loopback50, together with adding loopback0 to EIGRP, permits the loopback addresses to be advertised.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html

✗ **A** Configure EIGRP 100 with a network statement for loopback 0 and configure and redistribute the static route for loopback 50.
Adding loopback0 to EIGRP addresses its omission, but the exhibit configures loopback50 as a connected interface, not as a static route. Configuring and redistributing a static route therefore does not inject the existing connected loopback50 route into EIGRP.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html

✓ **B** Configure EIGRP 100 with a network statement for loopback 0 and redistribute the connected route for loopback 50.
The EIGRP network statements shown do not include loopback0, so adding a network statement brings it into the process. Loopback50 is a connected route and is selected by the leak-map prefix list; redistributing the connected route makes it available for EIGRP advertisement and summary leaking.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html

✗ **C** Configure EIGRP 100 with a network statement for loopback 0 and remove the leak map for loopback 50.
Adding loopback0 to EIGRP is appropriate, but removing the leak map does not solve the advertisement problem for loopback50. The configured summary uses the leak map to permit the selected /32 route to be advertised along with the summary.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/command/ire-xe-3se-3850-cr-book.pdf
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-eigrp-00.html

✗ **D** Configure 100.1.1.1 permit statement in the prefix list LOPPBACK50 and redistribute the connected route for loopback 50.
The prefix list is intended to match loopback50 at 10.100.1.50/32, not the loopback0 address 100.1.1.1; adding that permit would not bring loopback0 into EIGRP. Although redistributing the connected route for loopback50 is relevant, this option does not add the required EIGRP network statement for loopback0.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/8606-redist.html
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/118974-technote-

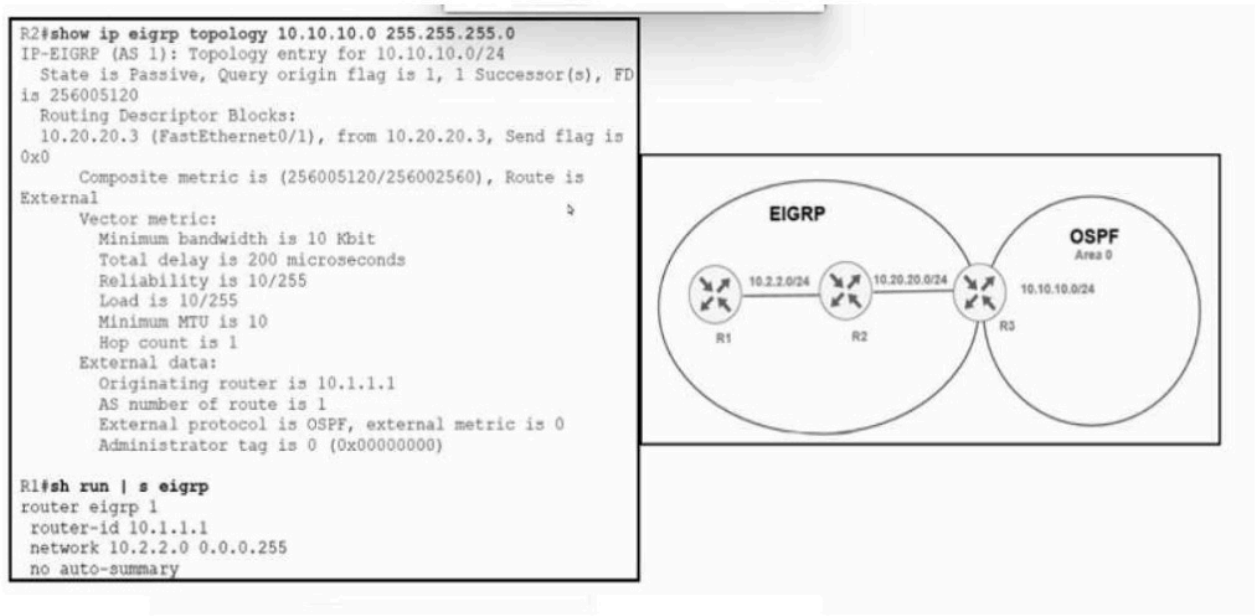
Free demo · full exam: 636 questions

eigrp-00.html

Question 34

Single choice · Domain: D1 Layer 3 Technologies › 1.4 Troubleshoot redistribution between any pair of routing protocols or routing sources

Refer to the exhibit.



Refer to the exhibit. An engineer configured router R3 to redistribute the prefix 10.10.10.0/24 from OSPF into EIGRP R1 has no connectivity to the prefix. Which action enables receipt of prefixes on R1?

A. R3 is advertising the 10.20.20 0'24 prefix with a TTL of 1, R3 must set the TTL to 2 for this prefix.

B. R1 docs not have a neighbor relationship with R2. The EIGRP process should be cleared on R1.

C. Duplicate router IDs on R1 and R3, R1 should modify its router ID.

D. R1 is not receiving the next-hop IP address of R3. R2 must enable the network 10 20.20.0V24 within EIGRP.

Answer: D. R1 is not receiving the next-hop IP address of R3. R2 must enable the network 10 20.20.0V24 within EIGRP.

Explanation

R1 learns the redistributed OSPF route through EIGRP across R2. For R1 to install an EIGRP route advertised by R3, the next-hop address must be reachable; R3's next hop lies across the 10.20.20.0/24 R2–R3 link. Enabling that link in R2's EIGRP process allows the required EIGRP connectivity and lets R1 receive the prefix. [3]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/command/ire-cr-book/ire-a1.html

Free demo · full exam: 636 questions

✗ **A** R3 is advertising the 10.20.20.0/24 prefix with a TTL of 1; R3 must set the TTL to 2 for this prefix.
EIGRP does not use IP TTL as a setting that must be raised to advertise this route. The issue is that the R2–R3 transit network must be included in EIGRP so R1 can resolve R3 as the route's next hop.

✗ **B** R1 does not have a neighbor relationship with R2. The EIGRP process should be cleared on R1.
Clearing EIGRP on R1 does not establish a missing neighbor relationship. The exhibit shows R1's EIGRP network statement, while the missing transit-network advertisement is on R2; configuring that network is the relevant fix.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/21324-trouble-eigrp.html

✗ **C** Duplicate router IDs on R1 and R3; R1 should modify its router ID.
The exhibit shows R1 and R3 using different router IDs: R1 is 10.1.1.1, while the route's originating router is also shown as 10.1.1.1, so this option suggests a potential ID conflict. However, changing R1's ID does not address the missing EIGRP advertisement of the R2–R3 transit network, which prevents next-hop resolution.

https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/18685-eigrp-dup-id.html

✓ **D** R1 is not receiving the next-hop IP address of R3. R2 must enable the network 10.20.20.0/24 within EIGRP.
The redistributed route uses R3 as its next hop, and R1 must be able to reach that next hop to use the route. Enabling 10.20.20.0/24 in R2's EIGRP process advertises the R2–R3 link into EIGRP, allowing R1 to resolve R3's address and receive the prefix.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/command/ire-cr-book/ire-a1.html
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/21324-trouble-eigrp.html

Dumpsgate

Question 35

Multiple choice · Domain: D2 VPN Technologies › 2.1 Describe MPLS operations (LSR, LDP, label switching, LSP)

Which two features are required for MPLS forwarding on which types of routers? (Choose two.)

A. MPLS on CE and core routers
B. LDP on PE and core routers
C. LDP on PE and CE routers
D. MPLS on PE and core routers
E. CEF on PE and CE routers

Answer

- B. **LDP on PE and core routers**
- D. **MPLS on PE and core routers**

Explanation

MPLS VPN forwarding requires the PE routers and routers in the provider core to forward labeled packets, so MPLS must be enabled on those routers. LDP distributes the labels used for that forwarding between the PE and core routers. CE routers connect to the provider edge and do not need MPLS or LDP for this design. [2]

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html

Free demo · full exam: 636 questions

✗ **A** MPLS on CE and core routers
MPLS is required on the provider-edge and core routers, not on the customer-edge routers in this design. This choice substitutes CE for PE, so it does not identify the required router types.
https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html

✓ **B** LDP on PE and core routers
LDP distributes labels used for MPLS forwarding within the provider network. It is required on the PE and core routers, rather than on the CE routers.
https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ias_and_csc_book/mp_carrier_ldp_igp.html

✗ **C** LDP on PE and CE routers
The CE does not participate in the provider's MPLS label-distribution domain. LDP is needed on the PE and core routers, so this choice names the wrong second router type.
https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html

✓ **D** MPLS on PE and core routers
PE routers impose or process VPN labels, while core routers forward labeled traffic across the provider network. MPLS therefore must be enabled on both PE and core routers.
https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/17-9/configuration_guide/mpls/b_179_mpls_9500_cg/multiprotocol_label__switching__mpls_.pdf

✗ **E** CEF on PE and CE routers
This does not specify the required MPLS and LDP features, and it includes CE routers instead of the provider core. It therefore does not meet the forwarding requirements described in the question.

Question 36

Single choice · Domain: D2 VPN Technologies › 2.2 Describe MPLS Layer 3 VPN

How are MPLS Layer 3 VPN services deployed?

A. The RD and RT values must match under the VRF.
B. The label switch path must be available between the local and remote PE routers.
C. The import and export RT values under a VRF must always be the same.
D. The RD and RT values under a VRF must match on the remote PE router.

Answer: B. The label switch path must be available between the local and remote PE routers.

Explanation

MPLS Layer 3 VPN forwarding requires a label-switched path (LSP) across the provider core so traffic can travel between the provider edge (PE) routers. Therefore, the required deployment condition among these choices is that an LSP be available between the local and remote PE routers. RD and RT values identify and control VPN route handling; they do not need to match in the ways the other options claim.
https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/74x/b-l3vpn-cg-ncs5500-74x/implementing-mpls-layer-3-VPNs.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-layer-3-vpn.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

✗ **A** The RD and RT values must match under the VRF.
An RD makes VPN routes unique, while RTs control which VPN routes are imported and exported. They are distinct attributes and are not required to match under a VRF for MPLS L3 VPN service to work.
https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/74x/b-l3vpn-cg-ncs5500-74x/implementing-mpls-layer-3-VPNs.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-layer-3-vpn.html

✓ **B** The label switch path must be available between the local and remote PE routers.
MPLS VPN traffic is carried across the provider network using labels, so a label-switched path must provide transport between the PE routers. This is the required condition stated by the correct choice.
https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/74x/b-l3vpn-cg-ncs5500-74x/implementing-mpls-layer-3-VPNs.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-layer-3-vpn.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

✗ **C** The import and export RT values under a VRF must always be the same.
Import and export route targets serve different purposes: export RTs tag routes advertised from a VRF, and import RTs determine which tagged routes the VRF accepts. They may be configured differently to support VPN connectivity and do not have to always match.
https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/74x/b-l3vpn-cg-ncs5500-74x/implementing-mpls-layer-3-VPNs.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-layer-3-vpn.html

✗ **D** The RD and RT values under a VRF must match on the remote PE router.
The RD and RT have different functions, and a remote PE does not need identical RD and RT values for routes to be exchanged. Route-target import/export policy and the provider-core label path determine how VPN routes and traffic are handled.
https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/74x/b-l3vpn-cg-ncs5500-74x/implementing-mpls-layer-3-VPNs.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-layer-3-vpn.html

Question 37

Single choice · Domain: D2 VPN Technologies › 2.2 Describe MPLS Layer 3 VPN

Which Layer 3 VPN attribute allows different customers to connect to the same MPLS network with overlapping IP ranges?

A. MP-BGP
B. VRF
C. RT
D. RD

Answer: B. VRF

Explanation

VRF (Virtual Routing and Forwarding) creates separate routing tables for customer VPNs on the same MPLS network. Because each customer's routes are kept in its own VRF, customers can use overlapping IP address ranges without those routes conflicting. [1][2]
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-layer-3-vpn.html

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/12-2sr/mp-l3-vpns-12-2sr-book.pdf

✗ **A** MP-BGP
MP-BGP distributes VPN routes between provider routers, but it is not the attribute that separates customers' routing tables. The isolation that permits overlapping customer IP ranges is provided by VRFs. [5][11]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/xe-3s/asr903/17-1-1/b-mpls-l3-vpns-xe-17-1-asr900/m_mp-bgp-mpls-vpn.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-s/mp-l3-vpns-15-s-book/mp-bgp-mpls-vpn.html

✓ **B** VRF
A VRF maintains a distinct routing table for a customer VPN on the shared MPLS network. Separate tables allow different customers to use the same IP prefixes without their routes colliding. [1][2]

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-layer-3-vpn.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/12-2sr/mp-l3-vpns-12-2sr-book.pdf

✗ **C** RT
A route target (RT) is an extended community used to control which VPN routes are imported into or exported from VRFs. It helps determine route membership, but the VRF's separate routing table provides the isolation for overlapping address ranges. [2][11]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/12-2sr/mp-l3-vpns-12-2sr-book.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-s/mp-l3-vpns-15-s-book/mp-bgp-mpls-vpn.html

✗ **D** RD
A route distinguisher (RD) makes otherwise identical customer IPv4 prefixes unique when represented as VPN-IPv4 routes. It supports route identification and distribution, while the VRF provides the customer-specific routing-table separation asked for here. [2][11]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/12-2sr/mp-l3-vpns-12-2sr-book.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-s/mp-l3-vpns-15-s-book/mp-bgp-mpls-vpn.html

Dumpsgate

Question 38

Single choice · Domain: D2 VPN Technologies › 2.2 Describe MPLS Layer 3 VPN

How does MPLS Layer 3 VPN function?

A. When a VPN route is learned from a CE device and injected into IGP, a VPN route distinguisher attribute is associated with it.
B. When a destination PE device receives a labeled packet, it pops the label and uses it to forward the packet to the correct CE device.
C. When an EIGRP internal route is redistributed into BGP by one PE and then back into EIGRP by another PE, the originating router ID for the route is changed to the router ID of the first PE
D. When a PE device forwards a packet received from a CE device across the provider network, it labels the packet with the label learned from the source PE device.

Answer: B. When a destination PE device receives a labeled packet, it pops the label and uses it to forward the packet to the correct CE device.

Explanation

Free demo · full exam: 636 questions

B is correct: the destination PE uses the VPN label to identify the appropriate VPN forwarding context and outgoing CE-facing route, then removes the VPN label before forwarding the IP packet to that CE. The other choices misstate when route distinguishers are added, how EIGRP router IDs are handled, or which PE supplies the forwarding label.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/29828-mplsvpnte.html
https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/26xx/configuration/guide/b-l3vpn-cg-ncs5500-26xx/implementing-mpls-layer-3-VPNs.pdf

✗ **A** When a VPN route is learned from a CE device and injected into IGP, a VPN route distinguisher attribute is associated with it.
A route distinguisher is used to make a VPN-IPv4 route unique when the PE advertises it through MP-BGP; it is not an attribute added when a CE route is injected into the provider IGP. The provider IGP carries reachability within the provider network, not customer VPN routes in this manner.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/29828-mplsvpnte.html
https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/26xx/configuration/guide/b-l3vpn-cg-ncs5500-26xx/implementing-mpls-layer-3-VPNs.pdf

✓ **B** When a destination PE device receives a labeled packet, it pops the label and uses it to forward the packet to the correct CE device.
The VPN label advertised by the destination PE identifies the VPN forwarding information and outgoing path for the packet. The destination PE removes that VPN label and forwards the resulting IP packet toward the appropriate CE device.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/29828-mplsvpnte.html
https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/26xx/configuration/guide/b-l3vpn-cg-ncs5500-26xx/implementing-mpls-layer-3-VPNs.pdf

✗ **C** When an EIGRP internal route is redistributed into BGP by one PE and then back into EIGRP by another PE, the originating router ID for the route is changed to the router ID of the first PE
The EIGRP-originating router ID is carried as route information across the PE-to-PE VPN route exchange, rather than being replaced with the first PE's router ID as stated. This option describes EIGRP route propagation incorrectly.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/xe-16-7/mp-l3-vpns-xe-16-7-book/mpls-vpn-support-for-eigrp-between-pe-and-ce.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-sy/mp-l3-vpns-15-sy-book/mp-vpn-support-eigrp-betw-pe-ce.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/xe-16/ire-xe-16-book/m_eigrp-mpls-vpn.html

✗ **D** When a PE device forwards a packet received from a CE device across the provider network, it labels the packet with the label learned from the source PE device.
The ingress PE uses the VPN label advertised by the destination (egress) PE for the route, not a label learned from the source PE. The provider network may also impose a transport label to carry the packet toward that egress PE.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/29828-mplsvpnte.html
https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/26xx/configuration/guide/b-l3vpn-cg-ncs5500-26xx/implementing-mpls-layer-3-VPNs.pdf

Free demo · full exam: 636 questions

Question 39

Single choice · Domain: D2 VPN Technologies › 2.1 Describe MPLS operations (LSR, LDP, label switching, LSP)

How are LDP neighbors discovered?

A. Unicast hellos are sent to directly connected neighbors IP addresses.
B. Multicast hellos are sent to the 224.0.0.2 group address.
C. Broadcasts hellos are sent to the 255.255.255.255 broadcast address.
D. Multicast hellos are sent to the 224.0.0.5 group address.

Answer: B. Multicast hellos are sent to the 224.0.0.2 group address.

Explanation

LDP discovers neighbors on a directly connected link by sending periodic multicast Hello messages to the all-routers group address 224.0.0.2. This allows LDP-enabled routers on the link to discover one another without knowing each other's unicast addresses in advance.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-label-distribution-protocol.pdf
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ldp_book/mp_ldp_overview.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-16-12/mp-ldp-xe-16-12-book/mp-ldp.pdf

✗ **A** Unicast hellos are sent to directly connected neighbors IP addresses.
LDP link discovery does not require a router to know each neighbor's IP address first. It uses multicast Hello messages on the link to discover neighboring LDP routers.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-label-distribution-protocol.pdf
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ldp_book/mp_ldp_overview.html

✓ **B** Multicast hellos are sent to the 224.0.0.2 group address.
LDP sends link Hello messages to the all-routers multicast address 224.0.0.2. LDP-enabled routers on the local link receive these messages and can discover one another.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-label-distribution-protocol.pdf
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ldp_book/mp_ldp_overview.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-16-12/mp-ldp-xe-16-12-book/mp-ldp.pdf

✗ **C** Broadcasts hellos are sent to the 255.255.255.255 broadcast address.
LDP neighbor discovery uses IP multicast, not an IP broadcast to 255.255.255.255. The broadcast address is not the LDP link Hello destination.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-label-distribution-protocol.pdf
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ldp_book/mp_ldp_overview.html

✗ **D** Multicast hellos are sent to the 224.0.0.5 group address.
Although LDP uses multicast Hellos, 224.0.0.5 is the AllSPFRouters group used by OSPF. LDP's link Hello destination is 224.0.0.2.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-label-distribution-protocol.pdf
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ldp_book/mp_ldp_overview.html

Free demo · full exam: 636 questions

Question 40

Single choice · Domain: D2 VPN Technologies › 2.2 Describe MPLS Layer 3 VPN

How many labels are present in an MPLS Layer 3 packet traversing through the network without traffic engineering?

A. 1
B. 2
C. 3
D. 4

Answer: B. 2

Explanation

An MPLS Layer 3 VPN packet normally carries two labels across the provider network: an outer transport label to reach the egress provider edge router and an inner VPN label to identify the destination VPN route. Without traffic engineering, there is no additional TE tunnel label, so the correct count is two.

https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/26xx/configuration/guide/b-l3vpn-cg-ncs5500-26xx/implementing-mpls-layer-3-VPNs.pdf
https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13734-mpls-vpn-tsh.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

✗ **A** 1
One label would not account for both functions in a Layer 3 VPN packet: the packet needs a transport label for forwarding through the provider core and a VPN label for delivery to the correct VPN route.

https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/26xx/configuration/guide/b-l3vpn-cg-ncs5500-26xx/implementing-mpls-layer-3-VPNs.pdf
https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13734-mpls-vpn-tsh.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

✓ **B** 2
The packet carries an outer transport label and an inner VPN label as it traverses the provider network. Because the scenario excludes traffic engineering, no additional TE label is present.

https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/26xx/configuration/guide/b-l3vpn-cg-ncs5500-26xx/implementing-mpls-layer-3-VPNs.pdf
https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13734-mpls-vpn-tsh.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

✗ **C** 3
Three labels would imply an additional label beyond the transport and VPN labels, such as one associated with a TE tunnel. The question explicitly says the packet is traversing without traffic engineering.

https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/26xx/configuration/guide/b-l3vpn-cg-ncs5500-26xx/implementing-mpls-layer-3-VPNs.pdf
https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13734-mpls-vpn-tsh.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

✗ **D** 4
Four labels would require multiple additional labels beyond the usual transport-and-VPN pair. No such extra label-stack requirements are specified, and traffic engineering is excluded.

https://www.cisco.com/c/en/us/td/docs/iosxr/ncs5500/vpn/26xx/configuration/guide/b-l3vpn-cg-ncs5500-26xx/implementing-mpls-layer-3-VPNs.pdf

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13734-mpls-vpn-tsh.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

Question 41

Single choice · Domain: D2 VPN Technologies › 2.2 Describe MPLS Layer 3 VPN

Which Layer 3 VPN attribute installs customer routes in the VRF?

A. extended-community
B. MPLS label
C. RD
D. RT

Answer: D. RT

Explanation

The correct answer is the route target (RT). An RT is an extended BGP community used to control which VPN routes are imported into a VRF, so it determines whether customer routes are installed there. [1]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/xe-16/mp-l3-vpns-xe-16-book/mpls-vpn-route-target-rewrite.pdf

✗ **A** extended-community
This is a broader BGP attribute category, not the specific attribute being asked for. The route target is an extended community, and that specific value controls importing routes into a VRF.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/xe-16/mp-l3-vpns-xe-16-book/mpls-vpn-route-target-rewrite.pdf

✗ **B** MPLS label
The MPLS label is used to forward traffic through the provider network toward the egress router or VPN context. It does not determine whether a customer route is imported into a VRF.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-mt/mp-l3-vpns-15-mt-book/mp-cfg-layer3-vpn.html

✗ **C** RD
A route distinguisher makes otherwise overlapping customer prefixes unique in the VPN routing information. It does not select routes for installation in a VRF; route targets perform that import function.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-mt/mp-l3-vpns-15-mt-book/mp-cfg-layer3-vpn.html

✓ **D** RT
A route target identifies VPN routes that a VRF is configured to import. Routes carrying a matching RT are eligible to be installed in that VRF.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/xe-16/mp-l3-vpns-xe-16-book/mpls-vpn-route-target-rewrite.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-mt/mp-l3-vpns-15-mt-book/mp-cfg-layer3-vpn.html

Free demo · full exam: 636 questions

Question 42

Single choice · Domain: D2 VPN Technologies › 2.2 Describe MPLS Layer 3 VPN

Which Layer 3 VPN attribute allows different customers to connect to the same MPLS network wrth overlapping IP ranges?

A. VRF
B. RT
C. MP-BGP
D. RD

Answer: D. RD

Explanation

The supplied answer is D, RD (Route Distinguisher). In an MPLS Layer 3 VPN, an RD makes otherwise identical customer routes unique when they are carried in the provider's VPN routing system, so customers can use overlapping IP address ranges. A VRF separates a customer's routing information, while the RD distinguishes the routes themselves.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

✗ **A** VRF
A VRF keeps a customer's routing table separate from other customers' routing tables on a provider edge router. It supports VPN isolation, but the route distinguisher is what makes overlapping customer prefixes unique in the VPN routing system.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

✗ **B** RT
A route target controls which VPN routes are imported into or exported from a VRF. It determines VPN membership and route distribution; it does not distinguish identical prefixes belonging to different customers.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/xe-16/mp-l3-vpns-xe-16-book/mpls-vpn-route-target-rewrite.pdf
https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html

✗ **C** MP-BGP
MP-BGP distributes VPN routes between provider edge routers and carries VPN-related attributes. It is the transport mechanism for exchanging those routes, not the attribute that makes overlapping prefixes unique.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

✓ **D** RD
An RD is added to an IPv4 prefix to create a unique VPN-IPv4 route, allowing the same IP range to be used by different customers without the routes colliding in the provider's VPN routing system. That is the capability required by the question.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/13733-mpls-vpn-basic.html
https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/l3vpn/l3vpn-config-cisco8000/mpls-l3vpn-fundamentals-and-core-deployment-w/mpls-layer-3-vpn-services.html

Free demo · full exam: 636 questions

Question 43

Multiple choice · Domain: D2 VPN Technologies › 2.1 Describe MPLS operations (LSR, LDP, label switching, LSP)

I .

