Cisco

(300-410)

Implementing Cisco Enterprise Advanced Routing and Services (ENARSI)

Total: **620 Questions**

Link:

Question: 1

Refer to the exhibit. Users in the branch network of 2001:db8:0:4::/64 report that they cannot access the Internet. Which command is issued in IPv6 router EIGRP 100 configuration mode to solve this issue?

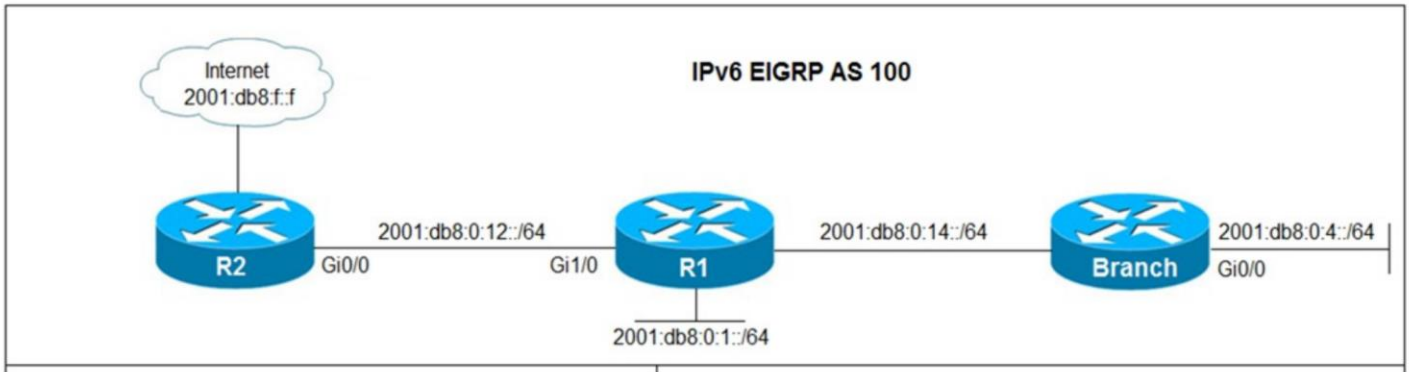


```
R1# show ipv6 eigrp topology
EIGRP-IPv6 Topology Table for AS(100)/ID(10.1.12.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
P 2001:DB8:0:4::/64, 1 successors, FD is 28416
       via FE80::C828:DFF:FEF4:1C (28416/2816), FastEthernet3/0
P 2001:DB8:0:1::/64, 1 successors, FD is 2816
       via Connected, GigabitEthernet0/0
P ::/0, 1 successors, FD is 2816
       via FE80::C821:17FF:FE04:8 (2816/256), GigabitEthernet1/0
P 2001:DB8:0:14::/64, 1 successors, FD is 28160
       via Connected, FastEthernet3/0
P 2001:DB8:0:12::/64, 1 successors, FD is 2816
       via Connected, GigabitEthernet1/0
```

```
Branch# show ipv6 eigrp topology
EIGRP-IPv6 Topology Table for AS(100)/ID(4.4.4.4)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
P 2001:DB8:0:4::/64, 1 successors, FD is 2816
       via Connected, GigabitEthernet0/0
P 2001:DB8:0:1::/64, 1 successors, FD is 28416
       via FE80::C820:17FF:FE04:54 (28416/2816), FastEthernet1/0
P 2001:DB8:0:14::/64, 1 successors, FD is 28160
       via Connected, FastEthernet1/0
P 2001:DB8:0:12::/64, 1 successors, FD is 28416
       via FE80::C820:17FF:FE04:54 (28416/2816), FastEthernet1/0
```

A. Issue the eigrp stub command on R1.
B. Issue the no eigrp stub command on R1.
C. Issue the eigrp stub command on R2.
D. Issue the no eigrp stub command on R2.

Answer: B

Explanation:

Issue the no eigrp stub command on R1.

Question: 2

Refer to the exhibit. Which configuration configures a policy on R1 to forward any traffic that is sourced from the 192.168.130.0/24 network to R2?

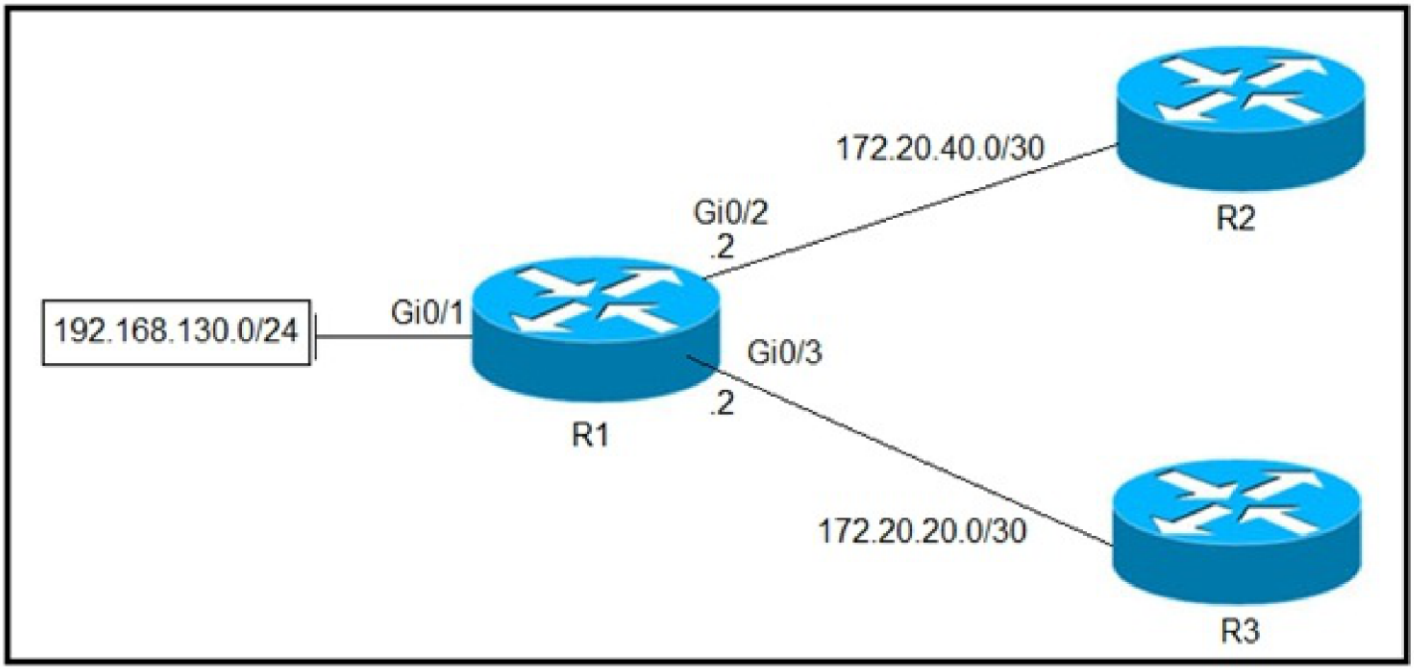

A.

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/2
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.20.2
```

B.

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/1
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.40.2
```

C.

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/2
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.20.1
```

D.

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/1
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.40.1
```

Answer: D

Explanation:

```
access-list 1 permit 192.168.130.0 0.0.0.255
!
interface Gi0/1
ip policy route-map test
!
route-map test permit 10
match ip address 1
set ip next-hop 172.20.40.1
```

Question: 3

R2 has a locally originated prefix 192.168.130.0/24 and has these configurations:

```
ip prefix-list test seq 5 permit 192.168.130.0/24
!
route-map OUT permit10
match ip address prefix-list test
set as-path prepend 65000
```

What is the result when the route-map OUT command is applied toward an eBGP neighbor R1 (1.1.1.1) by using the

neighbor 1.1.1.1 route-map OUT out command?

A. R1 sees 192.168.130.0/24 as two AS hops away instead of one AS hop away.
B. R1 does not accept any routes other than 192.168.130.0/24
C. R1 does not forward traffic that is destined for 192.168.30.0/24
D. Network 192.168.130.0/24 is not allowed in the R1 table

Answer: A

Explanation:

R1 sees 192.168.130.0/24 as two AS hops away instead of one AS hop away.

Question: 4

Which method changes the forwarding decision that a router makes without first changing the routing table or influencing the IP data plane?

A. nonbroadcast multiaccess
B. packet switching
C. policy-based routing
D. forwarding information base

Answer: C

Explanation:

The correct answer is **C. policy-based routing (PBR)**. PBR is a mechanism that allows network administrators to define routing policies based on criteria other than the destination IP address. This means traffic can be forwarded along specific paths depending on factors like source IP address, application type, or packet size, without requiring modifications to the main routing table. Traditional routing decisions rely on the longest prefix match in the routing table, directly affecting the IP data plane. PBR, however, operates outside of this normal routing process. It examines incoming packets and applies pre-defined policies. If a policy rule matches, the packet is forwarded based on that policy instead of the routing table's best path. This enables greater flexibility and control over network traffic flow. For instance, specific types of traffic can be routed through a particular ISP connection or security appliance, even if another path would have been chosen by standard routing. The Forwarding Information Base (FIB), option D, is a hardware-accelerated representation of the routing table used to quickly make forwarding decisions, not a tool to influence routing without changing the table. Nonbroadcast multiaccess (NBMA), option A, is a network type that requires specific configurations, not a means to alter forwarding policy without changing the routing table. Packet switching (option B) is a generic term for how data is routed, not a specific routing policy mechanism.

Further Research:

Cisco Documentation on Policy-Based Routing:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pbr/configuration/15-sy/irp-pbr-15-sy-book.html
Wikipedia on Policy-Based Routing:https://en.wikipedia.org/wiki/Policy-based_routing

Question: 5

Refer to the exhibits. The output of the trace route from R5 shows a loop in the network. Which configuration prevents this loop?

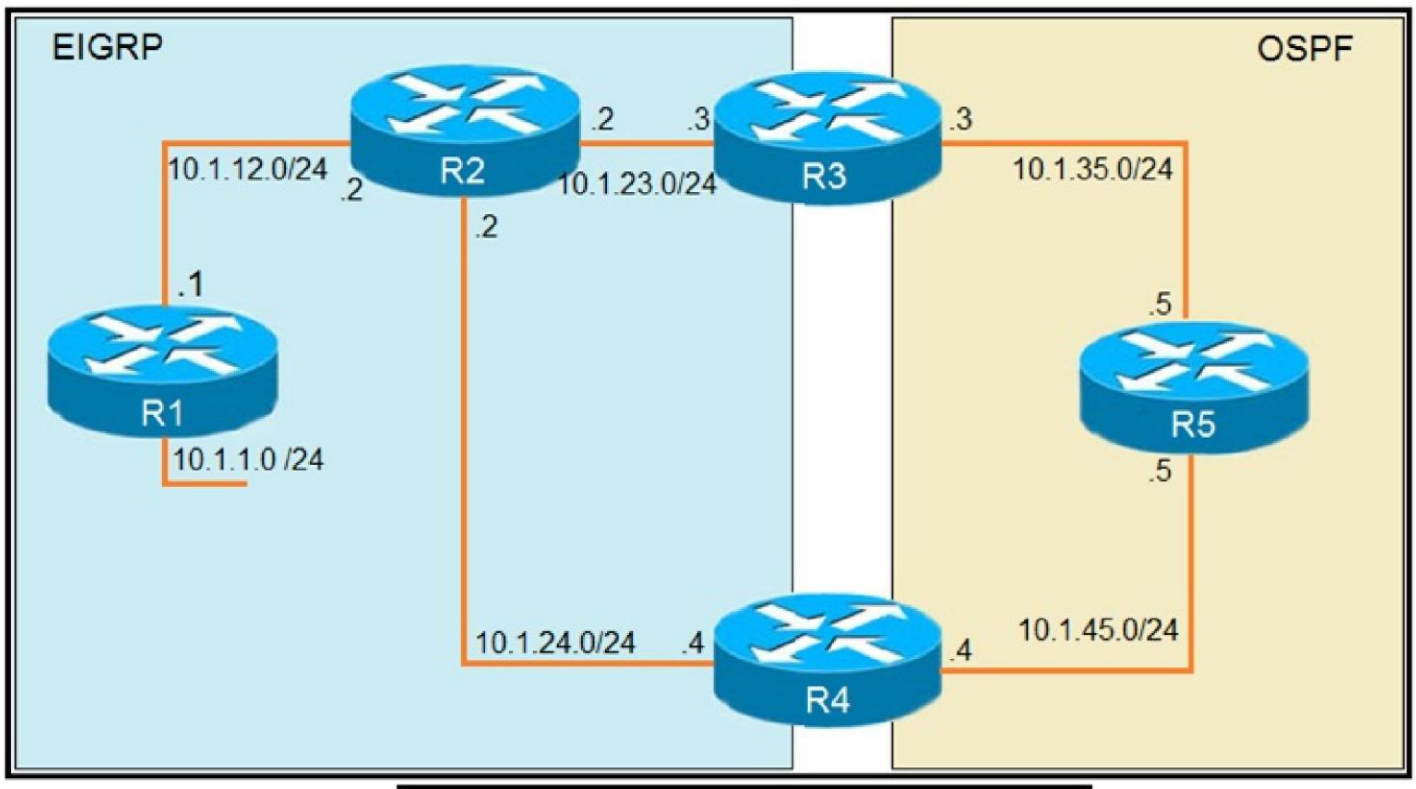


```
R1
router eigrp 1
  redistribute connected
  network 10.1.12.1 0.0.0.0

R3
router ospf 1
  redistribute eigrp 1 subnets
  network 10.1.35.3 0.0.0.0 area 0

R4
router eigrp 1
  redistribute ospf 1 metric 2000000 1 255 1 1500
!
router ospf 1
  network 10.1.45.4 0.0.0.0 area 0

R5#traceroute 10.1.1.1

Type escape sequence to abort.
Tracing the route to 10.1.1.1

 1 10.1.35.3 80 msec 44 msec 20 msec
 2 10.1.23.2 44 msec 104 msec 64 msec
 3 10.1.24.4 44 msec 64 msec 40 msec
 4 10.1.45.5 24 msec 40 msec 20 msec
 5 10.1.35.3 92 msec 144 msec 148 msec
 6 10.1.23.2 108 msec 76 msec 80 msec
       <output truncuated>
```

A.

```
R3
router ospf 1
  redistribute eigrp 1 subnets route-map SET-TAG
!
route-map SET-TAG permit 10
  set tag 1

R4
router eigrp 1
  redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG
!
route-map FILTER-TAG deny 10
  match tag 1
!
route-map FILTER-TAG permit 20
```

B.

```
R3
router eigrp 1
  redistribute OSPF 1 route-map SET-TAG
!
route-map SET-TAG permit 10
  set tag 1

R4
router eigrp 1
  redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG
  network 10.1.24.4 0.0.0.0
!
route-map FILTER-TAG deny 10
  match tag 1
!
route-map FILTER-TAG permit 20
```

C.

```
R3
router ospf 1
  redistribute eigrp 1 subnets route-map SET-TAG
!
route-map SET-TAG permit 10
  set tag 1

R4
router eigrp 1
  redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG
!
route-map FILTER-TAG permit 10
  match tag 1
```

D.

```
R3
router ospf 1
  redistribute eigrp 1 subnets route-map SET-TAG
!
route-map SET-TAG deny 10
  set tag 1

R4
router eigrp 1
  redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG
!
route-map FILTER-TAG deny 10
  match tag 1
```

Answer: A

Explanation:

```
R3
router ospf 1
 redistribute eigrp 1 subnets route-map SET-TAG
!
route-map SET-TAG permit 10
 set tag 1

R4
router eigrp 1
 redistribute ospf 1 metric 2000000 1 255 1 1500 route-map FILTER-TAG
!
route-map FILTER-TAG deny 10
 match tag 1
!
route-map FILTER-TAG permit 20
```

Question: 6

Refer to the exhibit. An engineer configures a static route on a router, but when the engineer checks the route to the destination, a different next hop is chosen.
What is the reason for this?

```
Router#show running-config | include ip route
ip route 192.168.2.2 255.255.255.255 209.165.200.225 130
Router#show ip route

<output omitted>

Gateway of last resort is not set

      192.168.1.0/32 is subnetted, 1 subnets
C        192.168.1.1 is directly connected, Loopback0
      192.168.2.0/32 is subnetted, 1 subnets
O        192.168.2.2[110/11] via 192.168.12.2, 00:52:09, Ethernet0/0
      192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.12.0/24 is directly connected, Ethernet0/0
L        192.168.12.1/32 is directly connected, Ethernet0/0
      209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C        209.165.200.0/24 is directly connected, Ethernet0/1
         209.165.200.226/32 is directly connected, Ethernet0/1
```

A. Dynamic routing protocols always have priority over static routes. B. The metric of the OSPF route is lower than the metric of the static route.

C. The configured AD for the static route is higher than the AD of OSPF. D. The syntax of the static route is not valid, so the route is not considered.

Answer: C

Explanation:

The configured AD for the static route is higher than the AD of OSPF.

Question: 7

Refer to the exhibit. An engineer is trying to generate a summary route in OSPF for network 10.0.0.0/8, but the summary route does not show up in the routing table.
Why is the summary route missing?

```
Router#show ip route
<output omitted>
Gateway of last resort is not set

      192.168.1.0/32 is subnetted, 1 subnets
O           192.168.1.1 [110/11] via 192.168.12.1, 16:56:40, Ethernet0/0
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C           192.168.2.0/24 is directly connected, Loopback0
L           192.168.2.2/32 is directly connected, Loopback0
      192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
C           192.168.3.0/24 is directly connected, Ethernet0/1
L           192.168.3.1/32 is directly connected, Ethernet0/1
      192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C           192.168.12.0/24 is directly connected, Ethernet0/0
L           192.168.12.2/32 is directly connected, Ethernet0/0
Router#show running-config | section ospf
router ospf 1
 summary-address 10.0.0.0 255.0.0.0
 redistribute static subnets
 network 192.168.3.0 0.0.0.255 area 0
 network 192.168.12.0 0.0.0.255 area 0
Router#
```

A. The summary-address command is used only for summarizing prefixes between areas.

B. The summary route is visible only in the OSPF database, not in the routing table.

C. There is no route for a subnet inside 10.0.0.0/8, so the summary route is not generated.

D. The summary route is not visible on this router, but it is visible on other OSPF routers in the same area.

Answer: C

Explanation:

There is no route for a subnet inside 10.0.0.0/8, so the summary route is not generated.

Question: 8

Refer to the exhibit. An engineer is trying to block the route to 192.168.2.2 from the routing table by using the configuration that is shown. The route is still present in the routing table as an OSPF route.
Which action blocks the route?

```
Router#show access-lists
Standard IP access list 1
        10 permit 192.168.2.2 (1 match)
Router#
Router#show route-map
route-map RM-OSPF-DL, permit, sequence 10
  Match clauses:
        ip address (access-lists): 1
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
Router#
Router#show running-config | section ospf
router ospf 1
  network 192.168.1.1 0.0.0.0 area 0
  network 192.168.12.0 0.0.0.255 area 0
  distribute-list route-map RM-OSPF-DL in
Router#|
```

A. Use an extended access list instead of a standard access list.

B. Change sequence 10 in the route-map command from permit to deny. C. Use a prefix list instead of an access list in the route map.

D. Add this statement to the route map: route-map RM-OSPF-DL deny 20.

Answer: B

Explanation:

Change sequence 10 in the route-map command from permit to deny.

Question: 9

What is a prerequisite for configuring BFD?

A. Jumbo frame support must be configured on the router that is using BFD.
B. All routers in the path between two BFD endpoints must have BFD enabled.
C. Cisco Express Forwarding must be enabled on all participating BFD endpoints.
D. To use BFD with BGP, the timers 3 9 command must first be configured in the BGP routing process.

Answer: C

Explanation:

The correct answer is **C. Cisco Express Forwarding must be enabled on all participating BFD endpoints.**

Bidirectional Forwarding Detection (BFD) is a low-overhead, fast-failure detection protocol designed to provide rapid detection of connectivity faults between two forwarding engines. It relies heavily on the underlying forwarding mechanism to quickly process and forward BFD control packets. Cisco Express Forwarding (CEF) is Cisco's primary mechanism for high-speed packet forwarding. It creates an optimized forwarding table which BFD leverages for its rapid processing.

Without CEF enabled, a router must rely on its slower process-switched forwarding, making BFD's fast failure detection capabilities ineffective. BFD requires the underlying infrastructure to respond rapidly to these control packets to detect link failures in milliseconds. Consequently, CEF's ability to quickly forward packets is essential for the rapid detection BFD aims to provide. This dependency makes CEF a mandatory prerequisite for configuring BFD on Cisco routers.

Option A, jumbo frame support, is not a prerequisite for BFD. While jumbo frames might improve the efficiency of data transfer, they don't directly affect BFD's functionality. Option B, requiring all routers in the path to have BFD enabled, is incorrect. BFD is a hop-by-hop protocol, meaning it's configured between two directly connected endpoints, not the entire path. Option D, concerning BGP timers, is also irrelevant. While BFD can be used with BGP to provide faster failover, setting BGP timers is not a prerequisite for configuring BFD. BFD is independent of BGP's configuration.

Authoritative Links:

Cisco Documentation on BFD:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-mt/iap-15-mt-book/iap-bfd.html
Cisco Documentation on CEF:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipswitch/configuration/15-mt/ips-15-mt-book/ips-cef.html

Question: 10

DRAG DROP -
Drag and drop the OSPF adjacency states from the left onto the correct descriptions on the right.
Select and Place:

Init	Each router compares the DBD packets that were received from the other router.
2-way	Routers exchange information with other routers in the multiaccess network.
Down	The neighboring router requests the other routers to send missing entries.
Exchange	The network has already elected a DR and a backup BDR.
ExStart	The OSPF router ID of the receiving router was not contained in the hello message.
Loading	No hellos have been received from a neighbor router.

Answer:

Init	Exchange
2-way	2-way
Down	Loading
Exchange	ExStart
ExStart	Init
Loading	Down

Explanation:

Reference:
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13685-13.html

Question: 11

Refer to the exhibit. R2 is a route reflector, and R1 and R3 are route reflector clients. The route reflector learns the

route to 172.16.25.0/24 from R1, but it does not advertise to R3.
What is the reason the route is not advertised?

```
R1 #show ip bgp summary
BGP router identifier 192.168.1.1, local AS number 65000
<output omitted>
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.2.2     4 65000      28      28       22    0    0 00:21:31        0
R1#show ip bgp
BGP table version is 22, local router ID is 192.168.1.1
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,
Origin codes: i – IGP, e – EGP, ? – incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
*>   172.16.25.0/24   209.165.200.225          0             32768 ?
R1#
```

```
R2 #show ip bgp summary
BGP router identifier 192.168.2.2, local AS number 65000
<output omitted>
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.1.1     4 65000      29      28        3    0    0 00:22:07        1
192.168.3.3     4 65000       7       8        3    0    0 00:02:55        0
R2#show ip bgp
BGP table version is 3, local router ID is 192.168.2.2
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,
Origin codes: i – IGP, e – EGP, ? – incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
* i  172.16.25.0/24    209.165.200.225          0    100      0 ?
R2#
```

```
R3 #show ip bgp summary
BGP router identifier 192.168.3.3, local AS number 65000
BGP table version is 4, main routing table version 4
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.2.2     4 65000       8       7        4    0    0 00:03:08        0
R3#
```

A.R2 does not have a route to the next hop, so R2 does not advertise the prefix to other clients. B.Route reflector setup requires full IBGP mesh between the routers.

C.In route reflector setup, only classful prefixes are advertised to other clients. D.In route reflector setups, prefixes are not advertised from one client to another.

Answer: A

Explanation:

R2 does not have a route to the next hop, so R2 does not advertise the prefix to other clients.

Question: 12

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

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

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

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

A.By default, only internal routes and external type 1 routes are redistributed into BGP B.Only classful networks are redistributed from OSPF to BGP
C.BGP convergence is slow, so the route will eventually be present in the BGP table D.By default, only internal OSPF routes are redistributed into BGP

Answer: D

Explanation:

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

Question: 13

```
R200#show ip bgp summary
BGP router identifier 10.1.1.1, local AS number 65000
BGP table version is 26, main routing table version 26
1 network entries using 132 bytes of memory
1 path entries using 52 bytes of memory
2/1 BGP path/bestpath attribute entries using 296 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Bitfield cache entries: current 1 (at peak 2) using 28 bytes of memory
BGP using 508 total bytes of memory
BGP activity 24/23 prefixes, 24/23 paths, scan interval 60 secs
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.0.2.2       4 65100   20335   20329        0    0    0 00:02:04 Idle (PfxCt)
R200#
```

Refer to the exhibit. In which circumstance does the BGP neighbor remain in the idle condition?

A.if prefixes are not received from the BGP peer
B.if prefixes reach the maximum limit
C.if a prefix list is applied on the inbound direction
D.if prefixes exceed the maximum limit

Answer: D

Explanation:

if prefixes exceed the maximum limit.

Question: 14

Which attribute eliminates LFAs that belong to protected paths in situations where links in a network are connected through a common fiber?

A.shared risk link group-disjoint
B.linecard-disjoint
C.lowest-repair-path-metric
D.interface-disjoint

Answer: A

Explanation:

The correct answer is **A. shared risk link group-disjoint (SRLG-disjoint)**. Here's why:

SRLG-disjointness is a crucial attribute in network resilience, specifically concerning situations where multiple links might fail due to a shared physical resource, like a common fiber optic cable. In the context of Loop-Free Alternates (LFAs), which provide backup paths for fast convergence after a failure, SRLG-disjointness ensures that the LFA path doesn't also rely on the same shared risk. When links share a fiber, a single fiber cut can take down all those links simultaneously. Without SRLG awareness, LFAs might inadvertently select paths that use the same shared resource as the primary path, making them ineffective in a real failure scenario.

By using SRLG-disjoint calculation, LFA algorithms avoid selecting paths that are part of the same SRLG as the primary path. This guarantees that in case of a fiber cut, the backup path is truly independent and available, ensuring high availability and business continuity. The other options are not relevant to solving the issue of common-fiber failures. Linecard-disjoint focuses on different line cards within a router, Lowest-repair-path-metric prioritizes the fastest backup route, and Interface-disjoint might choose interfaces within a single failure group. Only SRLG-disjoint explicitly addresses the scenario described in the question.

In summary, shared risk link groups (SRLGs) are used to identify links that are likely to fail together, and SRLG-disjointness prevents the selection of LFAs that might be affected by the same failures, thereby improving the overall robustness of the network. This method ensures more robust protection against correlated failures stemming from shared infrastructure.

Authoritative Links for further research:

1. Cisco Documentation on SRLG: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/15-mt/iro-ospf-15-mt-book/ospf-srlg.html
2. IETF RFC 5340 - OSPF for IPv6: (This describes some background and concept behind the solution) https://datatracker.ietf.org/doc/html/rfc5340

Question: 15

```
* Jun 28 14:41:57: %BGP-5-ADJCHANGE: neighbor 192.168.2.2 Down User reset
* Jun 28 14:41:57: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.2.2 IPv4 Unicast
topology base removed from session  User reset
* Jun 28 14:41:57: %BGP-5-ADJCHANGE: neighbor 192.168.2.2 Up
R1#show clock
*15:42:00.506 CET Fri Jun 28 2019
```

Refer to the exhibit. An engineer is troubleshooting BGP on a device but discovers that the clock on the device does not correspond to the time stamp of the log entries.
Which action ensures consistency between the two times?

A.Configure the service timestamps log uptime command in global configuration mode.
B.Configure the logging clock synchronize command in global configuration mode.
C.Configure the service timestamps log datetime localtime command in global configuration mode.
D.Make sure that the clock on the device is synchronized with an NTP server.

Answer: C

Explanation:

Configure the service timestamps log datetime localtime command in global configuration mode.

Question: 16

Refer to the exhibit. What is the result of applying this configuration?

```
R1#show policy-map control-plane
  Control Plane
        Service-policy input: CoPP-BGP
          Class-map: BGP (match all)
            2716 packets, 172071 bytes
            5 minute offered rate 0000 bps, drop rate 0000 bps
            Match: access-group name BGP
            drop

          Class-map: class-default (match-any)
            5212 packets, 655966 bytes
            5 minute offered rate 0000 bps, drop rate 0000 bps
            Match: any
```

A.The router can form BGP neighborships with any other device.
B.The router cannot form BGP neighborships with any other device.
C.The router cannot form BGP neighborships with any device that is matched by the access list named BGP.
D.The router can form BGP neighborships with any device that is matched by the access list named BGP.

Answer: C

Explanation:

The router cannot form BGP neighborships with any device that is matched by the access list named BGP.

Question: 17

Which command displays the IP routing table information that is associated with VRF-Lite?

A.show ip vrf
B.show ip route vrf
C.show run vrf
D.show ip protocols vrf

Answer: B

Explanation:

The correct answer is **B. show ip route vrf**.

Here's why:

The command show ip route vrf is specifically designed to display the IPv4 routing table associated with a particular Virtual Routing and Forwarding (VRF) instance within a Cisco device. VRF-Lite, also known as IP VRF, allows you to create multiple independent routing tables within the same router. Each VRF acts like a

separate logical router, maintaining its own routing information. Therefore, to view the routing table associated with a specific VRF, you must use the show ip route vrf command, followed by the VRF name.

Option A, show ip vrf, displays VRF information like the VRF name, route distinguisher, and the associated interface(s), but it does not show the actual IP routing table entries. Option C, show run vrf, is not a valid command. show running-config displays the router's full configuration, including VRF settings, but it will not display the actual routes in the routing table. Option D, show ip protocols vrf, will display routing protocols configured within a specific VRF, but does not display the IP routing table.

In essence, if you wish to see the prefixes, network masks and next hop for packets destined towards specific IP networks inside a VRF, the show ip route vrf command is indispensable. It allows for verification of routing reachability and troubleshooting within separate VRF contexts.

Authoritative Link:

Cisco: Understanding and Configuring VRF (Specifically, refer to the "Verification" section for show commands)

Question: 18

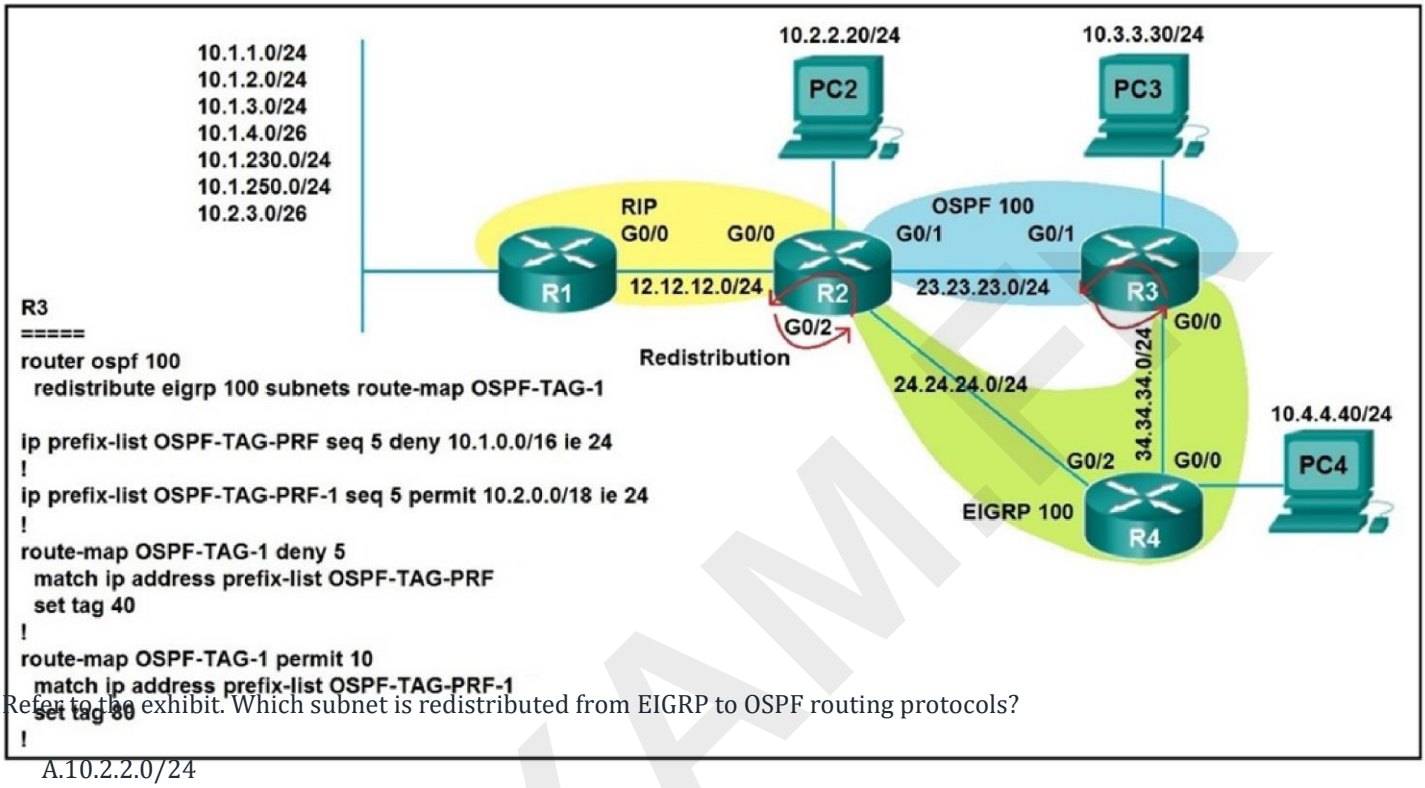

Refer to the exhibit. Which subnet is redistributed from EIGRP to OSPF routing protocols?

