

➤ **Vendor: Cisco**

➤ **Exam Code: 300-510**

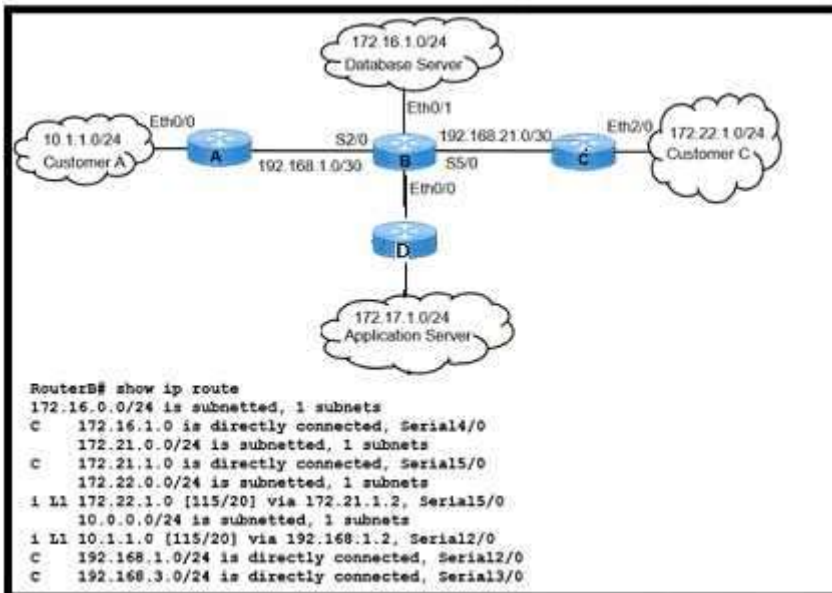
➤ **Exam Name: Implementing Cisco Service Provider Advanced Routing Solutions (SPRI)**

➤ **New Updated Questions from [Braindump2go](https://www.braindump2go.com) (Updated in [August/2022](#))**

Visit Braindump2go and Download Full Version 300-510 Exam Dumps

QUESTION 147

Refer to the exhibit. Customers A and C are experiencing packet drops when connecting to the application server. While troubleshooting the problem the network engineer confirms that the IS-IS Level-1/2 adjacency is up between routers A, B, and D and both customers can communicate with the database server without packet loss. Which action must the engineer take to resolve the issue?



- A. Advertise the application server subnet in the router D IS-IS database
- B. Advertise a static default route to the router B IS-IS database.
- C. Leak the 172.17.1.0/24 route in the IS-IS databases on routers A and C.
- D. Leak the customer A and customer C subnets in the router A IS-IS database.