Which two labet distribution methods are used by routers in MPLS? (Choose two)

A. targeted hello message
B. LDP discovery hello message
C. LDP session protection message
D. downstream unsolicited
E. downstream on demand

Answer

- D. **downstream unsolicited**
- E. **downstream on demand**

Explanation

The question asks for the two label distribution methods used by routers in MPLS. These are downstream unsolicited (DU) and downstream on demand (DoD), which describe how a router requests or receives label mappings from its downstream peer. The other options refer to LDP discovery or session mechanisms, not label distribution methods.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/mpls/mpls/m_mp-ldp-overview.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-16/mp-ldp-xe-16-book.pdf
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/12_2sr/mp_12_2sr_book/mp_ldp_overview.html

✗ **A** targeted hello message
A targeted Hello is an LDP discovery message used to establish discovery with a specifically identified peer. It is not a method for distributing label mappings.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/mpls/mpls/m_mp-ldp-overview.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-16/mp-ldp-xe-16-book.pdf

✗ **B** LDP discovery hello message
LDP Discovery Hello messages let routers discover LDP peers on a link or through targeted discovery. They do not describe the method by which label mappings are distributed.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/mpls/mpls/m_mp-ldp-overview.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-16/mp-ldp-xe-16-book.pdf

✗ **C** LDP session protection message
Session protection is an LDP resiliency feature that helps preserve an LDP session. It is not one of the two label distribution methods.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/mpls/mpls/m_mp-ldp-overview.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-16/mp-ldp-xe-16-book.pdf

✓ **D** downstream unsolicited
With downstream unsolicited distribution, a router advertises label mappings to its upstream peers without first receiving a request for each mapping. This is one of the two label distribution methods used with MPLS LDP.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/mpls/mpls/m_mp-ldp-overview.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-16/mp-ldp-xe-16-book.pdf
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/12_2sr/mp_12_2sr_book/mp_ldp_overview.html

✓ **E** downstream on demand
With downstream on-demand distribution, a router requests a label mapping from its downstream peer

Free demo · full exam: 636 questions

when needed. This is the other label distribution method used with MPLS LDP.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/mpls/mpls/m_mp-ldp-overview.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/xe-16/mp-ldp-xe-16-book.pdf
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/12_2sr/mp_12_2sr_book/mp_ldp_overview.html

Question 44

Single choice · Domain: D2 VPN Technologies › 2.1 Describe MPLS operations (LSR, LDP, label switching, LSP)

Which rouler takes an active role between two LDP neighbors when initiating LDP session negotiation and LDP TCP connection establishment?

A. with the higher IP address
B. with the larger number of LDP TCP neighbors
C. with the lowest IP address
D. with one interface in the MPLS backbone

Answer: A. with the higher IP address

Explanation

The router with the higher IP address takes the active role in LDP session negotiation and TCP connection establishment between two LDP neighbors. The active router initiates the TCP connection; the other neighbor responds. Therefore, option A is correct.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/15-s/mp-ldp-15-s-book/mp-ldp-overview.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-label-distribution-protocol.html
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ldp_book/mp_ldp_overview.html

✓ **A** with the higher IP address
The LDP neighbor with the higher IP address initiates the TCP connection and takes the active role in establishing the LDP session. This matches the question's description of initiating session negotiation and TCP connection establishment.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/15-s/mp-ldp-15-s-book/mp-ldp-overview.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-label-distribution-protocol.html
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ldp_book/mp_ldp_overview.html

✗ **B** with the larger number of LDP TCP neighbors
The number of other LDP TCP neighbors a router has does not determine which neighbor initiates the connection between this pair. The selection is based on the neighbors' IP addresses, not their neighbor counts.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/15-s/mp-ldp-15-s-book/mp-ldp-overview.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-label-distribution-protocol.html
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ldp_book/mp_ldp_overview.html

✗ **C** with the lowest IP address
This reverses the IP-address rule: the higher-IP-address neighbor is the active initiator, not the one with the lowest IP address.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ldp/configuration/15-s/mp-ldp-15-s-book/mp-ldp-overview.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/mpls/mpls-configuration-guide/mpls-label-distribution-

Free demo · full exam: 636 questions

protocol.html
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/convert/mp_ldp_book/mp_ldp_overview.html

✗ **D** with one interface in the MPLS backbone
Having an interface in the MPLS backbone is not the criterion that determines which of two LDP neighbors initiates session negotiation or TCP establishment. The active role is determined by the IP-address comparison.

Question 45

Single choice · Domain: D2 VPN Technologies › 2.2 Describe MPLS Layer 3 VPN

Which routing protocol is used by the PE router to advertise routes to a CE router without redistribution or static after removing the RD tag from the P router?

A. IS-IS
B. OSPF
C. BGPIPv4
D. MP-BGP

Answer: C. BGPIPv4

Explanation

The PE uses BGP for the PE–CE exchange in this scenario, so the correct choice is BGP IPv4. MP-BGP carries VPNv4 routes between PE routers; when a PE advertises a customer route to a CE, the route is presented as an IPv4 route rather than a VPNv4 route with its route distinguisher. [5]
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/12_2sr/mp_12_2sr_book/mp_cfg_layer3_vpn.html

✗ **A** IS-IS
IS-IS is an interior routing protocol and is not the protocol identified for advertising the customer routes from the PE to the CE here. The question's requirement points to the BGP PE–CE exchange.
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/12_2sr/mp_12_2sr_book/mp_cfg_layer3_vpn.html

✗ **B** OSPF
OSPF can be used as a PE–CE routing protocol in some MPLS VPN designs, but it is not the specified choice for this question. The correct protocol for the described advertisement is BGP IPv4.
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/12_2sr/mp_12_2sr_book/mp_cfg_layer3_vpn.html

✓ **C** BGPIPv4
BGP IPv4 is used for the PE-to-CE advertisement in this scenario. The PE advertises the customer route as an IPv4 route, without the VPN route distinguisher used for VPNv4 routes between PEs.
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/12_2sr/mp_12_2sr_book/mp_cfg_layer3_vpn.html

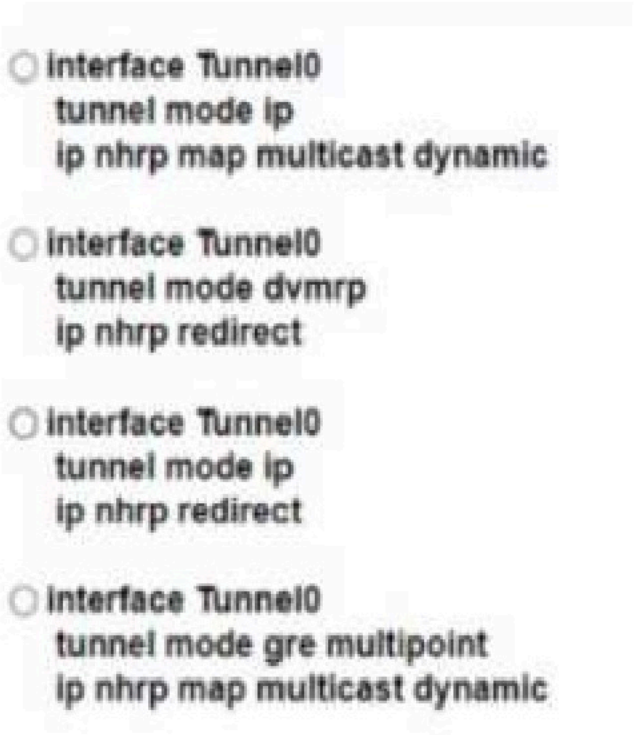
✗ **D** MP-BGP
MP-BGP is used to exchange VPNv4 routes between PE routers across the provider network. It is not the protocol used for this PE-to-CE advertisement as an IPv4 route.
https://www.cisco.com/c/en/us/td/docs/ios/mpls/configuration/guide/12_2sr/mp_12_2sr_book/mp_cfg_layer3_vpn.html

Free demo · full exam: 636 questions

Question 46

Single choice · Domain: D2 VPN Technologies › 2.3 Configure and verify DMVPN (single hub)

A company Is redesigning WAN infrastructure so that all branch sites must communicate via the head office and the head office can directly communicate with each site independently. The network engineer must configure the head office router by considering zero-touch technology when adding new sites in the same WAN infrastructure. Which configuration must be applied to the head office router to meet this requirement?



A. Option A
B. Option B
C. Option C
D. Option D

Answer: D. Option D

Explanation

Option D configures the head-office tunnel as multipoint GRE and enables dynamic NHRP multicast mappings. This supports a hub-and-spoke WAN in which branch routers register with the head office, allowing new sites to be added without manually configuring a separate tunnel or mapping for each one. The other choices either do not configure multipoint GRE or use NHRP redirect rather than the required dynamic multicast mapping. [12]

https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.html

✗ **A** interface Tunnel0; tunnel mode ip; ip nhrp map multicast dynamic
This uses IP-in-IP tunnel mode, not multipoint GRE, so it does not provide the multipoint tunnel required for the hub to accept dynamically added spokes. The dynamic multicast mapping command alone does not make this the appropriate hub configuration.

https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.html

✗ **B** interface Tunnel0; tunnel mode dvmrp; ip nhrp redirect
DVMRP is not the multipoint GRE tunnel mode used for this DMVPN hub-and-spoke configuration. NHRP redirect is associated with directing traffic toward direct spoke-to-spoke paths, rather than enabling new spokes to register through dynamic multicast mappings.

https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.html

✗ **C** interface Tunnel0; tunnel mode ip; ip nhrp redirect
This also uses IP-in-IP rather than multipoint GRE, so it does not establish the required multipoint hub tunnel. It configures NHRP redirect instead of the dynamic multicast mapping needed for zero-touch addition of spokes.

https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.html

✓ **D** interface Tunnel0; tunnel mode gre multipoint; ip nhrp map multicast dynamic
Multipoint GRE lets the head-office router use a single tunnel interface for multiple branch spokes. The dynamic NHRP multicast mapping supports spokes registering with the hub without requiring a static per-site mapping, meeting the zero-touch requirement.

https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.html

Dumpsgate

Question 47

Single choice · Domain: D2 VPN Technologies › 2.1 Describe MPLS operations (LSR, LDP, label switching, LSP)

Which technique removes the outermost label of an MPLS-tagged packet before the packet is forwarded to an adjacent LER?

A. label swap
B. explicit-null
C. label imposition
D. PHP

Answer: D. PHP

Explanation

The correct technique is PHP (Penultimate Hop Popping). The penultimate LSR removes the outermost MPLS label before forwarding the packet to the egress LER, which avoids requiring the egress LER to perform that label-removal operation. This is distinct from swapping or adding a label. [11][12]

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html?referring_site=bodynav
https://community.cisco.com/t5/mpls/implicit-null/m-p/1059417/highlight/true

✗ **A** label swap
A label swap replaces the top label with another label as the packet is forwarded. It does not describe removing the outermost label before forwarding to the adjacent LER.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html?

referring_site=bodynav

✗ **B** explicit-null
Explicit-null uses a specific label value that is retained in the label stack and delivered to the egress LER. It is not the operation that pops the outermost label before forwarding to that LER.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html?referring_site=bodynav

✗ **C** label imposition
Label imposition adds an MPLS label to a packet, typically as it enters an MPLS path. The question asks about removing the outermost label, which is the opposite operation.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html?referring_site=bodynav

✓ **D** PHP
PHP means Penultimate Hop Popping: the router immediately before the egress LER removes the outermost MPLS label and forwards the packet to that LER. This matches the specified label-removal step.

https://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html?referring_site=bodynav
https://community.cisco.com/t5/mpls/implicit-null/m-p/1059417/highlight/true

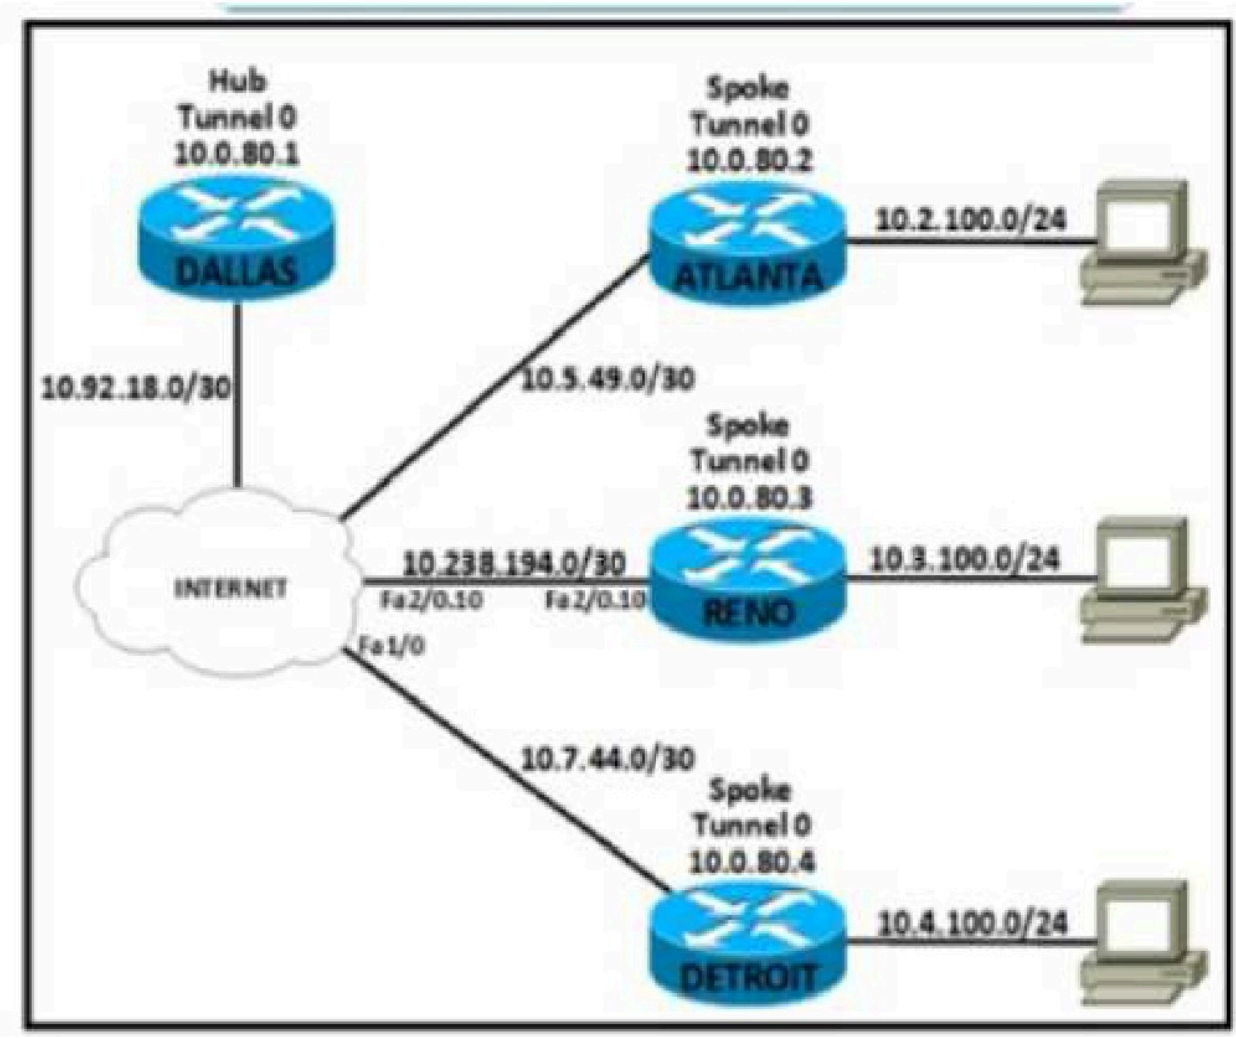
Question 48

Single choice · Domain: D2 VPN Technologies › 2.3.e Spoke-to-spoke

Refer to the exhibit.

Dumpsgate

Free demo · full exam: 636 questions



Dallas

```
interface Tunnel0
 ip address 10.0.80.1 255.255.255.0
 ip nhrp authentication cisco123
 ip nhrp map multicast dynamic
 ip nhrp network-id 5
 tunnel source Serial0/0
 tunnel mode gre multipoint
```

Reno

```
interface Tunnel0
 ip address 10.0.80.3 255.255.255.0
 ip nhrp authentication cisco321
 ip nhrp map multicast 10.92.18.2
 ip nhrp map 10.0.80.1 10.92.18.2
 ip nhrp network-id 5
```


Free demo · full exam: 636 questions

```
ip nhrp nhs 10.0.80.1
tunnel source 10.238.194.2
tunnel mode gre multipoint
```

Detroit

```
interface Tunnel0
 ip address 10.0.80.4 255.255.255.0
 ip nhrp authentication cisco321
 ip nhrp map 10.0.80.1 10.92.18.2
 ip nhrp map multicast 10.92.18.2
 ip nhrp network-id 5
 ip nhrp nhs 10.0.80.1
 tunnel source 10.7.44.2
 tunnel mode gre multipoint
```

Reno

```
interface Tunnel0
 ip address 10.0.80.3 255.255.255.0
 ip nhrp authentication cisco123
 ip nhrp map multicast 10.92.18.2
 ip nhrp map 10.92.18.2 10.0.80.1
 ip nhrp network-id 5
 ip nhrp nhs 10.0.80.1
 tunnel source 10.238.194.2
 tunnel mode gre multipoint
```

Detroit

```
interface Tunnel0
 ip address 10.0.80.4 255.255.255.0
 ip nhrp authentication cisco123
 ip nhrp map 10.92.18.2 10.0.80.1
 ip nhrp map multicast 10.92.18.2
 ip nhrp network-id 5
 ip nhrp nhs 10.0.80.1
 tunnel source 10.7.44.2
 tunnel mode gre multipoint
```

○ Reno

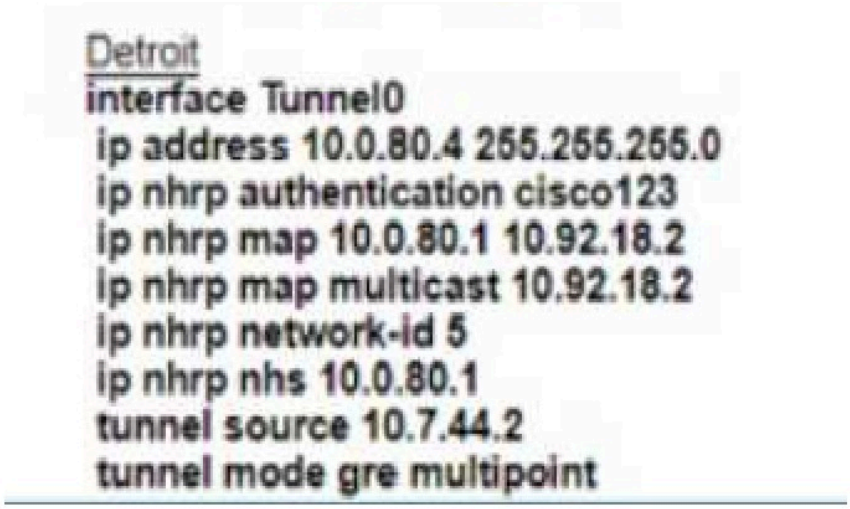
```
interface Tunnel0
 ip address 10.0.80.3 255.255.255.0
 ip nhrp authentication cisco123
 ip nhrp map broadcast 10.92.18.2
 ip nhrp map 10.0.80.1 10.92.18.2
 ip nhrp network-id 5
 ip nhrp nhs 10.0.80.1
 tunnel source 10.238.194.2
 tunnel mode gre multipoint
```

Detroit

```
interface Tunnel0
 ip address 10.0.80.4 255.255.255.0
 ip nhrp authentication cisco123
 ip nhrp map 10.0.80.1 10.92.18.2
 ip nhrp map broadcast 10.92.18.2
 ip nhrp network-id 5
 ip nhrp nhs 10.0.80.1
 tunnel source 10.7.44.2
 tunnel mode gre multipoint
```

○ Reno

```
interface Tunnel0
 ip address 10.0.80.3 255.255.255.0
 ip nhrp authentication cisco123
 ip nhrp map multicast 10.92.18.2
 ip nhrp map 10.0.80.1 10.92.18.2
 ip nhrp network-id 5
 ip nhrp nhs 10.0.80.1
 tunnel source 10.238.194.2
 tunnel mode gre multipoint
```



Refer to the exhibit An engineer must connect the Reno and Detroit spokes using DMVPN phase 2 Hub tunnel configuration is Which configuration accomplishes the task?

A. Option A
B. Option B
C. Option C
D. Option D

Answer: C. Option C

Explanation

Option C is the correct choice for connecting Reno and Detroit as DMVPN Phase 2 spokes. The spoke configurations must use the same NHRP authentication and network ID as the hub, identify the hub as the NHS, and map the hub's tunnel address to its reachable NBMA address. The multicast mapping also lets the spokes use the hub for NHRP and routing control traffic while Phase 2 allows spoke-to-spoke tunnels.

https://www.cisco.com/c/en/us/support/docs/security/dynamic-multipoint-vpn-dmvpn/221665-troubleshoot-dmvpn-phase-2-spoke-to-spok.html
https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.pdf

✗ **A** Option A
This configuration does not provide the required consistent NHRP authentication between the spokes and the hub: the displayed spoke authentication differs from the hub's `cisco123` setting. NHRP authentication must match for the spokes to participate in the hub's DMVPN network.

https://www.cisco.com/c/en/us/support/docs/security/dynamic-multipoint-vpn-dmvpn/111976-dmvpn-troubleshoot-00.html
https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.pdf

✗ **B** Option B
The Reno static NHRP mapping shown reverses the tunnel and NBMA addresses: it associates the hub's tunnel address with the wrong address field. A spoke needs to map the hub's tunnel IP to the hub's reachable NBMA address so it can resolve and reach the hub.

https://www.cisco.com/c/en/us/support/docs/security/dynamic-multipoint-vpn-dmvpn/111976-dmvpn-troubleshoot-00.html
https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.pdf

✓ **C** Option C

This is the configuration identified as correct for the required Phase 2 spoke-to-spoke setup. It provides the spokes with the hub mapping and multicast reachability needed for NHRP, while the matching authentication, network ID, NHS, and multipoint GRE settings allow the spokes to join the same DMVPN cloud.

https://www.cisco.com/c/en/us/support/docs/security/dynamic-multipoint-vpn-dmvpn/221665-troubleshoot-dmvpn-phase-2-spoke-to-spok.html
https://community.cisco.com/kxiwq67737/attachments/kxiwq67737/4561-docs-security/7611/1/DMVPN%20Phase%202%20and%203%20with%20OSPF%20and%20EIGRP.pdf
https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.pdf

✗ **D** Option D

This configuration does not correctly establish the required NHRP multicast mapping to the hub. Without a valid multicast mapping, the spokes cannot use the hub for the control traffic needed to form and maintain the DMVPN network.