A.10.2.2.0/24
B.10.1.4.0/26
C.10.1.2.0/24
D.10.2.3.0/26

Answer: A

Explanation:

Correct answer is A:10.2.2.0/24.

Question: 19

Which configuration adds an IPv4 interface to an OSPFv3 process in OSPFv3 address family configuration?

A.router ospfv3 1 address-family ipv4
B.Router(config-router)#ospfv3 1 ipv4 area 0
C.Router(config-if)#ospfv3 1 ipv4 area 0
D.router ospfv3 1 address-family ipv4 unicast

Answer: C

Explanation:

The correct answer is **C. Router(config-if)#ospfv3 1 ipv4 area 0**. This configuration command is applied under an interface configuration, not under a router configuration. OSPFv3, being a link-state routing protocol, needs to be enabled on an interface for it to participate in the routing process. The ospfv3 1 part indicates the OSPFv3 process ID and ipv4 specifies we're enabling it for IPv4 routing, even within OSPFv3 which generally deals with IPv6. The area 0 associates the interface with a particular OSPF area. Option A would be incorrect as it initiates the address-family subconfiguration under the OSPF router configuration, used for more general address family settings, but not interface activation. Option B attempts to enable OSPFv3 for IPv4 at the router level, which is incorrect. While router-level configurations are used for general OSPFv3 settings, interface participation is enabled at the interface level. Option D is also incorrect as it enters the address family configuration but does not enable the protocol on an interface. Thus, the correct placement of the command is directly under the interface context (config-if).

Authoritative Links for Further Research:

Cisco Documentation on OSPFv3:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/15-sy/iro-15-sy-book/ospf-ipv6.html (Focus on interface configuration aspects within OSPFv3)
Cisco Command Reference:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/command/iro-cr-book/ospf-i1.html#wp3220652266 (Search for "ospfv3" under interface commands to see proper use)

Question: 20

```
R1(config)#route-map ADD permit 20
R1(config-route-map)#set tag 1

R1(config)#router ospf1
R1(config-router)#redistribute rip subnets route-map ADD
```

Refer to the exhibit. Which statement about R1 is true?

A.OSPF redistributes RIP routes only if they have a tag of one.
B.RIP learned routes are distributed to OSPF with a tag value of one.
C.R1 adds one to the metric for RIP learned routes before redistributing to OSPF.
D.RIP routes are redistributed to OSPF without any changes.

Answer: B

Explanation:

RIP learned routes are distributed to OSPF with a tag value of one.

Question: 21

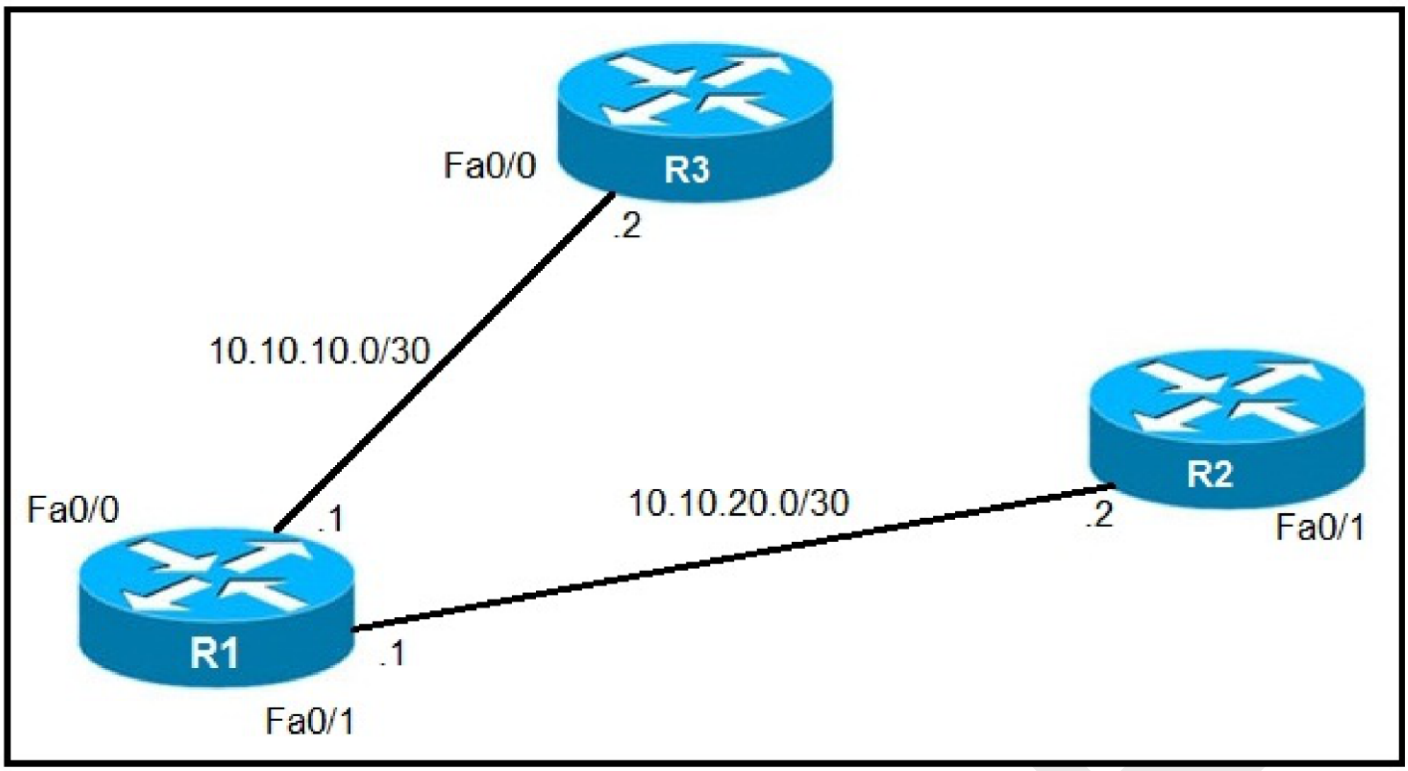

Refer to the exhibit. An IP SLA was configured on router R1 that allows the default route to be modified in the event that Fa0/0 loses reachability with the router R3
Fa0/0 interface. The route has changed to flow through router R2.
Which debug command is used to troubleshoot this issue?

A. debug ip flow
B. debug ip sla error
C. debug ip routing
D. debug ip packet

Answer: C

Explanation:

Correct answer is C:debug ip routing.

Question: 22

Which configuration enables the VRF that is labeled `Inet` on FastEthernet0/0?

A. R1(config)# ip vrf Inet R1(config-vrf)#ip vrf FastEthernet0/0
B. R1(config)#ip vrf Inet FastEthernet0/0
C. R1(config)# ip vrf Inet R1(config-vrf)#interface FastEthernet0/0 R1(config-if)#ip vrf forwarding Inet D. R1(config)#router ospf 1 vrf Inet R1(config-router)#ip vrf forwarding FastEthernet0/0

Answer: C

Explanation:

The correct answer is **C**. Let's break down why.

VRF (Virtual Routing and Forwarding) allows multiple routing tables to exist within a single router. To associate an interface with a VRF, you need to: 1) define the VRF, and 2) link the interface to the VRF. Option A incorrectly attempts to apply the VRF to an interface during VRF definition which is syntactically wrong. Option B incorrectly attempts to associate the VRF to the interface at the global config mode level rather than interface configuration. Option D incorrectly configures OSPF routing under a VRF context.

Option C correctly accomplishes VRF association. First, R1(config)# ip vrf Inet defines the VRF named "Inet." Then, the configuration moves to interface mode with R1(config-vrf)#interface FastEthernet0/0. Finally, R1(config-if)#ip vrf forwarding Inet assigns the FastEthernet0/0 interface to the previously defined "Inet" VRF. The command ip vrf forwarding is the specific command used to associate an interface to a given VRF. This is the correct and necessary step to enable the specified VRF on the chosen interface.

Here are some resources for further study:

1. **Cisco VRF Configuration Guide:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_vrf/configuration/15-mt/iad-vrf-15-mt-book.html
2. **VRF on Cisco Routers**: https://www.geeksforgeeks.org/virtual-routing-and-forwarding-vrf-in-networking/
3. **Implementing VRF:**https://www.practicalnetworking.net/stand-alone/vrf/

Question: 23

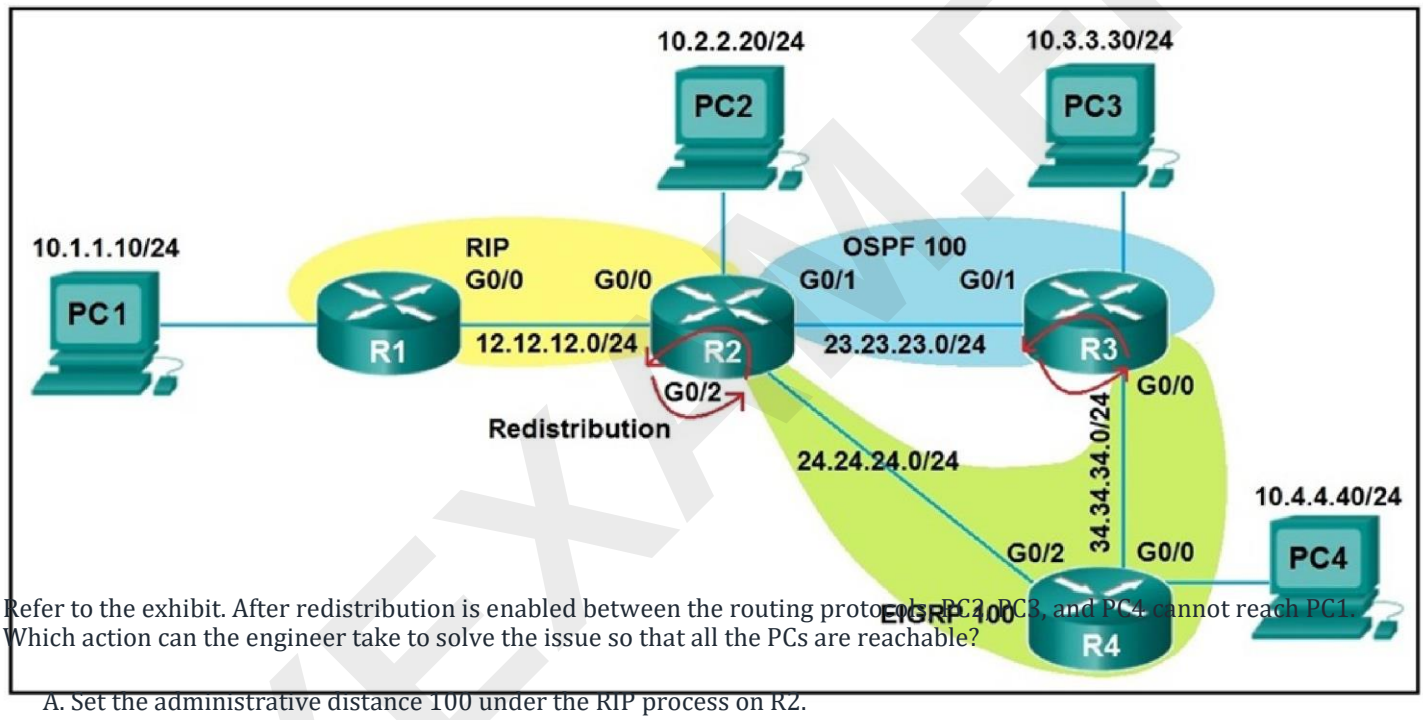

Refer to the exhibit. After redistribution is enabled between the routing protocols, PC2, PC3, and PC4 cannot reach PC1. Which action can the engineer take to solve the issue so that all the PCs are reachable?

A. Set the administrative distance 100 under the RIP process on R2.

B. Filter the prefix 10.1.1.0/24 when redistributed from OSPF to EIGRP.

C. Filter the prefix 10.1.1.0/24 when redistributed from RIP to EIGRP.

D. Redistribute the directly connected interfaces on R2.

Answer: A

Explanation:

Set the administrative distance 100 under the RIP process on R2.

Question: 24

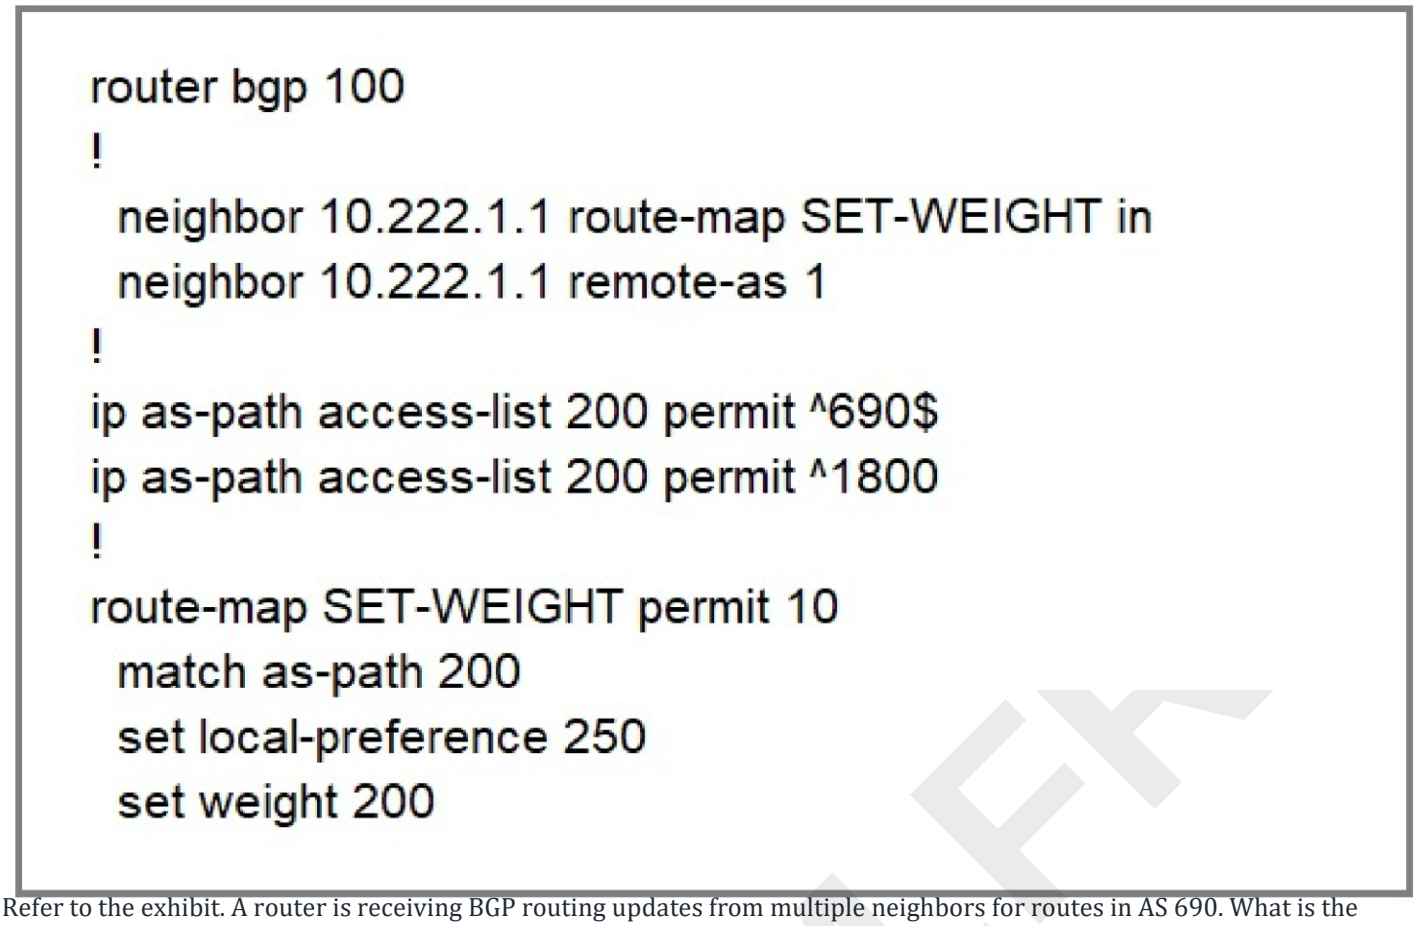

Refer to the exhibit. A router is receiving BGP routing updates from multiple neighbors for routes in AS 690. What is the reason that the router still sends traffic that is destined to AS 690 to a neighbor other than 10.222.1.1?

A. The local preference value in another neighbor statement is higher than 250.

B. The local preference value should be set to the same value as the weight in the route map. C. The route map is applied in the wrong direction.

D. The weight value in another neighbor statement is higher than 200.

Answer: D

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3se/3850/irg-xe-3se-3850-book/irg-prefix-filter.html

Question: 25

```
R1
interface Loopback0
    ip address 172.16.1.1 255.255.255.255
interface FastEthernet0/0
    ip address 192.168.12.1 255.255.255.0
router eigrp 100
    no auto-summary
    network 192.168.12.0
    network 172.16.0.0
    neighbor 192.168.12.2 FastEthernet0/0

R2
interface Loopback0
    ip address 172.16.2.2 255.255.255.255
interface FastEthernet0/0
    ip address 192.168.12.2 255.255.255.0
router eigrp 100
    network 192.168.12.0
    network 172.16.0.0
    neighbor 192.168.12.1 FastEthernet0/0
    passive-interface FastEthernet0/0
```

Refer to the exhibit. R1 and R2 cannot establish an EIGRP adjacency. Which action establishes EIGRP adjacency?

A. Remove the current autonomous system number on one of the routers and change to a different value.

B. Add the passive-interface command to the R1 configuration so that it matches the R2 configuration.

C. Remove the passive-interface command from the R2 configuration so that it matches the R1 configuration.

D. Add the no auto-summary command to the R2 configuration so that it matches the R1 configuration.

Answer: C

Explanation:

Remove the passive-interface command from the R2 configuration so that it matches the R1 configuration.

Question: 26

An engineer configured policy-based routing for a destination IP address that does not exist in the routing table. How is the packet treated through the policy for configuring the set ip default next-hop command?

A. Packets are not forwarded to the specific next hop.

B. Packets are forwarded based on the routing table.

C. Packets are forwarded based on a static route.

D. Packets are forwarded to the specific next hop.

Answer: D

Explanation:

Here's a detailed justification for why option D is the correct answer:

Policy-Based Routing (PBR) allows network administrators to override standard routing table lookups for specific traffic flows. When using the set ip default next-hop command within a PBR policy, you are instructing the router to forward packets matching the policy's criteria to a specified next-hop IP address, regardless of whether that destination IP address exists in the main routing table. This behavior is fundamental to PBR's function: it prioritizes the policy's defined forwarding path over the route selected by the normal routing process.

When a packet matches a PBR policy that includes the set ip default next-hop command, the router directly forwards the packet to that specified next-hop, even if the destination network isn't in the routing table. This is because PBR is meant to enforce policies that may deviate from normal routing paths. The default keyword indicates that even if a more specific route exists within the routing table, it will be ignored in favor of the defined next-hop in the policy. The destination IP address in the packet is not considered when the set ip default next-hop command is present in a PBR policy; rather, the packet is sent to the pre-defined IP. This is the key point of how PBR overrides normal routing decisions. If no PBR policy is matched then routing based on normal routing table lookup process would occur.

Therefore, even if the destination IP address isn't in the routing table, the packet will be forwarded based on the policy, directed to the specified next-hop, making option D ("Packets are forwarded to the specific next hop.") the correct answer. The set ip default next-hop instruction bypasses normal routing lookups, effectively creating a forced forwarding path. Options A, B, and C would be valid if normal routing rules applied.

Authoritative Links:

Cisco Documentation on Policy-Based Routing:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_pbr/configuration/15-mt/irp-pbr-15-mt-book/irp-pbr-oview.html
Cisco Configuration Examples with PBR:https://www.cisco.com/c/en/us/support/docs/ip/policy-routing/116202-technote-pbr-00.html

Question: 27

```
ip prefix-list DefaultRouteOnly seq 5 deny 0.0.0.0/0 le 32
ip prefix-list DefaultRouteOnly seq 10 permit 0.0.0.0/0

router eigrp ccnp
 address-family ipv4 unicast autonomous-system 1
  topology base
   distribute-list prefix DefaultRouteOnly out Tunnel0
```

Refer to the exhibit. The administrator configured route advertisement to a remote low resources router to use only the default route to reach any network but failed.
Which action resolves this issue?

A. Remove the prefix keyword from the distribute-list command. B. Remove the line with the sequence number 10 from the prefix list. C. Change the direction of the distribute-list command from out to in. D. Remove the line with the sequence number 5 from the prefix list.

Answer: D

Explanation:

Remove the line with the sequence number 5 from the prefix list.

Question: 28

```
Ipv6 unicast-routing
!
Router ospfv3 4
    Router-id 192.168.1.1
  !
Interface E 0/0
 Ipv6 enable
 Ip address 10.1.1.1 255.255.255.0
 Ospfv3 4 area 0 ipv4
 No shut
!
Interface Loopback0
 Ipv6 enable
 Ipv4 172.16.1.1 255.255.255.0
 Ospfv3 4 area 0 ipv4
```

Refer to the exhibit. The network administrator configured the branch router for IPv6 on the E 0/0 interface. The neighboring router is fully configured to meet requirements, but the neighbor relationship is not coming up. Which action fixes the problem on the branch router to bring the IPv6 neighbors up?

A. Disable OSPF for IPv4 using the no ospfv3 4 area 0 ipv4 command under the E 0/0 interface.

B. Enable the IPv4 address family under the router ospfv3 4 process by using the address-family ipv4 unicast command.

C. Disable IPv6 on the E 0/0 interface using the no ipv6 enable command.

D. Enable the IPv4 address family under the E 0/0 interface by using the address-family ipv4 unicast command.

Answer: B

Explanation:

Enable the IPv4 address family under the router ospfv3 4 process by using the address-family ipv4 unicast command.

Question: 29

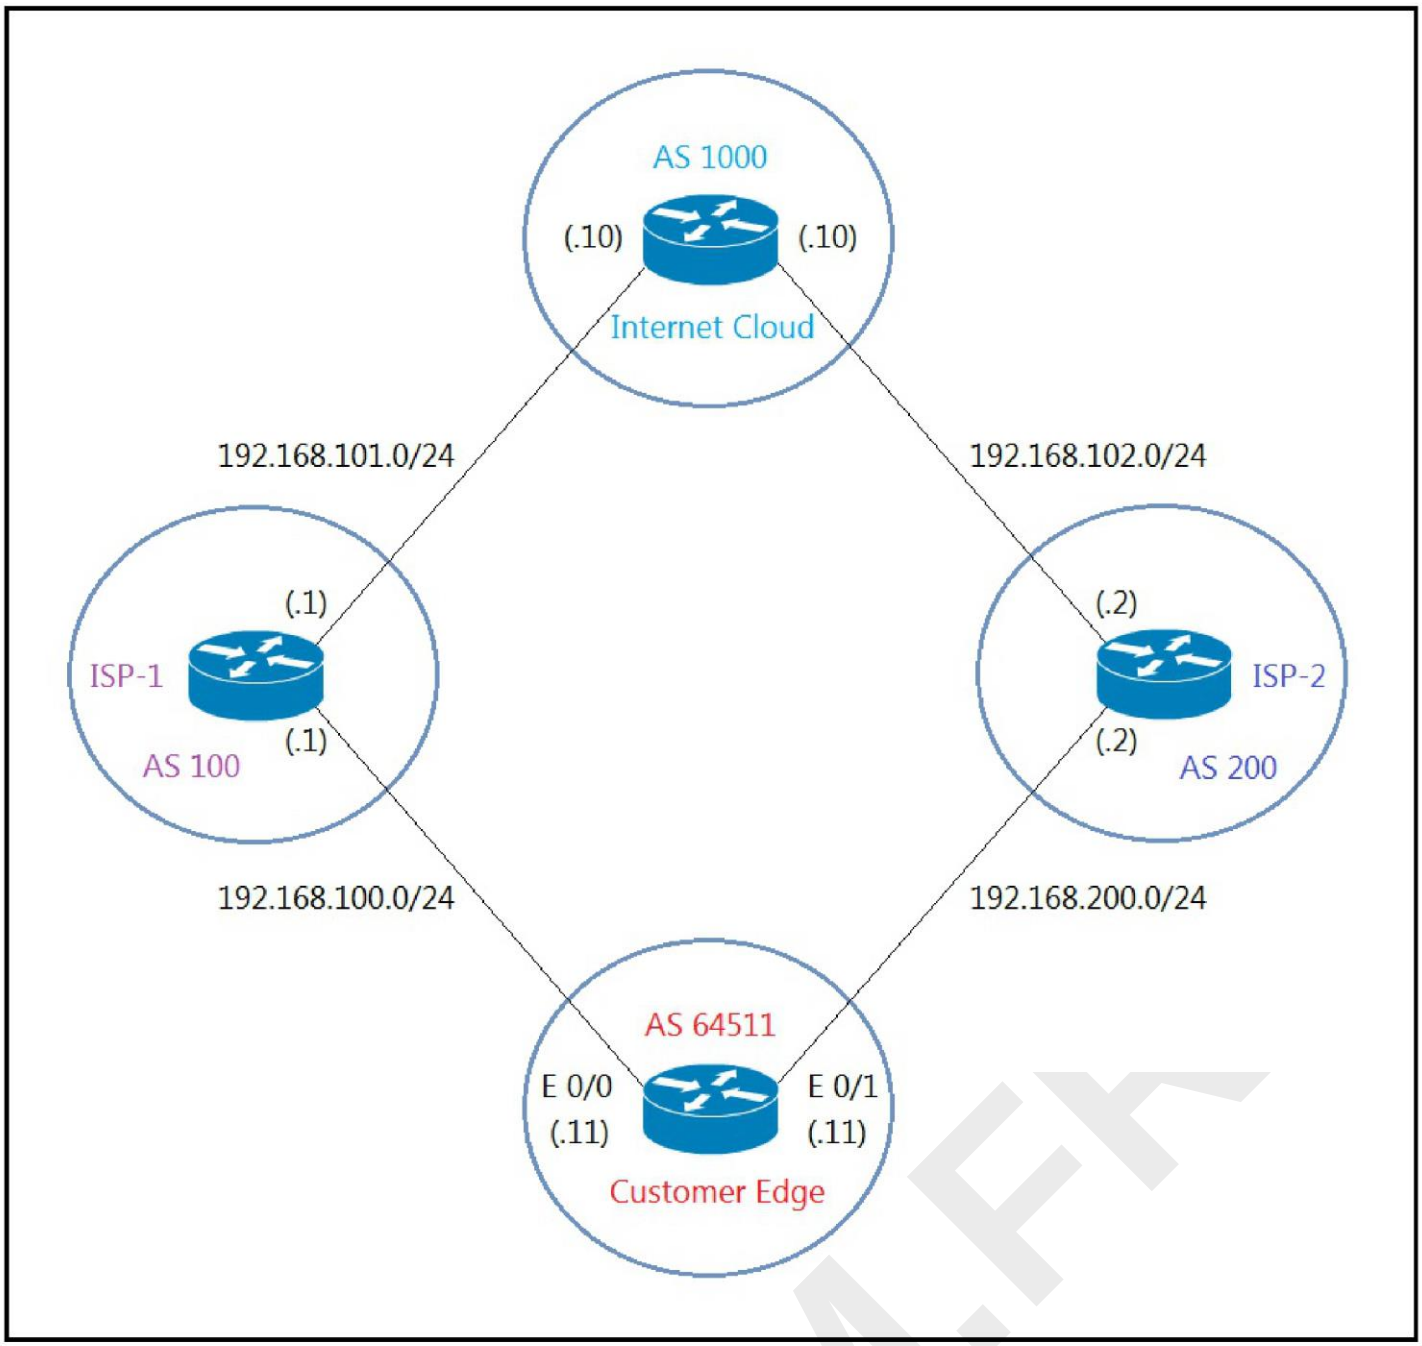

Refer to the exhibit. The network administrator has configured the Customer Edge router (AS 64511) to send only summarized routes toward ISP-1 (AS 100) and
ISP-2 (AS 200).

```
router bgp 64511
 network 172.16.20.0 mask 255.255.255.0
 network 172.16.21.0 mask 255.255.255.0
 network 172.16.22.0 mask 255.255.255.0
 network 172.16.23.0 mask 255.255.255.0
 aggregate-address 172.16.20.0 255.255.252.0
```

After this configuration, ISP-1 and ISP-2 continue to receive the specific routes and the summary route. Which configuration resolves the issue?

A.

```
router bgp 64511
 aggregate-address 172.16.20.0 255.255.252.0 summary-only
```

B.

```
router bgp 64511
 neighbor 192.168.100.1 summary-only
 neighbor 192.168.200.2 summary-only
```

C.

```
ip prefix-list PL_BLOCK_SPECIFIC deny 172.16.20.0/22 ge 22
ip prefix-list PL_BLOCK_SPECIFIC permit 172.16.20.0/22
!
route-map BLOCK_SPECIFIC permit 10
 match ip address prefix-list PL_BLOCK_SPECIFIC
!
router bgp 64511
 aggregate-address 172.16.20.0 255.255.252.0 suppress-map BLOCK_SPECIFIC
```

D.

```
interface E 0/0
 ip bgp suppress-map BLOCK_SPECIFIC
!
interface E 0/1
 ip bgp suppress-map BLOCK_SPECIFIC
!
 ip prefix-list PL_BLOCK_SPECIFIC permit 172.16.20.0/22 ge 24
!
route-map BLOCK_SPECIFIC permit 10
 match ip address prefix-list PL_BLOCK_SPECIFIC
```

Answer: A

Explanation:

```
router bgp 64511
 aggregate-address 172.16.20.0 255.255.252.0 summary-only
```

Question: 30

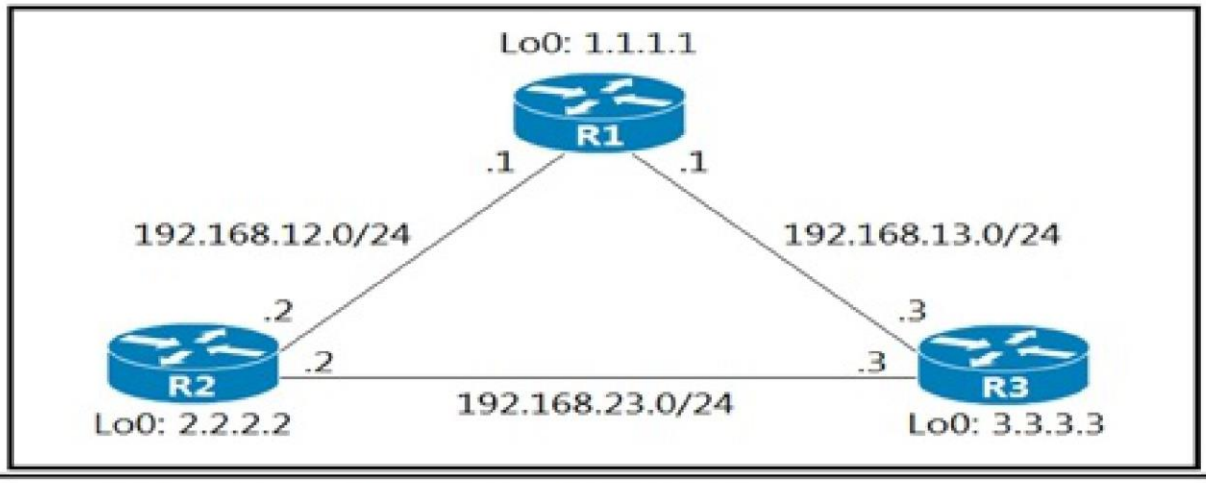


```
R2#show ip protocols | include eigrp|Maximum
Routing Protocol is "eigrp 1"
    Maximum path: 4
    Maximum hopcount 100
    Maximum metric variance 1

R2#show ip eigrp topology 192.168.13.0/24
EIGRP-IPv4 Topology Entry for AS(1)/ID(2.2.2.2) for 192.168.13.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s) FD is 1075200
  Descriptor Blocks
  192.168.23.3 (FastEthernet0/1), from 192.168.23.3, Send flag is 0x0
    Composite metric is (1075200/281600), route is internal
    Vector metric
      Minimum bandwidth is 2500 Kbit
      Total delay is 2000 microseconds
      Reliability is 255/255
      Load is 255/255
      Minimum MTU is 1500
      Hop count is 1
      Originating router is 3.3.3.3
  192.168.12.1 (FastEthernet0/0), from 192.168.12.1, Send flag is 0x0
    Composite metric is (2611200/281600), route is internal
    Vector metric
      Minimum bandwidth is 1000 Kbit
      Total delay is 2000 microseconds
      Reliability is 255/255
      Load is 1/255
      Minimum MTU is 1500
      Hop count is 1
```

Refer to the exhibit. R2 has two paths to reach 192.168.13.0/24, but traffic is sent only through R3. Which action allows traffic to use both paths?