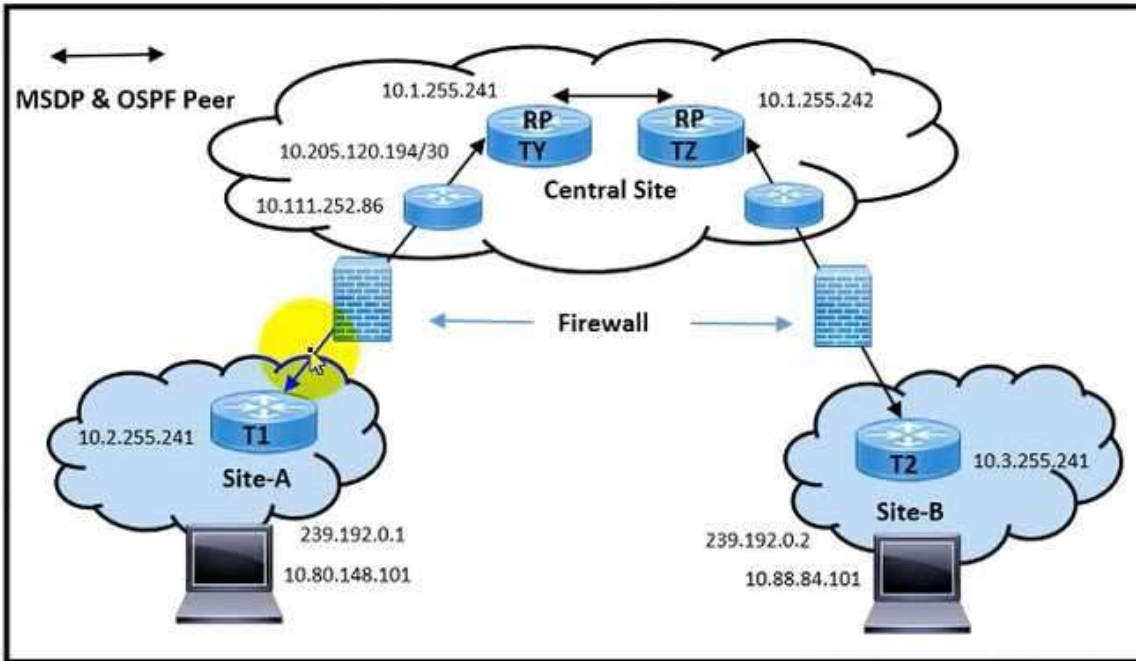
Answer: A

QUESTION 148

Refer to the exhibit. Multicast traffic destined from T1 and T2 routers to RP routers works well. A network engineer observes problems with multicast traffic flows between Site-A and Site-B. Site-A users fail to receive multicast stream on Site-B via RPTY site, while Site-B users fail to receive multicast stream on Site-A via RPTZ site. Which action must be implemented to resolve the issues?

[300-510 Exam Dumps](#) [300-510 Exam Questions](#) [300-510 PDF Dumps](#) [300-510 VCE Dumps](#)

<https://www.braindump2go.com/300-510.html>



```
TZ# show ip msdp sa-cache rejected-SA det read-only <snip>
86854209.328, (10.80.148.101, 239.192.0.1), RP: 10.2.255.241, Peer:
10.1.255.241, Reason: rpf-fail -> learned from central site RT1 but not
accepted (originated from site A RT1)
86854209.328, (10.88.84.101, 239.192.0.2), RP: 10.3.255.241, Peer:
10.1.255.241, Reason: rpf-fail -> learned from central site RT1 but not
accepted (originated from site B RT1)

TZ# show ip rpf 10.1.255.241
RPF information for ? (10.1.255.241)
RPF interface: Vlan10
RPF neighbor: ? (10.111.254.9)
RPF route/mask: 10.1.255.241/32
RPF type: unicast (ospf 15)
Doing distance-preferred lookups across tables
RPF topology: ipv4 multicast base, originated from ipv4 unicast base

TZ# show ip route 10.1.255.241
Routing Table: CENT1
Routing entry for 10.1.255.241/32
Known via "ospf 15", distance 110, metric 3, type intra area
Last update from 10.111.254.9 on Vlan10, 1d22h ago
Routing Descriptor Blocks:
* 10.111.254.9, from 10.205.0.197, 1d22h ago, via Vlan10
Route metric is 3, traffic share count is 1
```

```

TY# sh ip mosp sa-cache
MSDP Source-Active Cache - 2 entries
(10.80.148.101, 239.192.0.1), RP 10.2.255.241, AS ?,ld23h/00:05:42, Peer
10.2.255.241 -> learned from RT1 at site A (which is 10.2.255.241)
(10.88.84.101, 239.192.0.2), RP 10.3.255.241, AS ?,ld21h/00:05:31, Peer
10.3.255.241 -> learned from RT1 at site B (which is 10.3.255.241)

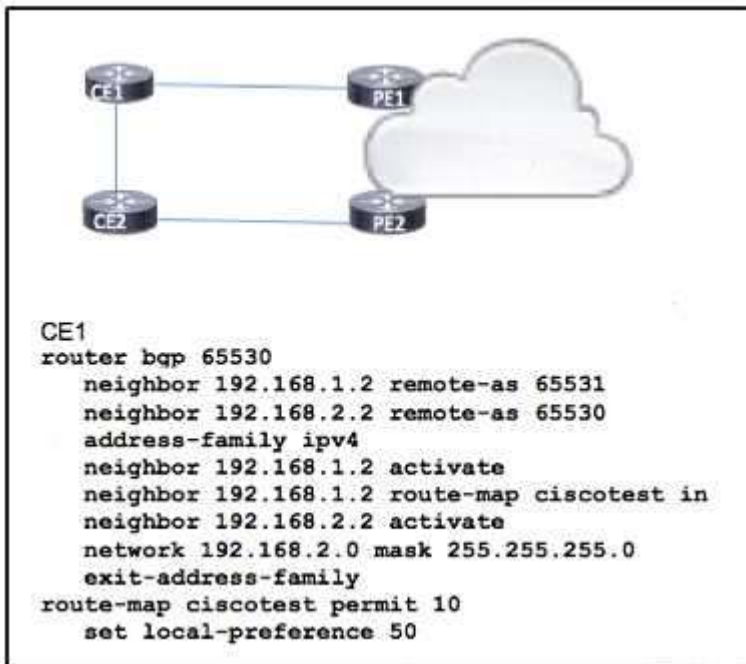
TY# sh ip rpf 10.2.255.241
RPF information for ? (10.2.255.241)
RPF interface: Fo9/1.1035
RPF neighbor: ? (10.111.252.86)
RPF route/mask: 10.2.255.241/32
RPF type: unicast (ospf 15)
Doing distance-preferred lookups across tables
RPF topology: ipv4 multicast base, originated from ipv4 unicast base
|
TY# sh ip route 10.2.255.241
Routing Table: CLNT1
Routing entry for 10.2.255.241/32
Known via "ospf 15", distance 110, metric 150, type extern 2, forward
metric 2
Last update from 10.111.252.86 on FortyGigabitEthernet9/1.1035, 04:06:26
ago
Routing Descriptor Blocks:
* 10.111.252.86, from 10.205.120.195, 04:06:26 ago, via
FortyGigabitEthernet9/1.1035
Route metric is 150, traffic share count is 1
  