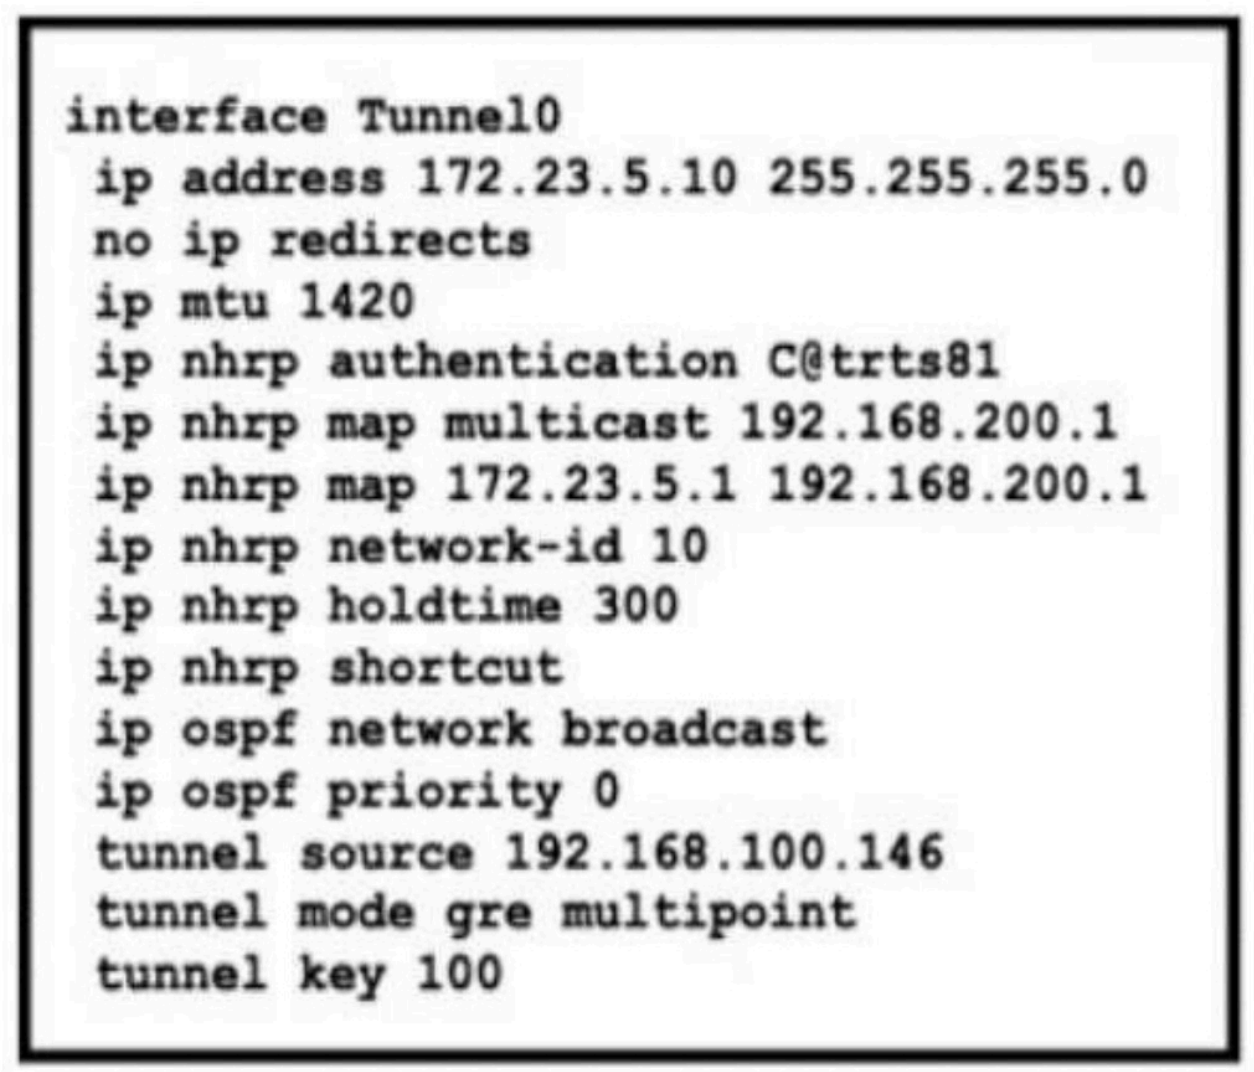
https://www.cisco.com/c/en/us/support/docs/security/dynamic-multipoint-vpn-dmvpn/111976-dmvpn-troubleshoot-00.html
https://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/41940-dmvpn.pdf

Dumpsgate

Question 49

Single choice · Domain: D2 VPN Technologies › 2.3.b NHRP

Refer to the exhibit.



```
interface Tunnel0
 ip address 172.23.5.10 255.255.255.0
 no ip redirects
 ip mtu 1420
 ip nhrp authentication C@trts81
 ip nhrp map multicast 192.168.200.1
 ip nhrp map 172.23.5.1 192.168.200.1
 ip nhrp network-id 10
 ip nhrp holdtime 300
 ip nhrp shortcut
 ip ospf network broadcast
 ip ospf priority 0
 tunnel source 192.168.100.146
 tunnel mode gre multipoint
 tunnel key 100
```

A network engineer is adding a new spoke router into an existing DMVPN Phase 3 tunnel with a hub router to provide secure communication between sites Which additional configuration must the engineer apply to enable the tunnel to come up?

A. ip nhrp registration no-unique
B. ip nhrp server-only non-caching
C. ip nhrp responder tunnel
D. ip nhrpnhs 172.23.5.1

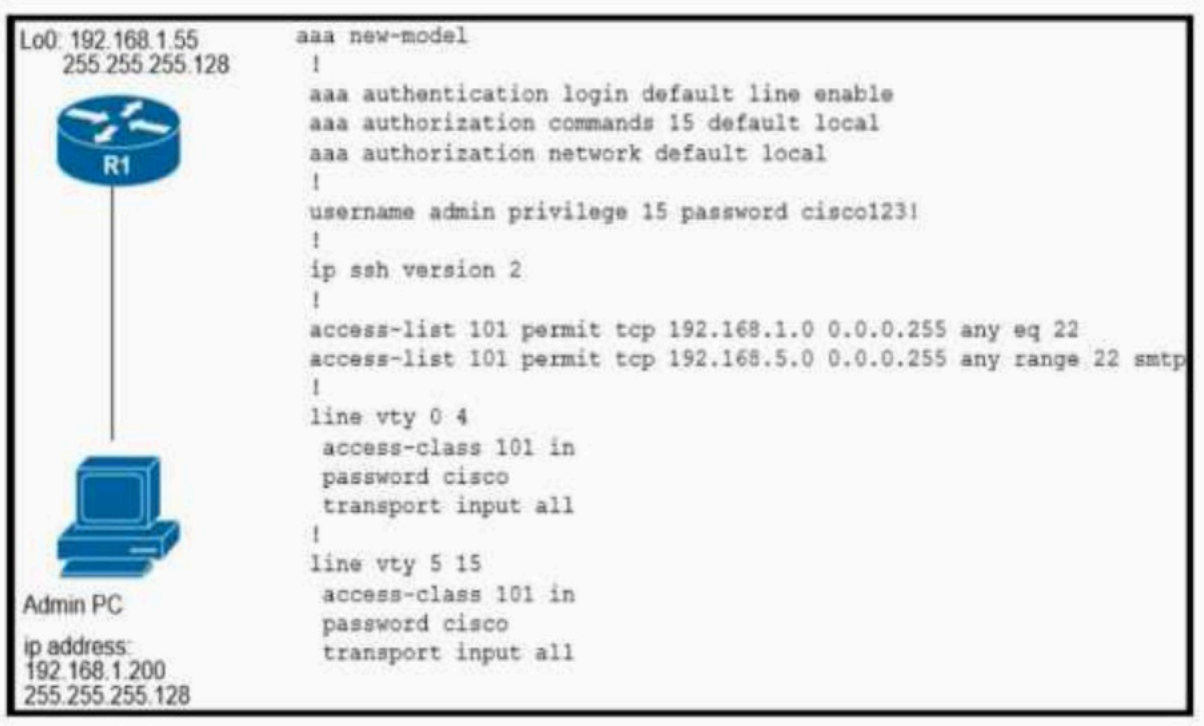
Answer: D. ip nhrpnhs 172.23.5.1

Dumpsgate

Question 50

Single choice · Domain: D3 Infrastructure Security › 3.1 Troubleshoot device security using IOS AAA (TACACS+, RADIUS, local database)

Refer to the exhibit.



Refer to the exhibit. The administrator successfully logs into R1 but cannot access privileged mode commands. What should be configured to resolve the issue?

A. aaa authorization reverse-access
B. enable secret or enable password commands to enter into privileged mode
C. Matching password on vty lines as cisco123!
D. secret cisco123! at the end of the username command instead of password cisco123!

Answer: B. enable secret or enable password commands to enter into privileged mode

Explanation

The administrator can authenticate to the VTY lines, but the configuration does not show an enable secret or enable password for entering privileged EXEC mode. Configure an enable secret or enable password so the administrator can move from the login session into privileged mode. The local username credential and the VTY line password serve different purposes from the enable credential.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-cfg-sec-4cli-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_cfg/configuration/15-s/sec-usr-cfg-15-s-book/sec-cfg-sec-4cli.html

✗ **A** aaa authorization reverse-access
This does not configure a credential for entering privileged EXEC mode. The issue is that the

configuration lacks an enable secret or enable password, not reverse-access authorization.

✓ **B** enable secret or enable password commands to enter into privileged mode
The exhibit configures login authentication and a privileged local username, but shows no enable secret or enable password. Configuring one supplies the credential needed to enter privileged EXEC mode after login.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-cfg-sec-4cli-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_cfg/configuration/15-s/sec-usr-cfg-15-s-book/sec-cfg-sec-4cli.html

✗ **C** Matching password on vty lines as cisco123!
The VTY lines already have a password configured, and the administrator successfully logs in. A VTY password controls access to the line; matching it to the username password does not provide the missing enable credential.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-cfg-sec-4cli-0.html

✗ **D** secret cisco123! at the end of the username command instead of password cisco123!
Changing the local username password to a secret changes how that account's credential is stored, but does not configure the enable credential required for privileged EXEC mode. The administrator already successfully logs in, so this does not address the reported problem.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-cfg-sec-4cli-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/a1/sec-a1-cr-book/sec-cr-a1.html

Dumpsgate

Question 51

Single choice · Domain: D3 Infrastructure Security › 3.2.c Unicast reverse path forwarding (uRPF)

Refer to the exhibit.



```
R3#show cef interface gi0/3
GigabitEthernet0/3 is up (if_number 5)
  Corresponding hwidb fast_if_number 5
  Corresponding hwidb firstsw->if_number 5
  Internet address is 172.16.4.253/30
  ICMP redirects are never sent
  Per packet load-sharing is disabled
  IP unicast RPF check is enabled
  Input features: uRPF
  IP policy routing is disabled
  BGP based policy accounting on input is disabled
  BGP based policy accounting on output is disabled
  Hardware idb is GigabitEthernet0/3
  Fast switching type 1, interface type 27
  IP CEF switching enabled
  IP CEF switching turbo vector
  IP prefix lookup IPv4 mtrie 8-8-8-8 optimized
  Input fast flags 0x4000, Output fast flags 0x0
  ifindex 5(5)
  Slot  Slot unit 3 VC -1
  IP MTU 1500
R3#show run int gi0/3
Building configuration...

Current configuration : 162 bytes
!
interface GigabitEthernet0/3
 ip address 172.16.4.253 255.255.255.252
 ip verify unicast source reachable-via rx
 duplex auto
 speed auto
 media-type rj45
end
```

Refer to the exhibit. An engineer implements uRPF to increase security and stop incoming spoofed IP packets. Some asymmetrically routed packets are also blocked after the configuration. Which command resolves the issue?

A. ip verify unicast source reachable-via rx
B. ip verify unicast reverse-path
C. ip verify unicast source reachable-via any
D. ip verify unicast reverse-path any

Answer: C. ip verify unicast source reachable-via any

Explanation

The exhibit shows uRPF configured with `ip verify unicast source reachable-via rx`, which checks that the packet's source is reachable through the interface on which it arrived. That strict check can drop legitimate packets when routing is asymmetric. Replacing it with `ip verify unicast source reachable-via any` allows a source to be reachable through any interface while retaining the source-reachability check, addressing the asymmetric-routing issue. [5][7]

Free demo · full exam: 636 questions

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/command/ir-cr-book/ir-i1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_urpf/configuration/xe-3se/3850/sec-data-urpf-xe-3se-3850-book.pdf

✗ **A** ip verify unicast source reachable-via rx
This is the command shown in the exhibit. The `rx` mode requires the route to the source to use the receiving interface, so valid packets arriving over an asymmetric path can be rejected.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/command/ir-cr-book/ir-i1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_urpf/configuration/xe-3se/3850/sec-data-urpf-xe-3se-3850-book.pdf

✗ **B** ip verify unicast reverse-path
This is not the command that changes the shown strict interface check to accept sources reachable through another interface. The option does not specify the `any` behavior needed for the asymmetric-routing scenario.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/command/ir-cr-book/ir-i1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/d1/sec-d1-cr-book/sec-cr-i3.html

✓ **C** ip verify unicast source reachable-via any
The `any` mode accepts a packet when its source is reachable through any interface, rather than requiring the route to point back through the receiving interface. This accommodates asymmetric routing while continuing to validate source reachability.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/command/ir-cr-book/ir-i1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_urpf/configuration/xe-3se/3850/sec-data-urpf-xe-3se-3850-book.pdf

✗ **D** ip verify unicast reverse-path any
This does not match the documented `source reachable-via any` command form used to replace the configuration shown in the exhibit. The required change is to allow source reachability through any interface, as in option C.

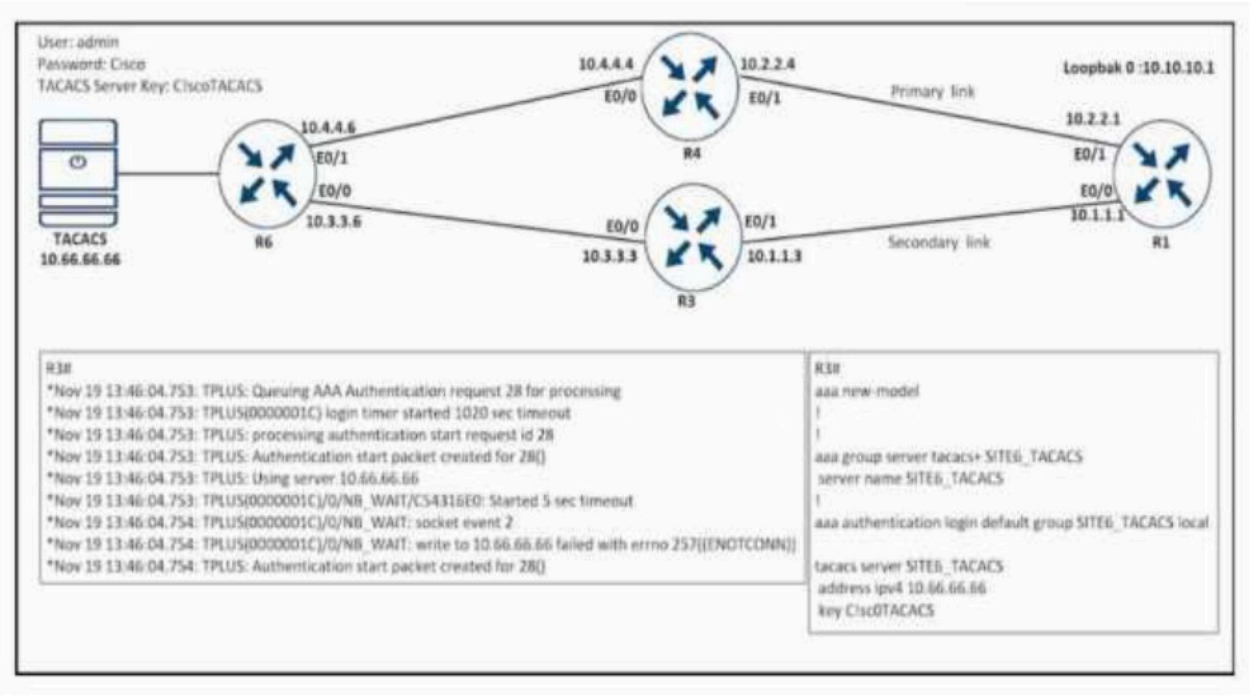
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/interface/command/ir-cr-book/ir-i1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/d1/sec-d1-cr-book/sec-cr-i3.html

Dumpsgate

Question 52

Single choice · Domain: D3 Infrastructure Security › 3.1 Troubleshoot device security using IOS AAA (TACACS+, RADIUS, local database)

Refer to the exhibit.



Refer to the exhibit. R3 cannot authenticate via TACACS. Which configuration resolves the issue?

A. tacacs server SITE6_TACACS key C!scOTACACS address ipv4 10.60.66.66
B. tacacs server SITE6_TACACS key C!scoTACACS address ipv4 10.60.66.66
C. tacacs server SITE6_TACACS key CiscoTACACS address ipv4 10.66.66.66
D. tacacs server SITE6_TACACS key CIscoTACACS

Answer: C. tacacs server SITE6_TACACS key CiscoTACACS address ipv4 10.66.66.66

Explanation

R3's TACACS server definition must use the TACACS server's address and the shared key shown in the exhibit. The exhibit identifies the server as 10.66.66.66 and its key as CiscoTACACS, which are both present in option C. Matching the configured address and key allows R3 to communicate with and authenticate to the TACACS server.

https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/200467-Troubleshoot-TACACS-Authentication-Issue.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/17-8/configuration_guide/sec/b_178_sec_9500_cg/configuring_tacacs_.pdf

✗ **A** tacacs server SITE6_TACACS key C!scOTACACS address ipv4 10.60.66.66
This does not match the exhibit's TACACS server details: the server address is shown as 10.66.66.66 and the key as CiscoTACACS. The incorrect address and key prevent R3 from using the configured server

successfully.

https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/200467-Troubleshoot-TACACS-Authentication-Issue.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/17-8/configuration_guide/sec/b_178_sec_9500_cg/configuring_tacacs_.pdf

✗ **B** tacacs server SITE6_TACACS key C!scoTACACS address ipv4 10.60.66.66
The option specifies 10.60.66.66, but the TACACS server in the exhibit is 10.66.66.66. Its key also differs from the displayed CiscoTACACS value, so this configuration does not match the server.

https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/200467-Troubleshoot-TACACS-Authentication-Issue.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/17-8/configuration_guide/sec/b_178_sec_9500_cg/configuring_tacacs_.pdf

✓ **C** tacacs server SITE6_TACACS key CiscoTACACS address ipv4 10.66.66.66
This option configures both values shown in the exhibit: the TACACS server address 10.66.66.66 and the shared key CiscoTACACS. Using the matching server address and key resolves the TACACS configuration mismatch for R3.

https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/200467-Troubleshoot-TACACS-Authentication-Issue.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/17-8/configuration_guide/sec/b_178_sec_9500_cg/configuring_tacacs_.pdf

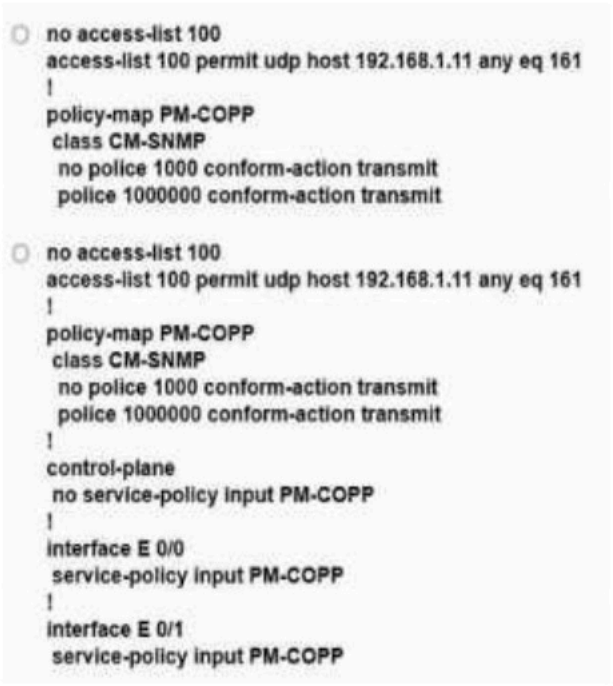
✗ **D** tacacs server SITE6_TACACS key CIscoTACACS
This key is not the CiscoTACACS value shown in the exhibit, and the option does not specify the TACACS server address 10.66.66.66. It therefore does not provide the matching server configuration R3 needs.

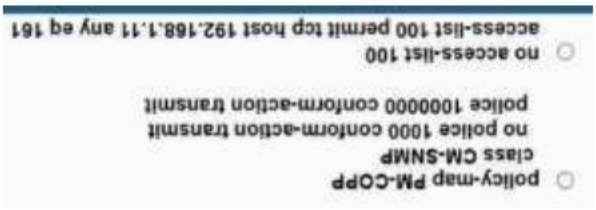
https://www.cisco.com/c/en/us/support/docs/security-vpn/terminal-access-controller-access-control-system-tacacs-/200467-Troubleshoot-TACACS-Authentication-Issue.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/17-8/configuration_guide/sec/b_178_sec_9500_cg/configuring_tacacs_.pdf

Question 53

Single choice · Domain: D3 Infrastructure Security › 3.3 Troubleshoot control plane policing (CoPP) (Telnet, SSH, HTTP(S), SNMP, EIGRP, OSPF, BGP)

The network administrator configured CoPP so that all SNMP traffic from Cisco Prime located at 192.168.1.11 toward the router CPU is limited to 1000 kbps. Any traffic that exceeds this limit must be dropped. access-list 100 permit udp any any eq 161 class-map CM-SNMP match access-group 100 policy-map PM-COPP class CM-SNMP police 1000 conform-action transmit ! control-plane service-policy input PM-COPP The network administrator is not getting the desired result for the SNMP traffic and SNMP traffic is getting dropped frequently. Which set of configurations resolves the issue?





A. Option A
B. Option B
C. Option C
D. Option D

Answer: A. Option A

Explanation

Option A corrects both causes of the problem: the ACL matches SNMP traffic from Cisco Prime at 192.168.1.11, and the policer is set to 1,000,000 bits per second, equivalent to 1000 kbps. The original ACL

matches SNMP from any source, and `police 1000` does not specify a rate of 1000 kbps. With the corrected rate and the existing conform action, traffic within the limit is transmitted and excess traffic is dropped by the policer's default behavior.

https://www.cisco.com/c/en/us/about/security-center/copp-best-practices.html
https://www.cisco.com/c/en/us/td/docs/ios/qos/command/reference/qos_book/qos_n1.html
https://www.cisco.com/en/US/docs/ios-xml/ios/qos_plcshp/configuration/xe-3s/qos-plcshp-ctrl-pln-plc.pdf

✓ **A** Replace the ACL with a permit for UDP from host 192.168.1.11 to destination port 161, and replace the 1000 policer rate with 1000000, retaining conform-action transmit.
This matches the specified SNMP source and sets the rate to 1,000,000 bits per second (1000 kbps). Traffic that conforms is transmitted, while traffic exceeding the policing rate is dropped by default.

https://www.cisco.com/c/en/us/about/security-center/copp-best-practices.html
https://www.cisco.com/c/en/us/td/docs/ios/qos/command/reference/qos_book/qos_n1.html
https://www.cisco.com/en/US/docs/ios-xml/ios/qos_plcshp/configuration/xe-3s/qos-plcshp-ctrl-pln-plc.pdf

✗ **B** Use the corrected host-specific ACL and 1000000 policer rate, remove the control-plane service policy, and apply the policy to interfaces E0/0 and E0/1.
Although the ACL and rate address the matching and policing-rate issues, moving the policy to interface input does not apply it as the configured control-plane policy for traffic headed to the router CPU. The scenario requires policing that traffic at the control plane.

https://www.cisco.com/c/en/us/about/security-center/copp-best-practices.html
https://www.cisco.com/en/US/docs/ios-xml/ios/qos_plcshp/configuration/xe-3s/qos-plcshp-ctrl-pln-plc.pdf

✗ **C** Option C.
The provided option text identifies this choice only as "Option C," and the image does not show enough legible configuration detail to establish that it both matches the specified source and applies the required 1000-kbps control-plane policing. Option A is the indicated configuration that resolves the reported issue.