A. Configure the variance 4 command under the EIGRP process on R2.

B. Configure the bandwidth 2000 command under interface FastEthernet0/0 on R2. C. Configure the delay 1 command under interface FastEthernet0/0 on R2.

D. Configure the variance 2 command under the EIGRP process on R2.

Answer: A

Explanation:

Configure the variance 4 command under the EIGRP process on R2.

Question: 31

```
OSPF: Send DBD to 10.100.1.2 on GigabitEthernet0/1 seq 0x9E6 opt
0x52 flag 0x7
   len 32
OSPF: Retransmitting DBD to 10.100.1.2 on GigabitEthernet0/1
[10]
OSPF: Send DBD to 10.100.1.2 on GigabitEthernet0/1 seq 0x9E6 opt
0x52 flag 0x7
   len 32
OSPF: Retransmitting DBD to 10.100.1.2 on GigabitEthernet0/1
[11]
%OSPF-5-ADJCHG: Process 1, Nbr 10.100.1.2 on GigabitEthernet0/1
from EXSTART to
   DOWN, Neighbor Down: Too many retransmissions
```

Refer to the exhibit. The OSPF neighbor relationship is not coming up. What must be configured to restore OSPF neighbor adjacency?

A.matching hello timers
B.OSPF on the remote router
C.use router ID
D.matching mtu values

Answer: D

Explanation:

Correct answer is D:matching mtu values.

Question: 32

An engineer configured two routers connected to two different service providers using BGP with default attributes. One of the links is presenting high delay, which causes slowness in the network.
Which BGP attribute must the engineer configure to avoid using the high-delay ISP link if the second ISP link is up?

A.AS-PATH
B.WEIGHT
C.MED
D.LOCAL_PREF

Answer: D

Explanation:

The correct answer is **D. LOCAL_PREF**. Here's why:

When two routers connect to different ISPs using BGP, and you need to influence path selection to avoid a high-delay link, the most effective attribute to manipulate is **LOCAL_PREF**. LOCAL_PREF is an attribute that is **local to the autonomous system (AS)** and isn't advertised to neighboring ASes. It's used to dictate the preferred path within your network. A higher LOCAL_PREF value means that path is more preferred within your AS.

In this scenario, you would assign a higher LOCAL_PREF value to routes learned from the desired (low-delay) ISP compared to the routes from the high-delay ISP. Therefore, even with default BGP attributes, this preference is only valid within your network. By setting a higher LOCAL_PREF for routes learned from the desired ISP, BGP will choose the more preferred routes when available within your AS.

While AS-PATH (A) might be useful for influencing routing based on the AS hops, and WEIGHT (B) is another local preference mechanism, and MED (C) is intended for use between ASes, but LOCAL_PREF remains the most appropriate solution to enforce path selection inside of your AS and avoids using a link with delay. The other options are not best to do route selection for the purpose of avoiding slowness caused by high delays, rather you should use local pref.

In summary: LOCAL_PREF is the best BGP attribute to use for influencing path selection within your autonomous system to prefer a low-delay ISP link over a high-delay one, as it directly determines internal route preference.

Authoritative Links:

Cisco: Understanding the BGP Path Selection Algorithm (Specifically, look for the "Local Preference" section) Juniper: BGP Route Selection Process (Focus on the section discussing LOCAL_PREF)

Question: 33

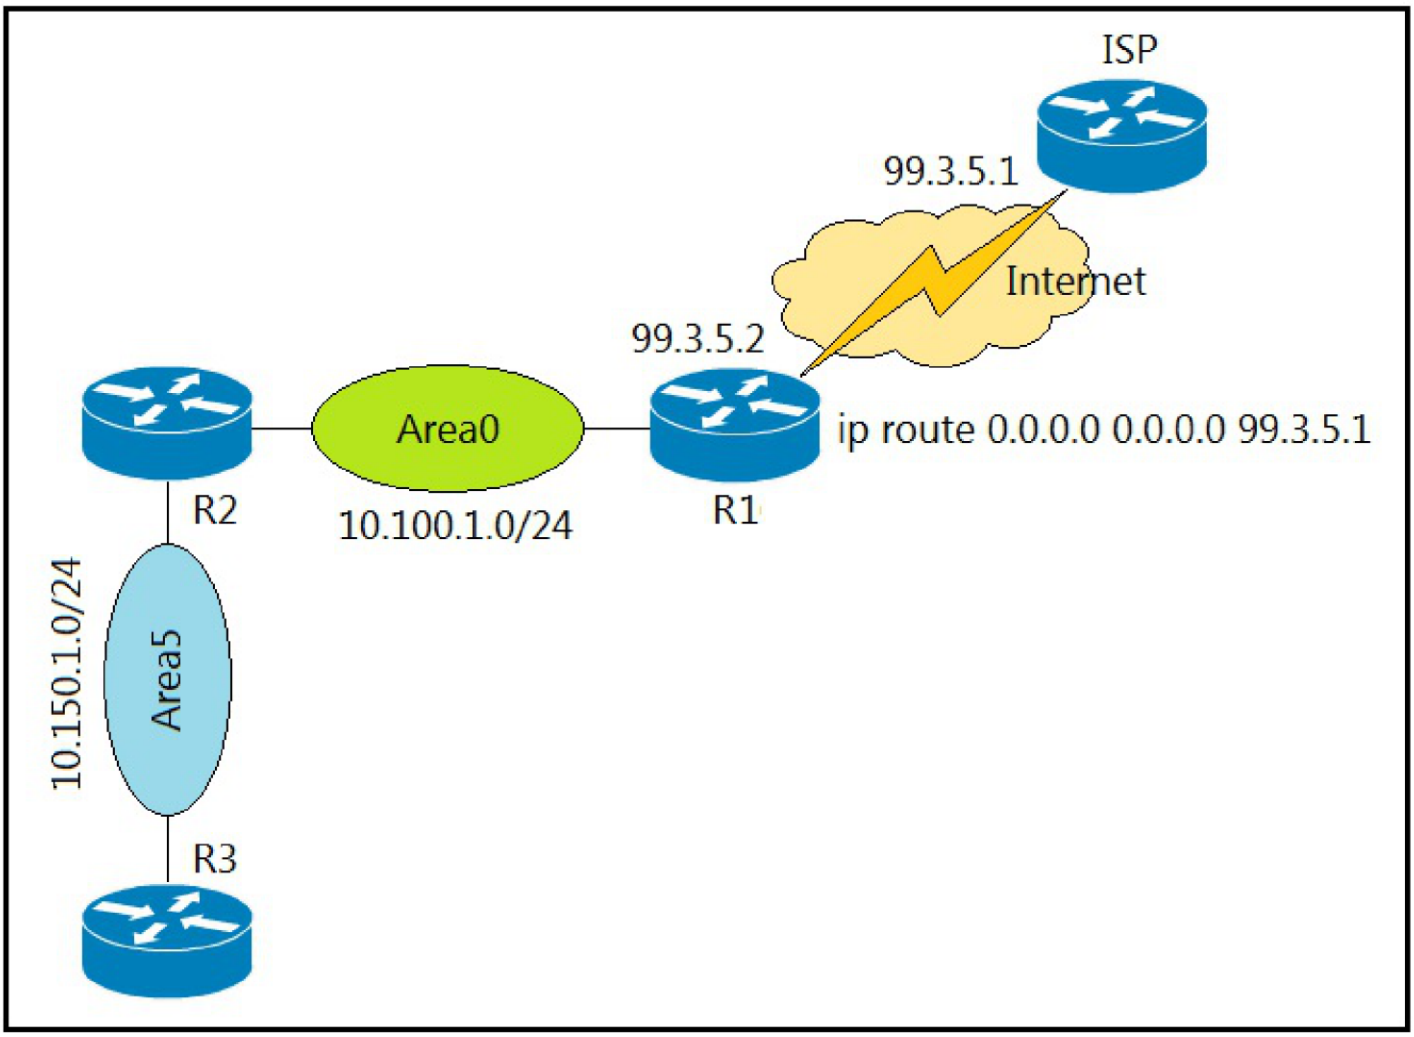

Refer to the exhibit. A network administrator redistributed the default static route into OSPF toward all internal routers to reach to Internet.
Which set of commands restores reachability to the Internet by internal routers?

A.router ospf 1 redistribute static subnets
B.router ospf 1 network 0.0.0.0 0.0.0.0 area 0
C.router ospf 1 redistribute connected 0.0.0.0
D.router ospf 1 default-information originate

Answer: D

Explanation:

router ospf 1 default-information originate.

Question: 34

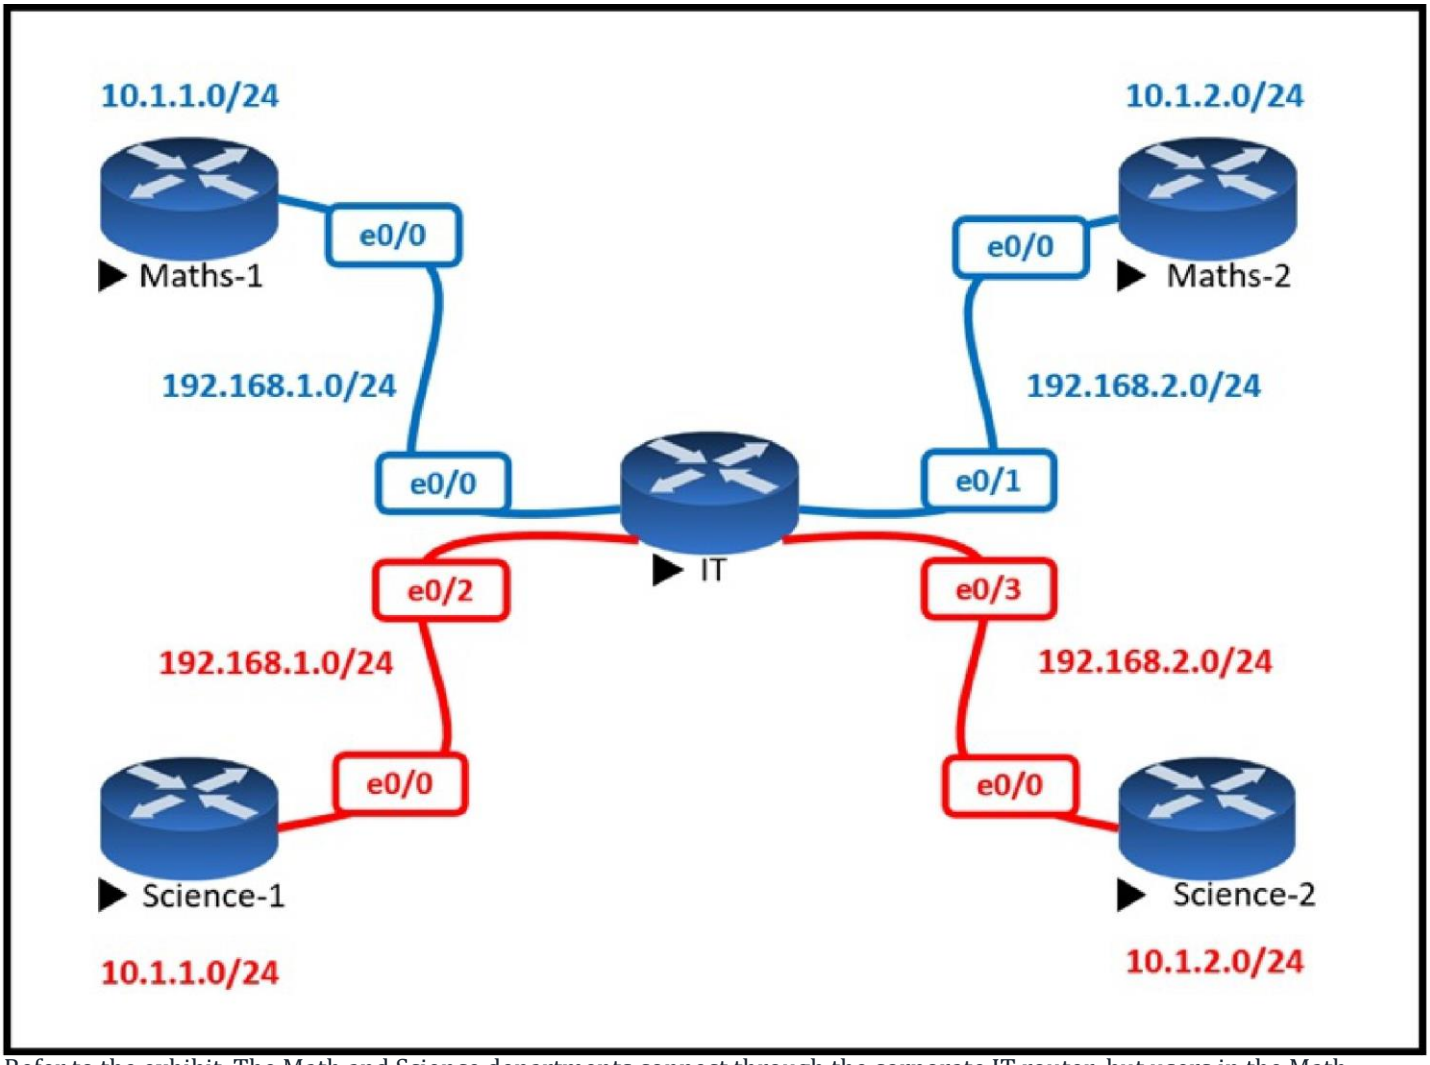

Refer to the exhibit. The Math and Science departments connect through the corporate IT router, but users in the Math department must not be able to reach the
Science department and vice versa.
Which configuration accomplishes this task?

A.vrf definition Science address-family ipv4 ! interface E 0/2 ip address 192.168.1.1 255.255.255.0 no shut ! interface E 0/3 ip address 192.168.2.1 255.255.255.0 no shut

B.vrf definition Science address-family ipv4 ! interface E 0/2 vrf forwarding Science ip address 192.168.1.1 255.255.255.0 no shut ! interface E 0/3 vrf forwarding Science ip address 192.168.2.1 255.255.255.0 no shut

C.vrf definition Science address-family ipv4 ! interface E 0/2 ip address 192.168.1.1 255.255.255.0 vrf forwarding Science no shut ! interface E 0/3 ip address 192.168.2.1 255.255.255.0 vrf forwarding Science no shut

D.vrf definition Science ! interface E 0/2 ip address 192.168.1.1 255.255.255.0 no shut ! interface E 0/3 ip address 192.168.2.1 255.255.255.0 no shut

Answer: B

Explanation:

vrf definition Science address-family ipv4 ! interface E 0/2 vrf forwarding Science ip address 192.168.1.1 255.255.255.0 no shut ! interface E 0/3 vrf forwarding Science ip address 192.168.2.1 255.255.255.0 no shut.

Question: 35

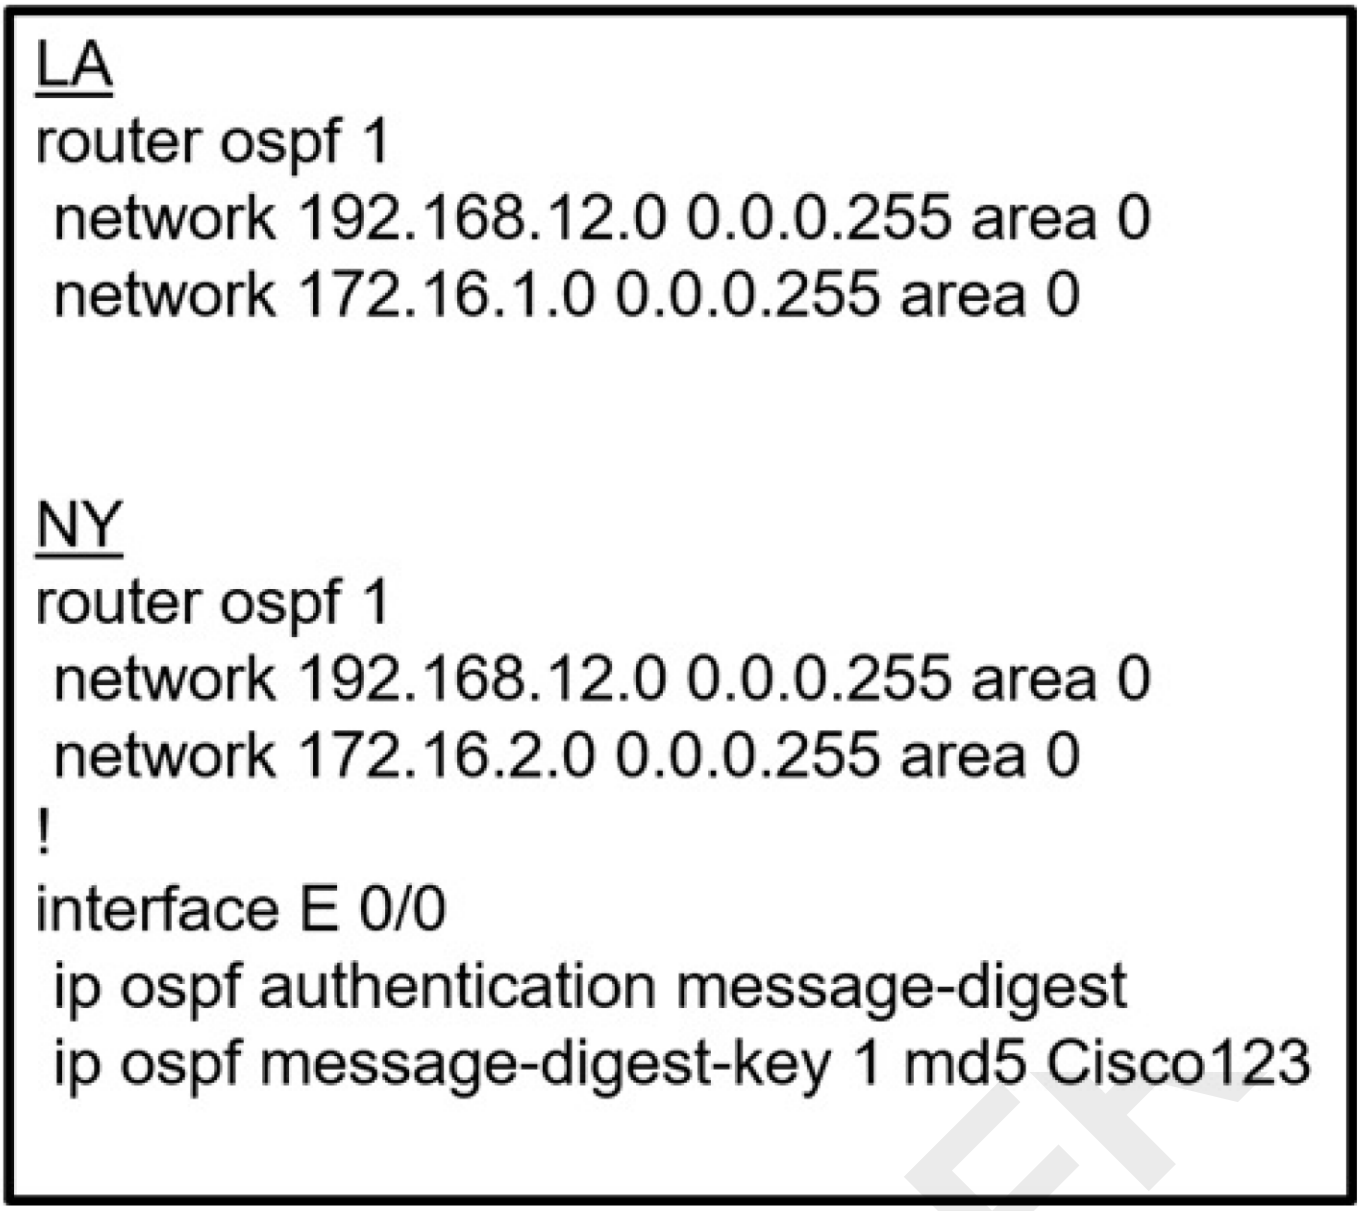

Refer to the exhibit. The neighbor relationship is not coming up. Which two configurations bring the adjacency up? (Choose two.)

A.LA interface E 0/0 ip ospf authentication-key Cisco123
B.NY interface E 0/0 no ip ospf message-digest-key 1 md5 Cisco123 ip ospf authentication-key Cisco123 C.LA interface E 0/0 ip ospf message-digest-key 1 md5 Cisco123
D.LA router ospf 1 area 0 authentication message-digest
E.NY router ospf 1 area 0 authentication message-digest

Answer: CD

Explanation:

C.LA interface E 0/0 ip ospf message-digest-key 1 md5 Cisco123.

D.LA router ospf 1 area 0 authentication message-digest.

Question: 36

```
router ospf 1
 redistribute eigrp 1 subnets route-map EIGRP->OSPF
!
router eigrp 1
 network 10.0.106.0 0.0.0.255
!
route-map EIGRP->OSPF permit 10
 match ip address WAN_PREFIXES
route-map EIGRP->OSPF permit 20
 match ip address LOCAL_PREFIXES
route-map EIGRP->OSPF permit 30
 match ip address VPN_PREFIXES
!
ip prefix-list LOCAL_PREFIXES seq 5 permit 172.16.0.0/12 le 24
ip prefix-list VPN_PREFIXES seq 5 permit 192.168.0.0/16 le 24
ip prefix-list WAN_PREFIXES seq 5 permit 10.0.0.0/8 le 24
!
```

Refer to the exhibit. The network administrator configured redistribution on an ASBR to reach to all WAN networks but failed. Which action resolves the issue?

A.The route map EIGRP->OSPF must have the 10.0.106.0/24 entry to exist in one of the three prefix lists to pass B.EIGRP must redistribute the 10.0.106.0/24 route instead of using the network statement
C.The OSPF process must have a metric when redistributing prefixes from EIGRP
D.The route map must have the keyword prefix-list to evaluate the prefix list entries

Answer: D

Explanation:

The route map must have the keyword prefix-list to evaluate the prefix list entries.

Question: 37

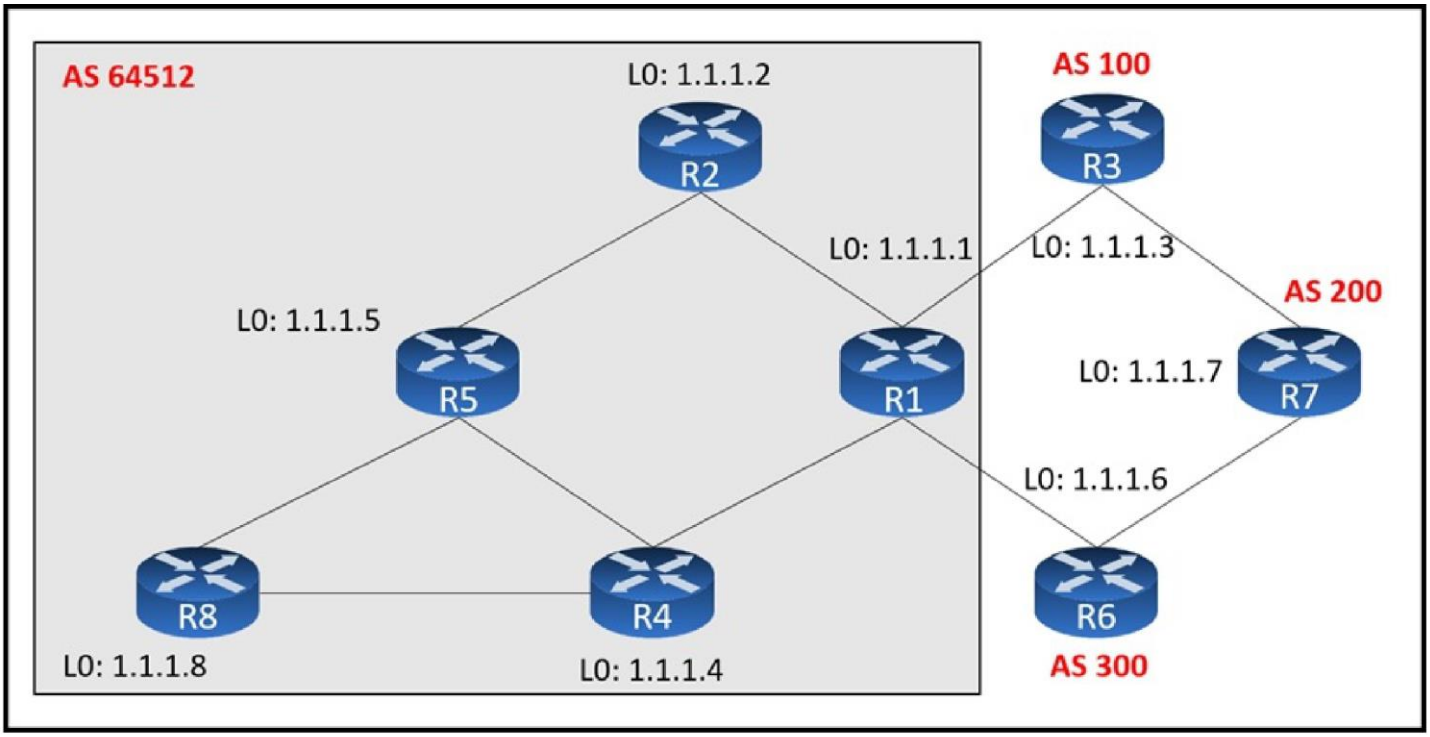

Refer to the exhibit. An engineer configured R2 and R5 as route reflectors and noticed that not all routes are sent to R1 to advertise to the eBGP peers.
Which iBGP routers must be configured as route reflectors to advertise all routes to restore reachability across all networks?

A.R1 and R4
B.R1 and R5
C.R4 and R5
D.R2 and R5

Answer: C

Explanation:

Correct answer is C:R4 and R5.

Question: 38

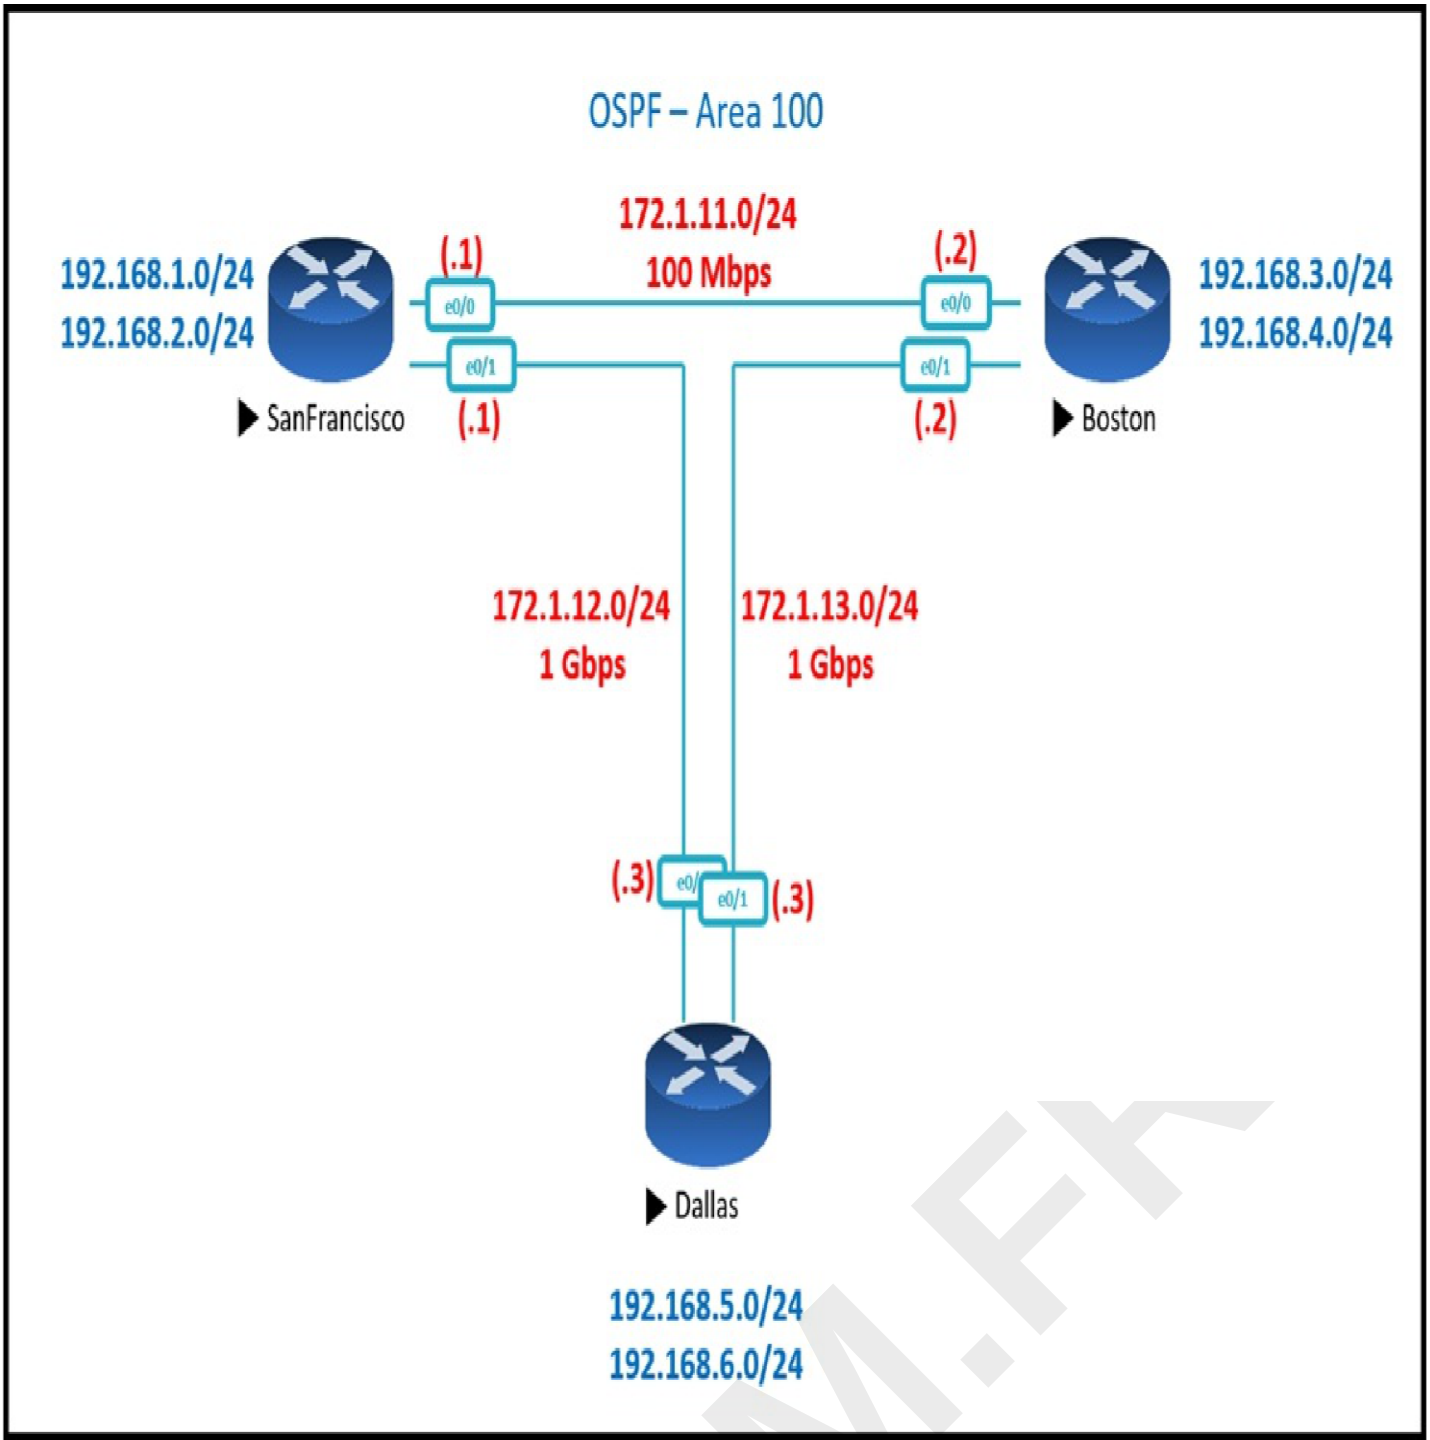

Show IP Route – San Francisco Router

Gateway of last resort is not set

```
     172.1.0.0/16 is variably subnetted, 5 subnets, 2 masks
C       172.1.11.0/24 is directly connected, Ethernet0/0
L     172.1.11.1/32 is directly connected, Ethernet0/0
C       172.1.12.0/24 is directly connected, Ethernet0/0
L     172.1.12.1/32 is directly connected, Ethernet0/0
O       172.1.13.0/24 [110/11] via 172.1.11.2, 00:02:34, Ethernet0/0
```

Refer to the exhibits. SanFrancisco and Boston routers are choosing slower links to reach each other despite the direct links being up.
Which configuration fixes the issue?