```

- A. Establish MSDP peering with interface IP subnet.
- B. Configure Site-A and Site-B in 10.80.14804
- C. Allow the OSPF and MSDP packets on the firewall.
- D. Configure direct OSPF peering between Site-A and Ste-B

Answer: C

QUESTION 149

Refer to the exhibit. Routers CE1 and CE2 are in AS 65530. which is multihomed for Internet access. An engineer expects inbound traffic to AS 65530 to arrive from PE1. But it is coming from PE2 instead PE1 and PE2 routers are connected with CE routers through the same bandwidth. Which action must be taken to correct the problem?



- A. On router CE2, configure inbound routes from PE2 to CE2 with a local-preference value of 50 or

[300-510 Exam Dumps](#) [300-510 Exam Questions](#) [300-510 PDF Dumps](#) [300-510 VCE Dumps](#)

<https://www.braindump2go.com/300-510.html>

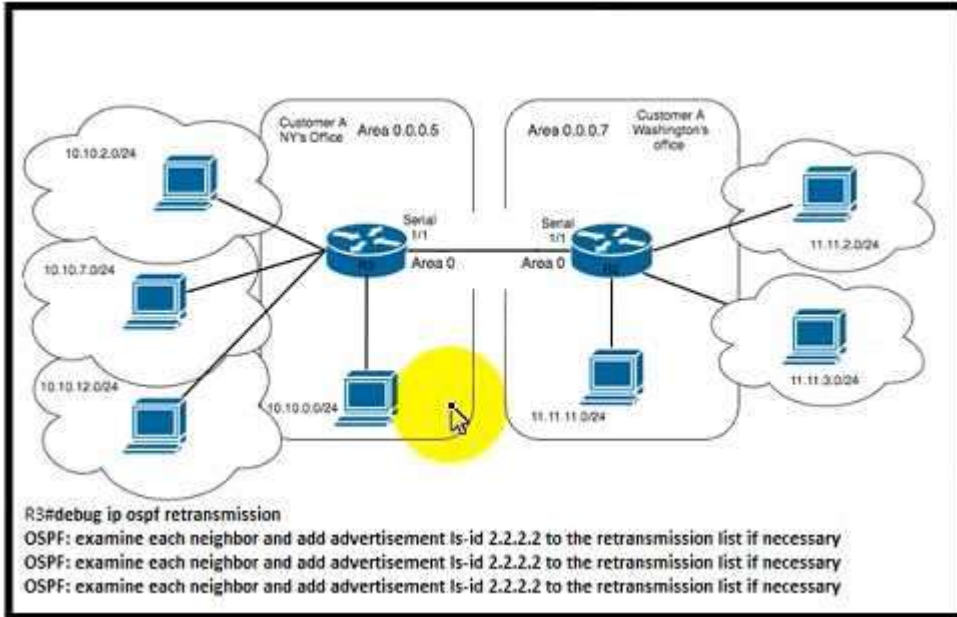
greater.