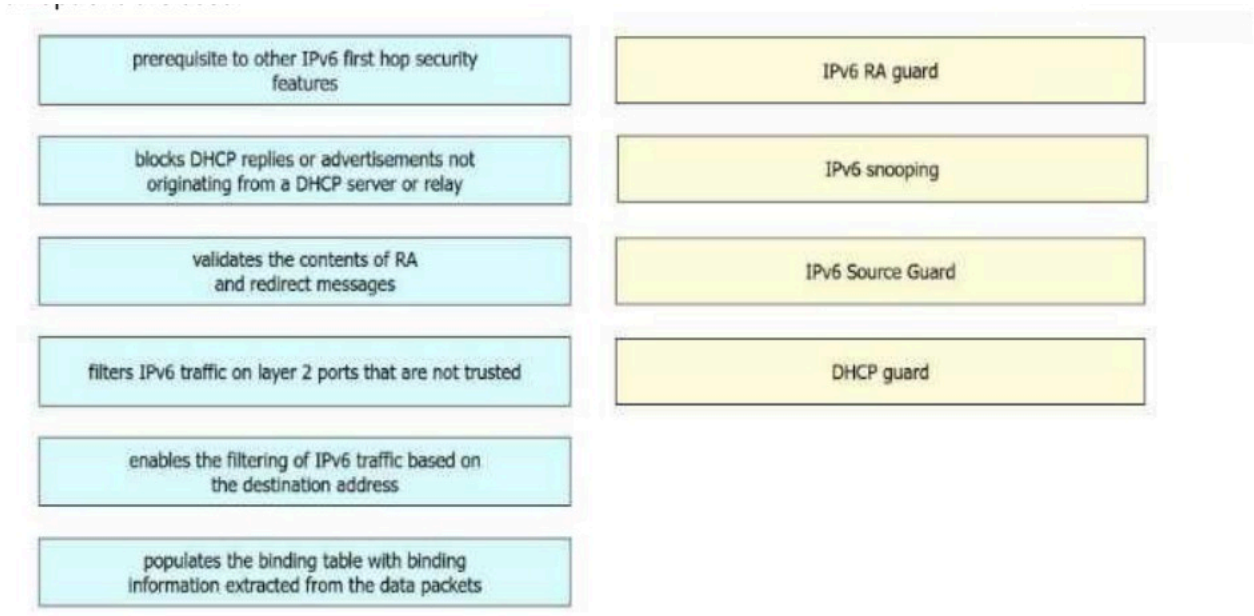
✗ **D** Option D.
The provided option text identifies this choice only as "Option D," and the image does not show enough legible configuration detail to establish that it corrects the source match and policing rate while retaining the control-plane attachment. Option A is the indicated configuration that meets those requirements.

Dumpsgate

Question 54

Drag and drop · Domain: D3 Infrastructure Security › 3.4 Describe IPv6 First Hop security features (RA guard, DHCP guard, binding table, ND inspection/snooping, source guard)

Drag and drop the descriptions from the left onto the IPv6 first hop security features on the right. Not all options are used.



Items:

- A. prerequisite to other IPv6 first hop security features
- B. blocks DHCP replies or advertisements not originating from a DHCP server or relay
- C. validates the contents of RA and redirect messages
- D. filters IPv6 traffic on layer 2 ports that are not trusted
- E. enables the filtering of IPv6 traffic based on the destination address
- F. populates the binding table with binding information extracted from the data packets

Place them on:

1. IPv6 RA guard
2. IPv6 snooping
3. IPv6 Source Guard
4. DHCP guard

Answer

- IPv6 RA guard → **C. validates the contents of RA and redirect messages**
- IPv6 snooping → **F. populates the binding table with binding information extracted from the data packets**
- IPv6 Source Guard → **D. filters IPv6 traffic on layer 2 ports that are not trusted**
- DHCP guard → **B. blocks DHCP replies or advertisements not originating from a DHCP server or relay**

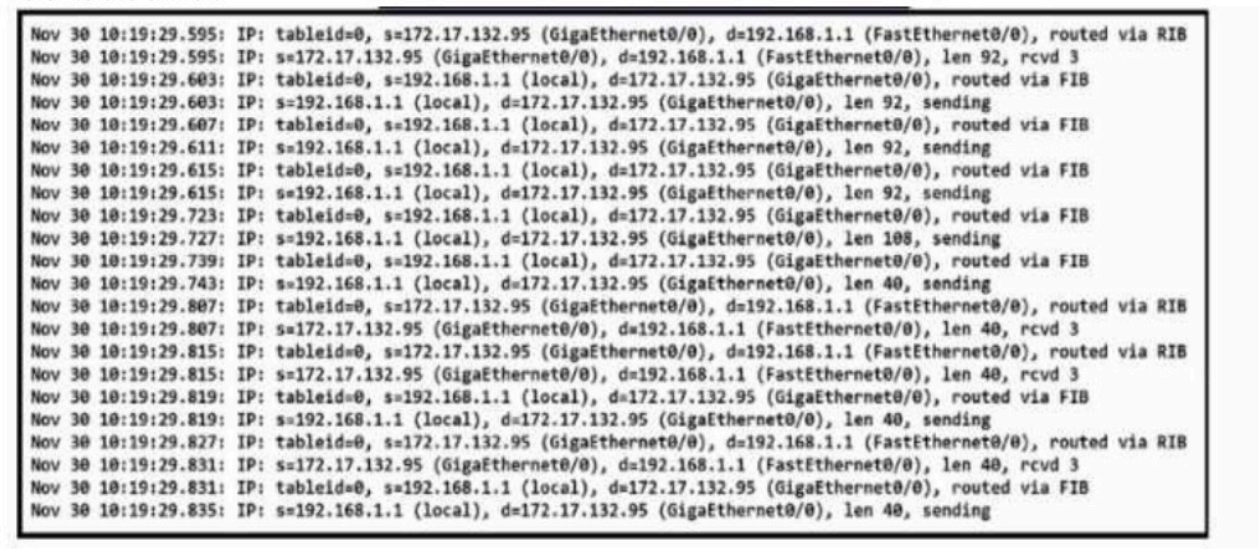
Free demo · full exam: 636 questions

Dumpsgate

Question 55

Single choice · Domain: D3 Infrastructure Security › 3.3 Troubleshoot control plane policing (CoPP) (Telnet, SSH, HTTP(S), SNMP, EIGRP, OSPF, BGP)

Refer to the exhibit.



```
Nov 30 10:19:29.595: IP: tableid=0, s=172.17.132.95 (GigaEthernet0/0), d=192.168.1.1 (FastEthernet0/0), routed via RIB
Nov 30 10:19:29.595: IP: s=172.17.132.95 (GigaEthernet0/0), d=192.168.1.1 (FastEthernet0/0), len 92, rcvd 3
Nov 30 10:19:29.603: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), routed via FIB
Nov 30 10:19:29.603: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), len 92, sending
Nov 30 10:19:29.607: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), routed via FIB
Nov 30 10:19:29.611: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), len 92, sending
Nov 30 10:19:29.615: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), routed via FIB
Nov 30 10:19:29.615: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), len 92, sending
Nov 30 10:19:29.723: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), routed via FIB
Nov 30 10:19:29.727: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), len 108, sending
Nov 30 10:19:29.739: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), routed via FIB
Nov 30 10:19:29.743: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), len 40, sending
Nov 30 10:19:29.807: IP: tableid=0, s=172.17.132.95 (GigaEthernet0/0), d=192.168.1.1 (FastEthernet0/0), routed via RIB
Nov 30 10:19:29.807: IP: s=172.17.132.95 (GigaEthernet0/0), d=192.168.1.1 (FastEthernet0/0), len 40, rcvd 3
Nov 30 10:19:29.815: IP: tableid=0, s=172.17.132.95 (GigaEthernet0/0), d=192.168.1.1 (FastEthernet0/0), routed via RIB
Nov 30 10:19:29.815: IP: s=172.17.132.95 (GigaEthernet0/0), d=192.168.1.1 (FastEthernet0/0), len 40, rcvd 3
Nov 30 10:19:29.819: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), routed via FIB
Nov 30 10:19:29.819: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), len 40, sending
Nov 30 10:19:29.827: IP: tableid=0, s=172.17.132.95 (GigaEthernet0/0), d=192.168.1.1 (FastEthernet0/0), routed via RIB
Nov 30 10:19:29.831: IP: s=172.17.132.95 (GigaEthernet0/0), d=192.168.1.1 (FastEthernet0/0), len 40, rcvd 3
Nov 30 10:19:29.831: IP: tableid=0, s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), routed via FIB
Nov 30 10:19:29.835: IP: s=192.168.1.1 (local), d=172.17.132.95 (GigaEthernet0/0), len 40, sending
```

Refer to the exhibit. Users are experiencing slow response times from critical application servers connected to the network. The engineer determined from the console log that there is a DOS attack on the interface GigabitEthernet0/0. Which action must the engineer take to resolve the issue?

A. Configure an access group on GigabitEthernet0/0 to block traffic from 192.168.1.1.
B. Configure an access group on GigabitEthernet0/0 to block traffic from 172.17.132.95.
C. Restrict CoPP traffic on the management interface from 172.17.132.95 and 192.168.1.1.
D. Configure QoS on FastEthernet0/0 to restrict bandwidth usage for 172.17.132.95 and 192.168.1.1.

Answer: B. Configure an access group on GigabitEthernet0/0 to block traffic from 172.17.132.95.

Explanation

The log repeatedly shows traffic sourced from 172.17.132.95 and destined for 192.168.1.1 arriving on GigabitEthernet0/0. Because the attack is entering on that interface, the engineer should apply an access group there to block the identified source, 172.17.132.95. This directly filters the offending traffic before it continues toward the network.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-create-ip-apply-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_acl/configuration/15-mt/sec-data-acl-15-mt-book/sec-create-ip-apply.html

✗ **A** Configure an access group on GigabitEthernet0/0 to block traffic from 192.168.1.1.
The log identifies 192.168.1.1 as the destination/local address, while 172.17.132.95 is the source of the traffic arriving on GigabitEthernet0/0. Blocking the destination does not target the identified attacking source.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-create-ip-apply-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_acl/configuration/15-mt/sec-data-acl-15-mt-book/sec-

create-ip-apply.html

✓ **B** Configure an access group on GigabitEthernet0/0 to block traffic from 172.17.132.95.
The console log shows repeated packets sourced from 172.17.132.95 entering GigabitEthernet0/0 and destined for 192.168.1.1. An access group on the affected interface that blocks this source addresses the DoS traffic at its point of entry.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-create-ip-apply-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_acl/configuration/15-mt/sec-data-acl-15-mt-book/sec-create-ip-apply.html

✗ **C** Restrict CoPP traffic on the management interface from 172.17.132.95 and 192.168.1.1.
The exhibit identifies the traffic on GigabitEthernet0/0, not a management interface, and the required action is to filter the attacking source on that affected interface. CoPP is not the interface-specific source block described by the log.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_plcshp/configuration/xe-3s/qos-plcshp-xe-3s-book/qos-plcshp-ctrl-pln-plc.pdf
https://www.cisco.com/c/en/us/td/docs/dcn/aci/apic/6x/security-configuration/cisco-apic-security-configuration-guide-62x/q_control_plane_traffic/information-about-copp.html
https://www.cisco.com/c/en/us/td/docs/ios/12_2sb/feature/guide/cpp.pdf
https://www.cisco.com/c/en/us/td/docs/ios/qos/configuration/guide/mgmt_plane_prot.pdf

✗ **D** Configure QoS on FastEthernet0/0 to restrict bandwidth usage for 172.17.132.95 and 192.168.1.1.
The attack is shown on GigabitEthernet0/0, not FastEthernet0/0, so configuring QoS on FastEthernet0/0 would not target the identified ingress path. The log also identifies 172.17.132.95 as the source to block rather than calling for bandwidth limits on both addresses.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/quality-of-service/quality-of-service/m_qos-plcshp-oview-0.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/quality-of-service/quality-of-service/m_qos-plcshp-pct-plc-0.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/qos/quality-of-service-configuration-guide/m-quality-of-service.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_plcshp/configuration/15-s/qos-plcshp-15-s-book.pdf
https://www.cisco.com/c/en/us/td/docs/routers/ios/config/17-x/qos/b-quality-of-service/m_qos-policing.html

Question 56

Single choice · Domain: D3 Infrastructure Security › 3.4 Describe IPv6 First Hop security features (RA guard, DHCP guard, binding table, ND inspection/snooping, source guard)

Which feature is required for IPv6 Source Guard to block traffic arriving on the server interface from unknown sources?

A. Dynamic ARP Inspection
B. IPv6 RA Guard
C. IPv6 ND Inspection
D. DHCPv6 Guard

Answer: C. IPv6 ND Inspection

Explanation

IPv6 Source Guard checks whether packets arriving on an interface use a source IPv6 address allowed for that interface. IPv6 ND Inspection provides the neighbor-discovery information used to validate those addresses, so it is the required feature for blocking traffic from unknown sources. The other choices

protect different protocols or functions.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_ipv6-src-guard-xe.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-16-8/ip6f-xe-16-8-book/ip6-src-guard.html

✗ **A** Dynamic ARP Inspection
Dynamic ARP Inspection validates IPv4 ARP traffic, not IPv6 source addresses. It therefore does not provide the IPv6 neighbor information needed for this Source Guard check.

✗ **B** IPv6 RA Guard
IPv6 RA Guard filters unauthorized Router Advertisement messages. It does not validate the source IPv6 address of ordinary traffic arriving on the server interface.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/sec-crypto/fhs-sisf/fhs-and-sisf-configuration-guide/ipv6-first-hop-security.html

✓ **C** IPv6 ND Inspection
IPv6 ND Inspection learns and validates IPv6 neighbor information that Source Guard can use to determine which source addresses are legitimate on an interface. This lets Source Guard block packets arriving from unknown sources.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_ipv6-src-guard-xe.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-16-8/ip6f-xe-16-8-book/ip6-src-guard.html

✗ **D** DHCPv6 Guard
DHCPv6 Guard controls DHCPv6 messages, such as preventing unauthorized DHCPv6 servers. The required feature for validating IPv6 source addresses in this scenario is IPv6 ND Inspection.

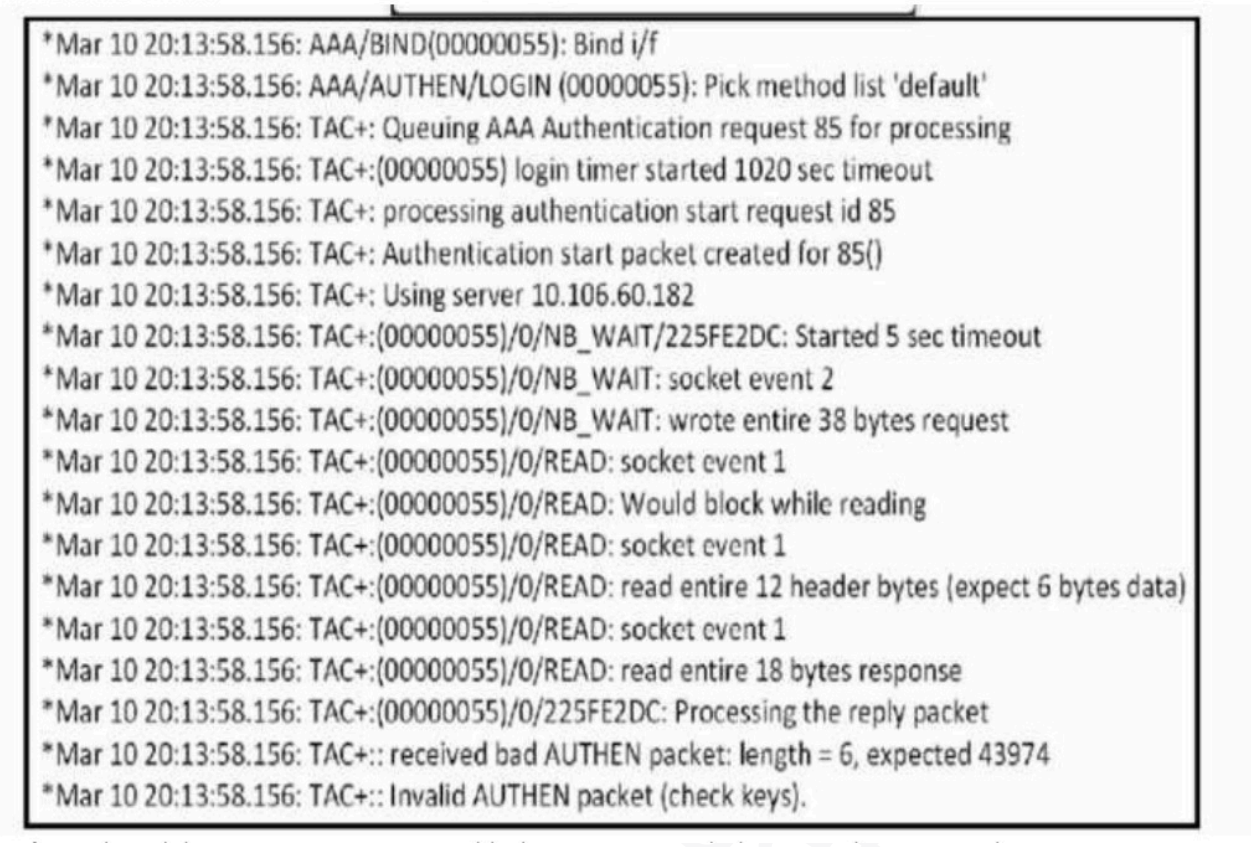
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/sec-crypto/fhs-sisf/fhs-and-sisf-configuration-guide/ipv6-first-hop-security.html

Dumpsgate

Question 57

Single choice · Domain: D3 Infrastructure Security › 3.1 Troubleshoot device security using IOS AAA (TACACS+, RADIUS, local database)

Refer to the exhibit.



Refer to the exhibit. An engineer must troubleshoot an issue with the aaa authentication that affected the user's login to router R1. Which command allows the configured user to authenticate?

A. aaa authentication login default group radius local
B. aaa authentication login default group radius tacacs+
C. aaa authentication login default group tacacs+
D. aaa authentication login default group radius

Answer: A. aaa authentication login default group radius local

Explanation

The exhibit shows that R1 selects the default login method list, sends the request to a TACACS+ server, and receives an invalid AUTHEN packet because the response length is unexpected. Configure RADIUS as the primary method and local authentication as the fallback: if the server-based method cannot authenticate the user, IOS can check the locally configured user credentials. This makes option A the command that allows the configured user to authenticate in this scenario.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-cfg-authentifcn-0.html

https://www.cisco.com/c/en/us/td/docs/ios/sec_user_services/configuration/guide/convert/aaa/sec_cfg_authentifcn.html
https://www.cisco.com/c/en/us/td/docs/ios/sec_user_services/configuration/guide/12_4/sec_securing_user_services_12-4_book/sec_cfg_authentifcn.html

✓ **A** aaa authentication login default group radius local
This uses RADIUS first and then falls back to the local user database when the remote method does not successfully authenticate the login. The exhibit indicates the remote authentication exchange is failing, so the local fallback can allow the configured user to log in.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-cfg-authentifcn-0.html
https://www.cisco.com/c/en/us/td/docs/ios/sec_user_services/configuration/guide/convert/aaa/sec_cfg_authentifcn.html
https://www.cisco.com/c/en/us/td/docs/ios/sec_user_services/configuration/guide/12_4/sec_securing_user_services_12-4_book/sec_cfg_authentifcn.html

✗ **B** aaa authentication login default group radius tacacs+
This configures RADIUS followed by TACACS+, but neither method checks the router's local user database. The exhibit already shows a failed TACACS+ authentication exchange, so this does not provide the needed local authentication fallback.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-cfg-authentifcn-0.html
https://www.cisco.com/c/en/us/td/docs/ios/sec_user_services/configuration/guide/convert/aaa/sec_cfg_authentifcn.html
https://www.cisco.com/c/en/us/td/docs/ios/sec_user_services/configuration/guide/12_4/sec_securing_user_services_12-4_book/sec_cfg_authentifcn.html

✗ **C** aaa authentication login default group tacacs+
This relies only on TACACS+ for login authentication. The exhibit shows the TACACS+ reply is invalid, and this command does not provide a local method to authenticate the configured user.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-cfg-authentifcn-0.html
https://www.cisco.com/c/en/us/td/docs/ios/sec_user_services/configuration/guide/convert/aaa/sec_cfg_authentifcn.html
https://www.cisco.com/c/en/us/td/docs/ios/sec_user_services/configuration/guide/12_4/sec_securing_user_services_12-4_book/sec_cfg_authentifcn.html

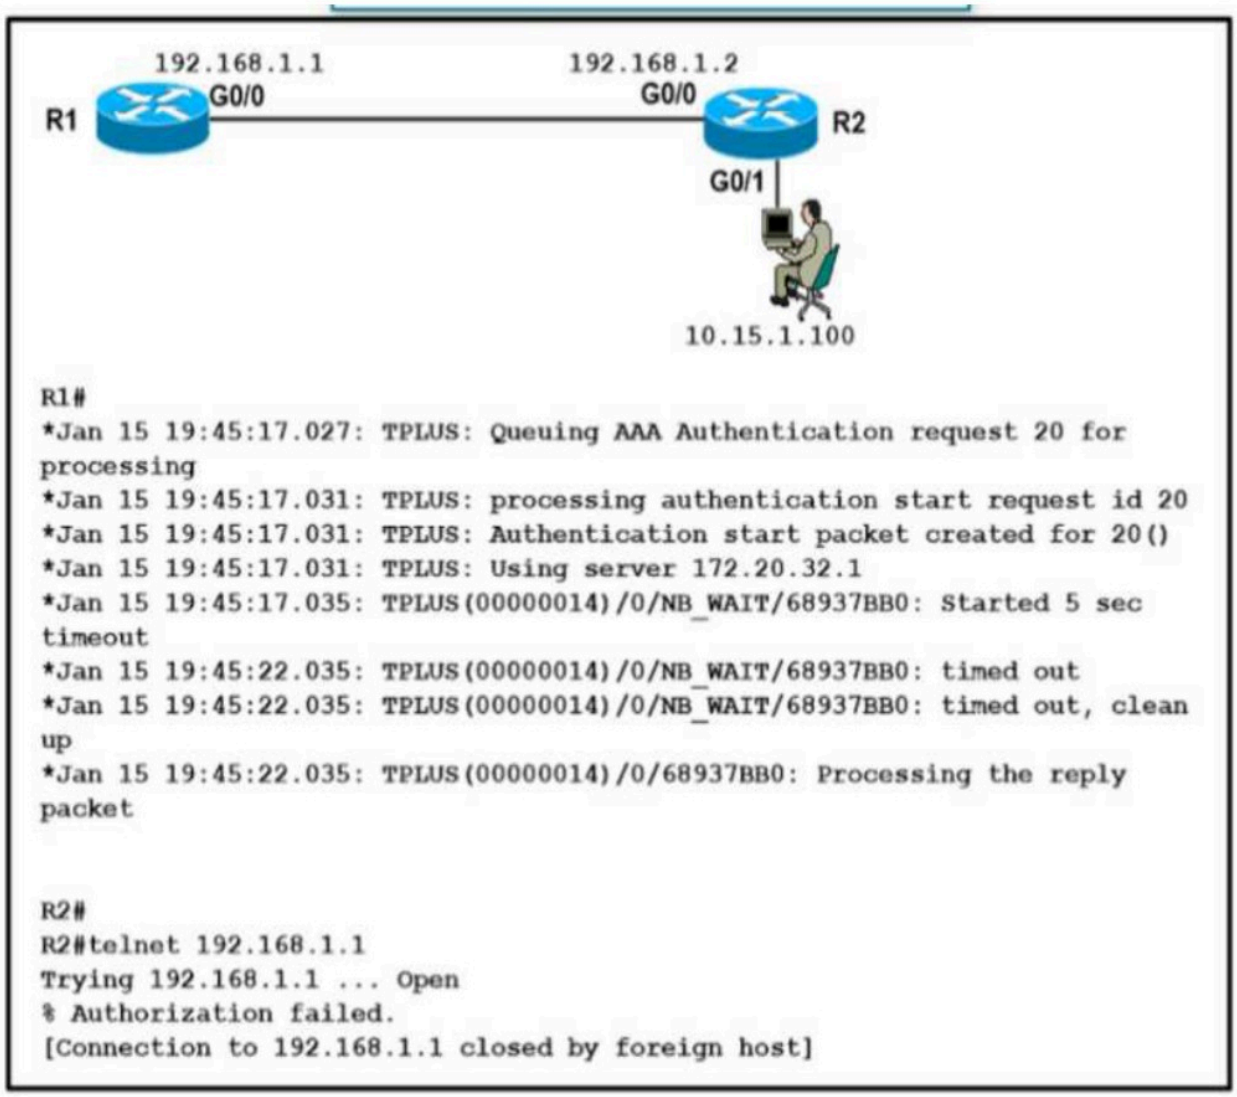
✗ **D** aaa authentication login default group radius
This relies only on RADIUS and has no local fallback. If the remote server method cannot authenticate the user, the configured local account will not be used by this method list.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-cfg-authentifcn-0.html
https://www.cisco.com/c/en/us/td/docs/ios/sec_user_services/configuration/guide/convert/aaa/sec_cfg_authentifcn.html
https://www.cisco.com/c/en/us/td/docs/ios/sec_user_services/configuration/guide/12_4/sec_securing_user_services_12-4_book/sec_cfg_authentifcn.html

Question 58

Single choice · Domain: D3 Infrastructure Security › 3.1 Troubleshoot device security using IOS AAA (TACACS+, RADIUS, local database)

Refer to the exhibit.