A.All Routers router ospf 1 auto-cost reference-bandwidth 100
B.SanFrancisco Router router ospf 1 auto-cost reference-bandwidth 1000
C.Boston Router router ospf 1 auto-cost reference-bandwidth 1000 D.All Routers router ospf 1 auto-cost reference-bandwidth 1000

Answer: D

Explanation:

All Routers router ospf 1 auto-cost reference-bandwidth 1000.

Question: 39

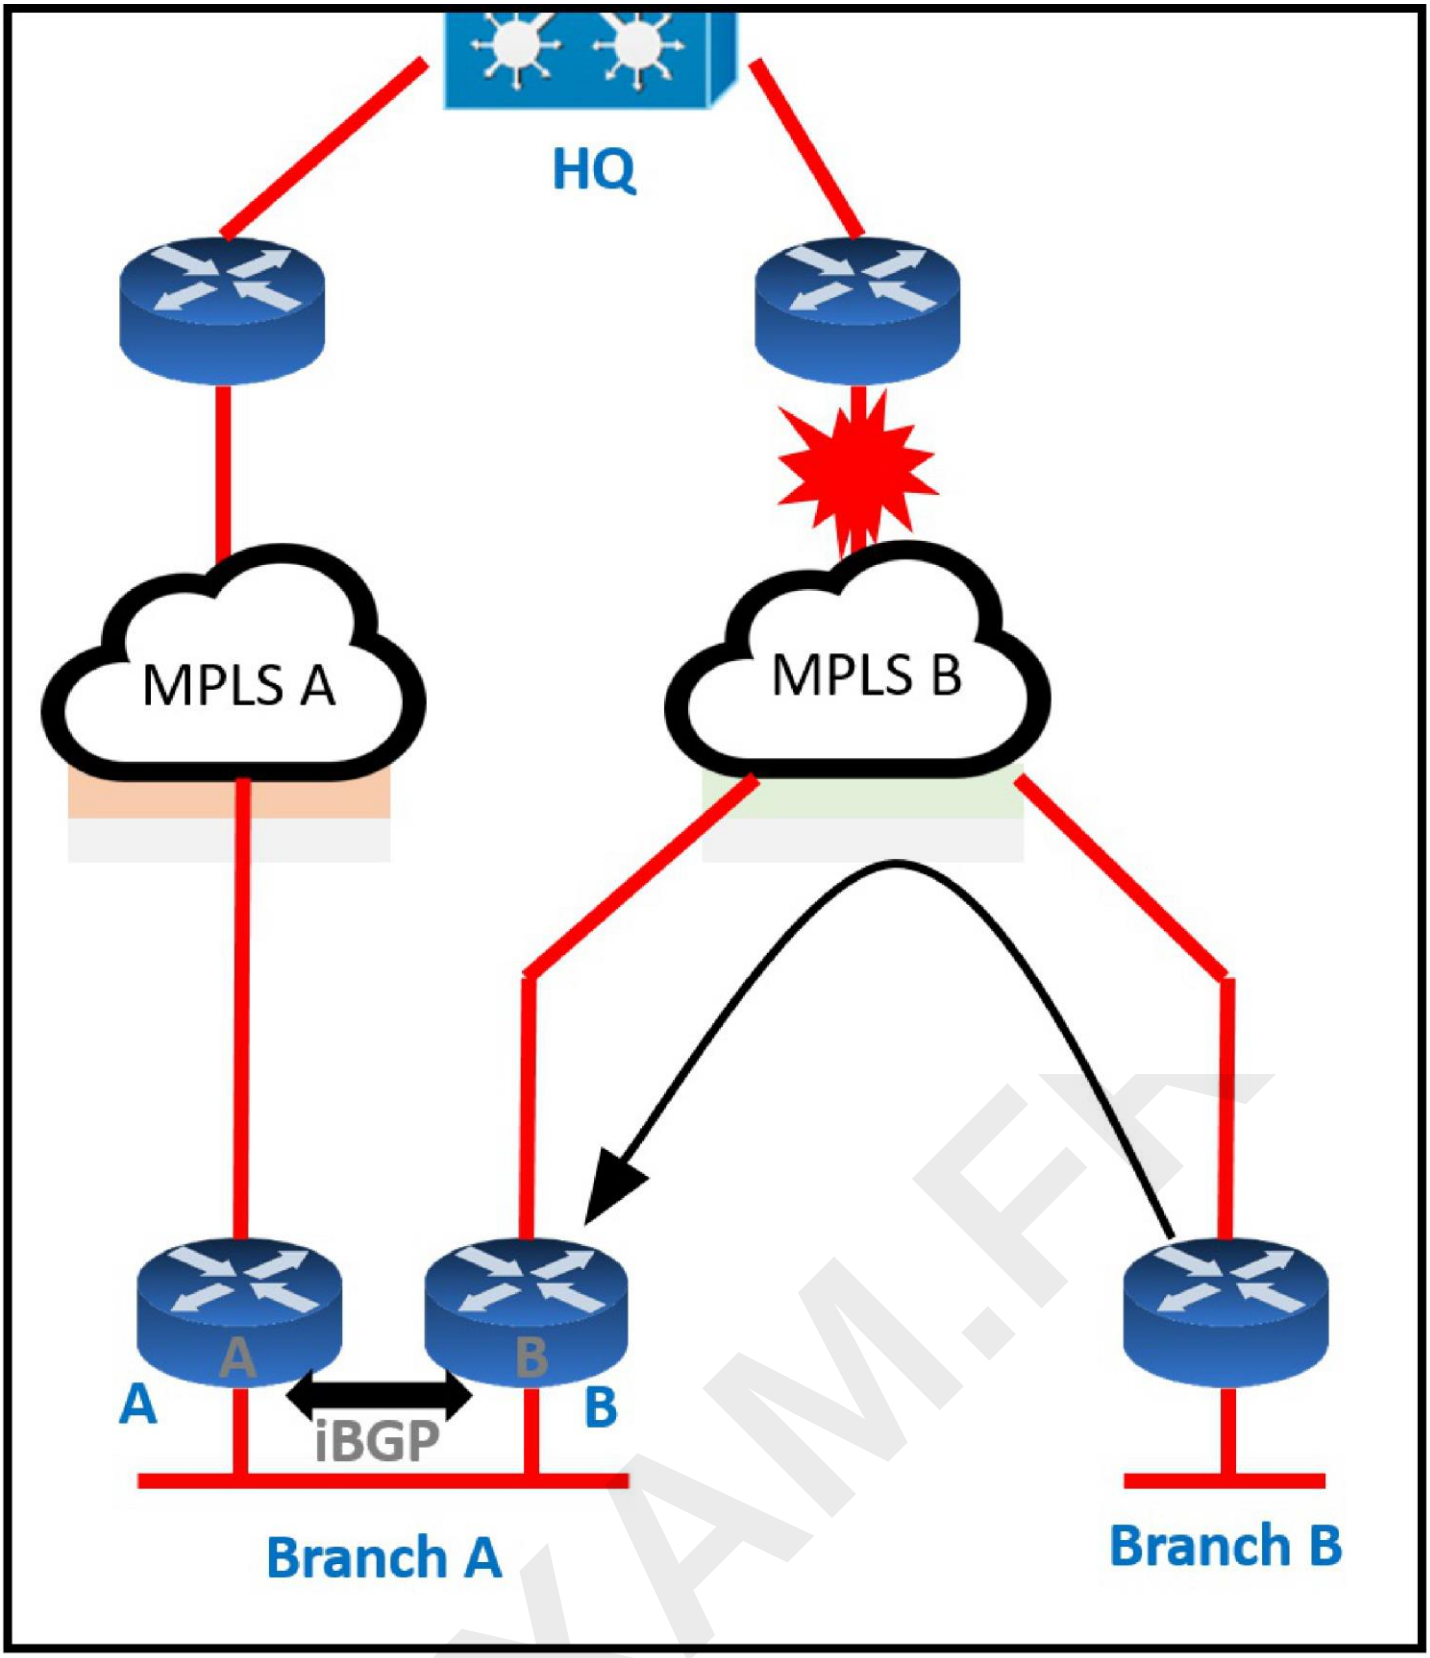

Refer to the exhibit. Troubleshoot and ensure that branch I' only ever uses the MPLS I' network to reach HQ. Which action achieves this requirement?

A.Introduce AS path prepending on the branch A MPLS I' network connection so that any HQ advertisements from branch A toward the MPLS I' network are prepended three times

B.Modify the weight of all HQ prefixes received at branch I' from the MPLS I' network to be higher than the weights used on the MPLS A network

C.Increase the local preference for all HQ prefixes received at branch I' from the MPLS I' network to be higher than the local preferences used on the MPLS A network

D.Introduce an AS path filter on branch A routers so that only local prefixes are advertised into BGP

Answer: D

Explanation:
Introduce an AS path filter on branch A routers so that only local prefixes are advertised into BGP.

Question: 40

Router Configuration:

```
router ospf 0.0.0.0
 network 2.0.0.0 0.255.255.255 area 0.0.0.0
!
router bgp 100
 redistribute ospf 0.0.0.0
!
neighbor 3.3.3.2 remote-as 200
!
end
```

Router# show ip route

```
     2.0.0.0/24 is subnetted, 1 subnets
C       2.2.2.0 is directly connected, Ethernet0/0
C   3.0.0.0/8 is directly connected, Serial1/0
O E2 200.1.1.0/24 [110/20] via 2.2.2.2, 00:16:17, Ethernet 0/0
O E1 200.2.2.0/24 [110/104] via 2.2.2.2, 00:00:41, Ethernet 0/0
     131.108.0.0/24 is subnetted, 2 subnets
O       131.108.2.0 [110/74] via 2.2.2.2, 00:16:17, Ethernet 0/0
O IA    131.108.1.0 [110/74] via 2.2.2.2, 00:16:17, Ethernet 0/0
```

Router# show ip bgp

```
  Network          Next Hop            Metric LocPrf Weight Path
*> 2.2.2.0/24      0.0.0.0                  0         32768 ?
*> 131.108.1.0/24  2.2.2.2                    84        32768 ?
*> 131.108.2.0/24  2.2.2.2                    74        32768 ?
```

Refer to the exhibit. The OSPF routing protocol is redistributed into the BGP routing protocol, but not all the OSPF routes are distributed into BGP.
Which action resolves the issue?

A.Include the word external in the redistribute command
B.Use a route-map command to redistribute OSPF external routes defined in an access list
C.Include the word internal external in the redistribute command
D.Use a route-map command to redistribute OSPF external routes defined in a prefix list

Answer: C

Explanation:

Include the word internal external in the redistribute command.

Question: 41

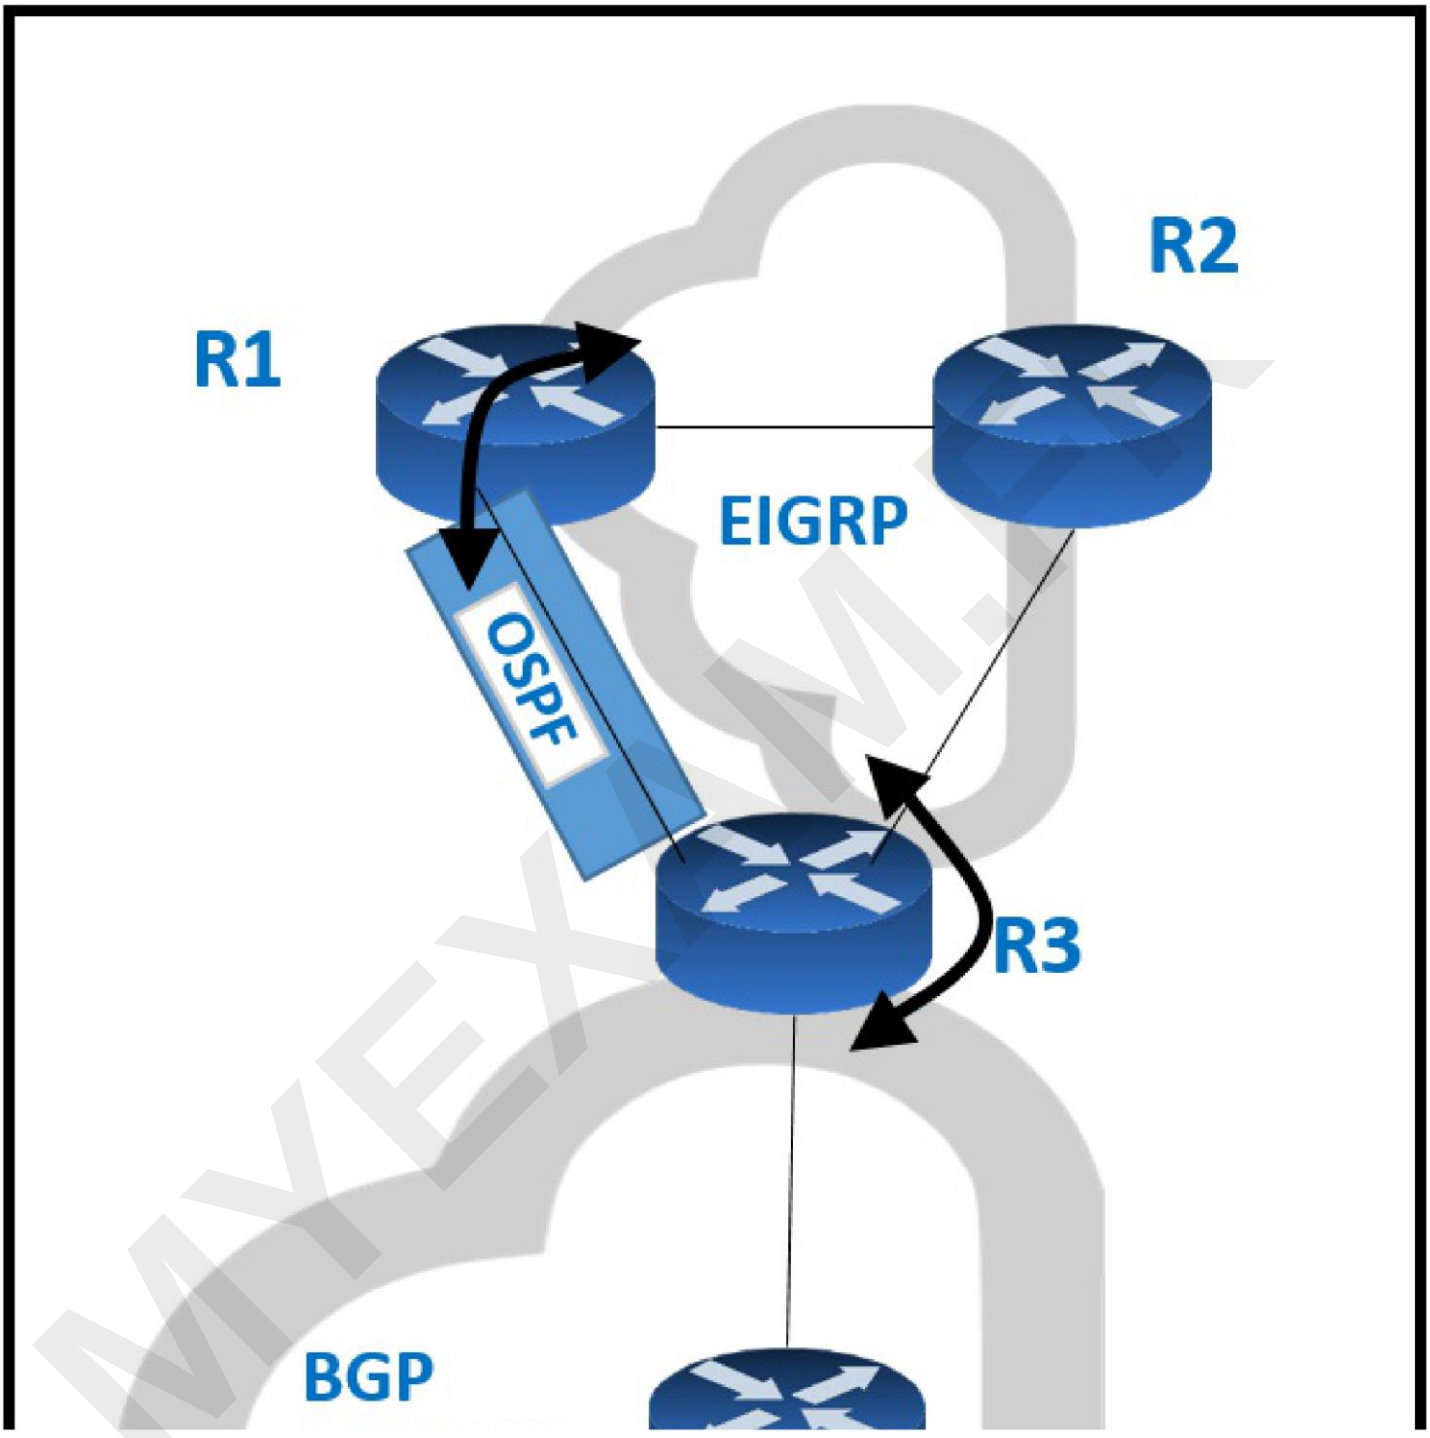

Refer to the exhibit. Routing protocols are mutually redistributed on R3 and R1. Users report intermittent connectivity to services hosted on the 10.1.1.0/24 prefix.
Significant routing update changes are noticed on R3 when the show ip route profile command is run. How must the services be stabilized?

A.The routing loop must be fixed by reducing the admin distance of OSPF from 110 to 80 on R3
B.The routing loop must be fixed by reducing the admin distance of iBGP from 200 to 100 on R3
C.The issue with using BGP must be resolved by using another protocol and redistributing it into EIGRP on R3 D.The issue with using iBGP must be fixed by running eBGP between R3 and R4

Answer: B

Explanation:

The routing loop must be fixed by reducing the admin distance of iBGP from 200 to 100 on R3.

Question: 42

When determining if a system is capable of support, what is the minimum time spacing required for a BFD control packet to receive once a control packet is arrived?

A. Desired Min TX Interval
B. Detect Mult
C. Required Min RX Interval
D. Required Min Echo RX Interval

Answer: C

Explanation:

The correct answer is C, Required Min RX Interval. Bidirectional Forwarding Detection (BFD) relies on the

regular exchange of control packets to detect link failures. When a BFD session is established, both endpoints negotiate parameters, one of which is the "Required Min RX Interval". This parameter specifies the minimum interval at which a BFD-enabled router is prepared to receive a BFD control packet from its peer. If a peer sends packets slower than this interval, it could be flagged as down. In contrast, "Desired Min TX Interval" (A) refers to the interval at which a router sends BFD packets, not how frequently it expects to receive them.

"Detect Mult" (B) is a multiplier that, when combined with the "Required Min RX Interval", determines the maximum time before a link is considered down if no BFD packets are received. "Required Min Echo RX Interval" (D) is relevant when BFD echo mode is used, which is different from the typical BFD control packets. It deals with the minimum interval when receiving echo packets, not standard control packets. Therefore, the 'Required Min RX Interval' is the critical parameter defining the minimum acceptable reception rate of BFD control packets, hence establishing if a system is meeting its necessary support conditions. It is fundamental in determining link stability and quickly detecting network outages. The system must receive BFD control packets as a minimum at this interval to be considered operationally healthy by its neighbor.

Authoritative links:

1. Cisco Documentation on BFD: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-mt/iap-15-mt-book/ipapp-bfd.html
2. Juniper Documentation on BFD: https://www.juniper.net/documentation/us/en/software/junos/routing/topics/concept/bfd-overview.html

Question: 43

An engineer is configuring a network and needs packets to be forwarded to an interface for any destination address that is not in the routing table.
What should be configured to accomplish this task?

A.set ip next-hop

B.set ip default next-hop

C.set ip next-hop recursive

D.set ip next-hop verify-availability

Answer: B

Explanation:

The correct answer is **B. set ip default next-hop**. This command is specifically designed to configure a default gateway, which is crucial for routing traffic destined for networks not explicitly defined in the routing table.

When a router receives a packet with a destination IP address that doesn't match any known routes, it consults the default gateway. If configured, the router forwards the packet to the interface/IP specified by set ip default next-hop. This ensures that traffic intended for unknown destinations doesn't get dropped, but instead, gets sent to a predetermined path, often a router connected to an internet or core network.

Option A, set ip next-hop, is used for configuring the next hop IP address for specific static routes, not for a default gateway applicable to all unknown destinations. Option C, set ip next-hop recursive, is used with static routes where the next hop IP address is not directly connected but is resolved through another hop, and also does not address a default route scenario. Option D, set ip next-hop verify-availability, is not a valid command in the context of routing configurations; while it implies verification of next-hop availability, such functionality is handled through mechanisms like routing protocol adjacencies or tracking objects. Therefore, only set ip default next-hop fulfills the requirement of forwarding packets for unknown destinations by setting up a default gateway.

For further research on default routing, refer to Cisco's documentation on IP routing configuration:

Cisco IP Routing Configuration Guide
Understanding IP Routing

Question: 44

What is an advantage of using BFD?

A.It detects local link failure at layer 1 and updates the routing table.
B.It detects local link failure at layer 3 and updates the routing protocols.
C.It has sub-second failure detection for layer 1 and layer 3 problems.
D.It has sub-second failure detection for layer 1 and layer 2 problems.

Answer: B

Explanation:

The correct answer is **B. It detects local link failure at layer 3 and updates the routing protocols.**

Bidirectional Forwarding Detection (BFD) is a network protocol designed for rapid failure detection between two forwarding engines. It operates at the application layer, facilitating quick identification of communication issues, primarily for layer 3 protocols like routing protocols (OSPF, EIGRP, BGP). BFD doesn't directly interact with layer 1 (physical) or layer 2 (data link) hardware; instead, it relies on layer 3 connectivity to establish and maintain sessions with neighboring devices. When a failure occurs, BFD swiftly signals the routing protocol, enabling faster convergence. This means the routing protocol can recalculate routes much quicker compared to relying solely on the protocol's timers alone, improving network resilience. Options A, C, and D are inaccurate because BFD doesn't directly detect layer 1 or layer 2 issues. While it might react to such issues indirectly as they impact layer 3 communication, its detection mechanism is centered around layer 3 reachability. BFD's primary benefit is not direct layer 1 detection or layer 2 detection, instead, it focuses on efficient layer 3 protocol fault notification.

Authoritative Links for further research:

1. **Cisco Documentation on BFD:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/15-mt/ipapp-15-mt-book/ipapp-bidir-forward.html
2. **RFC 5880 - Bidirectional Forwarding Detection (BFD):**https://datatracker.ietf.org/doc/html/rfc5880

Question: 45

An engineer needs dynamic routing between two routers and is unable to establish OSPF adjacency. The output of the show ip ospf neighbor command shows that the neighbor state is EXSTART/EXCHANGE.
Which action should be taken to resolve this issue?

A. match the passwords
B. match the hello timers
C. match the MTUs
D. match the network types

Answer: C

Explanation:

The OSPF neighbor state stuck in EXSTART/EXCHANGE indicates a problem with Database Description (DBD) packet negotiation. This phase occurs after the initial Hello packets are exchanged successfully, meaning basic connectivity and OSPF process configuration are likely correct. During EXSTART/EXCHANGE, routers establish a master/slave relationship and exchange DBD packets to identify which Link State Advertisements (LSAs) they need. Mismatched MTU (Maximum Transmission Unit) sizes often prevent the routers from successfully exchanging these larger DBD packets, leading to this stuck state. Larger DBD packets are fragmented or not properly reassembled if they exceed the interface MTU. Passwords and hello timers influence the initial neighbor discovery; if they are mismatched, the neighbor state wouldn't even progress to EXSTART. Similarly, incorrect network types might prevent the neighbor relationship from forming initially.

Therefore, matching MTUs on both interfaces involved in the OSPF adjacency is the most pertinent action. Resolving this MTU issue will allow successful DBD negotiation and move the OSPF neighbor relationship to a more advanced state like FULL.

Further Research:

Cisco Documentation on OSPF Neighbor States:https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html
Understanding OSPF MTU:https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/117777-technote-ospf-00.html

Question: 46

```
*Jun 24 08:54:51.530: IF-EvD(GigabitEthernet0/0): IP Routing reports state transition from DOWN to DOWN
*Jun 24 08:54:52.525: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/0, changed state to down
*Jun 24 08:54:52.528: IF-EvD(GigabitEthernet0/0): IP Routing reports state transition from DOWN to DOWN
*Jun 24 08:54:53.215: IF-EvD(GigabitEthernet0/0): IP Routing reports state transition from DOWN to DOWN
*Jun 24 08:54:54.998: %LINK-3-UPDOWN: Interface GigabitEthernet0/0, changed state to up
*Jun 24 08:54:55.006: IF-EvD(GigabitEthernet0/0): IP Routing reports state transition from DOWN to UP
*Jun 24 08:54:55.998: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/0, changed state to up
```

Refer to the exhibit. R1 is connected with R2 via GigabitEthernet0/0, and R2 cannot ping R1. What action will fix the issue?

A.Fix route dampening configured on the router.
B.Replace the SFP module because it is not supported.
C.Fix IP Event Dampening configured on the interface.
D.Correct the IP SLA probe that failed.

Answer: C

Explanation:

Fix IP Event Dampening configured on the interface.

Question: 47

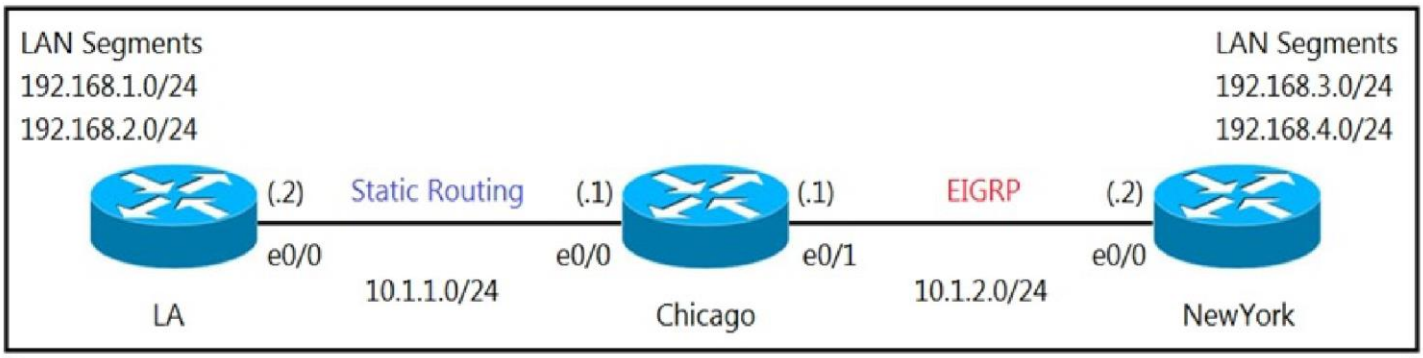


```
Chicago Router
ip route 192.168.1.0 255.255.255.0 10.1.1.2
ip route 192.168.2.0 255.255.255.0 10.1.1.2
!
router eigrp 100
 redistribute static

LA Router
ip route 0.0.0.0 0.0.0.0 10.1.1.1
```

Refer to the exhibits. A user on the 192.168.1.0/24 network can successfully ping 192.168.3.1, but the administrator cannot ping 192.168.3.1 from the LA router.
Which set of configurations fixes the issue?
A.

```
Chicago Router

ip route 192.168.3.0 255.255.255.0 10.1.2.2
ip route 192.168.4.0 255.255.255.0 10.1.2.2
```

B.

```
LA Router

ip route 192.168.3.0 255.255.255.0 10.1.1.1
ip route 192.168.4.0 255.255.255.0 10.1.1.1
```

C.

```
Chicago Router

router eigrp 100
 redistribute static metric 10 10 10 10 10
```

D.

```
Chicago Router

router eigrp 100
   redistribute connected
```

Answer: D

Explanation:

```
Chicago Router

router eigrp 100
   redistribute connected
```

Question: 48

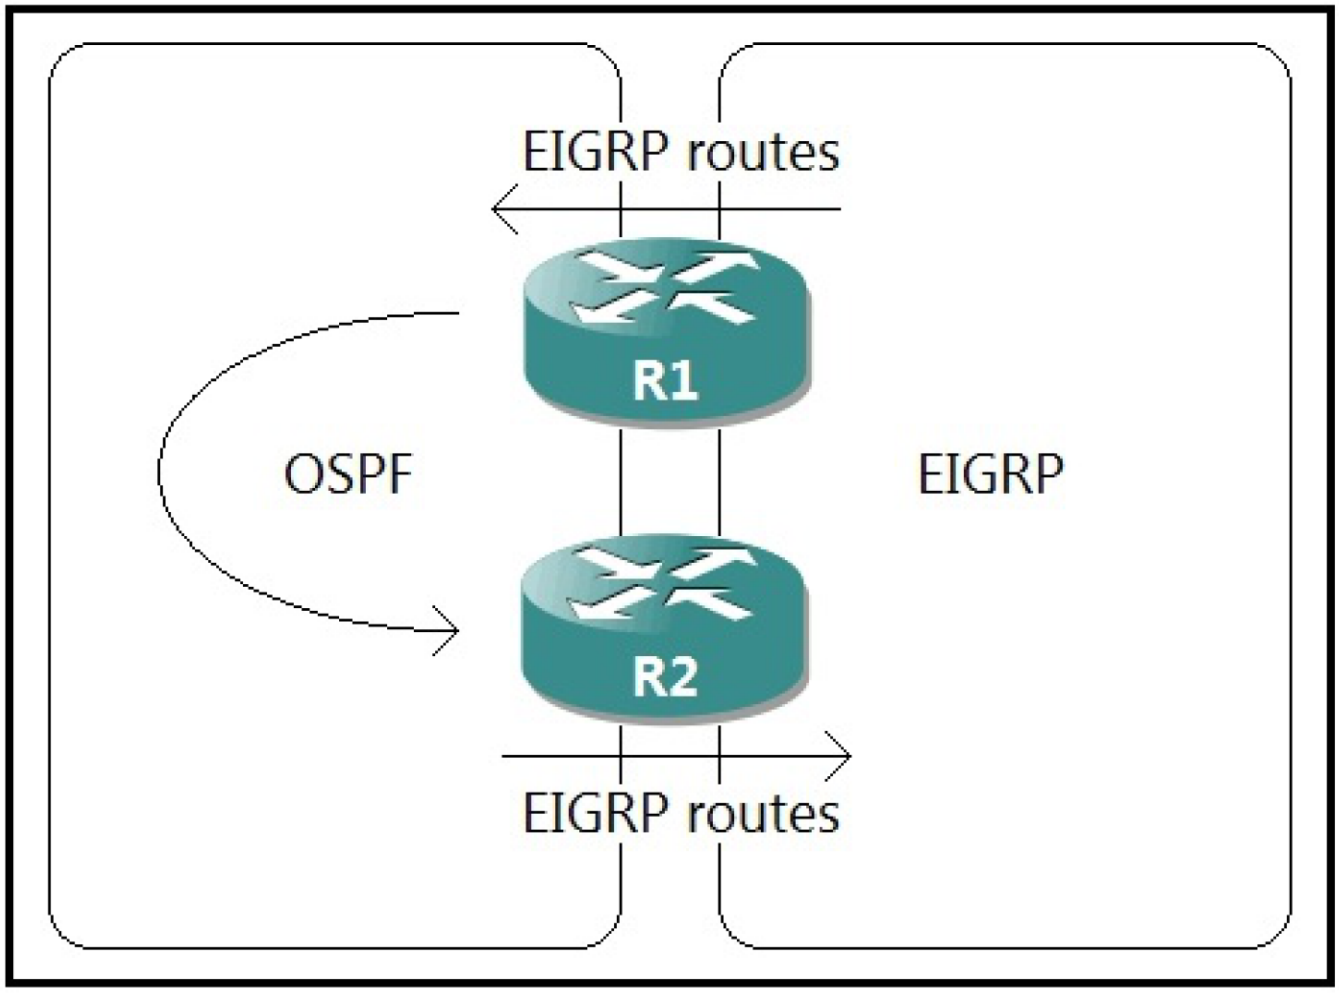

Refer to the exhibit. A network administrator configured mutual redistribution on R1 and R2 routers, which caused instability in the network.
Which action resolves the issue?

A. Set a tag in the route map when redistributing EIGRP into OSPF on R1, and match the same tag on R2 to deny when redistributing OSPF into EIGRP.

B. Set a tag in the route map when redistributing EIGRP into OSPF on R1, and match the same tag on R2 to allow when redistributing OSPF into EIGRP.

C. Apply a prefix list of EIGRP network routes in OSPF domain on R1 to propagate back into the EIGRP routing domain.

D. Advertise summary routes of EIGRP to OSPF and deny specific EIGRP routes when redistributing into OSPF.

Answer: A

Explanation:

Set a tag in the route map when redistributing EIGRP into OSPF on R1, and match the same tag on R2 to deny when redistributing OSPF into EIGRP.