- B. Configure router CE1 to prepend the AS path to routes it receives from PE 1.
- C. Set the local-preference value on router CE1 to 100 or greater
- D. On router PE1, change the origin for routes that are redistributed from CE1 to CE2.

Answer: C

QUESTION 150

Refer to the exhibit. Customer A is a small media company with two offices connected by a 512 Kbps line. Their NY office is connected to several external partners by static routes on router R3. VoIP services use VoIP codec G729. Users reported poor voice quality and slow data transfer between the offices. A network engineer configured ip tcp header-compression iphc-format on R2 and R3 routers. Which additional action must the engineer take to fix the issue?

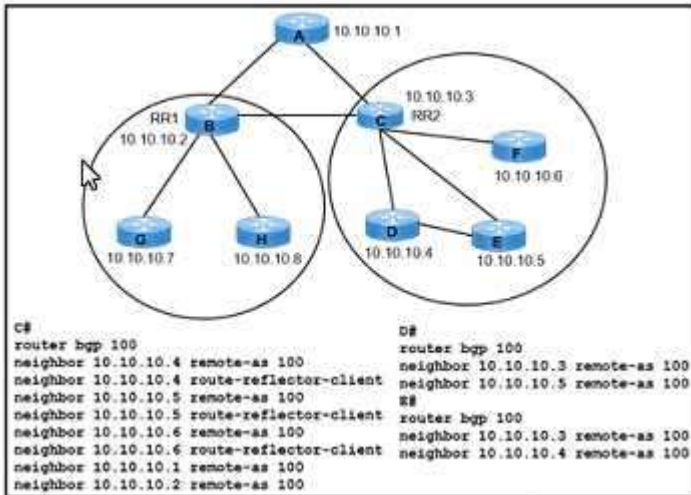


- A. Configure the ip ospf l area O command under Serial 1/1 interfaces on R2 and R3 to avoid routing loops
- B. Change the OSPF router ID on either router so that the router IDs are unique.
- C. Configure the summary-address 10.10.0.0 255.255.240.0 command on R3 to optimize OSPF communication
- D. Configure the BGP routing protocol between R2 and R3 to control route propagation.

Answer: C

QUESTION 151

Refer to the exhibit. While troubleshooting a networking issue an engineer identified a suboptimal communication issue on route reflector RR2. In the current environment Router A is a non-route-reflector client for RR1 and RR2. Routers D and E are directly connected iBGP peers. Router F is not an iBGP peer of routers D and E. Which action resolves the issue?

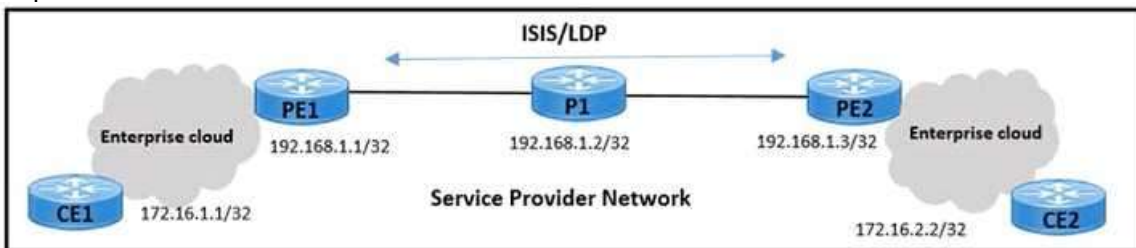


- A. Disable BGP Client-to-Client reflection on router RR2.
- B. Enable next-hop-self for BGP peering on router C.
- C. Remove the route-reflector configuration on router RR2.
- D. Enable next-hop-self for BGP peering on router D.