Refer to the exhibit A network engineer is troubleshooting an AAA authentication issue for R1 from R2 When an engineer tries to open a telnet connection to R1 it opens the connection but shows a %Authorization failed error message on the terminal and closes the connection silently Which action resolves the issue?

A. Resolve tacacs+ server host IP authentication miss configuration on the R1 router
B. Resolve tacacs+ server reachability from the R1 router.
C. Configure the tacacs+ server host IP on the R1 router
D. Configure authorization commands in the tacacs* server for the R1 router.

Answer: D. Configure authorization commands in the tacacs* server for the R1 router.

Dumpsgate

Question 59

Single choice · Domain: D3 Infrastructure Security › 3.3 Troubleshoot control plane policing (CoPP) (Telnet, SSH, HTTP(S), SNMP, EIGRP, OSPF, BGP)

The network administrator configured CoPP so that all SNMP traffic from Cisco Prime located at 192.168.1.11 toward the router CPU is limited to 1000 kbps. Any traffic that exceeds this limit must be dropped. access-list 100 permit udp any any eq 161 ! class-map CM-SNMP match access-group 100 ! policy-map PM-COPP class CM-SNMP police 1000 conform-action transmit ! control-plane service-policy input PM-COPP The network administrator is not getting the desired result for the SNMP traffic and SNMP traffic is getting dropped frequently. Which set of configurations resolves the issue?

A. no access-list 100access-list 100 permit tcp host 192.168.1.11 any eq 161

B. no access-list 100access-list 100 permit udp host 192.168.1.11 any eq 161!policy-map PM-COPPclass CM-SNMPno police 1000 conform-action transmitpolice 1000000 conform-action transmit! control-planeno service-policy input PM-COPP!interface E 0/0service-policy input PM- COPP!interface E 0/1service-policy input PM-COPP

C. no access-list 100access-list 100 permit udp host 192.168.1.11 any eq 161!policy-map PM-COPPclass CM-SNMPno police 1000 conform-action transmitpolice 1000000 conform-action transmit

D. policy-map PM-COPPclass CM-SNMPno police 1000 conform-action transmitpolice 1000000 conform-action transmit

Answer: C. no access-list 100access-list 100 permit udp host 192.168.1.11 any eq 161!policy-map PM- COPPclass CM-SNMPno police 1000 conform-action transmitpolice 1000000 conform-action transmit

Explanation

Option C corrects both issues in the configuration: it narrows the ACL to UDP SNMP traffic sourced from Cisco Prime at 192.168.1.11, and raises the policer rate from 1000 to 1,000,000 bps (1000 kbps). The original ACL matches UDP destination port 161 from any source, so unrelated SNMP traffic can consume the allowance. The original policer rate is also 1000 bps, not 1000 kbps, which explains the frequent drops. The existing control-plane policy attachment is retained, so the policer continues to apply to traffic destined for the router CPU.

https://www.cisco.com/c/en/us/td/docs/ios/qos/command/reference/qos_book/qos_n1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/command/qos-cr-book/qos-n1.html
https://www.cisco.com/c/en/us/td/docs/routers/ios/config/17-x/qos/b-quality-of-service/m_qos-plcshp-ctrl-pln-plc-0.pdf

✗ **A** Replace the ACL with `permit tcp host 192.168.1.11 any eq 161`.
SNMP traffic in this scenario uses UDP, so an ACL permitting TCP would not match the SNMP packets that need to be policed. It also leaves the policer's 1000 bps rate unchanged.

https://www.cisco.com/c/en/us/support/docs/ip/simple-network-management-protocol-snmp/118901-technote-snmp-00.html

✗ **B** Restrict the ACL to UDP from 192.168.1.11 and increase the policer rate, then remove the control-plane policy and apply the policy inbound on interfaces E0/0 and E0/1.
The ACL and rate changes address the traffic classification and overly low rate, but moving the policy to physical interfaces and removing it from the control plane does not meet the requirement to police SNMP traffic toward the router CPU. CoPP is applied to traffic destined for the control plane; the additional interface attachments are not the needed correction here.

https://www.cisco.com/c/en/us/about/security-center/copp-best-practices.html
https://www.cisco.com/c/en/us/td/docs/routers/ios/config/17-x/qos/b-quality-of-service/m_qos-plcshp-ctrl-pln-

Free demo · full exam: 636 questions

plc-0.pdf

✓ **C** Restrict the ACL to UDP from 192.168.1.11 and change the policer to 1,000,000 bps.
The ACL now matches SNMP packets from the specified Cisco Prime host rather than SNMP traffic from any source. The policer's rate is changed to 1,000,000 bps, equivalent to 1000 kbps, while the existing control-plane service-policy attachment remains in place.

https://www.cisco.com/c/en/us/td/docs/ios/qos/command/reference/qos_book/qos_n1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/command/qos-cr-book/qos-n1.html
https://www.cisco.com/c/en/us/td/docs/routers/ios/config/17-x/qos/b-quality-of-service/m_qos-plcshp-ctrl-pln-plc-0.pdf

✗ **D** Change the policer to 1,000,000 bps but leave ACL 100 unchanged.
This fixes the rate mismatch, but the ACL still matches UDP destination-port-161 traffic from any source. As a result, other SNMP senders can still share or consume the configured policing allowance instead of limiting the match to Cisco Prime at 192.168.1.11.

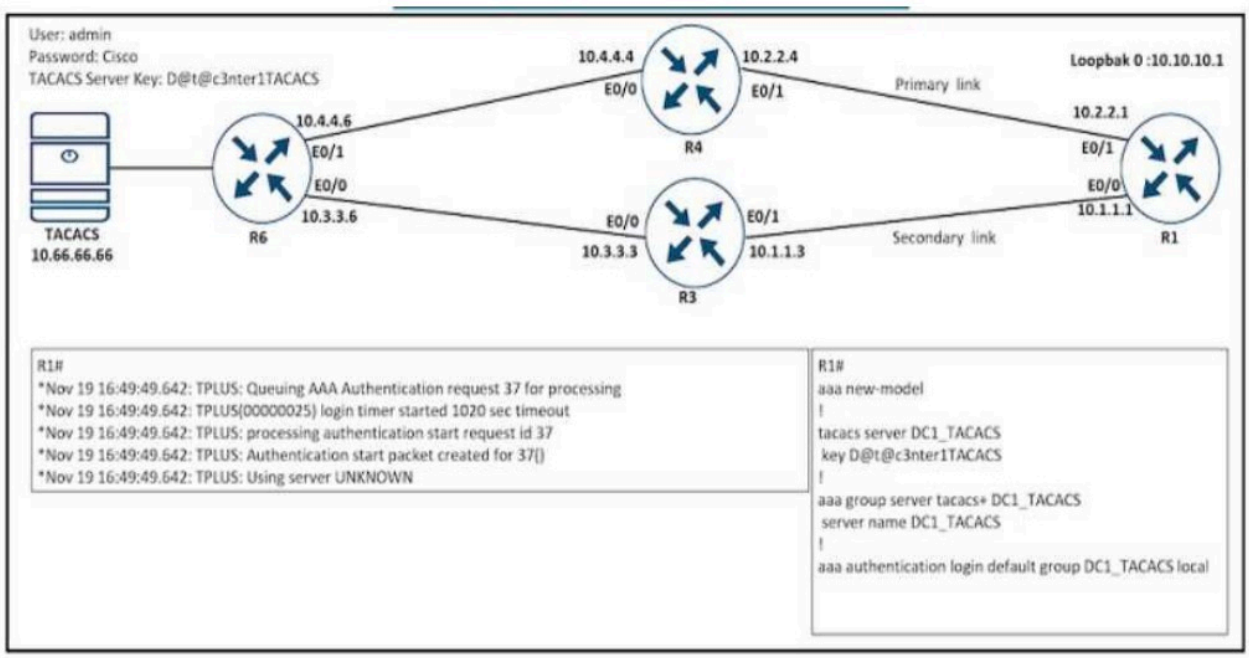
https://www.cisco.com/c/en/us/td/docs/ios/qos/command/reference/qos_book/qos_n1.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/command/qos-cr-book/qos-n1.html

Question 60

Single choice · Domain: D3 Infrastructure Security › 3.1 Troubleshoot device security using IOS AAA (TACACS+, RADIUS, local database)

Refer to the exhibit.

Dumpsgate



- aaa group server tacacs+ DC_TACACS
  server name DC_TACACS

- tacacs server DC1_TACACS
  address ipv4 10.66.66.66
  key D@t@c3nter1TACACS

- aaa group server tacacs+ DC1_TACACS
  server name DC_TACACS

- tacacs server DC1_TACACS
  address ipv4 10.60.66.66
  key D@t@c3nter1TACACS

Refer to the exhibit R1 cannot authenticate via TACACS Which configuration resolves the issue?

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B. Option B

Dumpsgate

Question 61

Single choice · Domain: D3 Infrastructure Security › 3.2.b IPv6 traffic filter

Refer to the exhibit. An engineer is trying to log in to R1 via R3 loopback address. Which action resolves the issue?

A. Add transport input SCP
B. Add transport input none
C. Remove the IPv6 traffic filter from R1, which is blocking the Telnet.
D. Remove the IPv6 traffic from R1, which is blocking the SSH

Answer: C. Remove the IPv6 traffic filter from R1, which is blocking the Telnet.

Explanation

The issue is the IPv6 traffic filter on R1: it is blocking the Telnet connection needed for the login attempt through R3's loopback address. Removing that filter resolves the stated blockage. Changing the allowed transport protocol on the VTY lines would not remove the IPv6 filtering that is preventing the connection.

https://www.cisco.com/c/en/us/support/docs/ip/ip-version-6/113126-ipv6-acl-00.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_nman/configuration/xe-3se/5700/ip6-telnet-acc.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ip6-15-2mt-book/ip6-telnet-acc.html

✗ **A** Add transport input SCP
This changes which protocols the VTY lines accept; it does not remove the IPv6 traffic filter blocking the connection. SCP is also not the Telnet access required by the scenario.

https://www.cisco.com/c/en/us/td/docs/ios/termserv/command/reference/tsv_book.pdf
https://www.cisco.com/c/en/us/td/docs/ios/termserv/command/reference/tsv_book/tsv_s1.html

✗ **B** Add transport input none
This would disable incoming access through the VTY lines rather than restore the blocked connection. The problem is the IPv6 filter on R1, not an absence of a transport-input setting.

https://www.cisco.com/c/en/us/td/docs/ios/termserv/command/reference/tsv_book.pdf
https://www.cisco.com/c/en/us/td/docs/ios/termserv/command/reference/tsv_book/tsv_s1.html

✓ **C** Remove the IPv6 traffic filter from R1, which is blocking the Telnet.
The filter is preventing the Telnet traffic needed to log in to R1 via R3's loopback address. Removing the filter addresses that specific blockage.

https://www.cisco.com/c/en/us/support/docs/ip/ip-version-6/113126-ipv6-acl-00.pdf
https://www.cisco.com/c/en/us/support/docs/ip/ip-version-6/113126-ipv6-acl-00.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_nman/configuration/xe-3se/5700/ip6-telnet-acc.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ip6-15-2mt-book/ip6-telnet-acc.html

✗ **D** Remove the IPv6 traffic from R1, which is blocking the SSH
This option refers to SSH, whereas the stated issue and the correct resolution concern an IPv6 filter blocking Telnet. The action that addresses the described problem is to remove the filter, not to remove IPv6 traffic generally.

https://www.cisco.com/c/en/us/support/docs/ip/ip-version-6/113126-ipv6-acl-00.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_nman/configuration/xe-3se/5700/ip6-telnet-acc.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ip6-15-2mt-book/ip6-telnet-acc.html

Free demo · full exam: 636 questions

Question 62

Single choice · Domain: D3 Infrastructure Security › 3.3 Troubleshoot control plane policing (CoPP) (Telnet, SSH, HTTP(S), SNMP, EIGRP, OSPF, BGP)

Refer to the exhibit.

```
ip access-list extended CoPP-ICMP
 permit icmp any any echo
!
ip access-list extended CoPP-BGP
 permit tcp any eq bgp any established
!
ip access-list extended CoPP-EIGRP
 permit eigrp any host 224.0.0.10
!
Class-map match-all CoPP-CLASS
 match access-group name CoPP-ICMP
 match access-group name CoPP-BGP
 match access-group name CoPP-EIGRP
!
```

Refer to the exhibit A CoPP policy is implemented to allow specific control traffic, but the traffic is not matching as expected and is getting unexpected behavior of control traffic. Which action resolves the issue?

A. Use match-any instruction in class-map
B. Create a separate class map against each ACL.
C. Create a separate class map lor ICMP traffic.
D. Use default-class to match ICMP traffic

Answer: A. Use match-any instruction in class-map

Explanation

Free demo · full exam: 636 questions

The class map is configured with `match-all` and lists three ACLs for different control protocols: ICMP, BGP, and EIGRP. With `match-all`, traffic must satisfy every listed ACL match, so the intended protocol-specific traffic will not match the class as expected. Changing the class map to `match-any` lets traffic match when it satisfies any one of the ACLs, resolving the mismatch. [2][3]

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus7000/sw/security/config/cisco_nexus7000_security_config_guide_8x/configuring_control_plane_policing.pdf
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x_chapter_010010.pdf

✓ **A** Use match-any instruction in class-map
The class map currently uses `match-all` for three ACLs that describe different kinds of traffic. Using `match-any` makes the class match traffic permitted by any one of those ACLs, which fits the intended CoPP policy.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus7000/sw/security/config/cisco_nexus7000_security_config_guide_8x/configuring_control_plane_policing.pdf
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x_chapter_010010.pdf

✗ **B** Create a separate class map against each ACL.
Separate class maps could classify each ACL independently, but the problem is that the existing class map requires all its ACL matches at once. Changing it to `match-any` directly fixes that issue without requiring separate class maps.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus7000/sw/security/config/cisco_nexus7000_security_config_guide_8x/configuring_control_plane_policing.pdf
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x_chapter_010010.pdf

✗ **C** Create a separate class map lor ICMP traffic.
The issue is not limited to ICMP: the class map includes ACL matches for ICMP, BGP, and EIGRP under `match-all`. A separate ICMP class map would not correct the overly restrictive matching behavior for the combined class.

✗ **D** Use default-class to match ICMP traffic
The exhibit defines an explicit ACL for ICMP echo traffic, so the issue is how the class map combines its ACL matches, not a need to classify ICMP as default traffic. `match-any` allows the ICMP ACL, as well as the other listed ACLs, to match independently.

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus7000/sw/security/config/cisco_nexus7000_security_config_guide_8x/configuring_control_plane_policing.pdf
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x_chapter_010010.pdf

Question 63

Single choice · Domain: D3 Infrastructure Security › 3.2.a IPv4 access control lists (standard, extended, time-based)

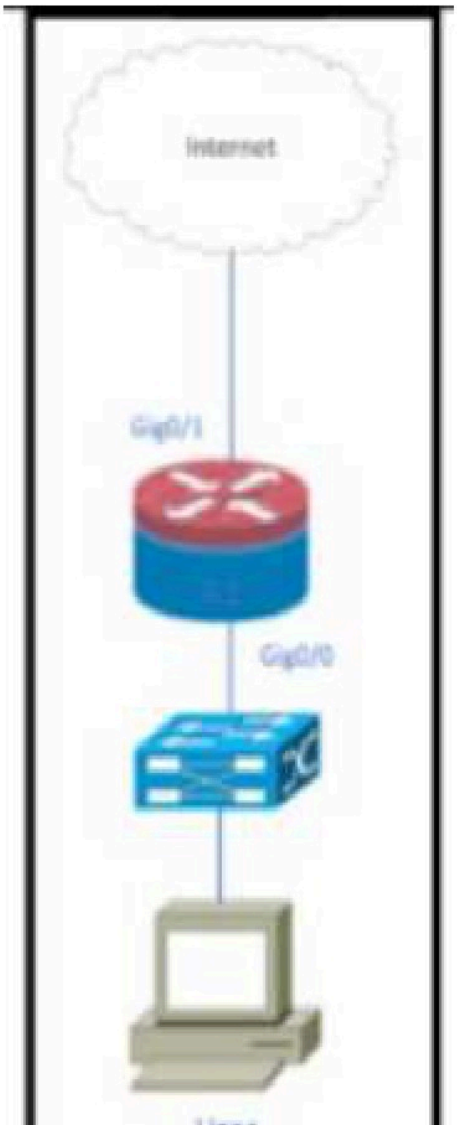
Refer to the exhibit.

```
R1#show time-range

time-range entry: timer (active)
   periodic weekend 9:00 to 17:00
   used in: IP ACL entry
   used in: IP ACL entry

R1#show ip access-list interface gig0/0

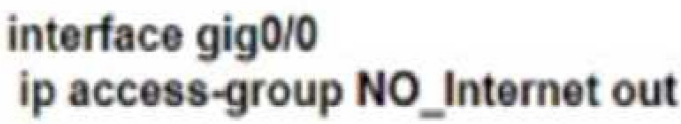
Extended IP access list NO_Internet in
   10 deny tcp any any eq www time-range timer (active)
   20 deny tcp any any eq 443 time-range timer (active)
   30 permit ip any any
```



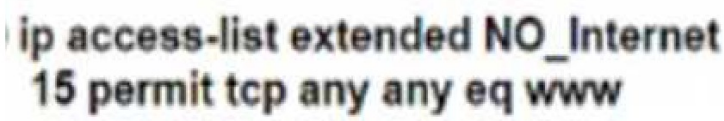


Refer to the exhibit. Users on a call center report that they cannot browse the internet on Saturdays during the afternoon. Which configuration resolves the issue?

A.



B.



C.



D.



Answer: D

Explanation

The ACL denies HTTP and HTTPS only while the referenced time range, `timer`, is active. That range currently includes Saturdays from 9:00 to 17:00, which covers the reported afternoon outage. Option D removes that interval and schedules the range for Saturdays from 17:00 to 23:59 instead, so the ACL's time-based denies no longer block browsing during the afternoon.

https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/time-range.htm
https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/periodic.htm

✗ **A** Apply `NO_Internet` outbound on interface gig0/0.
The exhibit shows `NO_Internet` applied inbound on gig0/0, where the user traffic enters the router. Changing the direction does not correct the Saturday time range causing the reported outage, and it may not filter the user traffic as intended.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/security/acls/acls-configuration-guide/access-control-lists.html

✗ **B** Add `permit tcp any any eq www` as sequence 15 in `NO_Internet`.
Sequence 15 comes after the existing sequence 10 deny for HTTP, so HTTP traffic matching that deny is rejected before this permit is reached. It also does not change the schedule that blocks traffic on Saturdays from 9:00 to 17:00.

https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/security/acls/acls-configuration-guide/access-control-lists.html

✗ **C** Remove the `timer` time-range definition.
The problem is that the existing time range includes Saturday afternoon; simply removing its definition does not configure the required replacement schedule. Option D explicitly changes the active weekend interval so the time-based denies do not apply during the reported afternoon hours.

https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/time-range.htm

✓ **D** Remove `periodic weekend 9:00 to 17:00` and add `periodic weekend 17:00 to 23:59` to time range `timer`.
The ACL denies HTTP and HTTPS while `timer` is active, and the current weekend interval includes the afternoon period when users report the outage. Replacing that interval with 17:00–23:59 means the time-based denies are not active during Saturday afternoon, resolving the reported browsing problem.

https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/time-range.htm
https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/periodic.htm

Question 64

Single choice · Domain: D3 Infrastructure Security › 3.4 Describe IPv6 First Hop security features (RA guard, DHCP guard, binding table, ND inspection/snooping, source guard)

What is a characteristic of IPv6 RA Guard?

A. RA messages are allowed from the host port to the switch
B. It is unable to protect tunneled traffic
C. It filters rogue RA broadcasts from connected hosts
D. It is supported on the egress direction of the switch

Answer: C. It filters rogue RA broadcasts from connected hosts

Explanation

IPv6 RA Guard helps prevent rogue Router Advertisement messages from connected hosts from being accepted on the network. Therefore, the characteristic that fits is filtering rogue RA broadcasts from connected hosts.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_ip6-ra-guard.html
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus3000/sw/security/7x/b_Cisco_n3k_Security_Config_7x/b_Cisco_n3k_Security_Config_7x_chapter_01111.pdf

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960l/software/15-2_7_e/configuration_guide/b_1527e_consolidated_2960l_cg/m_sec_ipv6_ra_guard_2960l.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-3s/ip6f-xe-3s-book/ip6-ra-guard.html

✗ **A** RA messages are allowed from the host port to the switch
This describes allowing RA messages from a host-facing port, which is contrary to RA Guard's purpose of restricting unauthorized Router Advertisement messages. Trusted router-facing ports may be configured to allow legitimate RAs, but this is not the stated characteristic.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_ip6-ra-guard.html
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus3000/sw/security/7x/b_Cisco_n3k_Security_Config_7x/b_Cisco_n3k_Security_Config_7x_chapter_01111.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960l/software/15-2_7_e/configuration_guide/b_1527e_consolidated_2960l_cg/m_sec_ipv6_ra_guard_2960l.pdf

✗ **B** It is unable to protect tunneled traffic
This is not the characteristic identified by the correct answer. RA Guard's defining function in this question is filtering rogue RAs from connected hosts; the cited material does not establish this statement as the applicable characteristic.

✓ **C** It filters rogue RA broadcasts from connected hosts
RA Guard is used to prevent unauthorized hosts from sending Router Advertisement messages that could mislead IPv6 clients. Filtering rogue RAs from connected hosts directly describes that protective function.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_ip6-ra-guard.html
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus3000/sw/security/7x/b_Cisco_n3k_Security_Config_7x/b_Cisco_n3k_Security_Config_7x_chapter_01111.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960l/software/15-2_7_e/configuration_guide/b_1527e_consolidated_2960l_cg/m_sec_ipv6_ra_guard_2960l.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xe-3s/ip6f-xe-3s-book/ip6-ra-guard.html

✗ **D** It is supported on the egress direction of the switch
This does not describe the characteristic being tested. RA Guard checks Router Advertisement traffic arriving from connected devices, rather than being characterized as an egress-direction feature.

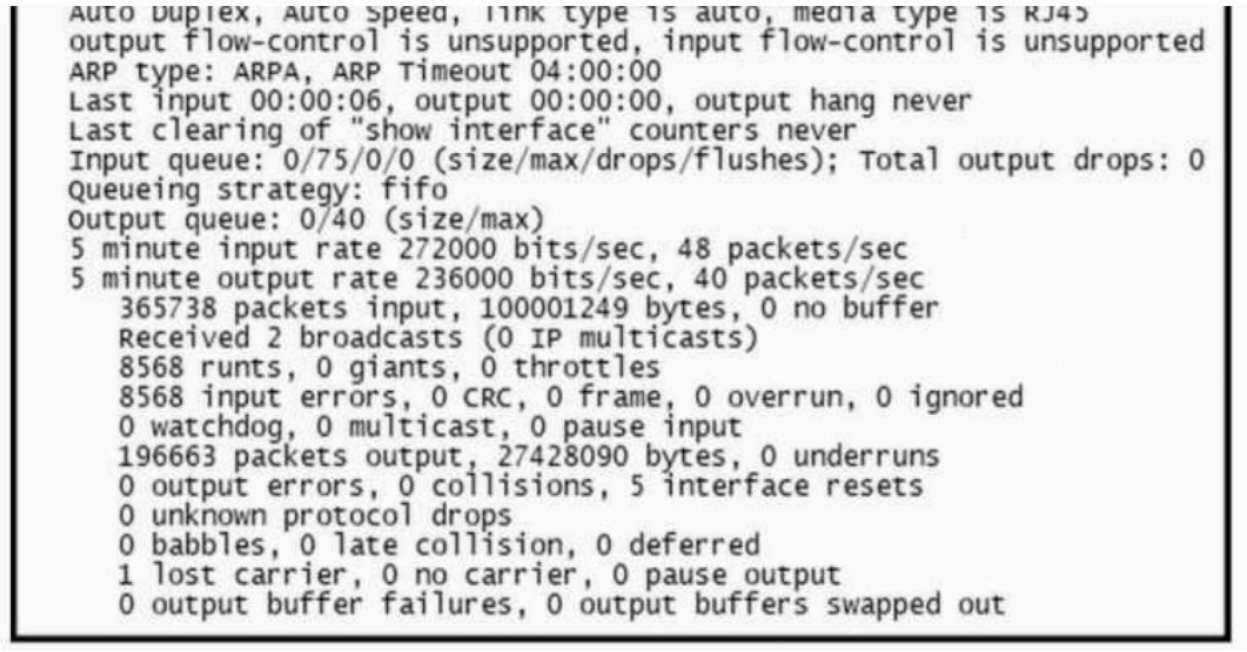
https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus3000/sw/security/7x/b_Cisco_n3k_Security_Config_7x/b_Cisco_n3k_Security_Config_7x_chapter_01111.pdf
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960l/software/15-2_7_e/configuration_guide/b_1527e_consolidated_2960l_cg/m_sec_ipv6_ra_guard_2960l.pdf

Dumpsgate

Question 65

Single choice · Domain: D4 Infrastructure Services

Refer to the exhibit.