Question: 49

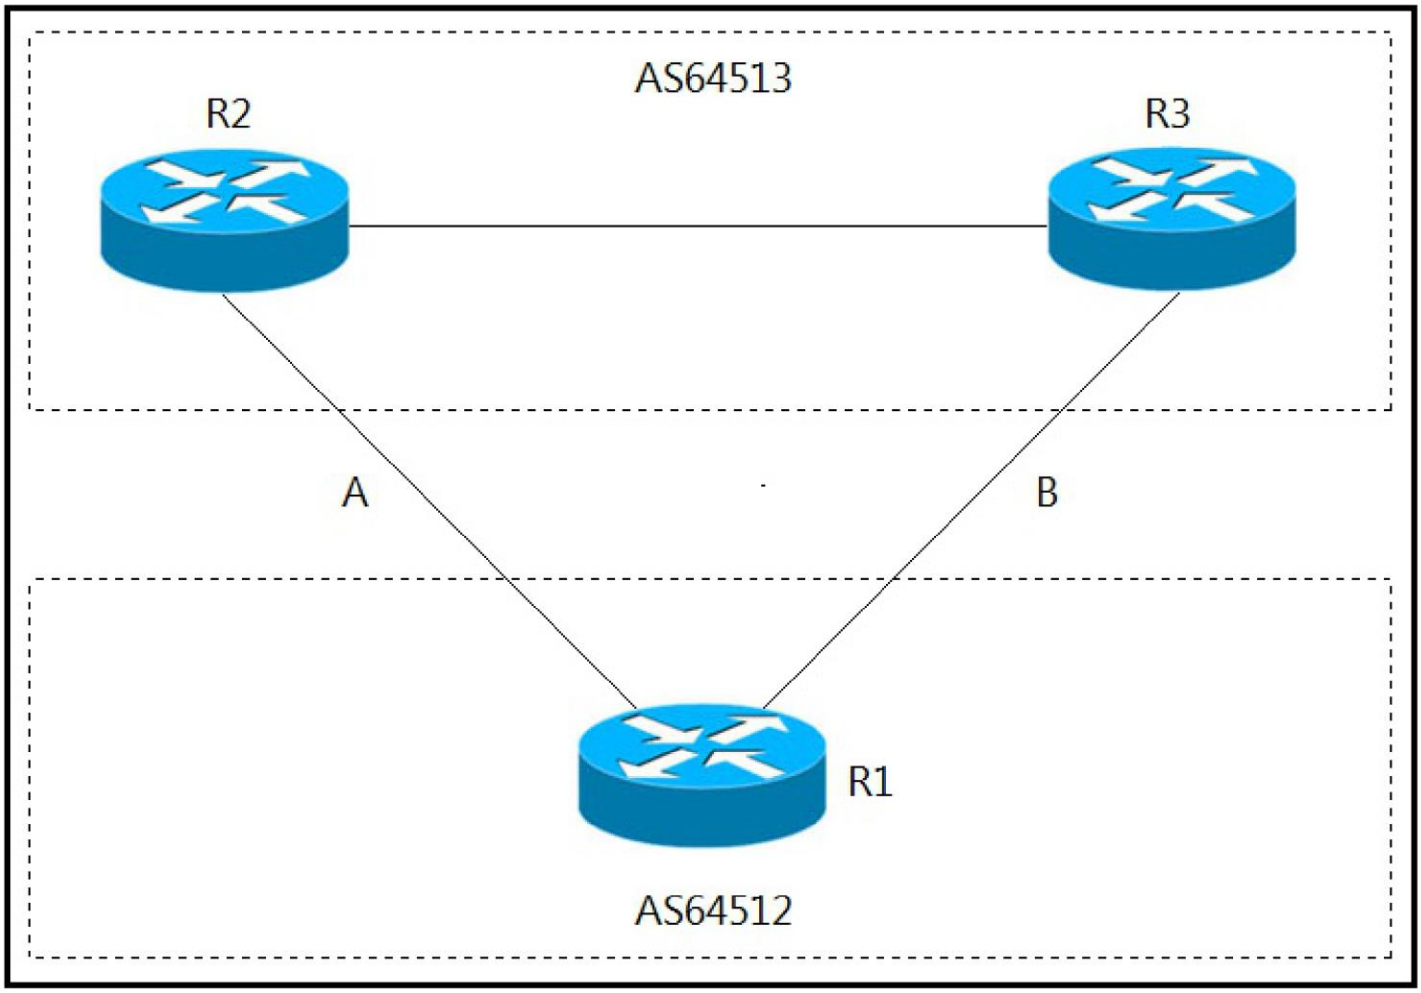

Refer to the exhibit. A network engineer for AS64512 must remove the inbound and outbound traffic from link A during maintenance without closing the BGP session so that there is still a backup link over link A toward the ASN. Which BGP configuration on R1 accomplishes this goal?

A.

```
route-map link-a-in permit 10
 set weight 200
route-map link-a-out permit 10
 set as-path prepend 64512
route-map link-b-in permit 10
 set weight 100
route-map link-b-out permit 10
```

B.

```
route-map link-a-in permit 10
 set weight 200
route-map link-a-out permit 10
route-map link-b-in permit 10
 set weight 100
route-map link-b-out permit 10
 set as-path prepend 64512
```

C.

```
route-map link-a-in permit 10
route-map link-a-out permit 10
 set as-path prepend 64512
route-map link-b-in permit 10
 set local-preference 200
route-map link-b-out permit 10
```

D.

```
route-map link-a-in permit 10
 set local-preference 200
route-map link-a-out permit 10
route-map link-b-in permit 10
route-map link-b-out permit 10
 set as-path prepend 64512
```

Answer: C

Explanation:

```
route-map link-a-in permit 10
route-map link-a-out permit 10
 set as-path prepend 64512
route-map link-b-in permit 10
 set local-preference 200
route-map link-b-out permit 10
```

Question: 50

An engineer configured access list NON-CISCO in a policy to influence routes.

```
route-map PBR, deny, sequence 5
 Match clauses:
  ip address (access-list): NON-CISCO
 Set clauses:
 Policy routing matches: 0 packets, 0 bytes
route-map PBR, permit, sequence 10
 Match clauses:
 Set clauses:
  ip next-hop 192.168.1.5
 Policy routing matches: 389362063 packets, 222009685077 bytes
```

What are the two effects of this route map configuration? (Choose two.)

A.Packets are forwarded using normal route lookup.
B.Packets are forwarded to the default gateway.
C.Packets are dropped by the access list.
D.Packets are evaluated by sequence 10.
E.Packets are not evaluated by sequence 10.

Answer: AD

Explanation:

A.Packets are forwarded using normal route lookup.

D.Packets are evaluated by sequence 10.

Question: 51

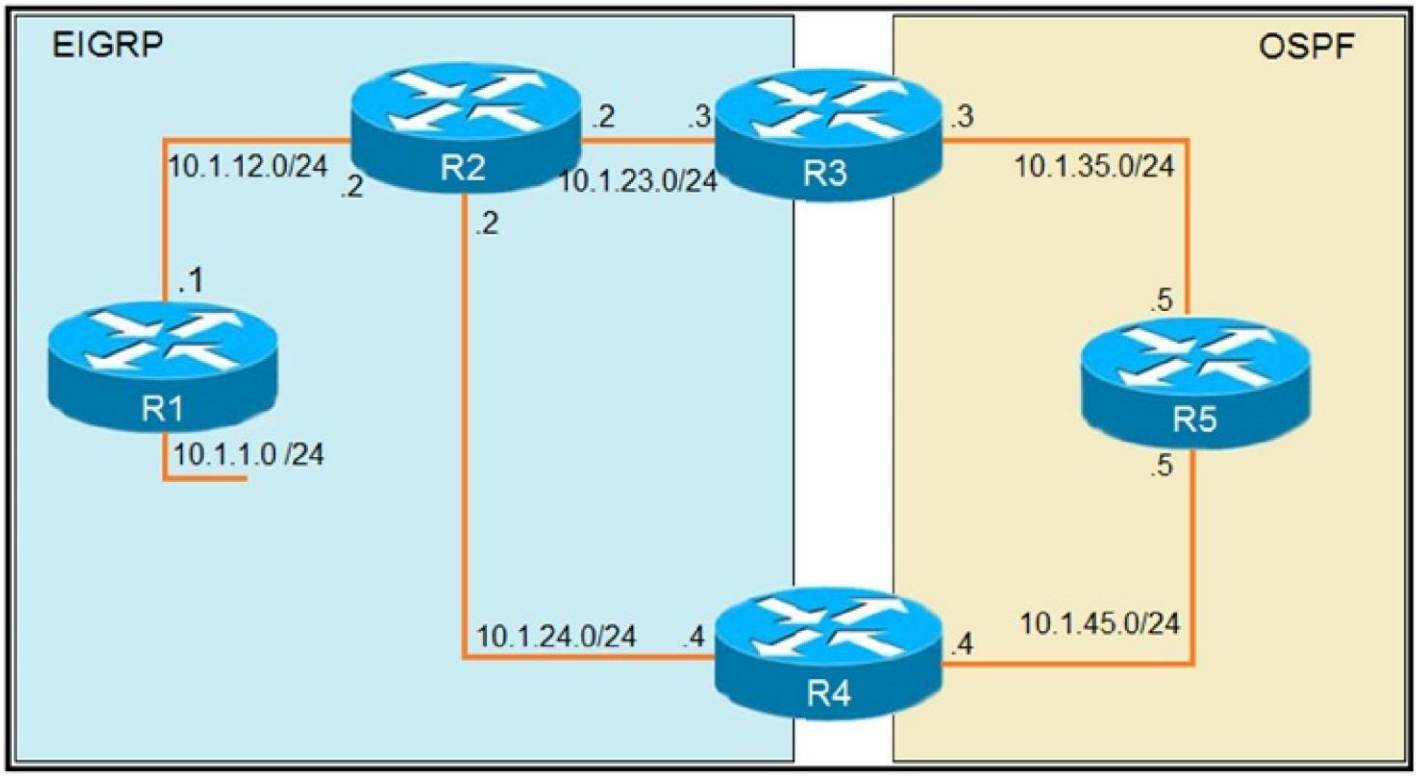

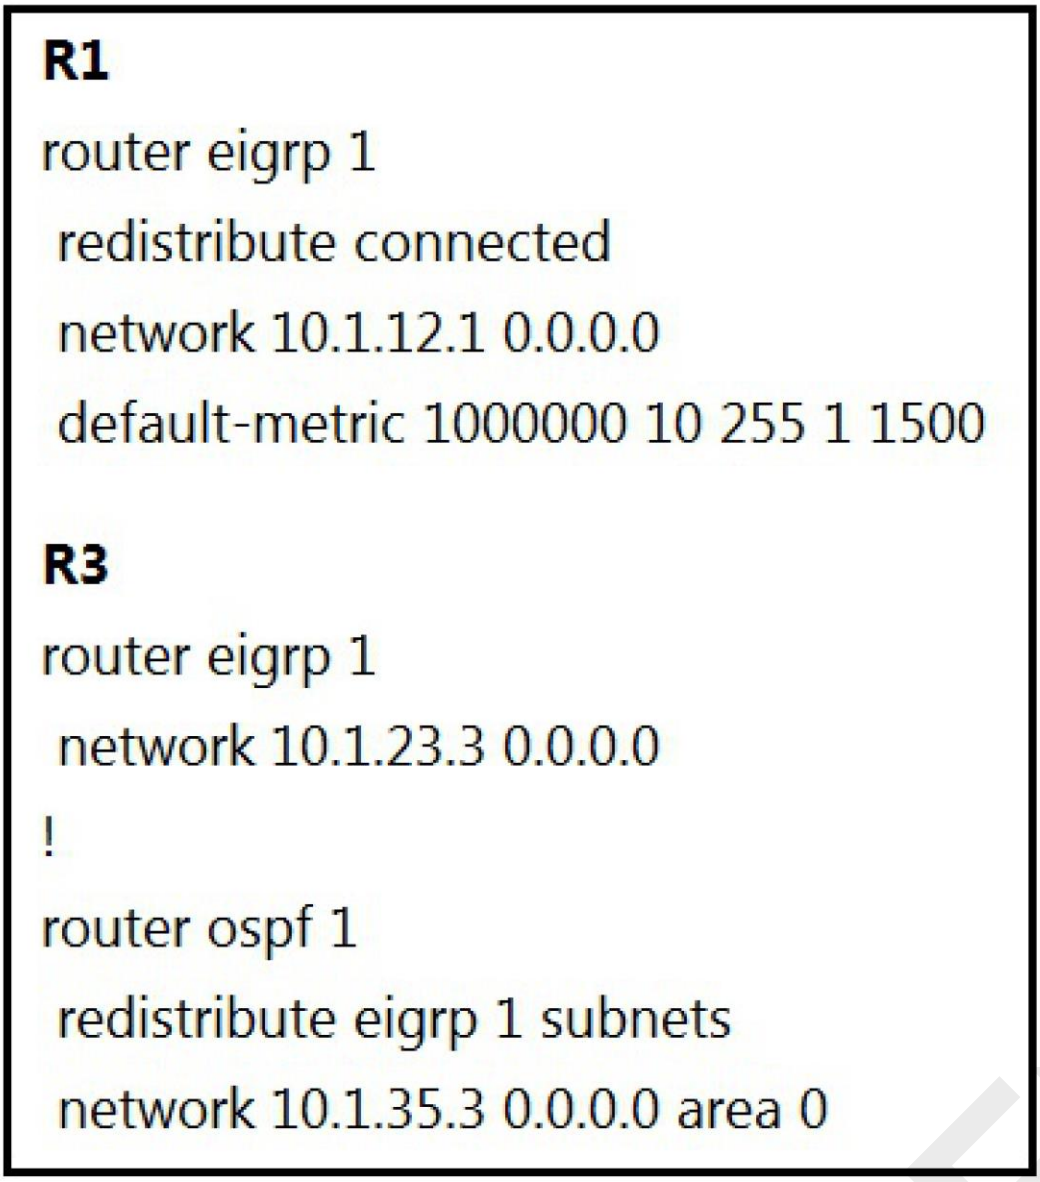

Refer to the exhibits. To provide reachability to network 10.1.1.0/24 from R5, the network administrator redistributes EIGRP into OSPF on R3 but notices that R4 is now taking a suboptimal path through R5 to reach 10.1.1.0/24 network. Which action fixes the issue while keeping the reachability from R5 to 10.1.1.0/24 network?

A.Change the administrative distance of the external EIGRP to 90.
B.Apply the outbound distribution list on R5 toward R4 in OSPF.
C.Change the administrative distance of OSPF to 200 on R5.
D.Redistribute OSPF into EIGRP on R4.

Answer: A

Explanation:

Change the administrative distance of the external EIGRP to 90.

Question: 52

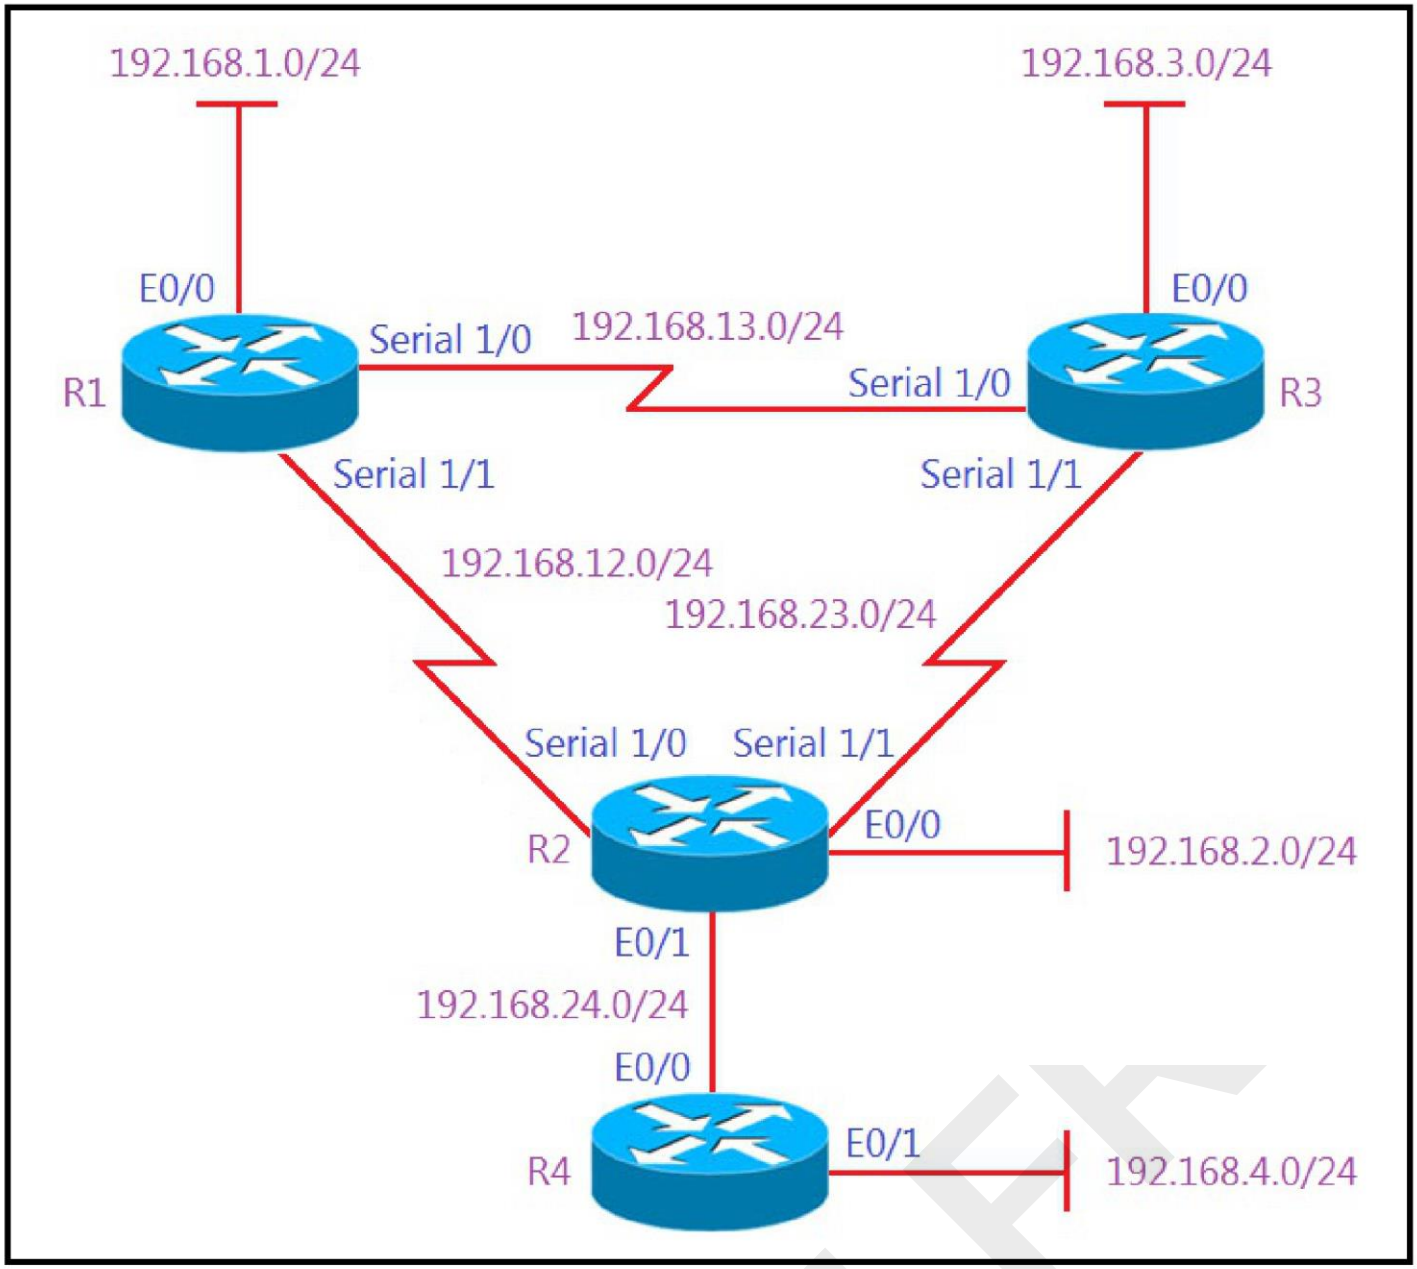
192.168.1.0/24
192.168.3.0/24
E0/0
E0/0
Serial 1/0
192.168.13.0/24
R1
Serial 1/0
R3
Serial 1/1
Serial 1/1
192.168.12.0/24
192.168.23.0/24
Serial 1/0
Serial 1/1
R2
E0/0
192.168.2.0/24
E0/1
192.168.24.0/24
E0/0
R4
E0/1
192.168.4.0/24

```
# Show IP route on R1
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, Ethernet0/0
L       192.168.1.1/32 is directly connected, Ethernet0/0
D       192.168.2.0/24 [90/2297856] via 192.168.12.2, 00:02:14, Serial1/1
S       192.168.3.0/24 [1/0] via 192.168.12.2
      192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.12.0/24 is directly connected, Serial1/1
L       192.168.12.1/32 is directly connected, Serial1/1
      192.168.13.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.13.0/24 is directly connected, Serial1/0
L       192.168.13.1/32 is directly connected, Serial1/0
D       192.168.23.0/24 [90/2681856] via 192.168.13.3, 00:06:38, Serial1/0
                        [90/2681856] via 192.168.12.2, 00:06:38, Serial1/1
D       192.168.24.0/24 [90/2195456] via 192.68.12.2, 00:06:38, Serial1/1
```

Refer to the exhibits. All the serial links between R1, R2, and R3 have the same bandwidth. Users on the 192.168.1.0/24 network report slow response times while they access resources on network 192.168.3.0/24. When a traceroute is run on the path, it shows that the packet is getting forwarded via R2 to R3 although the link between R1 and R3 is still up.
What must the network administrator do to fix the slowness?

A.Add a static route on R1 using the next hop of R3.
B.Remove the static route on R1.
C.Change the Administrative Distance of EIGRP to 5.
D.Redistribute the R1 static route to EIGRP.

Answer: B

Explanation:

Remove the static route on R1.

Question: 53

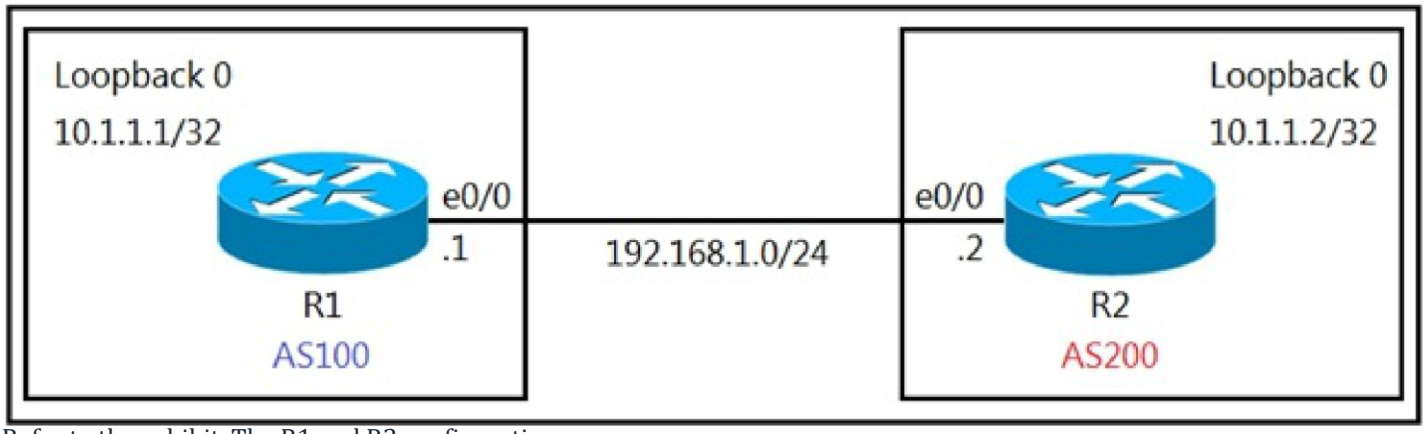

Refer to the exhibit. The R1 and R2 configurations are:

```
R1
router bgp 100
 neighbor 10.1.1.2 remote-as 200

R2
router bgp 200
 neighbor 10.1.1.1 remote-as 100
```

The neighbor relationship is not coming up.
Which two sets of configurations bring the neighbors up? (Choose two.) A.

```
 R1

 ip route 10.1.1.2 255.255.255.255 192.168.1.2
 !
 router bgp 100
  neighbor 10.1.1.1 ttl-security hops 1
  neighbor 10.1.1.2 update-source loopback 0
```

B.

R2

```
ip route 10.1.1.2 255.255.255.255 192.168.1.2
!
router bgp 100
 neighbor 10.1.1.2 ttl-security hops 1
 neighbor 10.1.1.2 update-source loopback 0
```

C.

R2

```
ip route 10.1.1.1 255.255.255.255 192.168.1.1
!
router bgp 200
neighbor 10.1.1.1 ttl-security hops 1
neighbor 10.1.1.1 update-source loopback 0
```

D.

R1

```
ip route 10.1.1.2 255.255.255.255 192.168.1.2
!
router bgp 100
 neighbor 10.1.1.2 disable-connected-check
 neighbor 10.1.1.2 update-source Loopback0
```

E.

R2

```
ip route 10.1.1.1 255.255.255.255 192.168.1.1
!
router bgp 200
 neighbor 10.1.1.1 disable-connected-check
 neighbor 10.1.1.1 update-source loopback 0
```

Answer: DE
Explanation:

R1

```
ip route 10.1.1.2 255.255.255.255 192.168.1.2
!
router bgp 100
 neighbor 10.1.1.2 disable-connected-check
 neighbor 10.1.1.2 update-source Loopback0
```

R2

```
ip route 10.1.1.1 255.255.255.255 192.168.1.1
!
router bgp 200
 neighbor 10.1.1.1 disable-connected-check
 neighbor 10.1.1.1 update-source loopback 0
```

Question: 54

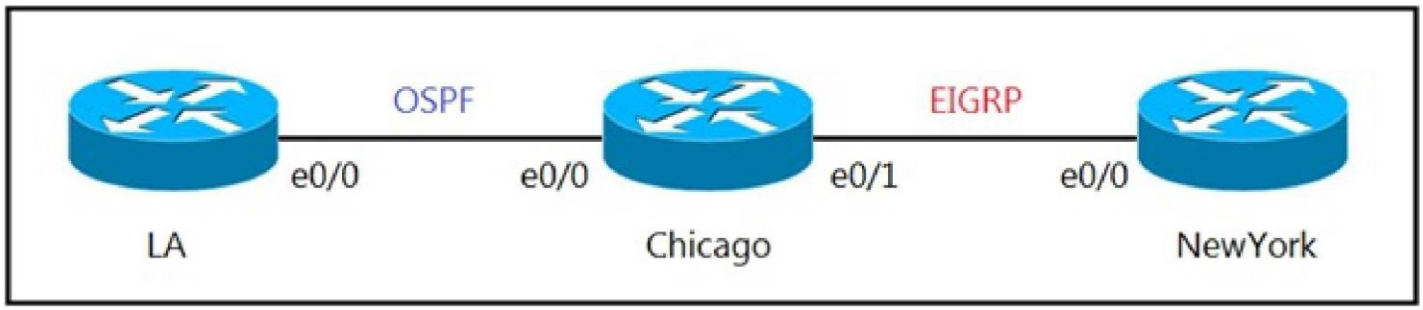

Refer to the exhibit. The network administrator must mutually redistribute routes at the Chicago router to the LA and NewYork routers. The configuration of the
Chicago router is this:

```
router ospf 1
 redistribute eigrp 100
router eigrp 100
 redistribute ospf 1
```

After the configuration, the LA router receives all the NewYork routes, but the NewYork router does not receive any LA routes.
Which set of configurations fixes the problem on the Chicago router? A.

```
router ospf 1
    redistribute eigrp 100 metric 20
```

B.

```
router eigrp 100
    redistribute ospf 1 metric 10 10 10 10 10
```

C.

```
router ospf 1
    redistribute eigrp 100 subnets
```

D.

```
router eigrp 100
    redistribute ospf 1 subnets
```

Answer: B

Explanation:

```
router eigrp 100
   redistribute ospf 1 metric 10 10 10 10 10
```

Question: 55

DRAG DROP -
Drag and drop the actions from the left into the correct order on the right to configure a policy to avoid following packet forwarding based on the normal routing path.
Select and Place:

Configure route map instances.	step 1
Configure set commands.	step 2
Configure fast switching for PBR.	step 3
Configure ACLs.	step 4
Configure match commands.	step 5
Configure PBR on the interface.	step 6

Answer:

Configure route map instances.	Configure ACLs.
Configure set commands.	Configure route map instances.
Configure fast switching for PBR.	Configure match commands.
Configure ACLs.	Configure set commands.
Configure match commands.	Configure PBR on the interface.
Configure PBR on the interface.	Configure fast switching for PBR.

Explanation:

Reference:
https://community.cisco.com/t5/networking-documents/how-to-configure-pbr/ta-p/3122774

Question: 56

```
R1

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

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

Refer to the exhibit. An engineer wanted to set a tag of 30 to route 10.1.80.65/32 but it failed. How is the issue fixed?

A.Modify route-map ospf-to-eigrp permit10 and match prefix-list ccnp2.
B.Modify prefix-list ccnp3 to add 10.1.64.0/20 ge 32.
C.Modify prefix-list ccnp3 to add 10.1.64.0/20 le 24.
D.Modify route-map ospf-to-eigrp permit 30 and match prefix-list ccnp2.

Answer: A

Explanation:

Modify route-map ospf-to-eigrp permit10 and match prefix-list ccnp2.

Question: 57

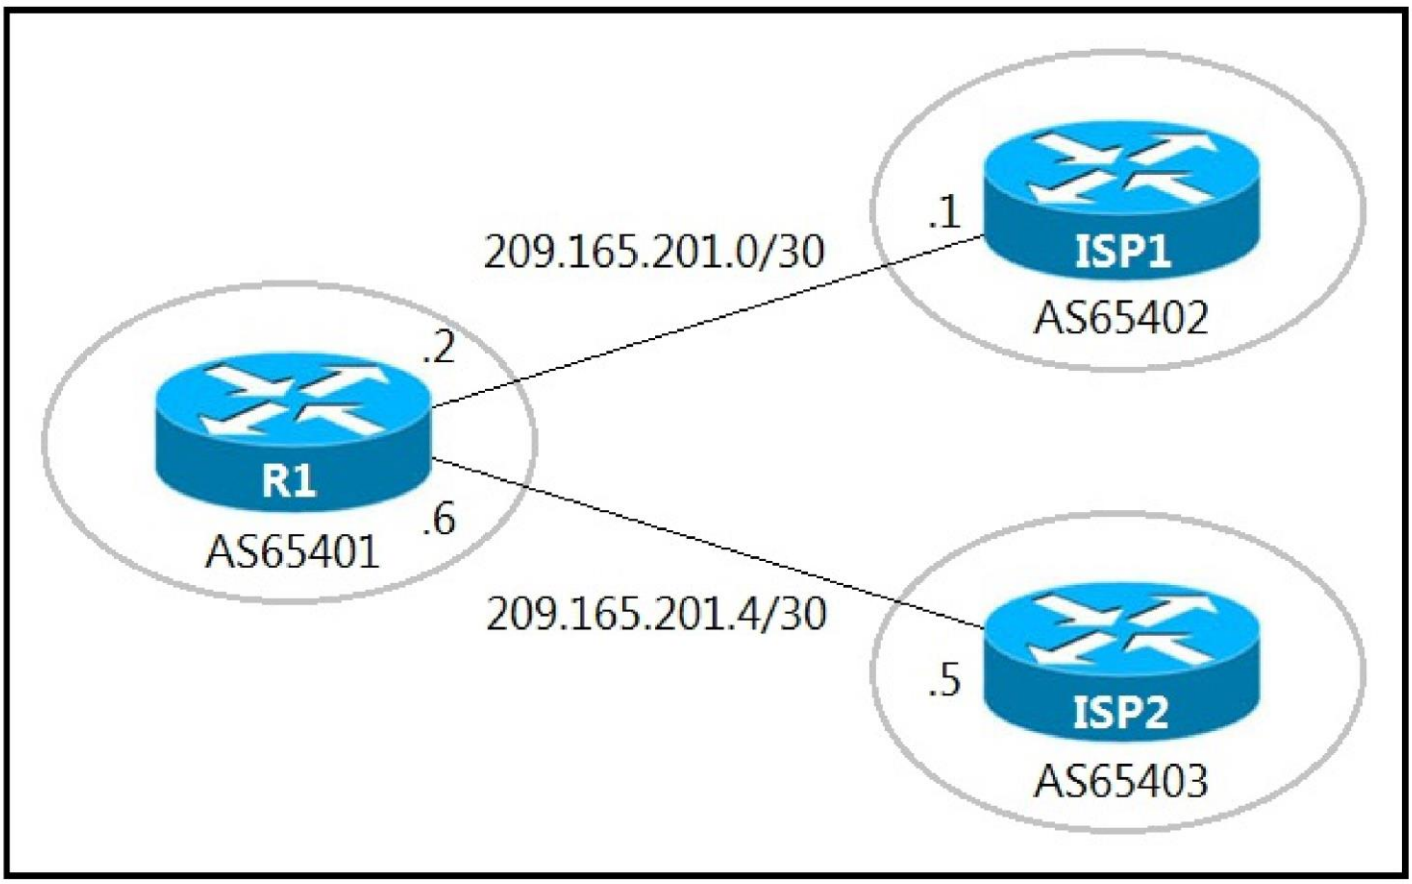


```
R1#
 interface GigabitEthernet0/0
  ip address 209.165.201.2 255.255.255.252
 !
 interface GigabitEthernet0/1
  ip address 209.165.201.6 255.255.255.252
 !
 router bgp 65401
  bgp log-neighbor-changes
  redistribute static
  neighbor 209.165.201.1 remote-as 65402
  neighbor 209.165.201.5 remote-as 65403
 !
 ip route 209.165.200.224 255.255.255.224 Null0
 ip route 209.165.202.128 255.255.255.224 Null0
 !
```

Refer to the exhibits. A company with autonomous system number AS65401 has obtained IP address block 209.165.200.224/27 from ARIN. The company needed more IP addresses and was assigned block 209.165.202.128/27 from ISP2. An engineer in ISP1 reports that they are receiving ISP2 routes from AS65401. Which configuration on R1 resolves the issue?

A.