Answer: B

QUESTION 152

Refer to the exhibit. An engineer working for a private telecommunication company with an employee id 4115 46 881 is enabling a segment routing solution with these requirements. A service provider is using the default range for prefix SID. PE1 must allocate the first SID from the default range for the loopback address PE1 and PE2 loopback SID allocation should have a minimum difference of 500. Which configuration must be implemented to meet the requirements?



- ☐ PE1(config-isis-if-af)# adjacency-sid absolute 16201
PE2(config-isis-if-af)# adjacency-sid absolute 16710
- ☐ PE1(config-isis-if-af)# prefix-sid absolute 16001
PE2(config-isis-if-af)# prefix-sid index 610
- ☐ PE1(config-isis-if-af)# prefix-sid absolute 16201
PE2(config-isis-if-af)# prefix-sid absolute 16710
- ☐ PE2(config-isis-if-af)# adjacency-sid absolute 16001
PE1(config-isis-if-af)# adjacency-sid index 610

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

Answer: B

QUESTION 153

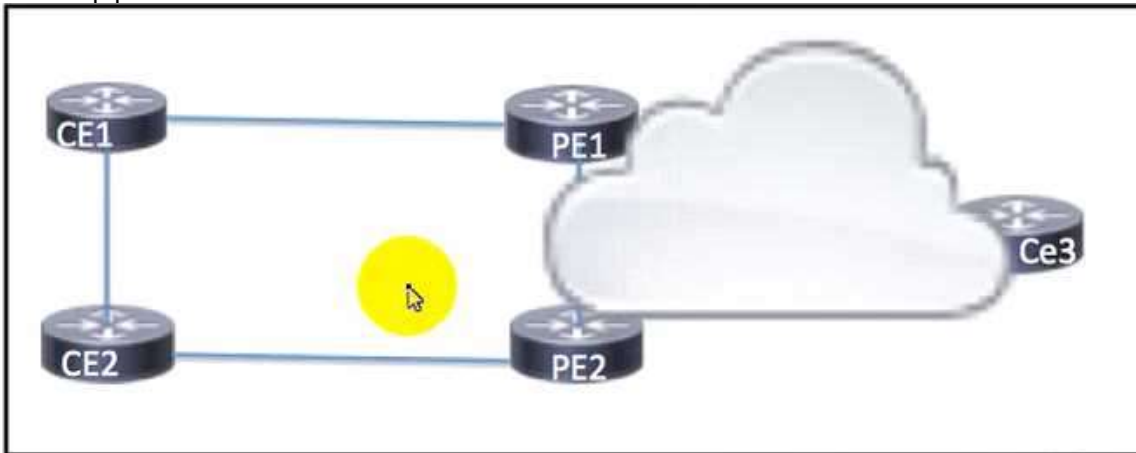
What is a requirement of PIM-SM?

- A. It requires Cisco Express Forwarding to be enabled.
- B. It must be enabled on loopback interfaces only
- C. It requires OSPF to be configured on the network.
- D. It must use an RP

Answer: D

QUESTION 154

Refer to the exhibit. CE1 and CE2 use connectivity over the service provider cloud to reach CE3. In the event of a link failure in the service provider cloud which BGP feature relies on IGP convergence to quickly assist in the installation of a backup path?

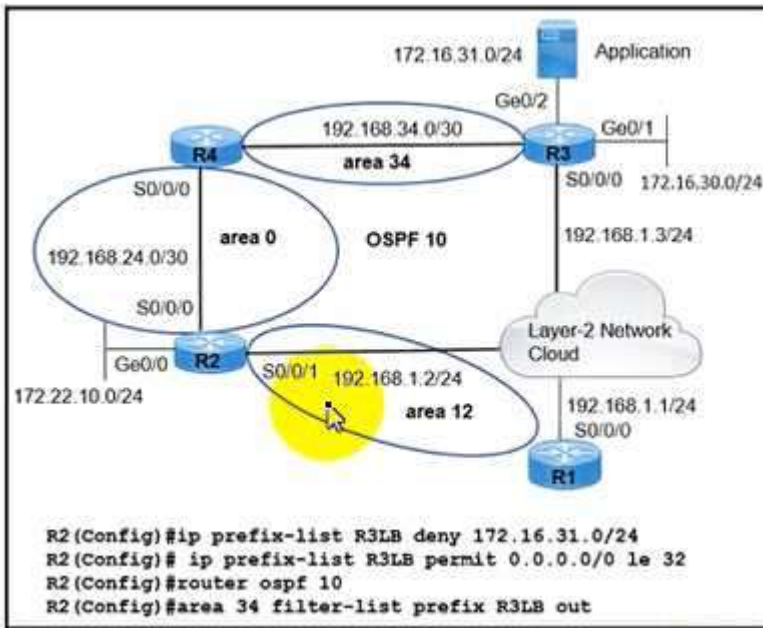


- A. BGP graceful restart
- B. BGP route dampening
- C. BGP PIC core
- D. BGP confederations