Refer to the exhibit. Which action resolves the issue?

A. Implement traffic policing to prevent the interface input traffic being exceeded.
B. Replace the faulty cable.
C. Upgrade the hardware to increase the interface input buffers.
D. Resolve the misconfigured QoS parameters.

Answer: B. Replace the faulty cable.

Explanation

The exhibit reports 8,568 runts and input errors, along with one lost carrier and five interface resets. Those counters point to a physical-layer problem, such as a faulty cable, rather than input congestion, insufficient buffers, or QoS misconfiguration. Replacing the cable addresses the indicated issue. [7]

https://www.cisco.com/en/US/docs/internetworking/troubleshooting/guide/tr1915.html

✗ **A** Implement traffic policing to prevent the interface input traffic being exceeded.
The exhibit does not show input queue drops or throttles that would indicate traffic exceeding the interface's capacity. The reported runts and lost-carrier event instead suggest a physical connection problem, which policing would not resolve.

https://www.cisco.com/c/en/us/td/docs/routers/ios/config/17-x/qos/b-quality-of-service/m_qos-policing.html

✓ **B** Replace the faulty cable.
The interface has many runts and input errors, as well as a lost-carrier event and multiple resets. These indicators are consistent with a faulty physical connection, so replacing the cable addresses the likely cause.

https://www.cisco.com/en/US/docs/internetworking/troubleshooting/guide/tr1915.html
https://www.cisco.com/c/en/us/support/docs/switches/catalyst-9500-series-switches/218397-troubleshoot-port-flaps-on-catalyst-9000.html

✗ **C** Upgrade the hardware to increase the interface input buffers.
The exhibit reports zero overruns and zero no-buffer errors, so it does not indicate that input buffers are being exhausted. Increasing buffer capacity would not address the runts or lost-carrier event.

https://www.cisco.com/c/en/us/support/docs/routers/asr-1000-series-ios-xe-sd-wan/224129-troubleshoot-interface-packet-drops-in.html

✗ **D** Resolve the misconfigured QoS parameters.
The output queue has no drops, and the exhibit does not identify QoS-related drops or congestion. The runts and lost carrier point instead to a physical-layer issue.

Question 66

Single choice · Domain: D4 Infrastructure Services › 4.1.b Telnet, HTTP, HTTPS, SSH, and SCP

Refer to the exhibit.

```
CPE# telnet 198.51.100.64
% telnet connections not permitted from this terminal
CPE# ssh -l admin 198.51.100.64
% ssh connections not permitted from this terminal
CPE#
```

Refer to the exhibit. After a security hardening was performed on a router, the administrator cannot access the command line of any remote device. Which action resolves the issue?

A. Remove the transport output none command from the line con 0 section.
B. Remove the egress ACL blocking Telnet and SSH on the router interfaces.
C. Move Telnet and SSH commands to a nonzero privilege level.
D. Modify the AAA policy to allow the user to run the Telnet and SSH commands.

Answer: A. Remove the transport output none command from the line con 0 section.

Explanation

The exhibit shows that both outbound Telnet and SSH attempts are rejected with “connections not permitted from this terminal.” The `transport output none` setting on the console line prevents outbound connections initiated from that terminal. Removing it allows the administrator to use the router to access remote devices. [2]

https://www.cisco.com/c/en/us/td/docs/ios/termserv/command/reference/tsv_book.pdf

✓ **A** Remove the transport output none command from the line con 0 section.
`transport output none` disables outbound connection protocols from the console line, matching the error shown for both Telnet and SSH. Removing the restriction allows the administrator to initiate those connections from the router. [2]

https://www.cisco.com/c/en/us/td/docs/ios/termserv/command/reference/tsv_book.pdf

✗ **B** Remove the egress ACL blocking Telnet and SSH on the router interfaces.
The exhibit identifies a terminal-level restriction, not an interface ACL denying traffic. Changing an egress ACL would not remove the console line's restriction on outbound connections.

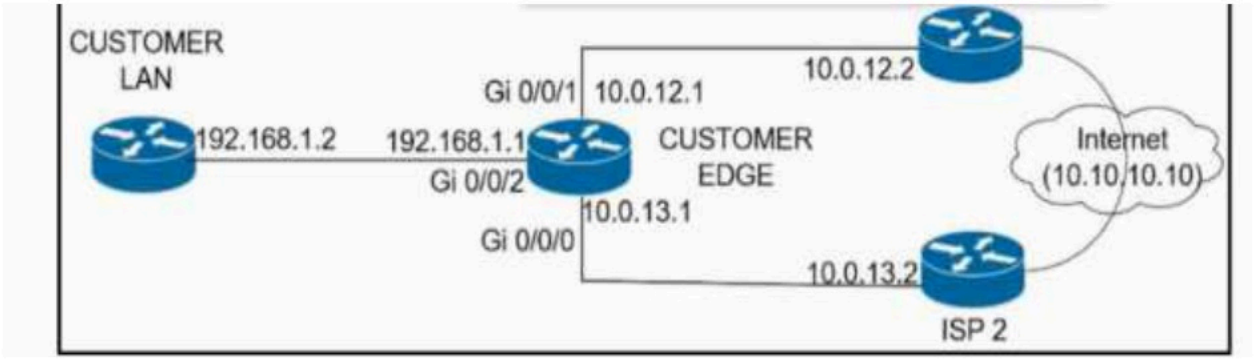
✗ **C** Move Telnet and SSH commands to a nonzero privilege level.
Privilege levels govern command authorization, but the displayed error says connections are not permitted from this terminal. The issue is the console line's outbound transport setting, not the privilege level of the commands.

✗ **D** Modify the AAA policy to allow the user to run the Telnet and SSH commands.
AAA authorization could affect whether a user may execute commands, but the exhibit reports a terminal transport restriction for both connection attempts. Allowing the commands through AAA would not remove that restriction.

Question 67

Single choice · Domain: D4 Infrastructure Services › 4.5 Troubleshoot network performance issues using IP SLA (jitter, tracking objects, delay, connectivity)

Refer to the exhibit.



Refer to the exhibit. ISP 1 and ISP 2 directly connect to the internet. A customer is tracking both ISP links to achieve redundancy but cannot see the Cisco IOS IP SLA tracking output on the router console. Which command is missing from the IP SLA configuration?

A. start-time 00:00
B. start-time now
C. start-time 0
D. start-time immediately

Answer: B. start-time now

Explanation

The IP SLA operation must be scheduled to run before it can generate results for the tracking configuration. The `start-time now` command starts the operation immediately, allowing the router to

Free demo · full exam: 636 questions

produce IP SLA tracking output for the ISP links. [1]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/command/sla-cr-book/sla_i1.html

✗ **A** start-time 00:00

This sets the operation's start time to midnight rather than starting it immediately. It does not meet the scenario's need to begin IP SLA tracking now. [1]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/command/sla-cr-book/sla_i1.html

✓ **B** start-time now

This schedules the IP SLA operation to start immediately. Once it runs, the router can produce the results needed for tracking the redundant ISP links. [1]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/command/sla-cr-book/sla_i1.html

✗ **C** start-time 0

This is not the command value used to start an IP SLA operation immediately. The supported immediate-start value is `now`. [1]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/command/sla-cr-book/sla_i1.html

✗ **D** start-time immediately

`immediately` is not the IP SLA start-time keyword. Use `now` to start the operation immediately. [1]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/command/sla-cr-book/sla_i1.html

Question 68

Single choice · Domain: D4 Infrastructure Services › 4.6 Troubleshoot NetFlow (Flexible NetFlow)

Refer to the exhibit.

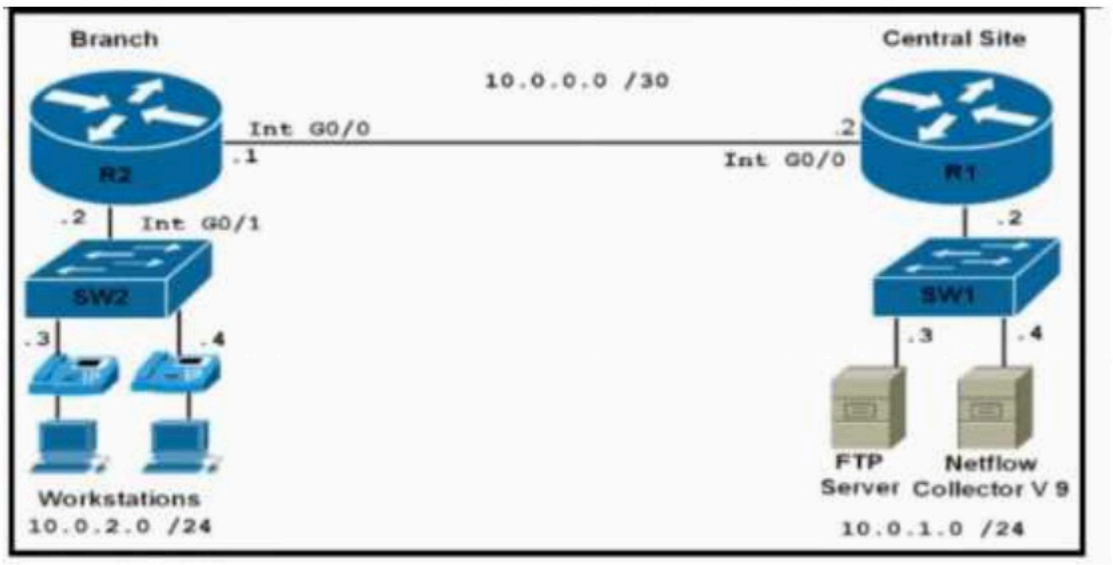
Dumpsgate

Free demo · full exam: 636 questions

```
R2#show flow monitor
Flow Monitor CCNPMON1:
  Description:       User defined
  Flow Record:       CCNPREC1
  Flow Exporter:     CCNPEXP1 (inactive)
  Cache:
    Type:            normal
    Status:          not allocated
    Size:            4096 entries / 0 bytes
    Inactive Timeout:  15 secs
    Active Timeout:    1800 secs
```

```
R2#show flow record CCNPREC1
flow record CCNPREC1:
  Description:       User defined
  No. of users:      1
  Total field space: 45 bytes
  Fields:
    match ipv4 protocol
    match ipv4 source address
    match ipv4 destination address
    match transport source-port
    match transport destination-port
    collect counter bytes long
    collect counter packets long
    collect timestamp absolute first
    collect timestamp absolute last
```

```
R2#show flow exporter
Flow Exporter CCNPEXP1:
  Description:            User defined
  Export protocol:        NetFlow Version 5
  Transport Configuration:
    Destination IP address: 10.0.1.4
    Source IP address:      10.0.2.2
    Source Interface:       GigabitEthernet0/1
    Transport Protocol:     UDP
    Destination Port:       2055
    Source Port:            57911
    DSCP:                   0x0
    TTL:                    255
    Output Features:        Not Used
```



Refer to the exhibit. An engineer uses NetFlow to troubleshoot a voice quality issue from the branch office and notices that the traffic is not arriving in the correct format to the collector. Which configuration resolves the issue?

Free demo · full exam: 636 questions

A. R2(config-flow-exporter)#export-protocol netflow-v5 R2(config)#flow exporter CCNPEXP1

B. R2(config-flow-exporter)#export-protocol netflow-v9 R2(config)#flow exporter CCNPEXP1

C. R2(config-flow-exporter)#export-protocol netflow-v9 R2(config-flow-monitor)#flow exporter CCNPEXP1 R2(config-flow-monitor)#no exporter CCNPEXP1 R2(config)#flow monitor CCNPMON1

D. R2(config-flow-exporter)#export-protocol ipfix R2(config-flow-monitor)#flow exporter CCNPEXP1 R2(config-flow-monitor)#no exporter CCNPEXP1 R2(config)#flow monitor CCNPMON1

Answer: C. R2(config-flow-exporter)#export-protocol netflow-v9 R2(config-flow-monitor)#flow exporter CCNPEXP1 R2(config-flow-monitor)#no exporter CCNPEXP1 R2(config)#flow monitor CCNPMON1

Explanation

The collector is identified as NetFlow Collector V9, so R2 must export flows using NetFlow Version 9 rather than the currently configured Version 5 or IPFIX. The flow monitor must also use the exporter so its records are sent to that collector. Option C makes the protocol and monitor-exporter configuration changes needed for this scenario.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fnetflow/configuration/xe-3s/fnf-xe-3s-book.pdf
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/network-services/network-services/m_fnf-v9-export.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fnetflow/configuration/xe-3se/fnf-xe-3e-book/fnf-v9-export.html

✗ **A** Set the exporter protocol to netflow-v5 and enter the exporter configuration.
The exhibit already shows NetFlow Version 5, while the destination is explicitly labeled Netflow Collector V9. Keeping Version 5 does not provide the format required by that collector.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/network-services/network-services/m_fnf-v9-export.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fnetflow/configuration/xe-3se/fnf-xe-3e-book/fnf-v9-export.html

✗ **B** Set the exporter protocol to netflow-v9 and enter the exporter configuration.
NetFlow Version 9 is the required format, but the shown `flow exporter CCNPEXP1` command is issued from global configuration and does not associate the exporter with the flow monitor. The monitor must reference the exporter for its flows to be sent.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fnetflow/configuration/xe-3s/fnf-xe-3s-book.pdf
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/network-services/network-services/m_fnf-v9-export.html

✓ **C** Set the exporter protocol to netflow-v9 and update the flow monitor's exporter association.
The collector requires NetFlow Version 9, so changing the exporter from Version 5 to Version 9 addresses the format mismatch. The monitor-side exporter configuration also ensures the flow monitor sends its records through CCNPEXP1.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fnetflow/configuration/xe-3s/fnf-xe-3s-book.pdf
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/network-services/network-services/m_fnf-v9-export.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fnetflow/configuration/xe-3se/fnf-xe-3e-book/fnf-v9-export.html

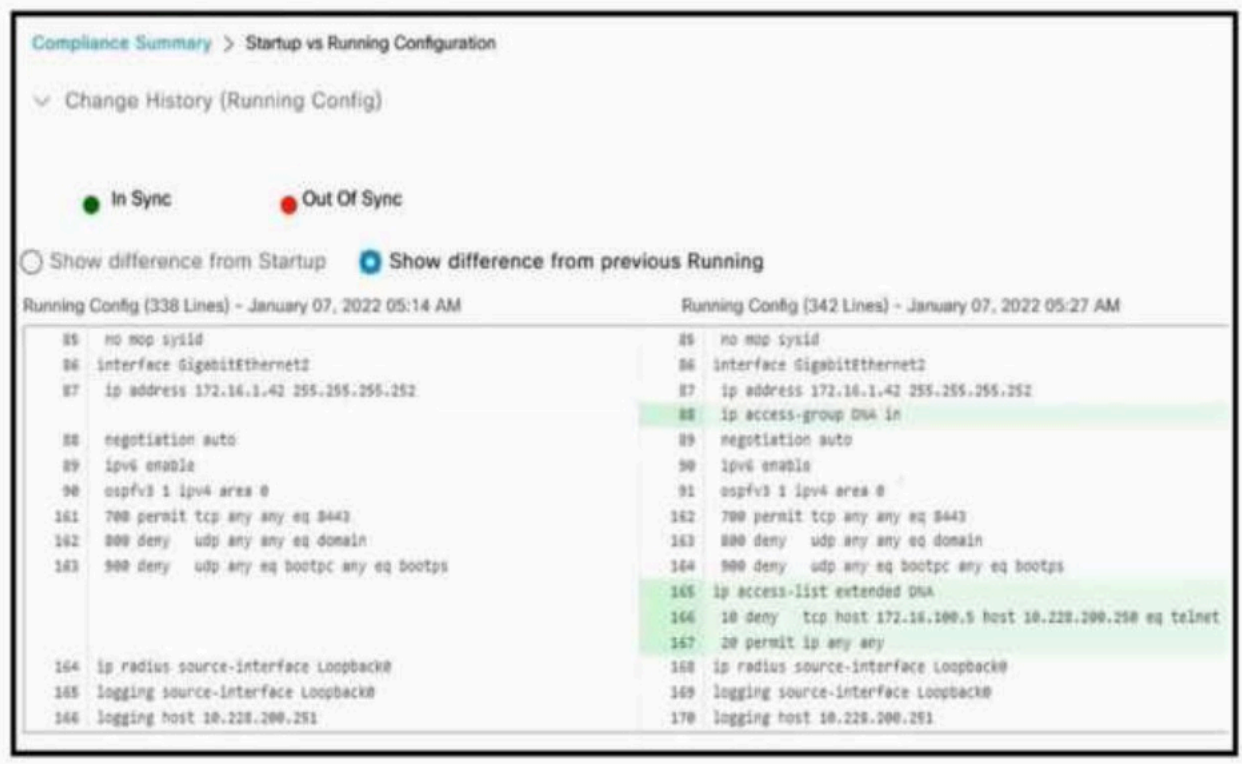
✗ **D** Set the exporter protocol to ipfix and update the flow monitor's exporter association.
Although associating an exporter with the monitor is relevant, IPFIX is not the format requested by the Netflow Collector V9 shown in the exhibit. The exporter must use NetFlow Version 9 instead.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fnetflow/configuration/xe-3s/fnf-xe-3s-book.pdf
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/network-services/network-services/m_fnf-v9-export.html

Question 69

Single choice · Domain: D4 Infrastructure Services › 4.1 Troubleshoot device management

Refer to the exhibit.



Refer to the exhibit. After an engineer updates the configuration on the device, they noticed unexpected behavior. Which command resolves the issue by completely replacing the startup configuration?

A. configure replace nvram:startup-config
B. copy system:running.config nvram:startup-config
C. configure replace nvram:private-config
D. copy running-config startup-config

Answer: A. configure replace nvram:startup-config

Explanation

The exhibit shows a difference between the running and startup configurations, and the question specifically requires completely replacing the startup configuration. `configure replace nvram:startup-config` replaces the running configuration with the contents of the startup configuration, rather than merging configuration commands as a copy operation can. This is the command that performs a configuration replacement using the startup-config file as its source. [1][3]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fundamentals/command/cf_command_ref/C_commands.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_cm-config-

files-0.html

✓ **A** configure replace nvram:startup-config
The `configure replace` command replaces the current running configuration with the configuration in the specified file; `nvram:startup-config` identifies the saved startup configuration as the source. That replacement behavior addresses the requirement to replace a configuration rather than simply copy or merge commands. [1][3]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fundamentals/command/cf_command_ref/C_commands.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_cm-config-files-0.html

✗ **B** copy system:running.config nvram:startup-config
This copies the running configuration to the startup-config file, so it saves the current running state as the startup configuration instead of replacing the running configuration from the saved startup configuration. The source and destination are opposite to what this scenario needs. [5][6]

https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/copy.htm
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/config-mgmt/configuration/15-sy/config-mgmt-15-sy-book/cm-config-files.html

✗ **C** configure replace nvram:private-config
Although `configure replace` is the relevant replacement command, `nvram:private-config` is not the startup configuration requested in the question. This would use a different configuration file as the replacement source. [1][3]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/fundamentals/command/cf_command_ref/C_commands.html
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_cm-config-files-0.html

✗ **D** copy running-config startup-config
This copies the current running configuration into startup-config, saving the running configuration for the next reload. It does not use the startup configuration to replace the running configuration, which is the needed direction and operation here. [5][6]

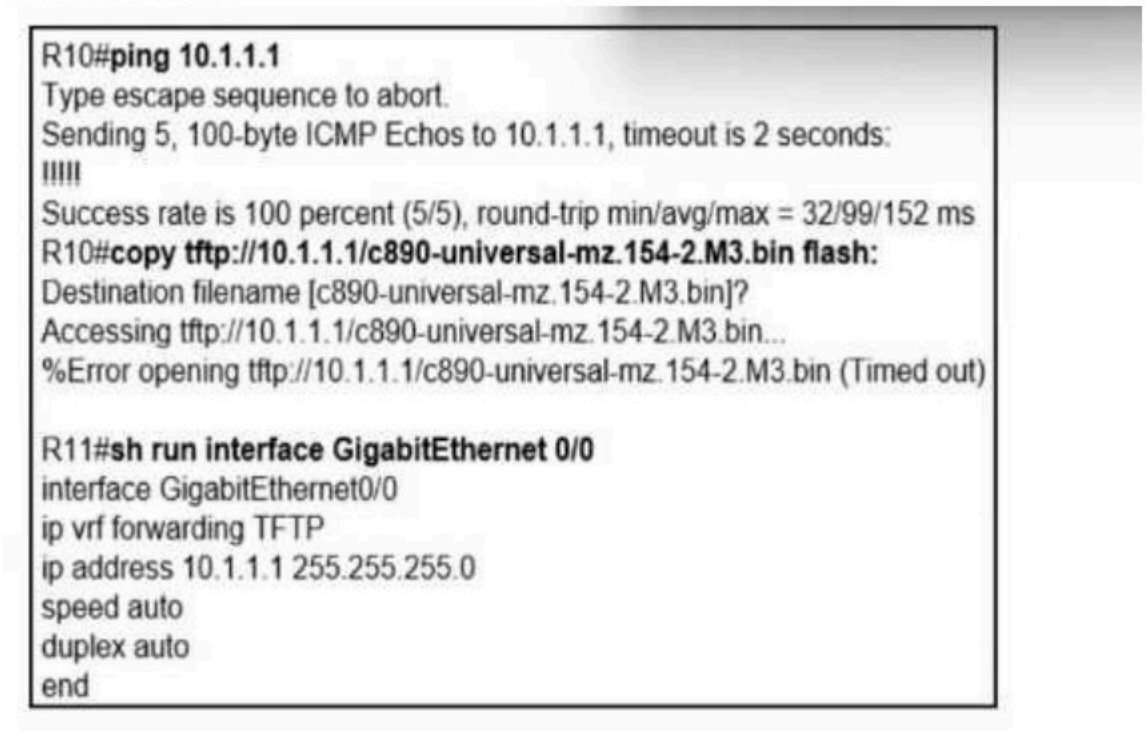
https://www.cisco.com/E-Learning/bulk/public/tac/cim/cib/using_cisco_ios_software/cmdrefs/copy.htm
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/config-mgmt/configuration/15-sy/config-mgmt-15-sy-book/cm-config-files.html

Dumpsgate

Question 70

Single choice · Domain: D4 Infrastructure Services › 4.1.c (T)FTP

Refer to the exhibit.



Refer to the exhibit. R10 attempts to copy an image file from R11 using TFTP, and the operation is timing out. Which action resolves the issue?

A. Change R10 duplex to auto to resolve the duplex mismatch between routers R10 and R11.
B. R10 must be configured to belong to the same VRF TFTP.
C. R11 requires the TFTP server source-interface to be set to GigabitEthernetO/0.
D. R11 requires the TFTP server command for the image to be set to the VRF TFTP.

Answer: C. R11 requires the TFTP server source-interface to be set to GigabitEthernetO/0.

Explanation

R11's GigabitEthernet0/0 interface is in VRF TFTP and owns the address 10.1.1.1, which R10 can reach with ping. The TFTP transfer times out because R11's TFTP server needs to use that VRF interface as its source interface. Configuring the TFTP server source interface as GigabitEthernet0/0 aligns the server's response traffic with the VRF-connected interface.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_cm-config-files-0.html
https://www.cisco.com/c/en/us/td/docs/ios/fundamentals/command/reference/cf_book.pdf
https://www.cisco.com/c/en/us/td/docs/ios/fundamentals/command/reference/cf_book/cf_f1.html

✗ **A** Change R10 duplex to auto to resolve the duplex mismatch between routers R10 and R11.
The exhibit shows a successful ping from R10 to 10.1.1.1, so basic connectivity is working; it provides no

Free demo · full exam: 636 questions

evidence of a duplex mismatch. The indicated issue is that R11's TFTP interface is in VRF TFTP, making the TFTP server's source-interface configuration relevant.