```
access-list 10 deny 209.165.202.128 0.0.0.31
access-list 10 permit any
!
router bgp 65401
 neighbor 209.165.201.1 distribute-list 10 out
```

B.

```
access-list 10 deny 209.165.202.128 0.0.0.31
 access-list 10 permit any
 !
router bgp 65401
  neighbor 209.165.201.1 distribute-list 10 in
```

C.

```
ip route 209.165.200.224 255.255.255.224 209.165.201.1
ip route 209.165.202.128 255.255.255.224 209.165.201.5
```

D.

```
ip route 0.0.0.0 0.0.0.0 209.165.201.1
ip route 0.0.0.0 0.0.0.0 100 209.165.201.5
```

Answer: A

Explanation:

```
access-list 10 deny 209.165.202.128 0.0.0.31
access-list 10 permit any
!
router bgp 65401
 neighbor 209.165.201.1 distribute-list 10 out
```

Question: 58

After some changes in the routing policy, it is noticed that the router in AS 45123 is being used as a transit AS router for several service providers.
Which configuration ensures that the branch router in AS 45123 advertises only the local networks to all SP neighbors?

A.

```
ip as-path access-list 1 permit ^45123$
!
router bgp 45123
 neighbor SP-Neighbors filter-list 1 out
```

B.

```
ip as-path access-list 1 permit ^45123
!
router bgp 45123
 neighbor SP-Neighbors filter-list 1 out
```

C.

```
ip as-path access-list 1 permit ^$
!
router bgp 45123
 neighbor SP-Neighbors filter-list 1 out
```

D.

```
ip as-path access-list 1 permit .*
!
router bgp 45123
 neighbor SP-Neighbors filter-list 1 out
```

Answer: C

Explanation:

```
ip as-path access-list 1 permit ^$
!
router bgp 45123
 neighbor SP-Neighbors filter-list 1 out
```

Question: 59

A network administrator is troubleshooting a high utilization issue on the route processor of a router that was reported by NMS. The administrator logged into the router to check the control plane policing and observed that the BGP process is dropping a high number of routing packets and causing thousands of routes to recalculate frequently.
Which solution resolves this issue?

A.Shape the pir for BGP, conform-action set-prec-transmit, and exceed action set-frde-transmit.

B.Police the pir for BGP, conform-action set-prec-transmit, and exceed action set-clp-transmit.

C.Shape the cir for BGP, conform-action transmit, and exceed action transmit.

D.Police the cir for BGP, conform-action transmit, and exceed action transmit.

Answer: D

Explanation:

The correct solution is **D. Police the cir for BGP, conform-action transmit, and exceed action transmit.** This addresses the observed issue of high route processor utilization due to excessive BGP packet processing, which leads to frequent route recalculations.

Here's a detailed breakdown:

The Problem: The route processor is overloaded due to a high volume of BGP routing packets. This excessive processing causes the BGP process to drop packets, triggering frequent route recalculations and potentially network instability. The issue suggests that the BGP process is being bombarded by more packets than it can handle, causing congestion at the control plane.

Why Policing is the Solution: Policing, unlike shaping, aims to limit traffic based on specified rates and actions. It's more appropriate when the traffic should never exceed a certain threshold. In this case, BGP route updates should be limited to avoid overwhelming the route processor. Shaping would hold traffic in a buffer, which isn't ideal for routing protocols that need a quicker response to network changes.

CIR (Committed Information Rate): The CIR defines the guaranteed average rate for traffic. Policing the CIR ensures that BGP packets are limited to a defined rate, avoiding bursts that can overwhelm the route processor.

Conform-action transmit: This action is taken for packets that are within the defined CIR. The router will forward these packets normally.

Exceed-action transmit: This means that packets exceeding the CIR will also be transmitted, but still, this approach will limit the overall traffic from the BGP. The router might mark these packets down (depending on configuration) and still transmit them for processing, preventing a total block and providing flexibility on managing the route propagation.

Why other options are wrong:

A and C: Shaping: Shaping delays packets that exceed the configured rate, typically by storing them in a buffer. This delay is unsuitable for real-time routing protocols like BGP, where timely updates are vital for network convergence. Delaying BGP updates could lead to prolonged outages or routing loops.

B: Policing with set-prec-transmit and set-clp-transmit: While these actions can prioritize some packets, they are not a direct solution for limiting BGP traffic to prevent overutilization of the router. They manipulate packet headers, which is not the required action in this scenario. Furthermore, the primary objective is to limit the number of routes being processed by the router, not to prioritize certain BGP update types.

Implementation: By implementing the solution, the route processor can effectively handle the BGP updates without getting overloaded. The constant route recalculation will be reduced, leading to improved network stability and performance.

Authoritative Links for Further Research:

Cisco IOS Quality of Service:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos/configuration/15-mt/qos-15-mt-book/qos-polic.html
BGP Overview:https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/16436-

bgp.html
Policing vs Shaping:https://www.omnisecu.com/tcpip/qos/traffic-shaping-vs-traffic-policing.php

Question: 60

Which mechanism must be chosen to optimize the reconvergence time for OSPF at company location 408817202 that is less CPU-intensive than reducing the hello and dead timers?

A.sso
B.BFD
C.Dead Peer Detection keepalives
D.OSPF demand circuit

Answer: B

Explanation:

The correct answer is B, Bidirectional Forwarding Detection (BFD). BFD is a lightweight, low-overhead protocol designed for rapid fault detection in network paths. Unlike relying on OSPF's hello and dead timers, which are process-intensive and can be slow to detect failures, BFD actively sends frequent, short packets between neighbors. If these packets are not received within a specific time, BFD quickly declares the link down. This proactive approach leads to significantly faster failure detection and, consequently, quicker reconvergence for routing protocols like OSPF. The other options are not as directly relevant for rapid failure detection in OSPF. SSO is related to high availability and redundant supervisors, not routing reconvergence.

Dead Peer Detection (DPD) keepalives are primarily used with VPNs, not OSPF, and are not as fast as BFD.

OSPF demand circuits are designed for on-demand links and do not directly optimize link failure detection.

Therefore, BFD offers a more efficient and less CPU-intensive solution for achieving quicker OSPF reconvergence than tuning hello and dead timers.Cisco BFD DocumentationBFD Explanation

Question: 61

Refer to the exhibit.

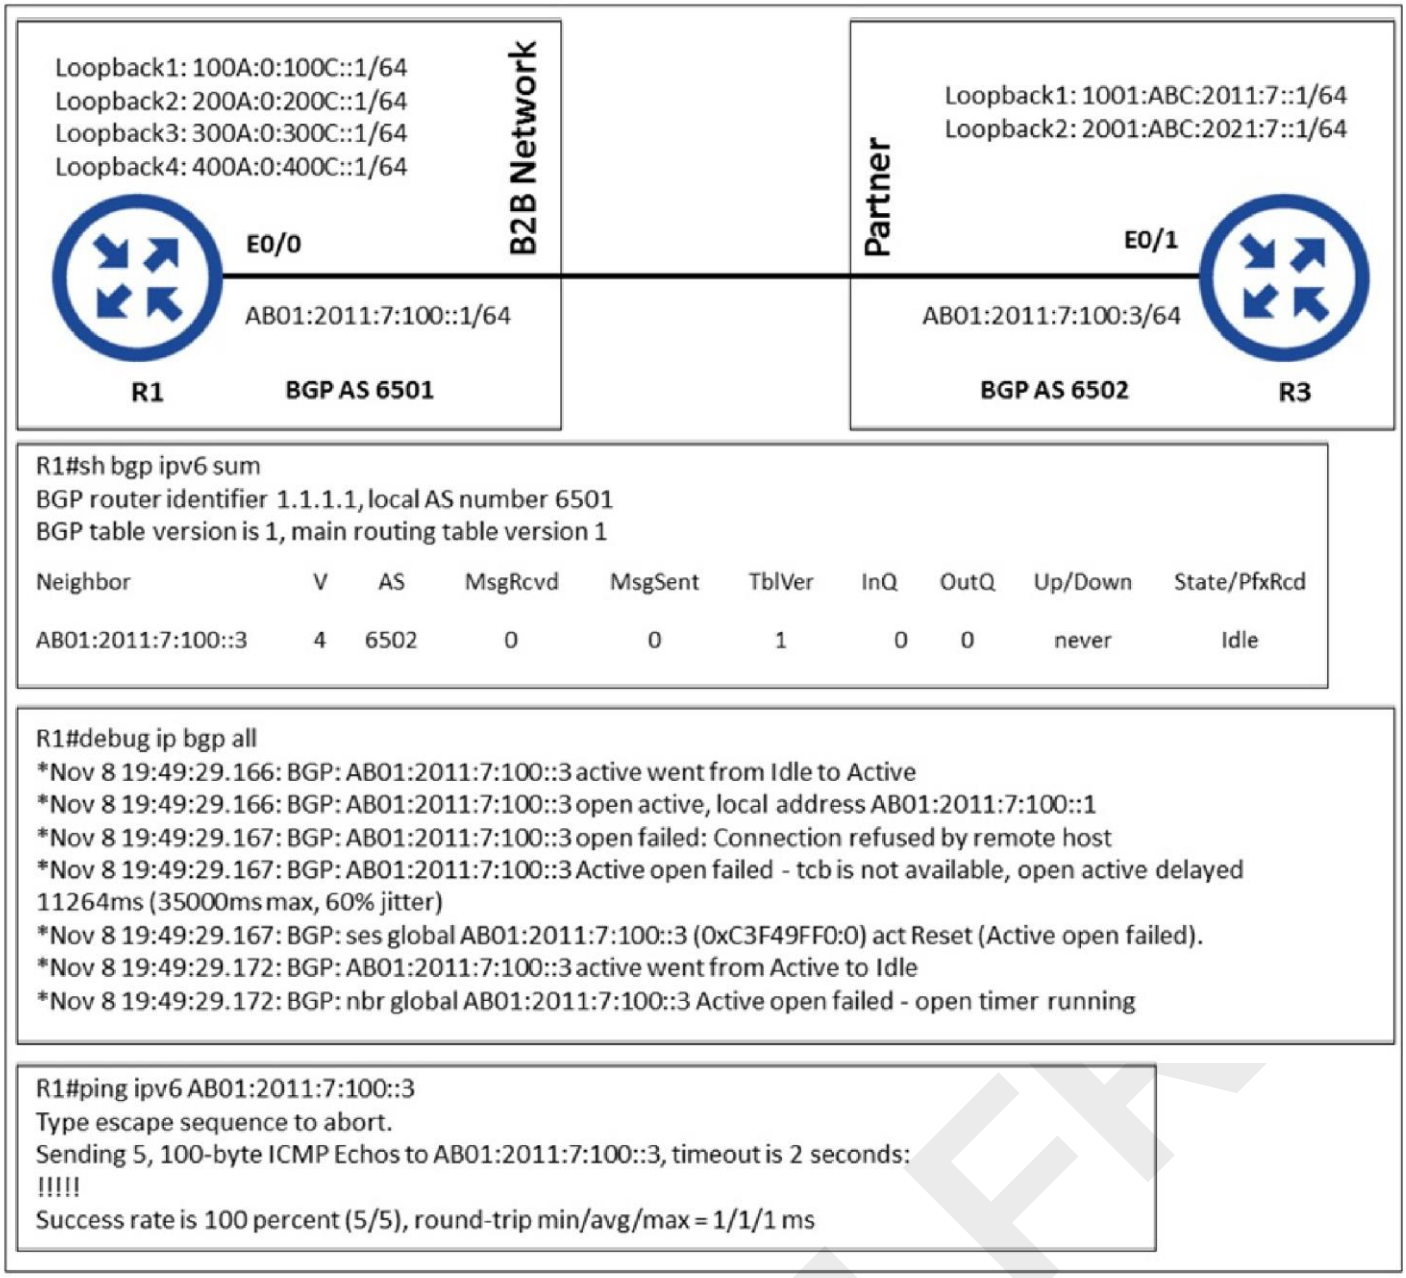

An engineer configured BGP between routers R1 and R3. The BGP peers cannot establish neighbor adjacency to be able to exchange routes.
Which configuration resolves this issue?

A. R1 router bgp 6501 address-family ipv6 neighbor AB01:2011:7:100::3 activate B. R3 router bgp 6502 address-family ipv6 neighbor AB01:2011:7:100::1 activate C. R1 router bgp 6501 neighbor AB01:2011:7:100::3 ebgp-multihop 255
D. R3 router bgp 6502 neighbor AB01:2011:7:100::1 ebgp-multihop 255

Answer: B

Explanation:

R3 router bgp 6502 address-family ipv6 neighbor AB01:2011:7:100::1 activate.

Question: 62

Refer to the exhibit.

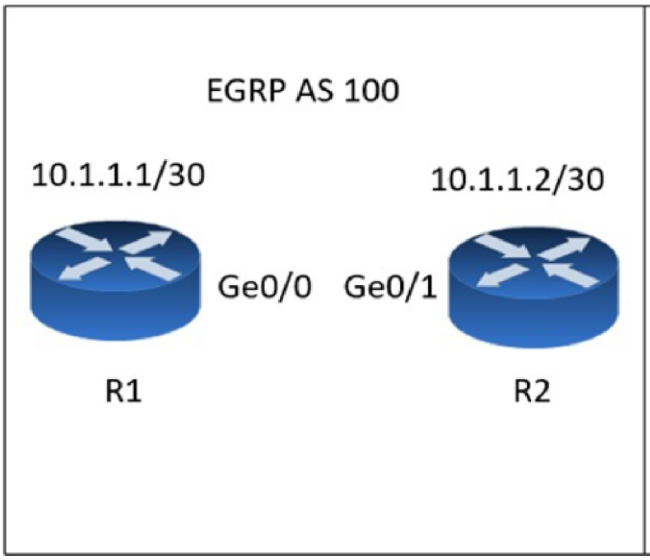


```
R1# debug eigrp packets
(UPDATE, REQUEST, QUERY, REPLY, HELLO,
UNKNOWN, PROBE, ACK, STUB, SIAQUERY,
SIAREPLY)
EIGRP Packet debugging is on
R1#
EIGRP: Sending HELLO on Gi0/0 - paklen 20
AS 100, Flags 0x0:(NULL), Seq 0/0 interfaceQ 0/0
iidbQ un/rely 0/0
R1#
EIGRP: Sending HELLO on Gi0/0 - paklen 20
AS 100, Flags 0x0:(NULL), Seq 0/0 interfaceQ 0/0
iidbQ un/rely 0/0
```

Which action resolves the adjacency issue?

A.Configure the same autonomous system numbers.
B.Match the hello interval timers.
C.Match the authentication keys.
D.Configure the same EIGRP process IDs.

Answer: A

Explanation:

Reference:
https://www.ciscopress.com/articles/article.asp?p=2999383&seqNum=2

Question: 63

Refer to the exhibit.

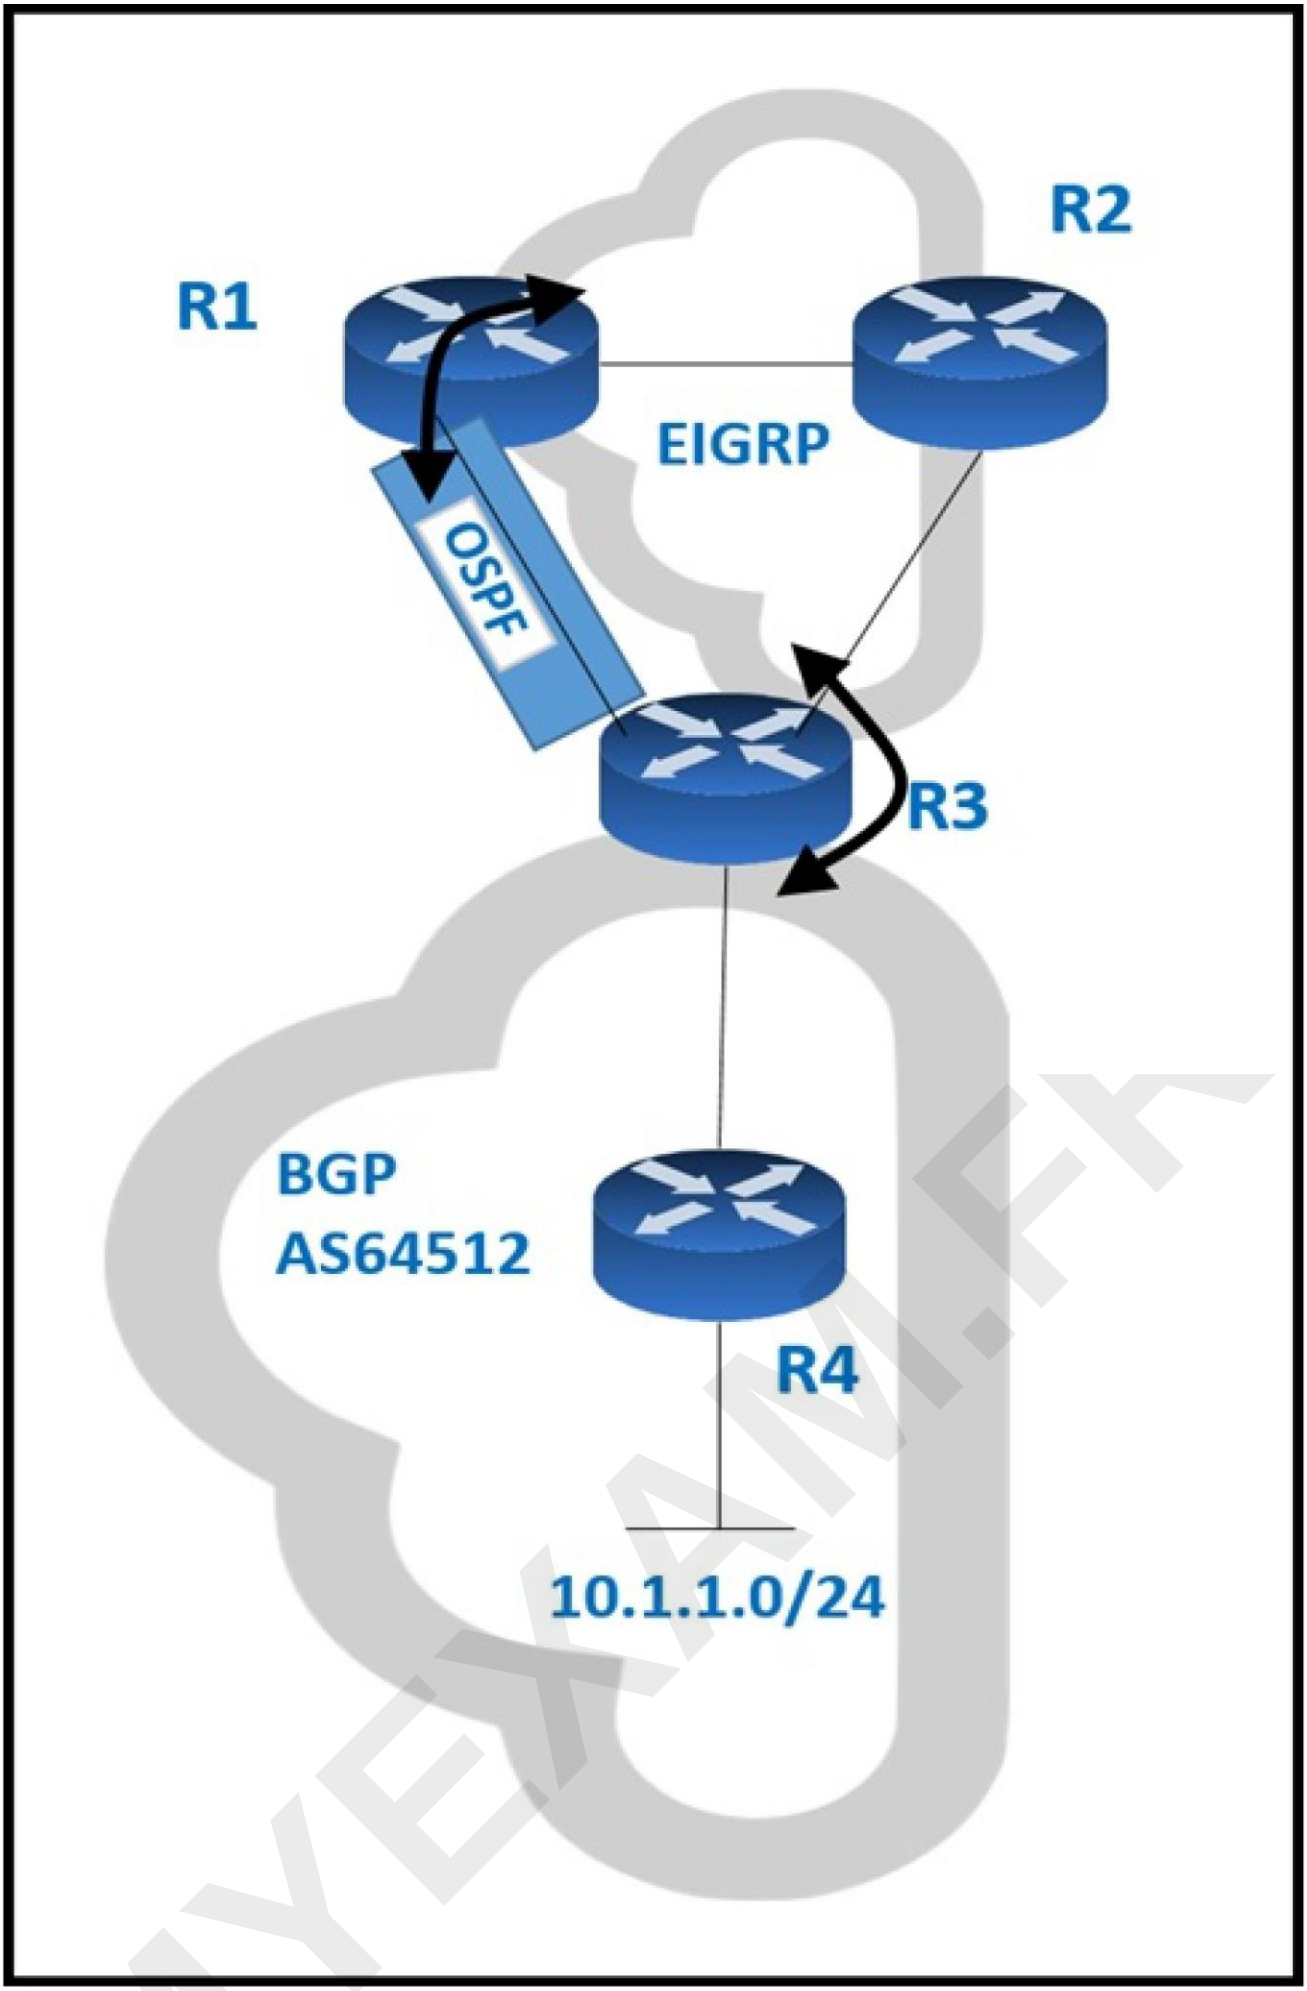

BGP and EIGRP are mutually redistributed on R3, and EIGRP and OSPF are mutually redistributed on R1. Users report packet loss and interruption of service to applications hosted on the 10.1.1.0/24 prefix. An engineer tested the link from R3 to R4 with no packet loss present but has noticed frequent routing changes on
R3 when running the debug ip route command.

Which action stabilizes the service?

A.Reduce frequent OSPF SPF calculations on R3 that cause a high CPU and packet loss on traffic traversing R3.

B.Tag the 10.1.1.0/24 prefix and deny the prefix from being redistributed into OSPF on R1.

C.Place an OSPF distribute-list outbound on R3 to block the 10.1.1.0/24 prefix from being advertised back to R3.

D.Repeat the test from R4 using ICMP ping on the local 10.1.1.0/24 prefix, and fix any Layer 2 errors on the host or switch side of the subnet.

Answer: B

Explanation:

Tag the 10.1.1.0/24 prefix and deny the prefix from being redistributed into OSPF on R1.

Question: 64

Refer to the exhibit. An engineer has configured policy-based routing and applied the configuration to the correct interface. How is the configuration applied to the traffic that matches the access list?

```
Route-map PBR, permit, sequence 10
  Match clauses:
    ip address (access lists): FILTER_ACL
  Set clauses:
    ip next-hop verify-availability 209.165.202.129 1 track 100 [down]
    ip next-hop verify-availability 209.165.202.131 2 track 200 [up]
  Policy routing matches: 0 packets, 0 bytes
route-map PBR, deny, sequence 20
  Match clauses:
  Set clauses:
    ip next-hop 209.165.201.30
  Policy routing matches: 275364861 packets, 12200235037 bytes
```

A. It is forwarded using the routing table lookup.

B. It is sent to 209.165.202.129.

C. It is dropped.

D. It is sent to 209.165.202.131.

Answer: D

Explanation:

The first next hop IP is down, so the second one will be used.

Question: 65

Refer to the exhibit.

```
Branch-Router#
"Nov 29 15.20.22.415: OSPF-1 HELLO Fa1/1: Rcv hello from 3.3.3.3 area 1 10.2.1.3
"Nov 29 15.20.23.195: OSPF-1 HELLO Fa1/1: Send hello to 224.0.0.5 area 1 from 10.2.1.1

Branch-Router#
"Nov 29 15.20:27.955: OSPF-1 HELLO Fa0/0: Rcv hello from 2.2.2.2 area 1 10.1.1.2
"Nov 29 15.20.27.955: OSPF-1 HELLO Fa0/0: Mismatched hello parameters from 10.1.1.2
"Nov 29 15.20.27.955: OSPF-1 HELLO Fa0/0: Dead R 40 C 40, Hello R 10 C 10 Mask R 255.255.255.0 C 255.255.255.240
"Nov 29 15.20.28.311: OSPF-1 HELLO Fa0/0: Send hello to 224.0.0.5 area 1 from 10.1.1.1
```

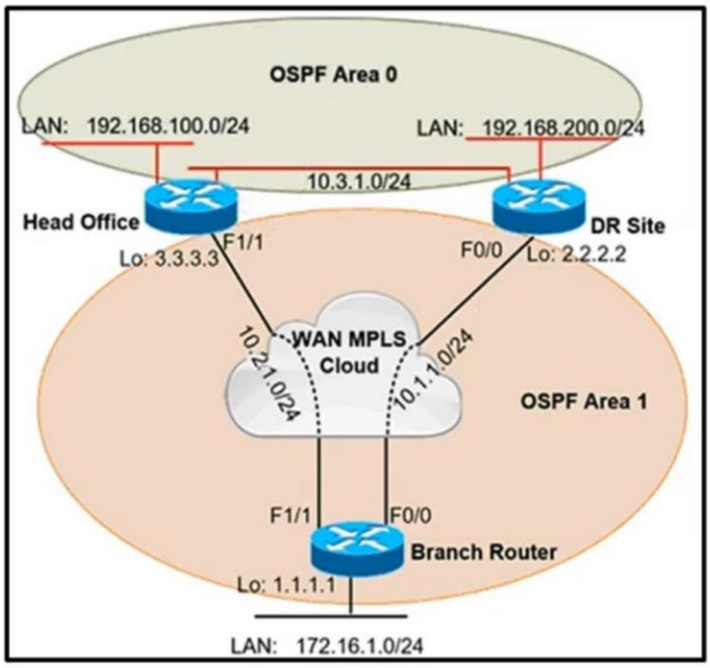

A network administrator reviews the branch router console log to troubleshoot the OSPF adjacency issue with the DR router.
Which action resolves this issue?

A.Stabilize the DR site flapping link to establish OSPF adjacency.
B.Advertise the branch WAN interface matching subnet for the DR site.
C.Configure the WAN interface for DR site in the related OSPF area.
D.Configure matching hello and dead intervals between sites.

Answer: B

Explanation:

Reference:
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13699-29.html

Question: 66

Refer to the exhibit.