Answer: C

QUESTION 155

Refer to the exhibit. Networks 172.16.31.0/24 and 172.16.30.0/24 are advertised in area 34. and network 172.22.10.0/24 is advertised in area 0. A recent security review discovered that users connected to routers R1 and R2 have been making unauthorized access to an application running on network 172.16.31.0/24. An engineer determined that routers R1 and R2 are receiving updates for network 172.16.31.0/24. Which action resolves the issue?

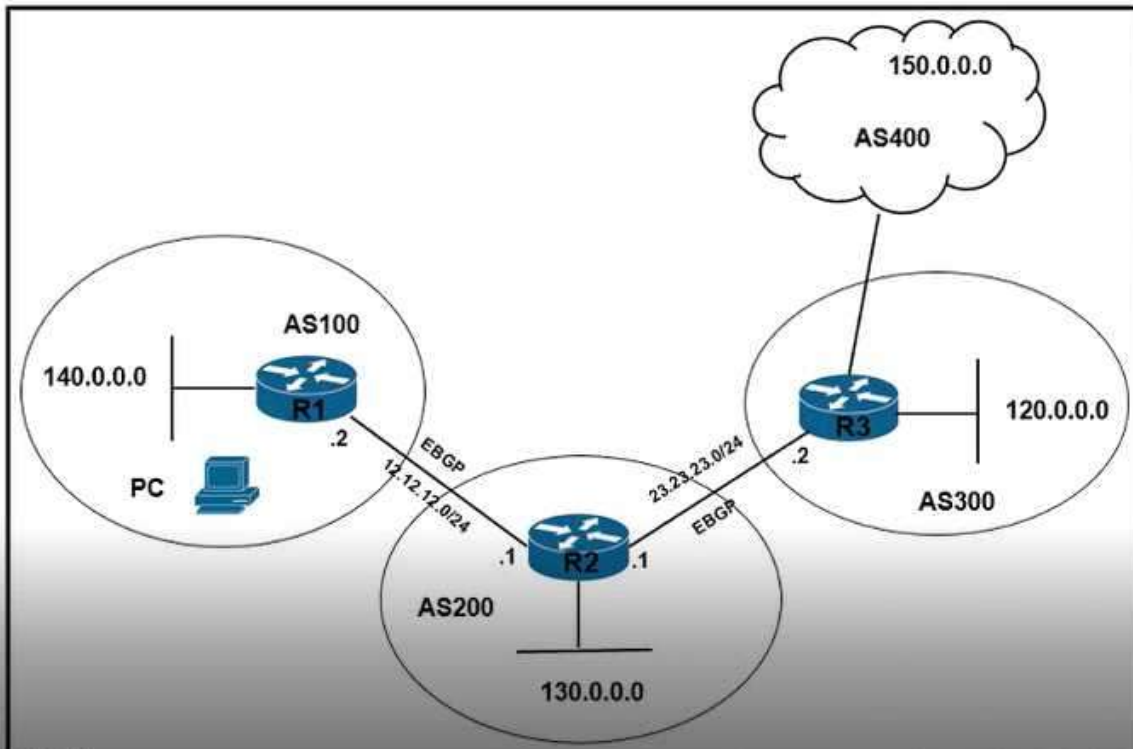


- A. Apply route filtering on routers R3 and R4.
- B. Apply route filtering on router R3 only.
- C. Apply route filtering on routers R1 and R2
- D. Apply route filtering on router R4 only.

Answer: D

QUESTION 156

Refer to the exhibit. Excessive routes are flooding from network 150.0.0.0 into AS100. Internet traffic between AS400 and AS300 is working normally. No route controlling mechanism is applied on incoming and outgoing traffic Which configuration resolves the issue?