✗ **B** R10 must be configured to belong to the same VRF TFTP.
R10 is already reaching R11's 10.1.1.1 address, and the exhibit identifies the problem on the TFTP server side of the transfer. The required fix is to have R11's TFTP server use its VRF TFTP interface as the source, not to move R10 into that VRF.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_cm-config-files-0.html
https://www.cisco.com/c/en/us/td/docs/ios/fundamentals/command/reference/cf_book.pdf

✓ **C** R11 requires the TFTP server source-interface to be set to GigabitEthernet0/0.
R11's GigabitEthernet0/0 has address 10.1.1.1 and is configured with `ip vrf forwarding TFTP`. Setting it as the TFTP server's source interface makes the server use the interface in the VRF where the file-serving endpoint resides, resolving the timeout.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_cm-config-files-0.html
https://www.cisco.com/c/en/us/td/docs/ios/fundamentals/command/reference/cf_book.pdf
https://www.cisco.com/c/en/us/td/docs/ios/fundamentals/command/reference/cf_book/cf_f1.html

✗ **D** R11 requires the TFTP server command for the image to be set to the VRF TFTP.
The exhibit shows R11's interface assigned to VRF TFTP, but it does not indicate that the image or TFTP file command must be assigned to a VRF. The needed change is to configure the TFTP server source interface, not the image's VRF.

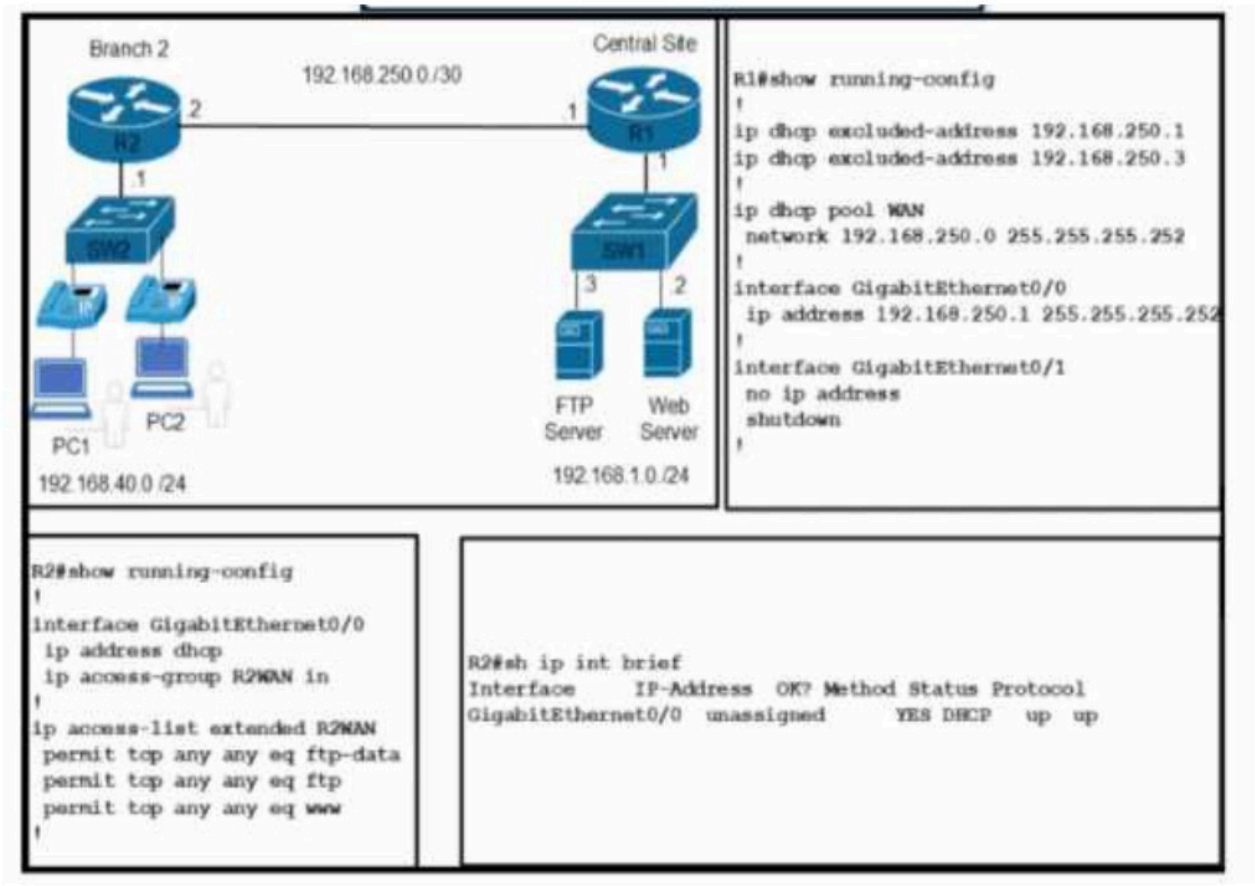
https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_cm-config-files-0.html
https://www.cisco.com/c/en/us/td/docs/ios/fundamentals/command/reference/cf_book.pdf

Dumpsgate

Question 71

Single choice · Domain: D4 Infrastructure Services › 4.4 Troubleshoot IPv4 and IPv6 DHCP (DHCP client, IOS DHCP server, DHCP relay, DHCP options)

Refer to the exhibit.



Refer to the exhibit. Which configuration is required for R2 to get the IP address from the DHCP server?

A. ip access-list extended R2WANpermit tcp any any eq 68
B. ip access-list extended R2WANpermit udp any any eq 68
C. ip access-list extended R2WANpermit udp any any eq 67
D. interface GigabitEthernet0/0ip access-group R2WAN out

Answer: B. ip access-list extended R2WANpermit udp any any eq 68

Explanation

R2's GigabitEthernet0/0 receives its address through DHCP, and the inbound ACL on that interface must permit the DHCP client traffic. DHCP uses UDP, and the client's source port is 68, so the matching ACL entry is `permit udp any any eq 68`. This allows R2 to obtain an address while retaining the existing ACL rules.

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html
https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.pdf

✗ **A** ip access-list extended R2WAN permit tcp any any eq 68
DHCP uses UDP, not TCP. This entry would not match the DHCP client traffic that R2 needs to send.

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html
https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.pdf

✓ **B** ip access-list extended R2WAN permit udp any any eq 68
A DHCP client sends its requests using UDP source port 68. Adding this permit entry allows that traffic through R2's inbound ACL so R2 can obtain an IP address.

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html
https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.pdf

✗ **C** ip access-list extended R2WAN permit udp any any eq 67
UDP port 67 is the DHCP server port, whereas the client request from R2 uses source port 68. This entry does not match the client traffic that must be permitted by the inbound ACL.

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html
https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.pdf

✗ **D** interface GigabitEthernet0/0 ip access-group R2WAN out
The exhibit already applies R2WAN inbound on GigabitEthernet0/0. Applying it outbound would not permit the DHCP request arriving at that interface, and the needed ACL rule is the UDP port 68 permit.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/security-vpn/security-vpn/m_sec-create-ip-apply-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_acl/configuration/xe-16-11/sec-data-acl-xe-16-11-book.pdf

Dumpsgate

Question 72

Single choice · Domain: D4 Infrastructure Services › 4.4 Troubleshoot IPv4 and IPv6 DHCP (DHCP client, IOS DHCP server, DHCP relay, DHCP options)

A switch has been configured to provide DHCP relay on VLAN 100 to a server with an IP address of 10.1.1.1. The DHCP server is sending syslog reports of multiple TFTP requests that also originate from the switch. As a result, the server CPU exceeded a configured threshold. Which action does the network administrator recommend to bring the server CPU threshold down?

A. Configure the switch with an access list on VLAN100 to deny TFTP.
B. Configure the switch with a VACL on VLAN100 to deny TFTP.
C. Configure the switch with ip forward-protocol udp 67 globally.
D. Configure the switch with no ip forward-protocol udp 69 on VLAN100.

Answer: D. Configure the switch with no ip forward-protocol udp 69 on VLAN100.

Explanation

The switch is relaying DHCP broadcasts and, by default, may also forward other UDP broadcasts such as TFTP to the configured DHCP server. Disabling forwarding for UDP port 69 on VLAN 100 stops those TFTP requests from reaching the server while leaving DHCP relay operational, reducing the load causing the server CPU threshold to be exceeded. [1]

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html

✗ **A** Configure the switch with an access list on VLAN100 to deny TFTP.

An access list applied on VLAN 100 is not the targeted control for stopping the switch's UDP broadcast forwarding to the DHCP server. The appropriate fix is to disable forwarding of the TFTP UDP port on the relay interface.

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/c9000/security/acls/acls-configuration-guide/access-control-lists.html

✗ **B** Configure the switch with a VACL on VLAN100 to deny TFTP.
A VLAN access map filters traffic switched within the VLAN; it does not directly disable the relay agent's forwarding of UDP broadcasts to the DHCP server. Removing TFTP from the relay's forwarded protocols addresses the source of the server requests.

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15-2_7_e/configuration_guide/b_1527e_consolidated_2960x_cg/m_sec_vlan_maps_cg_cauvery.html

✗ **C** Configure the switch with ip forward-protocol udp 67 globally.
This command adds UDP port 67 to the set of protocols forwarded globally, rather than suppressing the TFTP requests sent to the server. It does not address the reported CPU load from TFTP.

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html
https://www.cisco.com/c/dam/en/us/td/docs/ios/ipapp/command/reference/iap_cr_book.pdf

✓ **D** Configure the switch with no ip forward-protocol udp 69 on VLAN100.
The command disables forwarding of UDP port 69, which is used by TFTP, on VLAN 100. This prevents the switch from relaying the TFTP requests to the DHCP server while preserving DHCP relay for the VLAN.

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html
https://www.cisco.com/c/dam/en/us/td/docs/ios/ipapp/command/reference/iap_cr_book.pdf

Question 73

Single choice · Domain: D4 Infrastructure Services › 4.5 Troubleshoot network performance issues using IP SLA (jitter, tracking objects, delay, connectivity)

Refer to the exhibit.

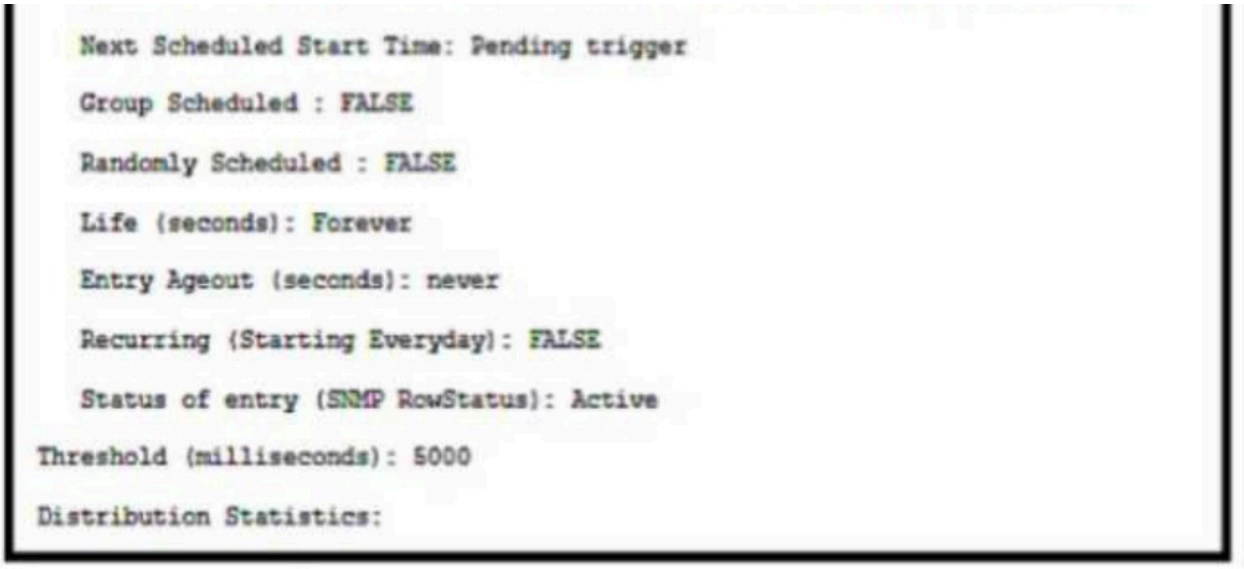
Dumpsgate

```
R1#sh track brief
Track Type        Instance                  Parameter        State Last Change
1     ip sla      10                        reachability     Down  00:03:52

R1#show ip sla configuration
IP SLAs Infrastructure Engine-III
Entry number: 10
Owner:
Tag:
Operation timeout (milliseconds): 5000
Type of operation to perform: icmp-echo
Target address/Source interface: 10.10.10.10/GigabitEthernet0/0
<->
Schedule:
   Operation frequency (seconds): 60  (not considered if randomly scheduled)
   Next Scheduled Start Time: Pending trigger
   Group Scheduled : FALSE
   Randomly Scheduled : FALSE
   Life (seconds): Forever
   Entry Ageout (seconds): never
   Recurring (Starting Everyday): FALSE
   Status of entry (SNMP RowStatus): Active
Threshold (milliseconds): 5000
Distribution Statistics:
```

```
Operation timeout (milliseconds): 5000
Type of operation to perform: icmp-echo
Target address/Source interface: 10.10.10.10/GigabitEthernet0/0
<->
Schedule:
   Operation frequency (seconds): 60  (not considered if randomly scheduled)
```


Free demo · full exam: 636 questions



Refer to the exhibit A network engineer notices that the configured track option is down Which configuration resolves the issue*?

A. ip sla schedule 10 start-time now
B. ip sla schedule 10 start-time pending life forever
C. ip sla schedule 10 no timeout
D. ip sla schedule 10 no threshold

Answer: A. ip sla schedule 10 start-time now

Explanation

The exhibit shows IP SLA operation 10 is configured and active, but its next scheduled start time is "Pending trigger," so it has not begun running. Scheduling it with a start time of now starts the probe immediately, allowing the associated track to reflect the operation's reachability result. [2]

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960cx_3650cx/software/release/15-2_6_e/configuration_guide/b_1526e_consolidated_3560cx_2960cx_cg/configuring_service_level_agreements.html

✓ **A** ip sla schedule 10 start-time now
This starts the configured IP SLA operation immediately. The exhibit shows its next start is still pending a trigger, which explains why the track is down; scheduling it to start now lets the probe run and provide a result for tracking. [2]

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960cx_3650cx/software/release/15-2_6_e/configuration_guide/b_1526e_consolidated_3560cx_2960cx_cg/configuring_service_level_agreements.html

✗ **B** ip sla schedule 10 start-time pending life forever
The exhibit already shows the next scheduled start time as "Pending trigger," so leaving the operation pending does not make the probe run. The operation needs to start now to produce a reachability result.

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960cx_3650cx/software/release/15-2_6_e/configuration_guide/b_1526e_consolidated_3560cx_2960cx_cg/configuring_service_level_agreements.html

✗ **C** ip sla schedule 10 no timeout
The configured timeout is a probe parameter, not the reason the operation has not started. The exhibit identifies the next start time as pending, so changing timeout behavior would not resolve that scheduling issue.

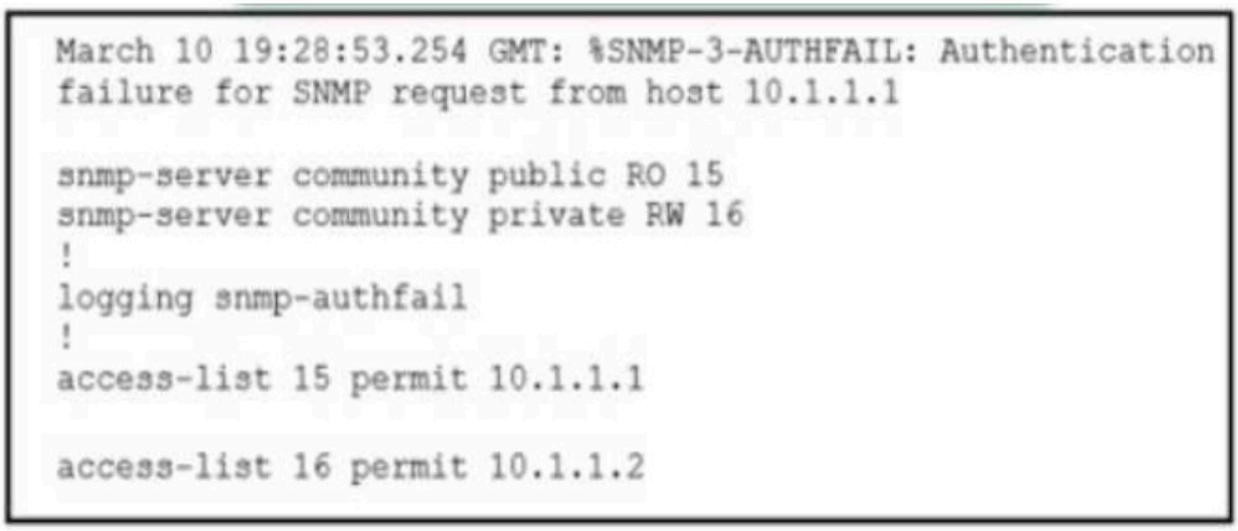
https://www.cisco.com/en/US/docs/ios-xml/ios/ipsla/configuration/15-1sg/sla_ftp.pdf

✗ **D** ip sla schedule 10 no threshold
The threshold affects how a probe's response time is assessed; it does not start the IP SLA operation. The exhibit shows the operation is still waiting for a start trigger, so it must be scheduled to start now.
https://www.cisco.com/en/US/docs/ios-xml/ios/ipsla/configuration/15-1sg/sla_ftp.pdf

Question 74

Single choice · Domain: D4 Infrastructure Services › 4.2 Troubleshoot SNMP (v2c, v3)

Refer to the exhibit.



Refer to the exhibit Which action resolves the issue?

A. Configure host IP address in access-list 16
B. Configure SNMPv3 on the router
C. Configure SNMP authentication on the router
D. Configure a valid SNMP community string

Answer: D. Configure a valid SNMP community string

Explanation

The log reports an SNMP authentication failure from host 10.1.1.1. The router has a read-only community named `public` protected by access list 15, which permits that host, but the failure indicates that the request did not present a matching community string. Configure the SNMP manager to use a valid community string configured on the router, such as `public`. [1][2]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/12-4/snmp-12-4-book.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/snmp-15-sy-book/nm-snmp-cfg-snmp-support.pdf

✗ **A** Configure host IP address in access-list 16
Access list 16 is attached to the `private` read-write community and currently permits 10.1.1.2. The failed request is from 10.1.1.1, but this does not explain the authentication failure for the `public` community, which is governed by access list 15 and already permits that host. [1][2]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/12-4/snmp-12-4-book.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/snmp-15-sy-book/nm-snmp-cfg-snmp-support.pdf

✗ **B** Configure SNMPv3 on the router
The exhibit shows SNMP community configuration, which is used for community-based SNMP rather than SNMPv3. Switching the router to SNMPv3 is not the necessary fix for a request that fails because its community string is not valid. [2][8]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/snmp-15-sy-book/nm-snmp-cfg-snmp-support.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/xe-3e/snmp-xe-3e-book/nm-snmp-snmpv3-comm-supp.pdf

✗ **C** Configure SNMP authentication on the router
The configuration already defines SNMP communities, and the log identifies an authentication failure for an incoming request. The issue is resolved by having the requester use a community string that matches a configured community, not by enabling a different authentication feature. [1][2]

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/12-4/snmp-12-4-book.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/snmp-15-sy-book/nm-snmp-cfg-snmp-support.pdf

✓ **D** Configure a valid SNMP community string
The router has configured communities `public` and `private`; the request from 10.1.1.1 is allowed by access list 15 for the `public` community. Configuring the SNMP manager to send the matching valid community string resolves the authentication failure. [1][2]

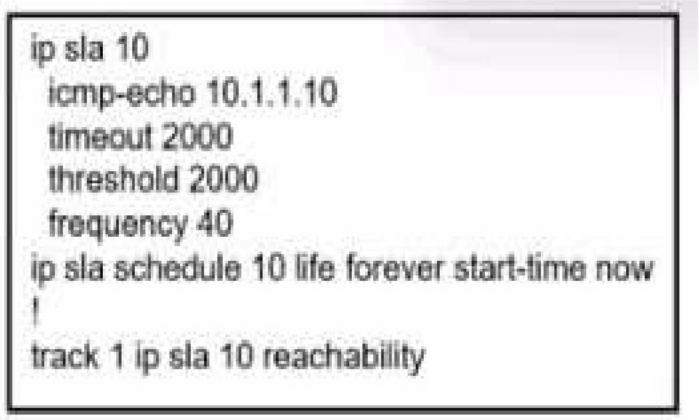
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/12-4/snmp-12-4-book.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/snmp-15-sy-book/nm-snmp-cfg-snmp-support.pdf

Dumpsgate

Question 75

Single choice · Domain: D4 Infrastructure Services › 4.5 Troubleshoot network performance issues using IP SLA (jitter, tracking objects, delay, connectivity)

Refer to the exhibit.



An engineer configured IP SLA to monitor a next hop on a router for reachability When the next hop is unreachable, the router is executing tracking and failing over another route, but packet loss is experienced because the reachability is flapping. Which action resolves the issue?

A. Append delay up 0 down Ovto the track command
B. Increase the timeout of the sla probe to 6000
C. Append delay up SO down 60vto the track command
D. Increase the frequency of the sla probe to 60.

Answer: C. Append delay up SO down 60vto the track command

Explanation

The reachability state is flapping, causing the tracked route to switch too quickly. Configure a delay on the tracking object so it waits 60 seconds before declaring the tracked object down, while allowing it to come back up after 50 seconds. This dampens transient probe failures and prevents unnecessary failover and packet loss. [1]

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-addressing/ip-addressing/m_iap-eot-xe.html

✗ **A** Append delay up 0 down 0 to the track command
A zero delay in both directions does not dampen state changes; the track object can still react immediately to short-lived probe failures. It therefore does not address the reported flapping. [1]

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-addressing/ip-addressing/m_iap-eot-xe.html

✗ **B** Increase the timeout of the SLA probe to 6000
The timeout controls how long the probe waits for a response, not how long the tracked state must remain changed before failover occurs. Increasing it does not provide the intended tracking-state delay and could make failure detection slower. [2]

https://www.cisco.com/c/en/us/td/docs/ios/ios_xe/ipapp/configuration/guide/2_xe/ipapp_xe_book.pdf

✓ **C** Append delay up 50 down 60 to the track command
The tracking delay holds a state change before it is acted on: 60 seconds for a transition to down and 50

seconds for a transition to up. This dampens brief reachability changes so they do not trigger rapid route failover and recovery. [1]

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-addressing/ip-addressing/m_iap-eot-xe.html

✗ **D** Increase the frequency of the SLA probe to 60
Changing the probe frequency changes how often the SLA operation runs; it does not add a stabilization delay to the tracked state. It is not the configuration that prevents flapping from causing repeated failover. [2]

https://www.cisco.com/c/en/us/td/docs/ios/ios_xe/ipapp/configuration/guide/2_xe/ipapp_xe_book.pdf

---

Question 76

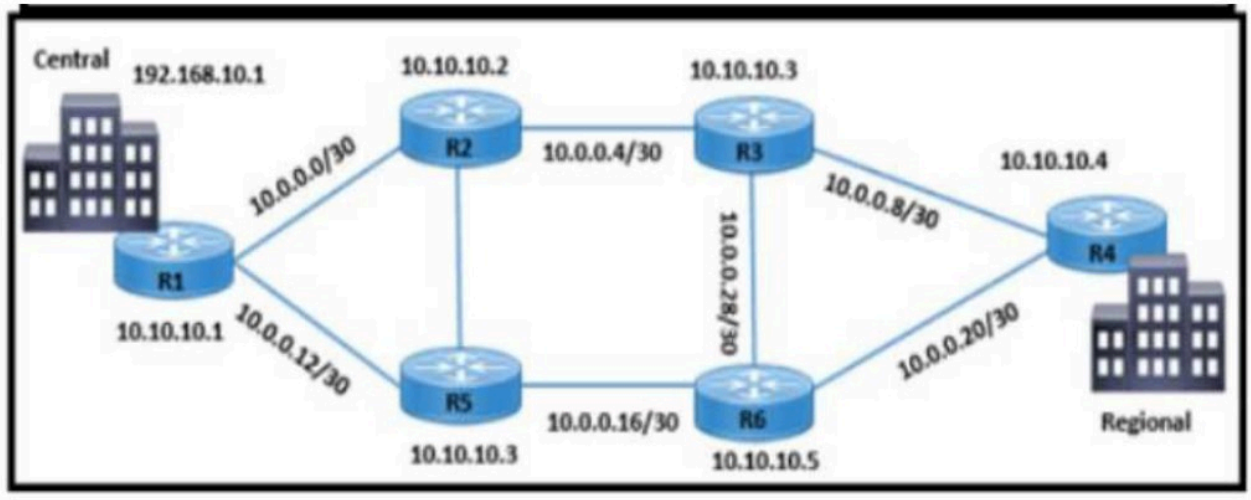
Single choice · Domain: D4 Infrastructure Services › 4.5 Troubleshoot network performance issues using IP SLA (jitter, tracking objects, delay, connectivity)

Refer to the exhibit.