```
P 172.29.0.0/16, 1 successors, FD is 307200, serno 2
        via 192.168.254.2 (307200/281600), FastEthernet0/1
        via 192.168.253.2 (410200/352300), FastEthernet0/0
```

When the FastEthernet0/1 goes down, the route to 172.29.0.0/16 via 192.168.253.2 is not installed in the RIB. Which action resolves the issue?

A.Configure feasible distance greater than the reported distance.

B.Configure feasible distance greater than the successor's feasible distance.
C.Configure reported distance greater than the successor's feasible distance.
D.Configure reported distance greater than the feasible distance.

Answer: A

Explanation:

Reference:

https://www.practicalnetworking.net/stand-alone/eigrp-feasibility-condition/

Question: 67

Refer to the exhibit.

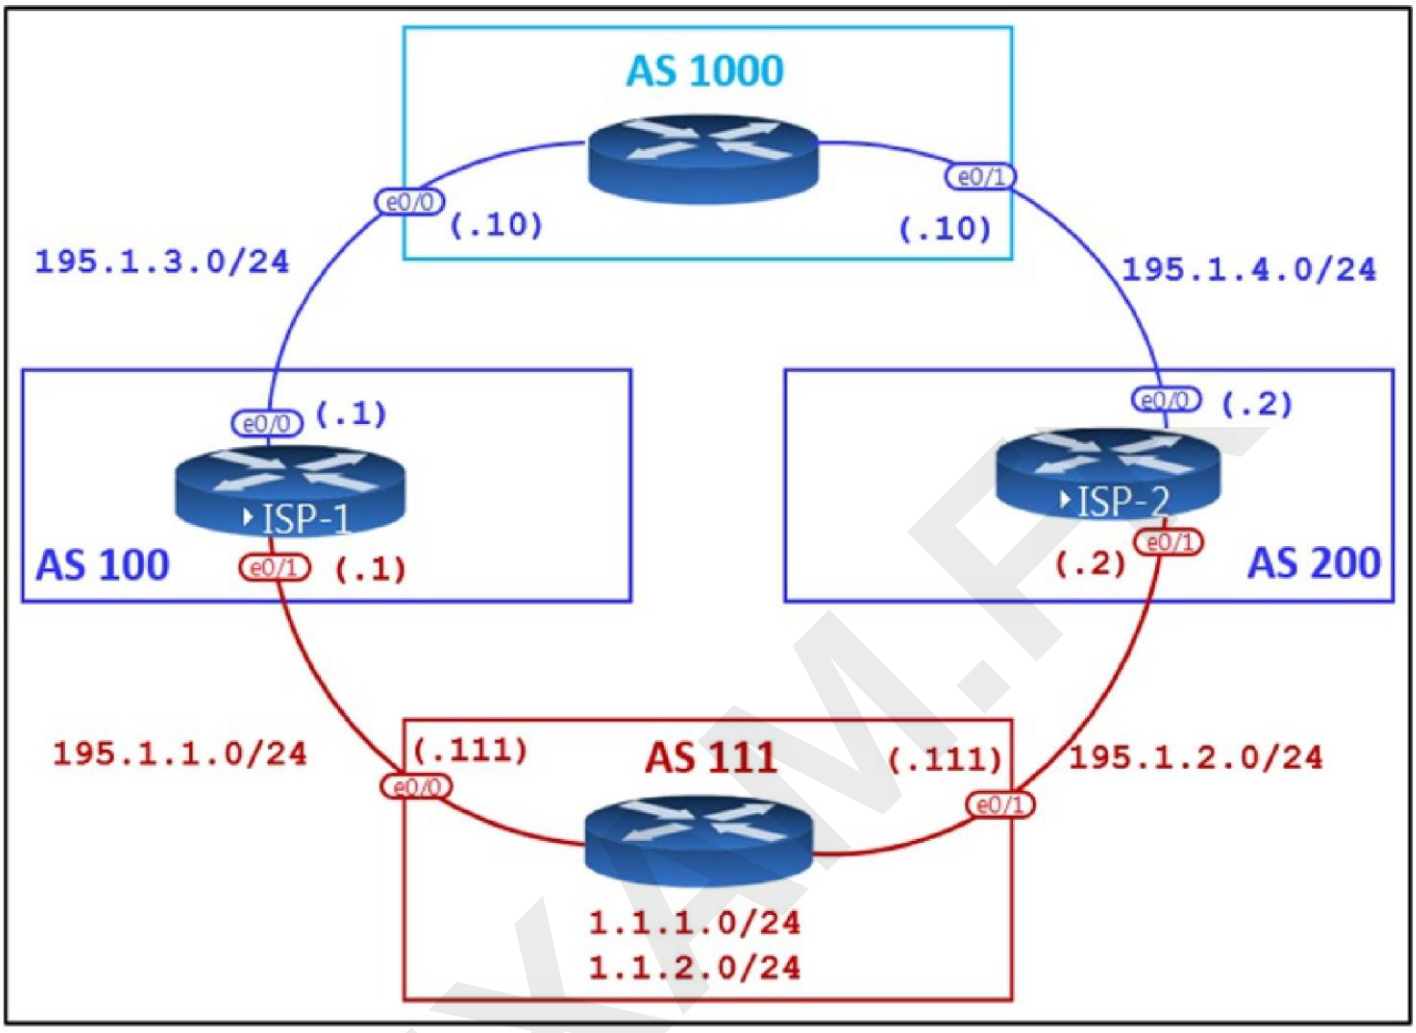

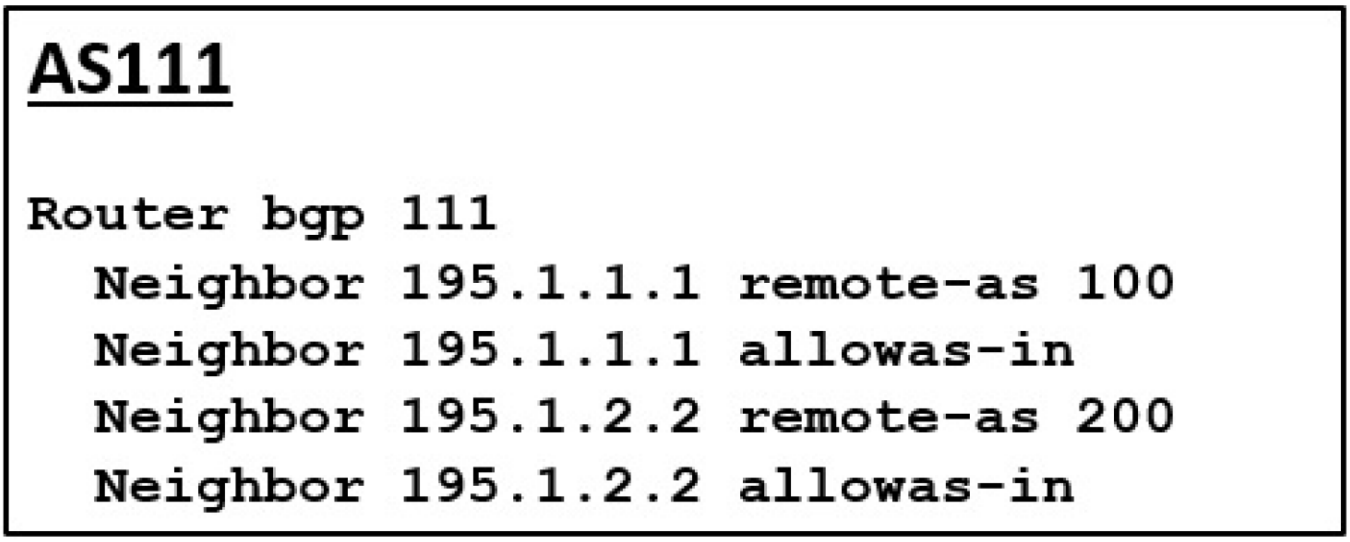

AS111 is receiving its own routes from AS200 causing a loop in the network.
Which configuration provides loop prevention?

A. router bgp 111 neighbor 195.1.1.1 as-override no neighbor 195.1.2.2 allowas-in B. router bgp 111 no neighbor 195.1.1.1 allowas-in no neighbor 195.1.2.2 allowas-in C. router bgp 111 neighbor 195.1.2.2 as-override no neighbor 195.1.1.1 allowas-in D. router bgp 111 neighbor 195.1.1.1 as-override neighbor 195.1.2.2 as-override

Answer: B

Explanation:

Reference:
https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/112236-allowas-in-bgp-config-example.html

Question: 68

Refer to the exhibit.

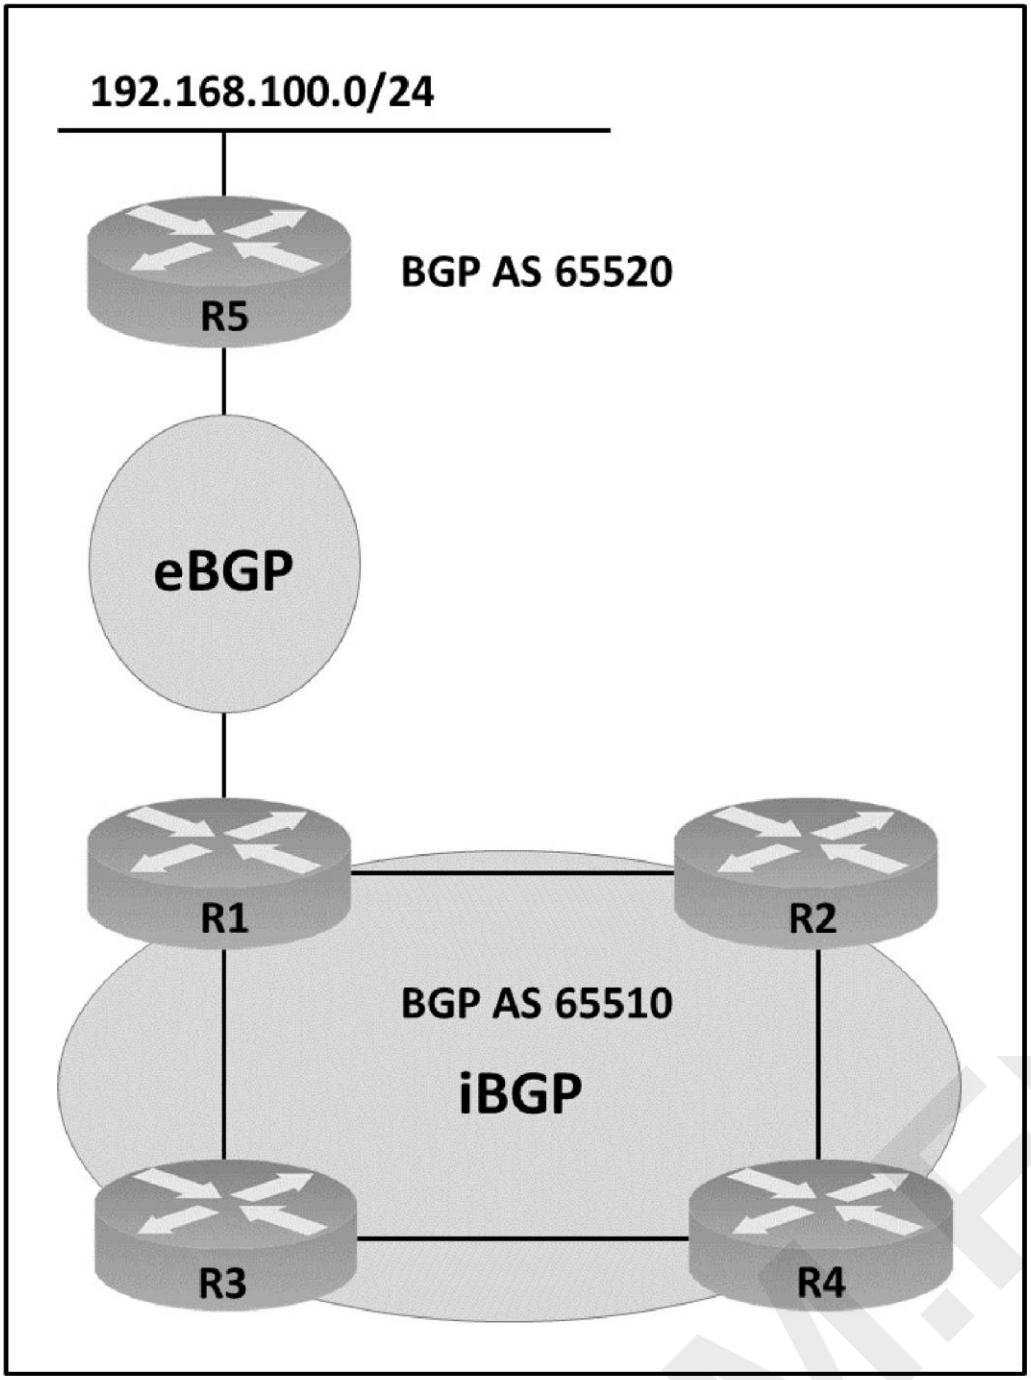

AS65510 iBGP is configured for directly connected neighbors. R4 cannot ping or traceroute network 192.168.100.0/24.
Which action resolves this issue?

A.Configure R1 as a route reflector server and configure R2 and R3 as route reflector clients.
B.Configure R4 as a route reflector server and configure R2 and R3 as route reflector clients.
C.Configure R4 as a route reflector server and configure R1 as a route reflector client.
D.Configure R1 as a route reflector server and configure R4 as a route reflector client.

Answer: B

Explanation:

Configure R4 as a route reflector server and configure R2 and R3 as route reflector clients.

Question: 69

Users report issues with reachability between areas as soon as an engineer configured summary routes between areas in a multiple area OSPF autonomous system.
Which action resolves the issue?

A.Configure the area range command on the ASBR.
B.Configure the summary-address command on the ASBR.
C.Configure the summary-address command on the ABR.
D.Configure the area range command on the ABR.

Answer: D

Explanation:

The correct answer is **D. Configure the area range command on the ABR.** Here's why:

The problem describes reachability issues arising after configuring summary routes between OSPF areas. This strongly suggests that the summarization is not being handled correctly at the Area Border Routers (ABRs). ABRs are the crucial points where routes between different OSPF areas are summarized and advertised.

The summary-address command is used on an Autonomous System Boundary Router (ASBR) to summarize external routes that are redistributed into OSPF, not routes between areas within OSPF. Conversely, the area range command is specifically used on ABRs to summarize routes within a specific area as they're being advertised into the backbone area or another non-backbone area. By using the area range command, the ABR condenses multiple network advertisements into a single summary route, reducing the size of the routing table and simplifying routing within the OSPF domain. When the area range command isn't configured correctly on the ABR, specific routes could be hidden within a broader range causing unreachable networks.

By correctly using the area range command on the ABR, we ensure that OSPF routes are summarized correctly when moving between OSPF areas, thus resolving the reachability issues reported by users.

Key Concepts

OSPF Areas: Dividing an OSPF network into areas improves scalability and efficiency.
Area Border Routers (ABRs): These routers connect different OSPF areas. They play a critical role in summarizing routes between the areas.
Autonomous System Boundary Routers (ASBRs): These routers connect an OSPF network to external networks.
Route Summarization: The process of condensing multiple smaller network advertisements into a single larger one.

Further Research:

Cisco Documentation on OSPF Route Summarization: https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13684-ospf-area-range.html
OSPF Area Configuration: https://www.oreilly.com/library/view/ccna-routing-and/9780134297306/ch10.html

Question: 70

Refer to the exhibit.

```
interface loopback0
ip address 4.4.4.4 255.255.255.0
!
interface FastEthernet1/0
Description **** WAN link ****
ip address 10.0.0.1 255.255.255.0
!
interface FastEthernet1/1
Description **** LAN Network ****
ip address 192.168.1.1 255.255.255.0
!
!
router ospf 1
router-id 4.4.4.4
log-adjacency-changes
network 4.4.4.4 0.0.0.0 area 0
network 10.0.0.1 0.0.0.0 area 0
network 192.168.1.1 0.0.0.0 area 10
!
```

Which set of commands restore reachability to loopback0?

A. interface loopback0 ip address 4.4.4.4 255.255.255.0 ip ospf network point-to-point B. interface loopback0 ip address 4.4.4.4 255.255.255.0 ip ospf interface area 10 C. interface loopback0 ip address 4.4.4.4 255.255.255.0 ip ospf network broadcast D. interface loopback0 ip address 4.4.4.4 255.255.255.0 ip ospf interface type network

Answer: A

Explanation:

Reference:
https://networkengineering.stackexchange.com/questions/13099/why-do-we-use-ospf-point-to-point-networ ks-for-loopbacks

Question: 71

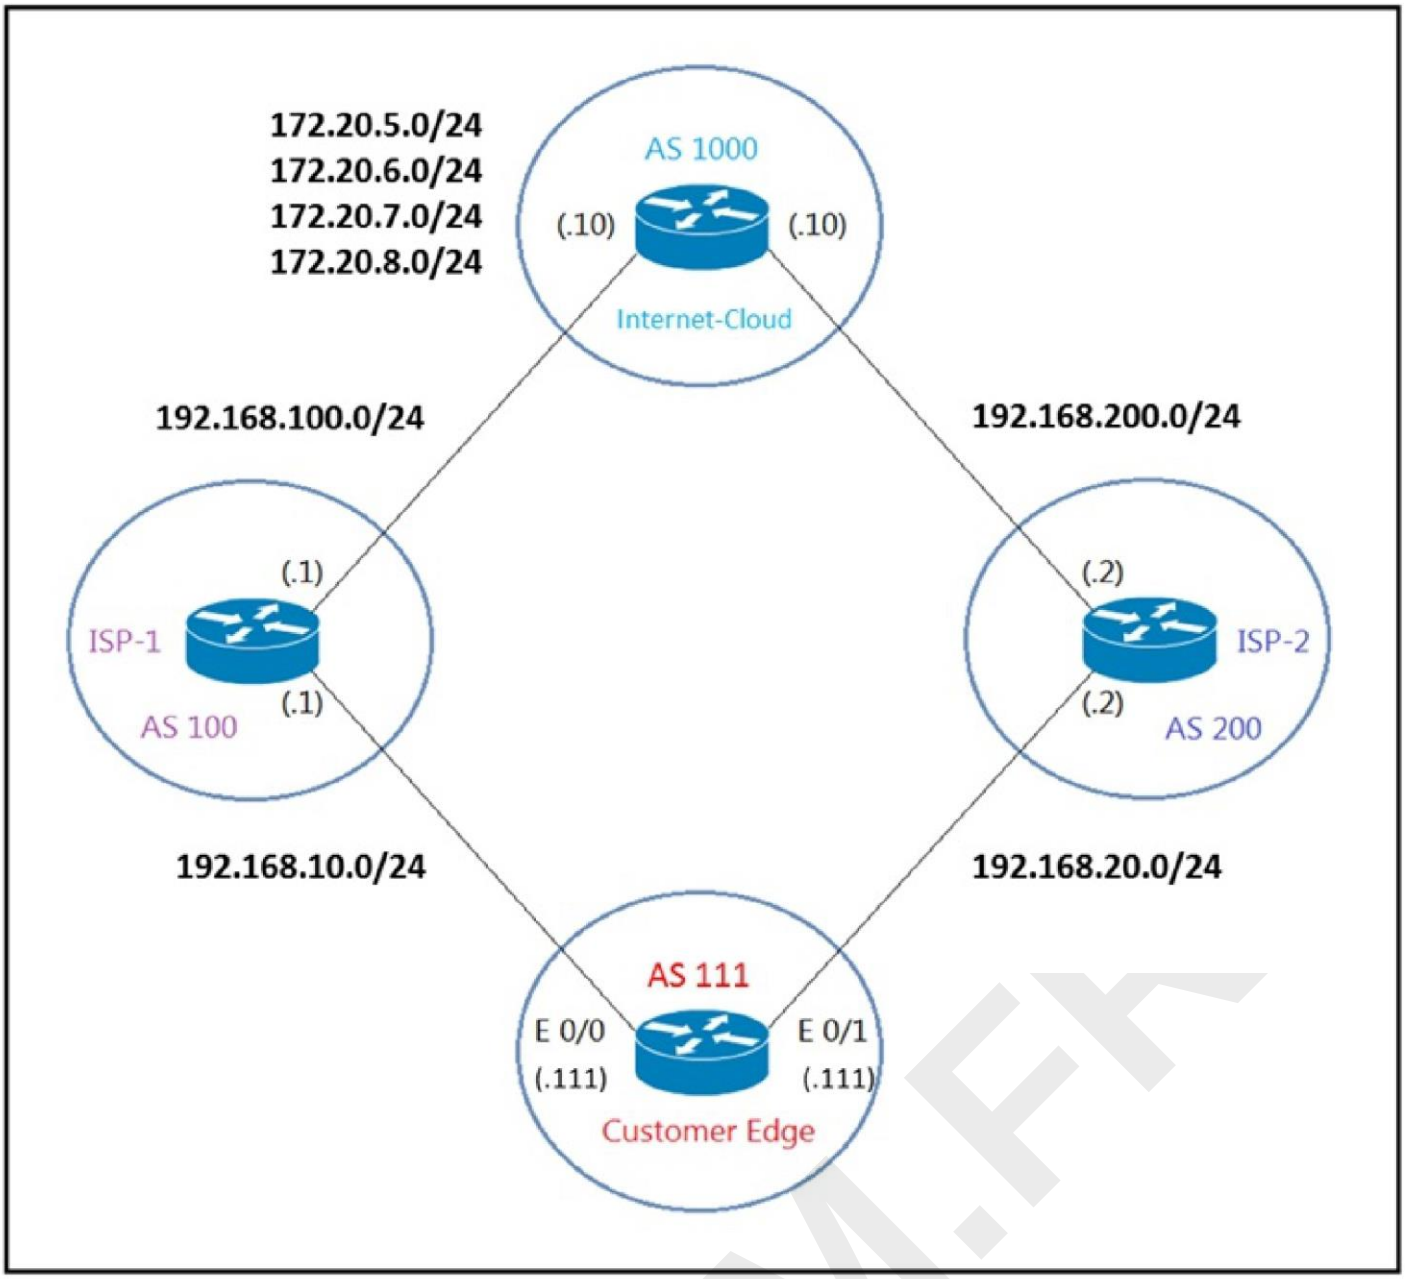

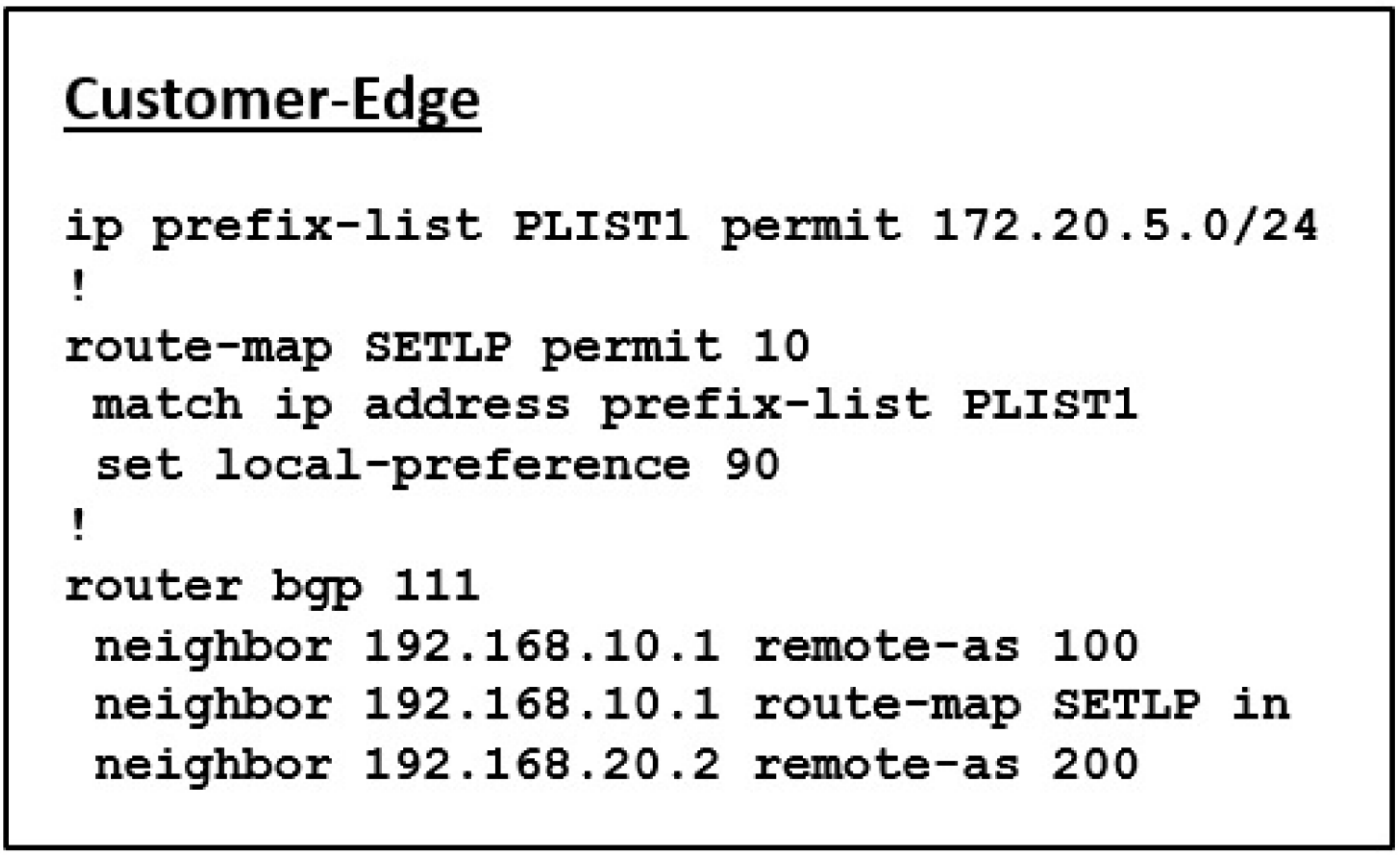

AS 111 wanted to use AS 200 as the preferred path for 172.20.5.0/24 and AS 100 as the backup. After the configuration, AS 100 is not used for any other routes.
Which configuration resolves the issue?

A.route-map SETLP permit 10 match ip address prefix-list PLIST1 set local-preference 99 route-map SETLP permit 20

B.router bgp 111 no neighbor 192.168.10.1 route-map SETLP in neighbor 192.168.20.2 route-map SETLP in

C.route-map SETLP permit 10 match ip address prefix-list PLIST1 set local-preference 110 route-map SETLP permit 20

D.router bgp 111 no neighbor 192.168.10.1 route-map SETLP in neighbor 192.168.10.1 route-map SETLP out

Answer: A

Explanation:
There is an implicit deny all at the end of any route-map so all other traffic that does not match 172.20.5.0/24 would be dropped. Therefore, we have to add a permit sequence at the end of the route-map to allow other traffic.

The default value of Local Preference is 100 and higher value is preferred so we have to set the local preference of AS100 lower than that of AS200.

Question: 72

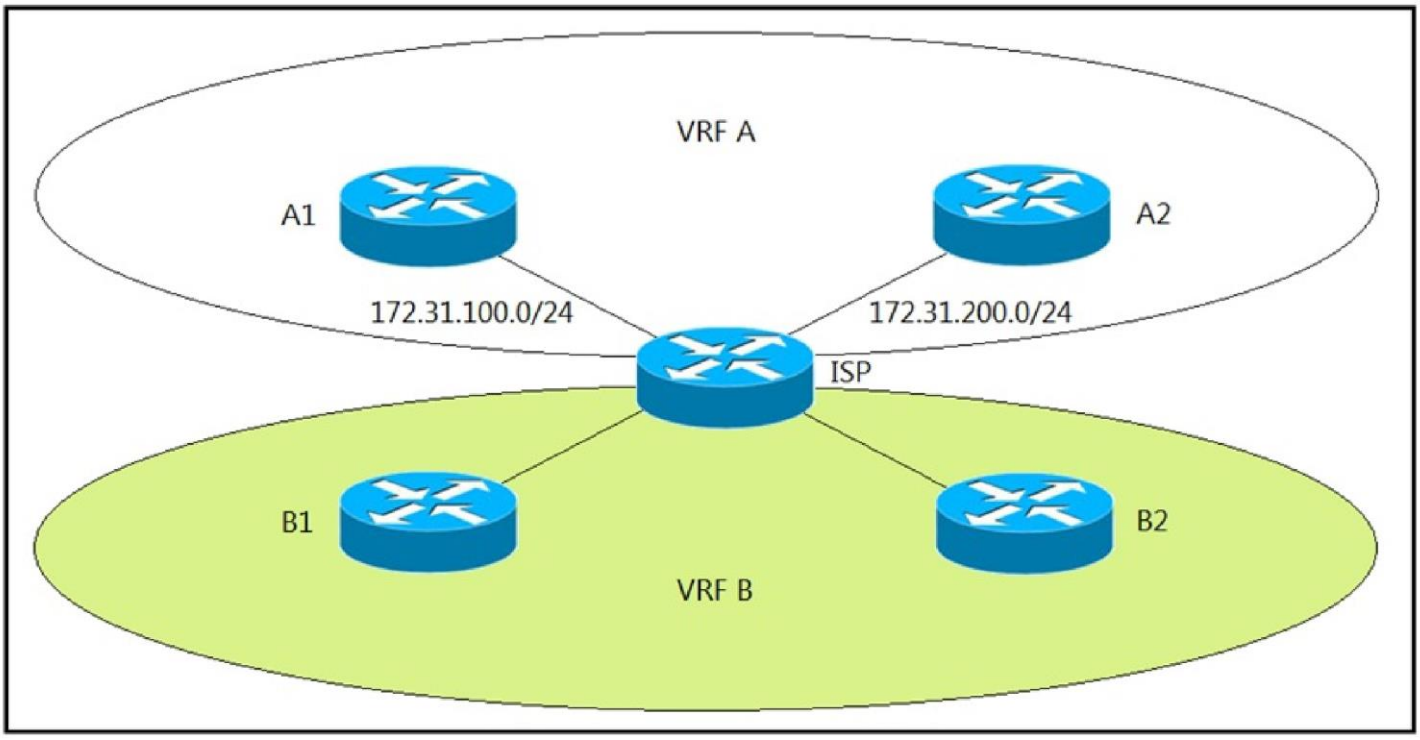

Refer to the exhibit. The ISP router is fully configured for customer A and customer B using the VRF-Lite feature. What is the minimum configuration required for customer A to communicate between routers A1 and A2?

A.A1 interface fa0/0 description To->ISP ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 net 172.31.100.1 0.0.0.255 area 0 A2 interface fa0/0 description To->ISP ip add 172.31.200.1 255.255.255.0 no shut ! router ospf 100 net 172.31.200.1 0.0.0.255 area 0

B.A1 interface fa0/0 description To->ISP ip vrf forwarding A ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 vrf A net 172.31.200.1 0.0.0.255 area 0 A2 interface fa0/0 description To->ISP ip vrf forwarding A ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 vrf A net 172.31.200.1 0.0.0.255 area 0

C.A1 interface fa0/0 description To->ISP ip vrf forwarding A ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 net 172.31.100.1 0.0.0.255 area 0 A2 interface fa0/0 description To->ISP ip vrf forwarding A ip add 172.31.200.1 255.255.255.0 no shut ! router ospf 100 net 172.31.200.1 0.0.0.255 area 0

D.A1 interface fa0/0 description To->ISP ip add 172.31.200.1 255.255.255.0 no shut ! router ospf 100 net 172.31.200.1 0.0.0.255 area 0 A2 interface fa0/0 description To->ISP ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 net 172.31.100.1 0.0.0.255 area 0

Answer: A

Explanation:

A1 interface fa0/0 description To->ISP ip add 172.31.100.1 255.255.255.0 no shut ! router ospf 100 net 172.31.100.1 0.0.0.255 area 0 A2 interface fa0/0 description To->ISP ip add 172.31.200.1 255.255.255.0 no shut ! router ospf 100 net 172.31.200.1 0.0.0.255 area 0

Question: 73

An engineer is implementing a coordinated change with a server team. As part of the change, the engineer must configure interface GigabitEthernet2 in an existing VRF "RED" then move the interface to an existing VRF "BLUE" when the server team is ready. The engineer configured interface GigabitEthernet2 in VRF "RED":

```
interface GigabitEthernet2
description Migration ID: B410A82D0935G35
vrf forwarding RED
ip address 10.0.0.0 255.255.255.254
negotiation auto
```

Which configuration completes the change?

A.interface GigabitEthernet2 no vrf forwarding RED vrf forwarding BLUE ip address 10.0.0.0 255.255.255.254

B.interface GigabitEthernet2 no ip address vrf forwarding BLUE

C.interface GigabitEthernet2 no vrf forwarding RED vrf forwarding BLUE

D.interface GigabitEthernet2 no ip address ip address 10.0.0.0 255.255.255.254 vrf forwarding BLUE

Answer: A

Explanation:

Okay, let's break down why option A is the correct approach for migrating interface GigabitEthernet2 from VRF "RED" to VRF "BLUE".

The initial configuration places GigabitEthernet2 in VRF "RED" with an IP address. To move it to "BLUE", we need to first remove it from "RED". This is achieved by using the no vrf forwarding RED command. Crucially, we then immediately configure the interface to use vrf forwarding BLUE to place it into the new VRF.

Option A uses no vrf forwarding RED to disassociate the interface from "RED", and then immediately associates the interface with vrf forwarding BLUE, while retaining the original IP address. This allows the interface to seamlessly switch to "BLUE" with no loss of network configuration. The IP address remains the same, as specified in the requirements, ensuring there is a minimal interruption to the connection when the server team is ready.

Option B removes the IP address, which is not desired. Option C does not specify the IP address requirement and will remove the IP address as well. Option D removes the IP address first then reconfigures it back to the same IP address but in a different VRF, which is also unnecessary. Only Option A successfully keeps the IP address through the change while changing to the correct VRF. This is the most efficient method that adheres to the prompt's parameters.

Key Concepts:

VRF (Virtual Routing and Forwarding): VRFs allow for multiple routing tables to exist on a single router. This is useful for network segmentation and overlapping IP addressing schemes.

Interface VRF Assignment: Interfaces are assigned to specific VRFs, dictating which routing table is used for their traffic.

no vrf forwarding <VRF_name>: Removes the interface from the specified VRF.

vrf forwarding <VRF_name>: Assigns the interface to the specified VRF.

Authoritative Links:

Cisco Documentation on VRF Configuration:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/configuration/15-mt/iad-15-mt-book/iad-vrf.html

Cisco's Command Reference Guide: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/iad-cr-book/iad-i1.html#wp1042310 (shows the vrf forwarding command).

In summary, option A provides the necessary commands to seamlessly transition the interface from one VRF to another while retaining the existing IP address which is crucial for this coordinated change.

Question: 74

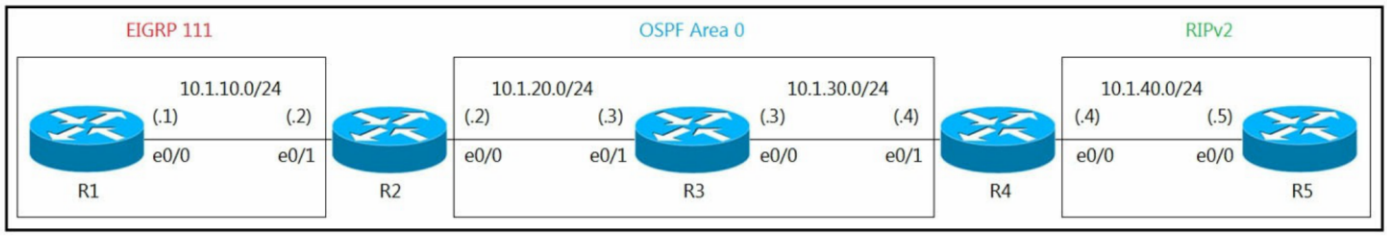


```
R2
route-map E20 permit 10
 set tag 111
!
router eigrp 111
 redistribute ospf 1 metric 10 10 10 10 10
!
router ospf 1
 redistribute eigrp 111 route-map E20 subnets

R4
router rip
 redistribute ospf 1 metric 1
!
router ospf 1
 redistribute rip subnets
```

Refer to the exhibit. R5 should not receive any routes originated in the EIGRP domain. Which set of configuration changes removes the EIGRP routes from the R5 routing table to fix the issue?

A.R4 route-map O2R deny 10 match tag 111 route-map O2R permit 20 ! router rip redistribute ospf 1 route-map O2R metric 1

B.R2 route-map E20 deny 20 R4 route-map O2R deny 10 match tag 111 ! router rip redistribute ospf 1 route-map O2R metric 1

C.R4 route-map O2R permit 10 match tag 111 route-map O2R deny 20 ! router rip redistribute ospf 1 route-map O2R metric 1

D.R4 route-map O2R deny 10 match tag 111 ! router rip redistribute ospf 1 route-map O2R metric 1

Answer: A

Explanation:

R4 route-map O2R deny 10 match tag 111 route-map O2R permit 20 ! router rip redistribute ospf 1 route-map O2R metric 1.

Question: 75

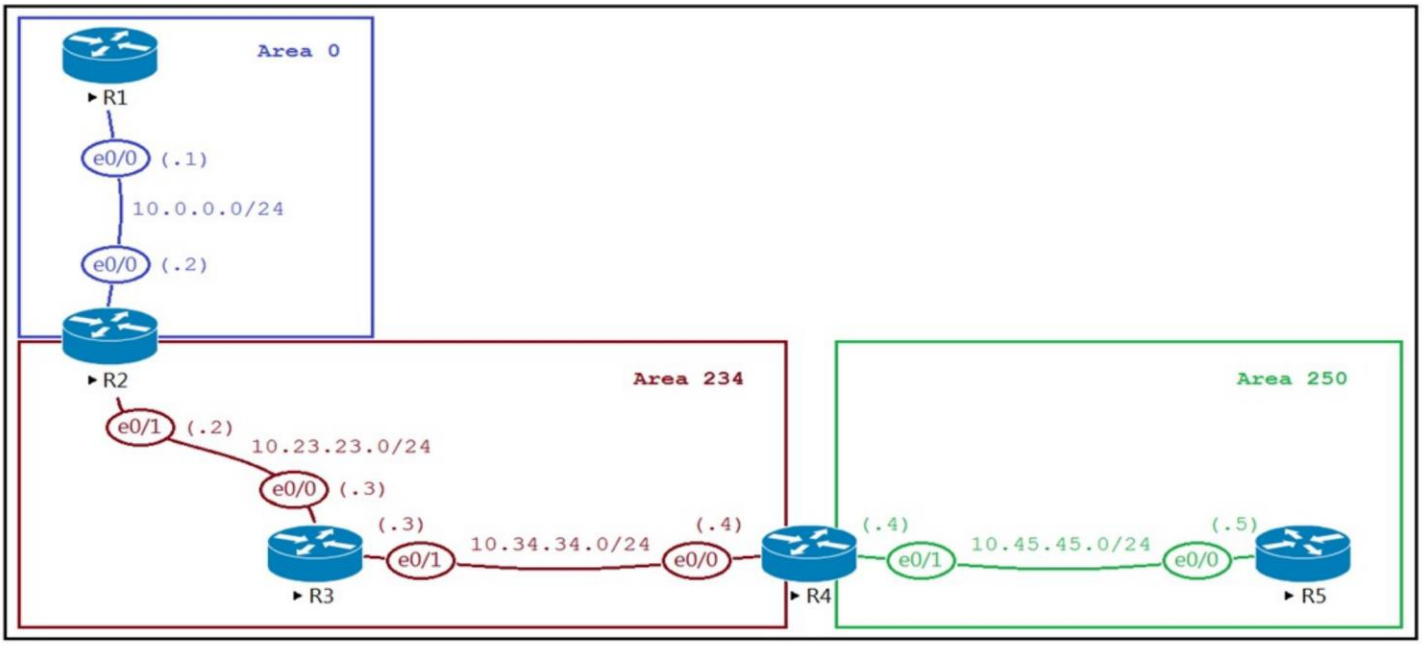

ABR Configurations

R2

```
router ospf 1
 router-id 0.0.0.22
 area 234 virtual-link 10.34.34.4
 network 10.0.0.0 0.0.0.255 area 0
 network 10.2.2.0 0.0.0.255 area 0
 network 10.22.22.0 0.0.0.255 area 234
 network 10.23.23.0 0.0.0.255 area 234
```

R4

```
router ospf 1
 router-id 0.0.0.44
 area 234 virtual-link 10.23.23.2
 network 10.34.34.0 0.0.0.255 area 234
 network 10.44.44.0 0.0.0.255 area 234
 network 10.45.45.0 0.0.0.255 area 250
```

Virtual Link Status

```
R4#sh ip ospf virtual-links
Virtual Link OSPF_VL0 to router 10.23.23.2 is down
Run as demand circuit
DoNotAge LSA allowed.
Transit area 234
Topology-MTID  Cost  Disabled  Shutdown  Topology Name
    0          65535  no        no        Base
Transmit Delay is 1 sec, State DOWN,
```

Refer to the exhibit. The network administrator configured the network to connect two disjointed networks and all the connectivity is up except the virtual link, which causes area 250 to be unreachable.
Which two configurations resolve this issue? (Choose two.)

A.R2 router ospf 1 no area 234 virtual-link 10.34.34.4 area 234 virtual-link 0.0.0.44

B.R2 router ospf 1 no area 234 virtual-link 10.34.34.4 area 0 virtual-link 0.0.0.44

C.R4 router ospf 1 no area 234 virtual-link 10.23.23.2 area 0 virtual-link 0.0.0.22

D.R2 router ospf 1 router-id 10.23.23.2
E.R4 router ospf 1 no area 234 virtual-link 10.23.23.2 area 234 virtual-link 0.0.0.22

Answer: AE

Explanation:

Reference:
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13703-8.html

Question: 76

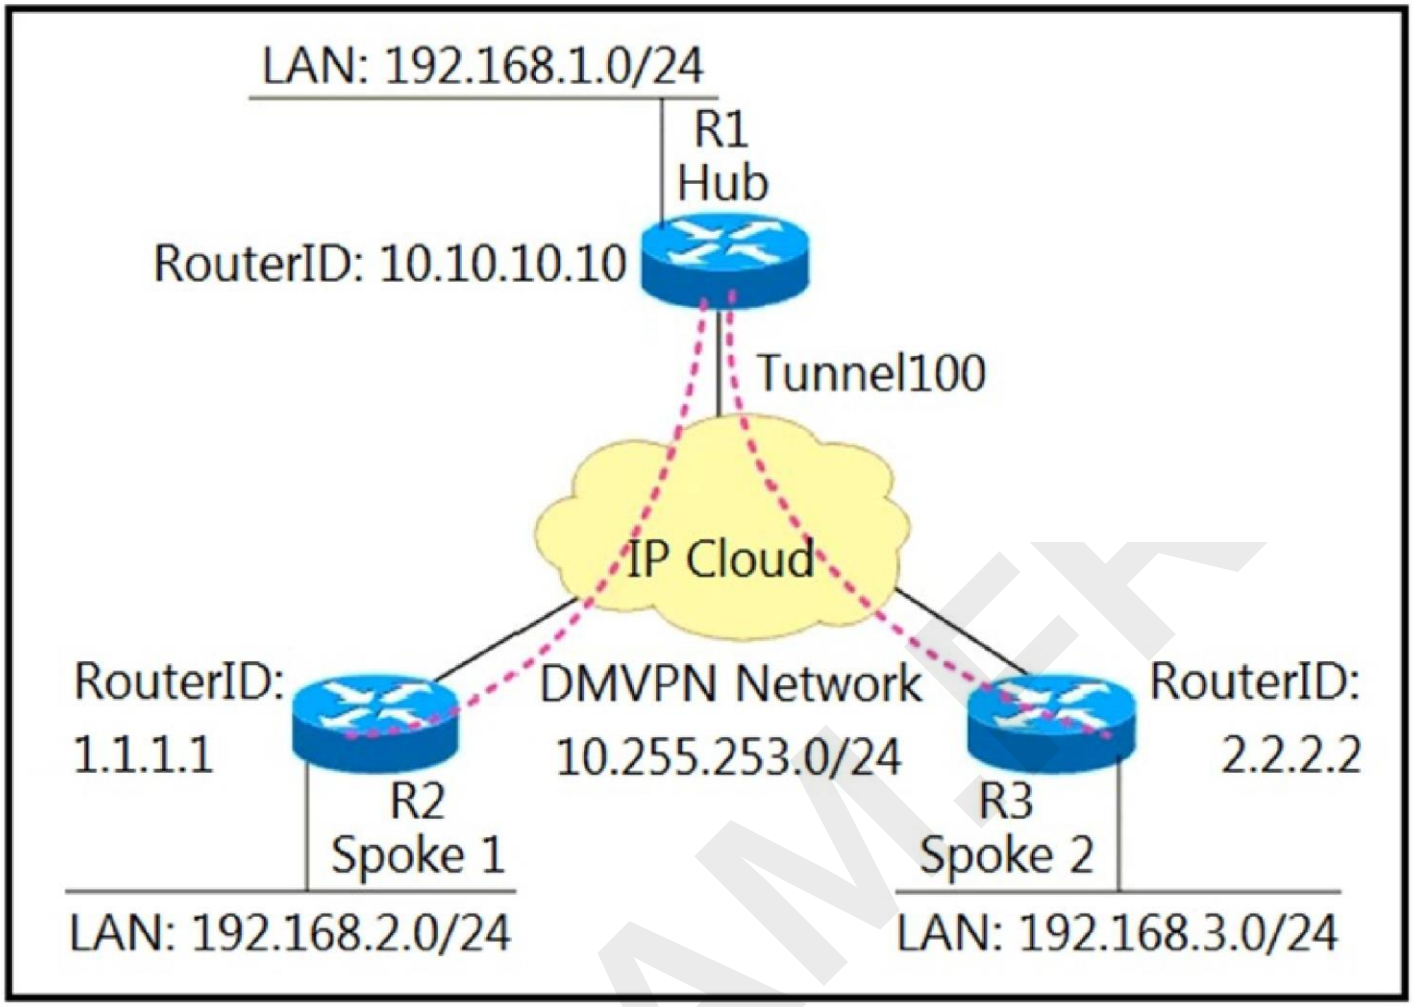


```
*Mar 1 17:19:04.051: %OSPF-5-ADJCHG: Process 100, Nbr 1.1.1.1 on Tunnel100 from LOADING to FULL, Loading Done
*Mar 1 17:19:06.375: %OSPF-5-ADJCHG: Process 100, Nbr 1.1.1.1 on Tunnel100 from FULL to DOWN, Neighbor Down: Adjacency forced to reset
*Mar 1 17:19:06.627: %OSPF-5-ADJCHG: Process 100, Nbr 2.2.2.2 on Tunnel100 from LOADING to FULL, Loading Done
*Mar 1 17:19:10.123: %OSPF-5-ADJCHG: Process 100, Nbr 2.2.2.2 on Tunnel100 from FULL to DOWN, Neighbor Down: Adjacency forced to reset
*Mar 1 17:19:14.499: %OSPF-5-ADJCHG: Process 100, Nbr 10.10.10.10 on Tunnel100 from LOADING to FULL, Loading Done
*Mar 1 17:19:19.139: %OSPF-5-ADJCHG: Process 100, Nbr 10.10.10.10 on Tunnel100 from EXSTART to DOWN, Neighbor Down: Interface down or detached
*Mar 1 17:01:51.975: %OSPF-4-NONEIGHBOR: Received database description from unknown neighbor 192.168.1.1
*Mar 1 17:01:57.783: OSPF: Rcv LS UPD from 192.168.1.1 on Tunnel100 length 88 LSA count 1
*Mar 1 17:01:57.155: OSPF: Send UPD to 10.255.253.1 on Tunnel100 length 100 LSA count 2
```

Refer to the exhibit. A network administrator sets up an OSPF routing protocol for a DMVPN network on the hub router. Which configuration command is required to establish a DMVPN tunnel with multiple spokes?

A.ip ospf network point-to-point on the hub router

B.ip ospf network point-to-multipoint on one spoke router
C.ip ospf network point-to-multipoint on both spoke routers
D.ip ospf network point-to-point on both spoke routers

Answer: C

Explanation:

ip ospf network point-to-multipoint on both spoke routers.

Question: 77

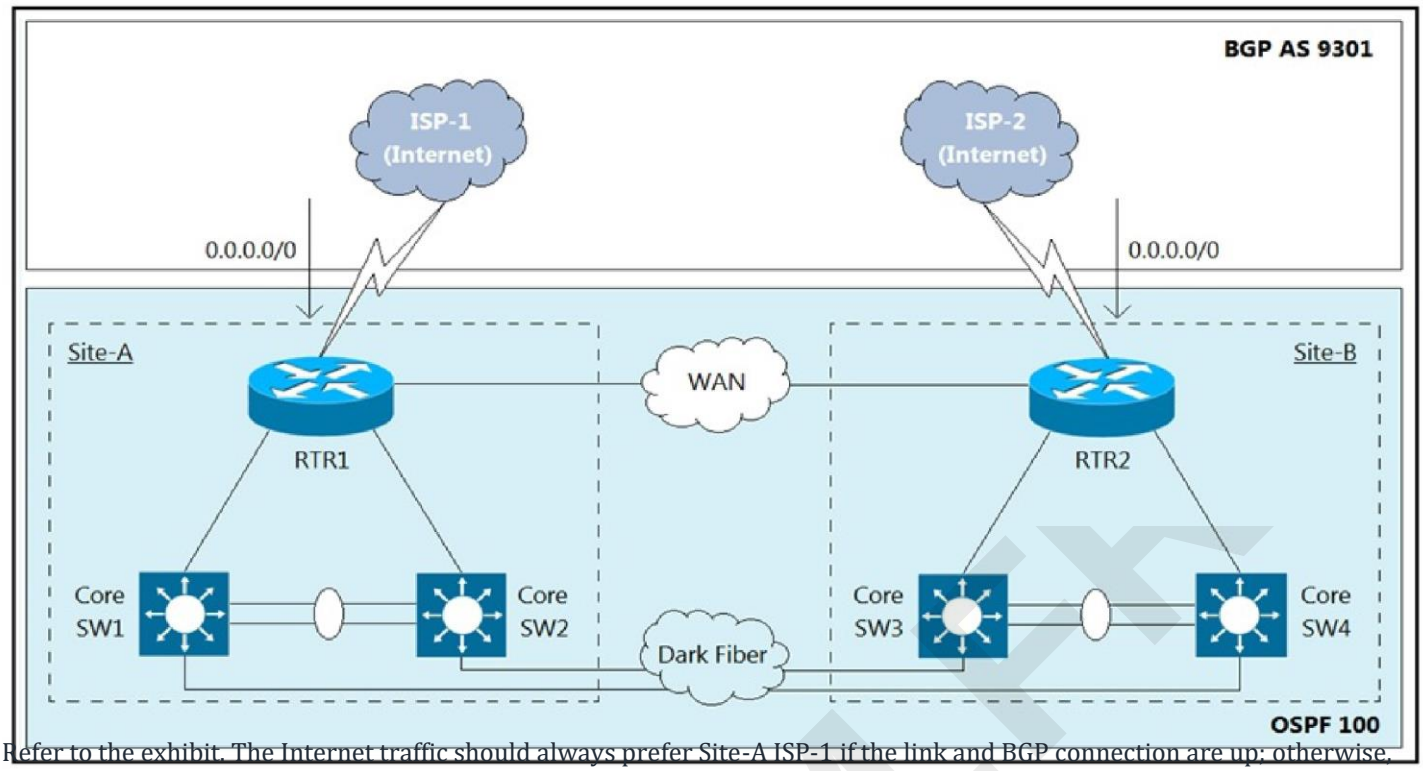

Refer to the exhibit. The Internet traffic should always prefer Site-A ISP-1 if the link and BGP connection are up; otherwise, all Internet traffic should go to ISP-2.
Redistribution is configured between BGP and OSPF routing protocols, and it is not working as expected. What action resolves the issue?

A.Set OSPF Cost 200 at Site-A RTR1, and set OSPF Cost 100 at Site-B RTR2.

B.Set metric-type 2 at Site-A RTR1, and set metric-type 1 at Site-B RTR2.

C.Set metric-type 1 at Site-A RTR1, and set metric-type 2 at Site-B RTR2. D.Set OSPF Cost 100 at Site-A RTR1, and set OSPF Cost 200 at Site-B RTR2.

Answer: C

Explanation:

Set metric-type 1 at Site-A RTR1, and set metric-type 2 at Site-B RTR2.

Question: 78

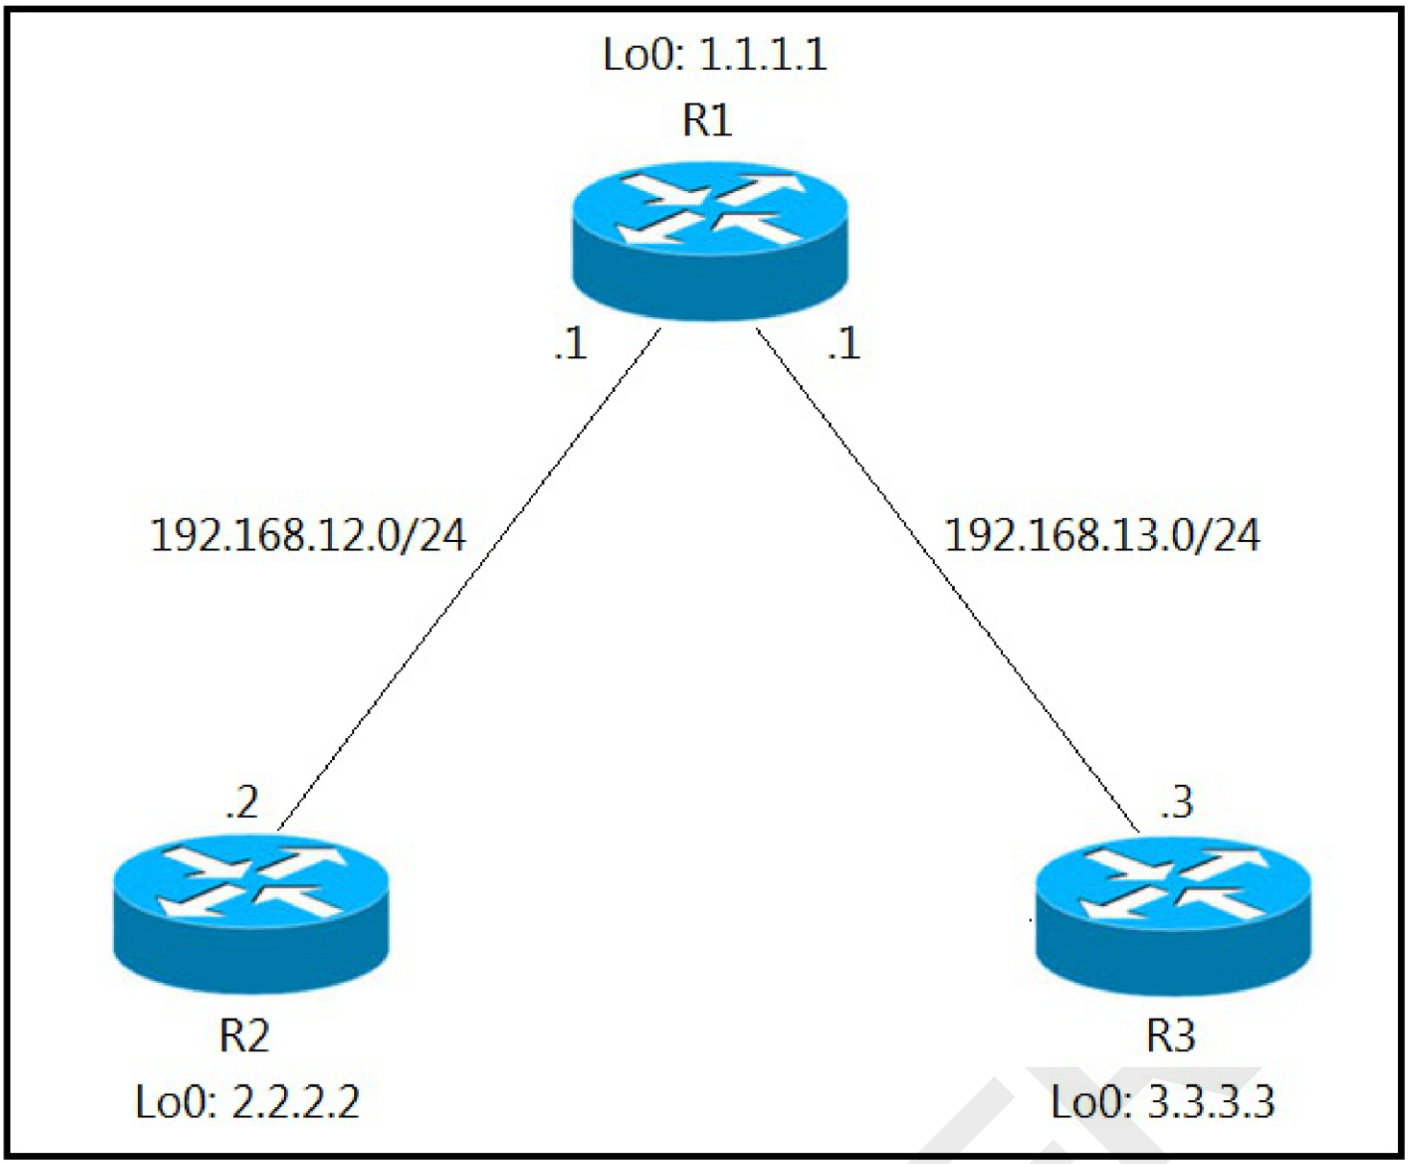

Refer to the exhibit. An engineer has configured R1 as EIGRP stub router. After the configuration, router R3 failed to reach to R2 loopback address.
Which action advertises R2 loopback back into the R3 routing table?

A.Add a static route for R2 loopback address in R1 and redistribute it to advertise to R3.

B.Use a leak map on R1 that matches the required prefix and apply it with the distribute list command toward R3.

C.Use a leak map on R3 that matches the required prefix and apply it with the EIGRP stub feature.

D.Add a static null route for R2 loopback address in R1 and redistribute it to advertise to R3.

Answer: B

Explanation:

Use a leak map on R1 that matches the required prefix and apply it with the distribute list command toward R3.

Question: 79

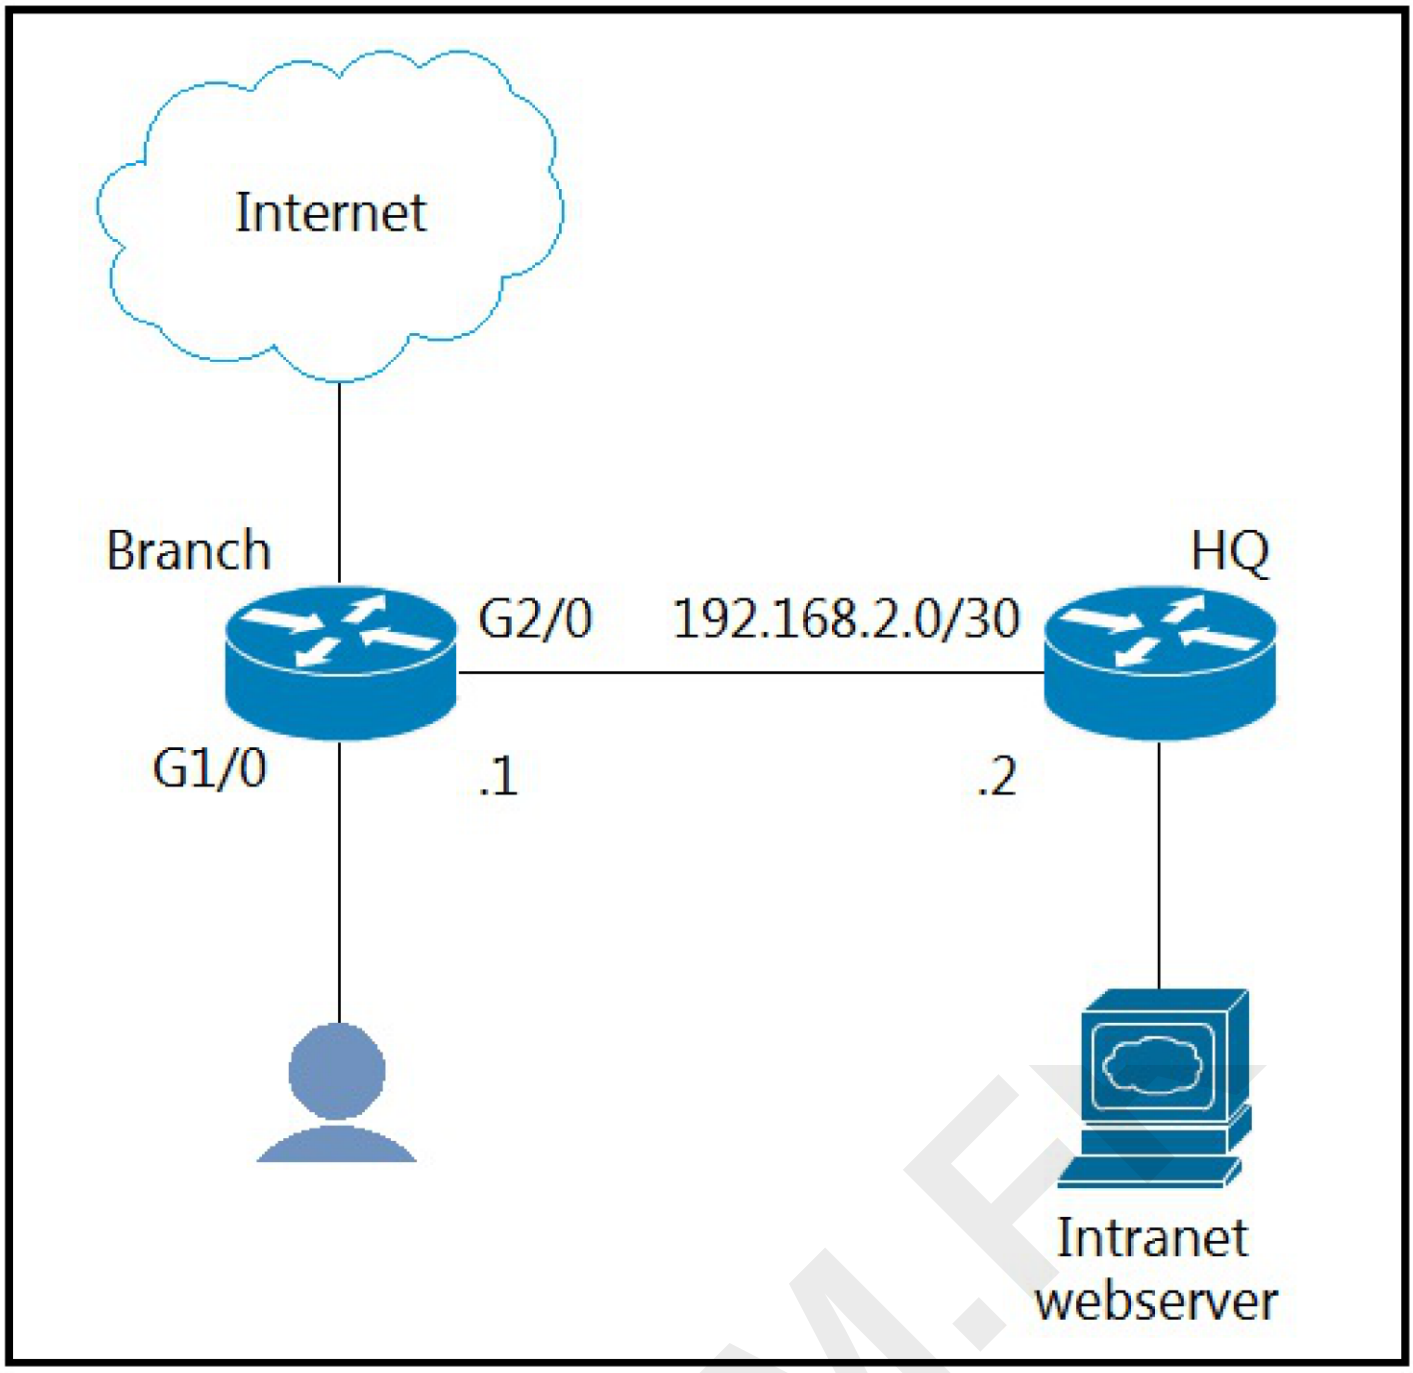

Refer to the exhibit. The branch router is configured with a default route toward the Internet and has no routes configured for the HQ site that is connected through interface G2/0. The HQ router is fully configured and does not require changes.
Which configuration on the branch router makes the intranet website (TCP port 80) available to the branch office users?

A.access-list 101 permit tcp any any eq 80 access-list 102 permit tcp any host intranet-webserver-ip ! route-map pbr permit 10 match ip address 101 set ip next-hop 192.168.2.2 route-map pbr permit 20 match ip address 102 set ip next-hop 192.168.2.2 ! interface G2/0 ip policy route-map pbr

B.access-list 100 permit tcp host intranet-webserver-ip eq 80 any ! route-map pbr permit 10 match ip address 100 set ip next-hop 192.168.2.2 ! interface G1/0 ip policy route-map pbr

C.access-list 100 permit tcp any host intranet-webserver-ip eq 80 ! route-map pbr permit 10 match ip address 100 set ip next-hop 192.168.2.2 ! interface G2/0 ip policy route-map pbr

D.access-list 101 permit tcp any any eq 80 access-list 102 permit tcp any host intranet-webserver-ip ! route-map pbr permit 10 match ip address 101 102 set ip next-hop 192.168.2.2 ! interface G1/0 ip policy route-map pbr

Answer: D

Explanation:

access-list 101 permit tcp any any eq 80 access-list 102 permit tcp any host intranet-webserver-ip ! route-map

pbr permit 10 match ip address 101 102 set ip next-hop 192.168.2.2 ! interface G1/0 ip policy route-map pbr.

Question: 80

R1 and R2 are configured as eBGP neighbors. R1 is in AS100 and R2 is in AS200. R2 is advertising these networks to R1:
172.16.16.0/20
172.16.3.0/24
172.16.4.0/24
192.168.1.0/24
192.168.2.0/24
172.16.0.0/16
The network administrator on R1 must improve convergence by blocking all subnets of 172.16.0.0/16 major network with a mask lower than 23 from coming in.
Which set of configurations accomplishes the task on R1?

A.ip prefix-list PL-1 deny 172.16.0.0/16 ge 23 ip prefix-list PL-1 permit 0.0.0.0/0 le 32 ! router bgp 100 neighbor 192.168.100.2 remote-as 200 neighbor 192.168.100.2 prefix-list PL-1 in

B.ip prefix-list PL-1 deny 172.16.0.0/16 le 23 ip prefix-list PL-1 permit 0.0.0.0/0 le 32 ! router bgp 100 neighbor 192.168.100.2 remote-as 200 neighbor 192.168.100.2 prefix-list PL-1 in

C.ip prefix-list PL-1 deny 172.16.0.0/16 ip prefix-list PL-1 permit 0.0.0.0/0 ! router bgp 100 neighbor 192.168.100.2 remote-as 200 neighbor 192.168.100.2 prefix-list PL-1 in

D.access-list 1 deny 172.16.0.0 0.0.254.255 access-list 1 permit any ! router bgp 100 neighbor 192.168.100.2 remote-as 200 neighbor 192.168.100.2 distribute-list 1 in

Answer: A

Explanation:

Okay, let's break down why option A is the correct configuration for filtering BGP routes on R1.

The goal is to block any subnet of 172.16.0.0/16 with a mask lower than /23 (meaning /16 to /22) from being received by R1 from R2. This is achieved using a prefix list.

Prefix List Logic

ip prefix-list PL-1 deny 172.16.0.0/16 ge 23: This is the core of the filtering. It says, "Deny any network that matches 172.16.0.0/16 AND has a subnet mask length that is greater than or equal to 23." In other words, any subnet with a mask of /23 or higher is not denied by this line. However, it denies anything matching /16, /17, /18, /19, /20, /21, or /22.
ip prefix-list PL-1 permit 0.0.0.0/0 le 32: This is a catch-all permit statement. It says "Permit any network with any mask length". It ensures that any other networks not explicitly denied are allowed. This line is crucial for the filtering functionality, preventing all other networks from being blocked. Without the permit all statement, implicit deny will exist at the end and nothing will pass.

BGP Configuration

router bgp 100: This enters the BGP configuration mode for AS 100.
neighbor 192.168.100.2 remote-as 200: This establishes a BGP peering relationship with R2 (192.168.100.2) which is in AS 200.
neighbor 192.168.100.2 prefix-list PL-1 in: This applies the previously defined prefix-list "PL-1" to incoming route updates from R2. This means all routes sent by R2 are evaluated against the PL-1 and filtered before entering the BGP routing table on R1.

Why Other Options Are Incorrect

Option B: This uses le 23 (less than or equal to 23). It would deny all subnets matching the 172.16.0.0/16 range, including /23 and /24 networks, which is not the goal.

Option C: This option is incorrect because it is missing "permit all" statement. If the permit all statement is not included, the prefix list will cause all routes to be filtered.

Option D: This uses an access list, which is less suitable for prefix matching based on variable mask lengths. Also, distribute-lists are not best practices with BGP.

Authoritative Links

Cisco Prefix Lists:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/15-sy/irg-15-sy-book/irg-pref-lists.html
BGP Route Filtering:https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13707-bgp-filter.html

In conclusion, option A precisely implements the required filtering policy by using a prefix list to deny the specific subnets and allow all the others, ultimately improving network convergence by reducing the number of advertisements being processed by R1.