Dumpsgate

```
R3#show ip sla statistics
IPSLAs Latest Operation Statistics
IPSLA operation id: 10
Type of operation: icmp-echo
        Latest RTT: 24 milliseconds
Latest operation start time: *21:26:43.211 UTC Sat Sep 18 2021
Latest operation return code: OK
Number of successes: 75
Number of failures: 0
Operation time to live: Forever

IPSLA operation id: 20
Type of operation: icmp-echo
        Latest RTT: NoConnection/Busy/Timeout
Latest operation start time: *21:26:47.499 UTC Sat Sep 18 2021
Latest operation return code: No connection
Number of successes: 128
Number of failures: 459
Operation time to live: Forever
```



Refer to me exhibit Traffic from R3 to the central site does not use alternate paths when R3 cannot reach 10 10 10 2 Traffic on R3 destined to R4 takes an alternate route via 10 10 10.6 when 10 10 10 4 is not accessible from R3 Which configuration switches traffic destined to 10 10 10 2 from R3 on the alternate path"

A. R3(config)#ip route 192.168.10.1 255.255.265.255 10.10.10.2 track 20

B. R2(config)#ip route 10.10 10 3 255 255.255 255 10.0.0.6

C. R3(config)#track(20 ip sla 20 reachability

D. R6(config)#ip route 10.10.10 3 255.255.255.255 10.0.0.30

Answer: A. R3(config)#ip route 192.168.10.1 255.255.265.255 10.10.10.2 track 20

Explanation

R3 needs a route to the central-site host 192.168.10.1 that is removed when tracking object 20 reports the primary path unavailable. Option A applies that tracking object to the static route via R2 (10.10.10.2), allowing another route to be selected when the tracked route is withdrawn. The mask shown in option A appears to contain a typo (255.255.265.255); the intended host mask is 255.255.255.255.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-addressing/ip-addressing/m_iap-eot-xe.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/12-2sx/iap-12-2sx-book.pdf

✓ **A** R3(config)#ip route 192.168.10.1 255.255.265.255 10.10.10.2 track 20
This is the R3 static route toward the central-site address, through R2, and ties the route to tracking object 20. When that tracking object goes down, the route is withdrawn so R3 can use an alternate route. The displayed mask appears mistyped; a host route to 192.168.10.1 would use 255.255.255.255.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/ip-addressing/ip-addressing/m_iap-eot-xe.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/12-2sx/iap-12-2sx-book.pdf

✗ **B** R2(config)#ip route 10.10 10 3 255 255.255 255 10.0.0.6
This configures a route on R2, not the required tracked route on R3 for traffic to the central site. It also does not include a tracking condition that would make R3 switch paths when its primary path fails.

✗ **C** R3(config)#track(20 ip sla 20 reachability
This defines a tracking relationship between object 20 and IP SLA operation 20, but it does not apply that object to a static route. By itself, it therefore does not cause R3 to withdraw the route to the central site when reachability fails.

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/command/sla-cr-book/sla_s2.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/iap-cr-book/iap-t1.pdf

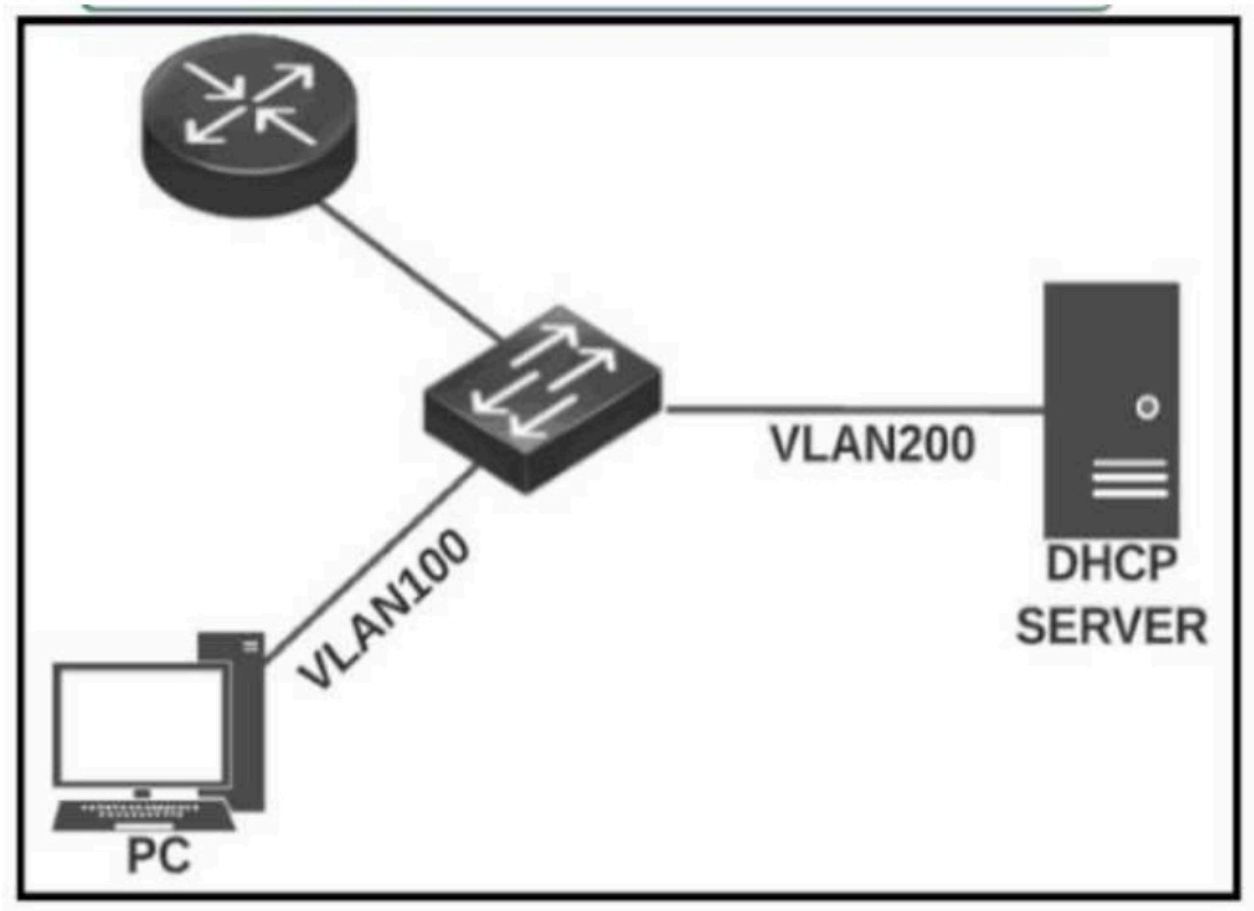
✗ **D** R6(config)#ip route 10.10.10 3 255.255.255.255 10.0.0.30
This configures a route on R6 rather than the tracked route on R3 that controls traffic to the central site. It does not associate the route with tracking object 20, so it cannot trigger the requested failover on R3.

Dumpsgate

Question 77

Single choice · Domain: D4 Infrastructure Services › 4.4 Troubleshoot IPv4 and IPv6 DHCP (DHCP client, IOS DHCP server, DHCP relay, DHCP options)

Refer to the exhibit.



Refer to lhe exhibit. APC is configured to obtain an IP address automatically, but it receives an IP address only from the 169.254.0.0 subnet The DHCP server logs contained no DHCPDISCOVER message from the MAC address of the PC. Which action resolves the issue?

A. Configure an ip helper-address on the router to forward DHCP messages to the server.
B. Configure DHCP Snooping on the switch to forward DHCP messages to the server.
C. Configure a DHCP reservation on the server for the PC.
D. Configure a static IP address on the PC and exclude it from the DHCP pool.

Answer: A. Configure an ip helper-address on the router to forward DHCP messages to the server.

Explanation

The PC is in VLAN 100, while the DHCP server is in VLAN 200. DHCPDISCOVER messages are broadcasts and are not normally routed between VLANs, which explains why the server logs show no discovery from

Free demo · full exam: 636 questions

the PC and the PC falls back to an APIPA address in 169.254.0.0/16. Configure an IP helper address on the router's VLAN 100 interface so it relays the DHCP request to the server in VLAN 200. [1][5]

https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/ipaddresses-and-services/configuration-guide/ip-addresses-and-services-config-cisco8000/implement-dhcp-overview/how-dhcp-relay-works.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-s/dhcp-15-s-book/dhcp-relay-agent.html

✓ **A** Configure an ip helper-address on the router to forward DHCP messages to the server.
The PC and DHCP server are on different VLANs, so the router must relay the PC's DHCP broadcast to the server. An IP helper address enables that relay; the absence of the PC's DHCPDISCOVER in the server logs is consistent with the request not reaching the server. [1][5]

https://www.cisco.com/c/en/us/td/docs/iosxr/cisco8000/ipaddresses-and-services/configuration-guide/ip-addresses-and-services-config-cisco8000/implement-dhcp-overview/how-dhcp-relay-works.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_dhcp/configuration/15-s/dhcp-15-s-book/dhcp-relay-agent.html

✗ **B** Configure DHCP Snooping on the switch to forward DHCP messages to the server.
DHCP snooping is a switch security feature that controls DHCP traffic; it does not perform the inter-VLAN relay needed here. The router must forward the DHCP request from VLAN 100 to the server in VLAN 200. [6]

https://www.cisco.com/c/en/us/td/docs/switches/lan/csbms/CBS_250_350/Administration-Guide/cbs-350/cbs_350_chapter_14.pdf

✗ **C** Configure a DHCP reservation on the server for the PC.
A reservation determines which address the server offers after it receives the client's DHCP request. The server has no DHCPDISCOVER from this PC, so a reservation would not make the request reach the server. [2]

https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/27470-100.html

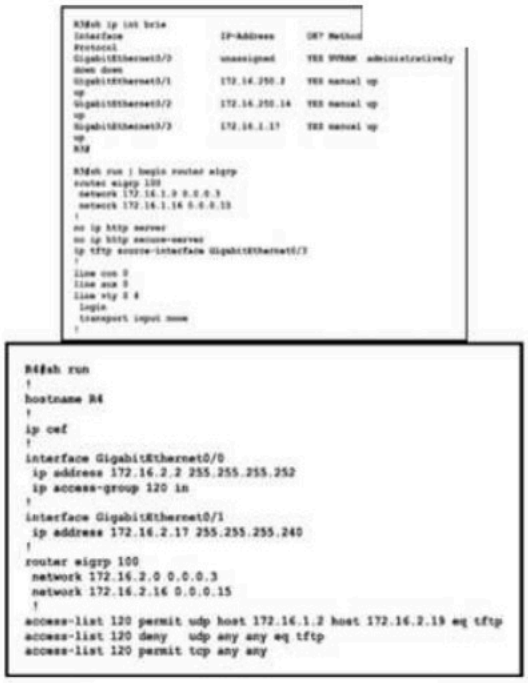
✗ **D** Configure a static IP address on the PC and exclude it from the DHCP pool.
A static address could avoid the DHCP process, but it does not resolve the failure to relay DHCP requests between VLAN 100 and VLAN 200. It also does not provide the requested automatic address assignment.

Question 78

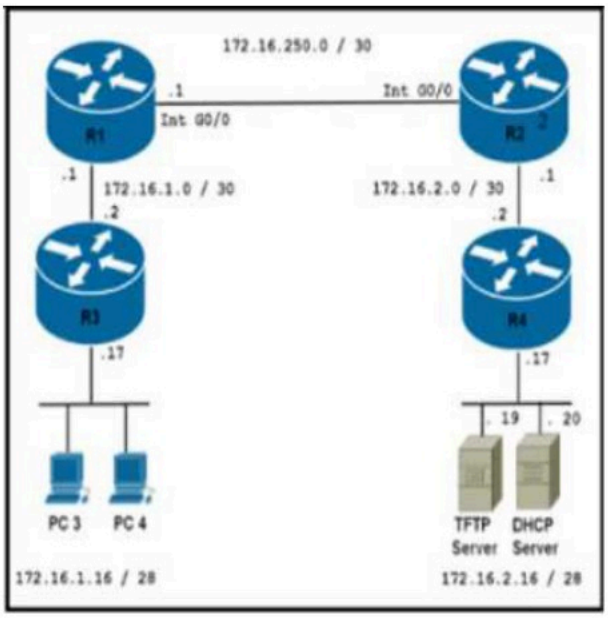
Single choice · Domain: D4 Infrastructure Services › 4.1.c (T)FTP

Refer to the exhibit.

Free demo · full exam: 636 questions



```
R4#sh run
!
hostname R4
!
ip cef
!
interface GigabitEthernet0/0
 ip address 172.16.2.2 255.255.255.252
 ip access-group 120 in
!
interface GigabitEthernet0/1
 ip address 172.16.2.17 255.255.255.240
!
router eigrp 100
 network 172.16.2.0 0.0.0.3
 network 172.16.2.16 0.0.0.15
 !
access-list 120 permit udp host 172.16.1.2 host 172.16.2.19 eq tftp
access-list 120 deny   udp any any eq tftp
access-list 120 permit tcp any any
```



- R4(config)#no access-list 120 permit udp host 172.16.1.2 host 172.16.2.19 eq 69
  R3(config)#no ip tftp source-interface GigabitEthernet0/3
- R4(config)#no access-list 120 permit udp host 172.16.1.2 host 172.16.2.19 eq 69
  R4(config)#access-list 120 permit udp host 172.16.1.17 host 172.16.2.19 eq 69
  R4(config)#access-list 120 permit tcp any any
- R4(config)#no access-list 120 permit udp host 172.16.1.2 host 172.16.2.19 eq 69
  R4(config)#no access-list 120 deny udp any any eq tftp
  R4(config)#access-list 120 permit tcp any any
- R4(config)#no access-list 120 permit udp host 172.16.1.2 host 172.16.2.19 eq 69
  R4(config)#access-list 120 permit tcp host 172.16.1.17 host 172.16.2.19 eq 69
  R4(config)#access-list 120 permit tcp any any

A. Option A
B. Option B
C. Option C
D. Option D

Answer: B. Option B

Explanation

R3's TFTP traffic must use its LAN interface address, 172.16.1.17, as the source when contacting the TFTP server at 172.16.2.19. Option B replaces the ACL entry that permits the wrong source address with one permitting UDP from 172.16.1.17 to the server's TFTP port, while retaining the TCP permit. This allows the transfer under the configuration shown.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_ifs-file-trans-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ifs/configuration/xe-16-7/ifs-xe-16-7-book/ifs-file-trans.pdf

✗ **A** Remove the UDP permit for host 172.16.1.2 to host 172.16.2.19 on TFTP port 69, and change R3's TFTP source interface to GigabitEthernet0/3.
This removes the existing TFTP permit but does not add a permit for R3's LAN source address, 172.16.1.17. Changing R3's TFTP source interface to GigabitEthernet0/3 does not fix the ACL mismatch shown in the exhibit.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_ifs-file-trans-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ifs/configuration/xe-16-7/ifs-xe-16-7-book/ifs-file-trans.pdf

✓ **B** Remove the UDP permit for host 172.16.1.2 to host 172.16.2.19 on TFTP port 69; permit UDP from host 172.16.1.17 to host 172.16.2.19 on port 69; permit TCP any any.
R3's LAN interface has address 172.16.1.17, which is the source address the ACL must allow for the TFTP request to the server at 172.16.2.19. Replacing the permit for 172.16.1.2 with a UDP permit for 172.16.1.17 allows the TFTP traffic; the TCP permit remains in place as shown.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_ifs-file-trans-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ifs/configuration/xe-16-7/ifs-xe-16-7-book/ifs-file-trans.pdf

✗ **C** Remove the UDP permit for host 172.16.1.2 to host 172.16.2.19 on TFTP port 69; remove the UDP deny for any-to-any TFTP; permit TCP any any.
Removing the deny does not add the needed permit for UDP traffic sourced from R3's LAN address, 172.16.1.17. The TCP permit does not authorize TFTP, which uses UDP.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_ifs-file-trans-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ifs/configuration/xe-16-7/ifs-xe-16-7-book/ifs-file-trans.pdf

✗ **D** Remove the UDP permit for host 172.16.1.2 to host 172.16.2.19 on TFTP port 69; permit TCP from host 172.16.1.17 to host 172.16.2.19 on port 69; permit TCP any any.
The source address 172.16.1.17 matches R3's LAN interface, but this option permits TCP rather than UDP. TFTP uses UDP, so the required transfer traffic would still not be permitted.

https://www.cisco.com/c/en/us/td/docs/routers/ios-xe/system-management/system-management/m_ifs-file-trans-0.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ifs/configuration/xe-16-7/ifs-xe-16-7-book/ifs-file-trans.pdf

Free demo · full exam: 636 questions

Question 79

Single choice · Domain: D4 Infrastructure Services › 4.7 Troubleshoot network problems using Cisco DNA Center assurance (connectivity, monitoring, device health, network health)

Refer to the exhibit.



Refer to the exhibit. Which action resolves the issue?

A. Establish connectivity between the NTP server and the switch.
B. Configure the local time on Cisco DNA Center
C. Configure the local time on the SW1 device
D. Establish connectivity between the NTP server and Cisco DNA Center.

Answer: C. Configure the local time on the SW1 device

Explanation

The alert says the time on Cisco DNA Center and SW1 has drifted by more than three minutes, preventing accurate processing of device data. The correct action is to configure the local time on SW1 so it aligns with Cisco DNA Center. Changing either clock or fixing NTP connectivity is not the identified remedy in the supplied answer.

https://www.cisco.com/c/en/us/support/docs/ip/network-time-protocol-ntp/221214-troubleshoot-ntp-on-dna-center.html
https://www.cisco.com/c/en/us/support/docs/ip/network-time-protocol-ntp/221214-troubleshoot-ntp-on-dna-center.pdf

✗ **A** Establish connectivity between the NTP server and the switch.
The exhibit identifies excessive time difference between Cisco DNA Center and SW1; it does not indicate that the switch cannot reach an NTP server. The specified correction is to configure SW1's local time, rather than to establish NTP connectivity.

Free demo · full exam: 636 questions

https://www.cisco.com/c/en/us/support/docs/ip/network-time-protocol-ntp/221214-troubleshoot-ntp-on-dna-center.html
https://www.cisco.com/c/en/us/support/docs/ip/network-time-protocol-ntp/221214-troubleshoot-ntp-on-dna-center.pdf

✗ **B** Configure the local time on Cisco DNA Center
Cisco DNA Center is one side of the time comparison, but the correct action for this issue is to correct the time on SW1. Changing Cisco DNA Center's local time would not address the specified device-side time configuration.

https://www.cisco.com/c/en/us/support/docs/ip/network-time-protocol-ntp/221214-troubleshoot-ntp-on-dna-center.html
https://www.cisco.com/c/en/us/support/docs/ip/network-time-protocol-ntp/221214-troubleshoot-ntp-on-dna-center.pdf

✓ **C** Configure the local time on the SW1 device
The alert explicitly reports that SW1's time has drifted too far from Cisco DNA Center and that a difference greater than three minutes prevents accurate processing. Configuring the local time on SW1 resolves the device-side time mismatch.

https://www.cisco.com/c/en/us/support/docs/ip/network-time-protocol-ntp/221214-troubleshoot-ntp-on-dna-center.html
https://www.cisco.com/c/en/us/support/docs/ip/network-time-protocol-ntp/221214-troubleshoot-ntp-on-dna-center.pdf

✗ **D** Establish connectivity between the NTP server and Cisco DNA Center.
The exhibit reports a time mismatch between Cisco DNA Center and SW1, not a connectivity failure between Cisco DNA Center and an NTP server. The correct action is to configure SW1's local time.

https://www.cisco.com/c/en/us/support/docs/ip/network-time-protocol-ntp/221214-troubleshoot-ntp-on-dna-center.html
https://www.cisco.com/c/en/us/support/docs/ip/network-time-protocol-ntp/221214-troubleshoot-ntp-on-dna-center.pdf

Dumpsgate

Question 80

Single choice · Domain: D4 Infrastructure Services › 4.1.b Telnet, HTTP, HTTPS, SSH, and SCP

Refer to the exhibit.

```
admin@linux:~$ telnet 198.51.100.64
Trying 198.51.100.64...
Connected to 198.51.100.64.
Escape character is '^]'.

User Access Verification

Password: admin
CPE> exit
Connection closed by foreign host.
admin@linux:~$ ssh 198.51.100.64
admin@198.51.100.64's password: admin
Permission denied, please try again.
admin@198.51.100.64's password: admin
Permission denied, please try again.
admin@198.51.100.64's password: admin
Connection closed by 198.51.100.64 port 22
admin@linux:~$
```

Refer to the exhibit. An administrator can log in to the device using Telnet but the attempts to log in to the same device using SSH with the same credentials fail Which action resolves this issue?

A. Configure SSH service on the router
B. Configure transport input all on the VTY lines to allow SSH
C. Configure to use the Telnet user database for SSH as well
D. Configure the VTY lines with login local

Answer: D. Configure the VTY lines with login local

Explanation

Telnet access works, but SSH authentication fails even when the same credentials are used. Configure the VTY lines with `login local` so incoming SSH sessions authenticate against the device's local username database. This addresses the authentication difference indicated by the failed SSH login, rather than merely enabling SSH transport.

https://www.cisco.com/c/en/us/support/docs/security-vpn/secure-shell-ssh/4145-ssh.pdf
https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/authentication-authorization-accounting-aaa/200173-Verify-AAA-behaviour-when-login-local.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_ssh/configuration/15-mt/sec-usr-ssh-15-mt-book/sec-ssh-term-line.pdf

✗ **A** Configure SSH service on the router
The device accepts an SSH connection far enough to request a password, so the failure shown is during authentication rather than because SSH service is unavailable. Enabling the service alone would not resolve the VTY login method issue.

https://www.cisco.com/c/en/us/support/docs/security-vpn/secure-shell-ssh/4145-ssh.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_ssh/configuration/15-mt/sec-usr-ssh-15-mt-book/sec-ssh-term-line.pdf

✗ **B** Configure transport input all on the VTY lines to allow SSH
The SSH session reaches a password prompt, which shows SSH is already permitted to connect to the VTY lines. Changing the permitted transport protocols would not make the device authenticate the supplied credentials successfully.

https://www.cisco.com/c/en/us/support/docs/security-vpn/secure-shell-ssh/212142-Configure-SSH-on-Tty-Lines-with-Menu-Opt.pdf
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_ssh/configuration/15-mt/sec-usr-ssh-15-mt-book/sec-ssh-term-line.pdf

✗ **C** Configure to use the Telnet user database for SSH as well
Telnet and SSH do not require separate user databases in this scenario; the VTY login configuration determines how the credentials are checked. Configure the VTY lines to use the local username database instead.

https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/authentication-authorization-accounting-aaa/200173-Verify-AAA-behaviour-when-login-local.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_ssh/configuration/15-mt/sec-usr-ssh-15-mt-book/sec-ssh-term-line.pdf

✓ **D** Configure the VTY lines with login local
`login local` makes the VTY lines authenticate users against the router's local username database. This provides the local-credential login method needed for SSH, resolving the failed authentication while retaining the same credentials.

https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/authentication-authorization-accounting-aaa/200173-Verify-AAA-behaviour-when-login-local.html
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_ssh/configuration/15-mt/sec-usr-ssh-15-mt-book/sec-ssh-term-line.pdf