- R2#router bgp 200
neighbor 12.12.12.2 remote-as 100
neighbor 23.23.23.2 remote-as 300
neighbor 12.12.12.12 filter-list 1 out
ip as-path access-list 1 deny ^400\$
ip as-path access-list 1 permit ."
- R2#router bgp 200
address-family ipv4 unicast
neighbor 12.12.12.2 remote-as 100
neighbor 12.12.12.2 activate
neighbor 12.12.12.2 route-map PREPEND out
exit-address-family
exit
route-map PREPEND permit 10
set as-path prepend 100 100
- R2#router bgp 200
neighbor 12.12.12.2 route-map FLOODING out
ip as-path access-list 1 permit ^400_
route-map FLOODING permit 10
match as-path 1
set metric 50000
- R1#router bgp 100
neighbor 12.12.12.1 remote-as 200
neighbor 12.12.12.1 route-map SET-LOCAL-PREF in
route-map SET-LOCAL-PREF permit 10
match ip address 2
set local-preference 700
route-map SET-LOCAL-PREF permit 20
access-list 2 permit 150.0.0.0 0.255.255.255
access-list 2 deny any

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

Answer: A

QUESTION 157

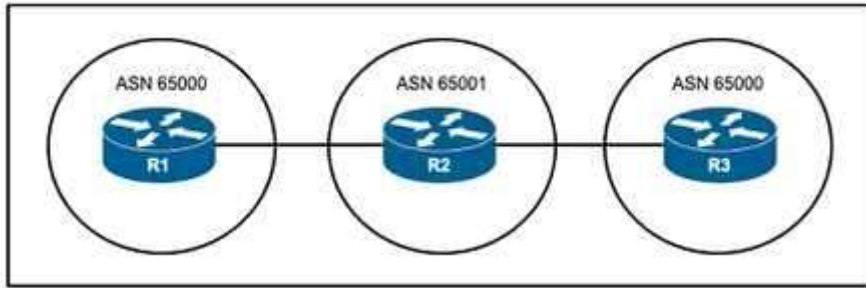
An engineer wants to map a multicast IP address to a multicast MAC.
How many bits are used to make the conversion?

- A. high-order 24 bits
- B. higher-order 23 bits
- C. low order 23 bits
- D. lower-order 24 bits

Answer: C

QUESTION 158

Refer to the exhibit. An engineer is troubleshooting an issue with this network and notices that prefixes from R3 are missing on the R1 routing table. Due to repeated ASN when the 10.0.0.0/8 prefix from R3 arrives at R1, BGP automatically rejects it. There is no prefix-list on R1 which blocks the traffic from R3. What should the engineer do to fix the problem so that BGP allows that prefix on R1?



- A. Configure R2 as a route reflector client of R1.
- B. Configure the allowas-in command on R1.
- C. Configure the next-hop-self command on R2.
- D. Configure identical confederation ASNs on R1 and R2.

Answer: B

QUESTION 159

Refer to the exhibit. A network engineer configured the redistribute connected subnets route-map filtering command on R1 to redistribute connected interfaces to the OSPF process. The engineer also wants to filter out IP address 10.0.1.0/24, but the prefix still appears in the routing tables of the other routers on the network. Which action corrects the problem?

```

R1#show route-map
route-map filtering, permit, sequence 10
  Match clauses:
    ip address (access-lists): 1
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
route-map filtering, deny, sequence 20
  Match clauses:
    ip address (access-lists): 2
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes
route-map filtering, permit, sequence 30
  Match clauses:
  Set clauses:
  Policy routing matches: 0 packets, 0 bytes

R1#show access-lists
Standard IP access list 1
  10 permit 10.0.0.0, wildcard bits 0.0.0.255 (8 matches)
Standard IP access list 2
  10 deny 10.0.1.0, wildcard bits 0.0.0.255 (1 match)
  
```

- A. Remove route-map sequence 30.
- B. Add a set statement to route-map sequence 20.
- C. Change the deny statement in access list 2 to permit
- D. Remove the subnets keyword from the redistribute connected subnets route-map filtering command.

Answer: C

QUESTION 160

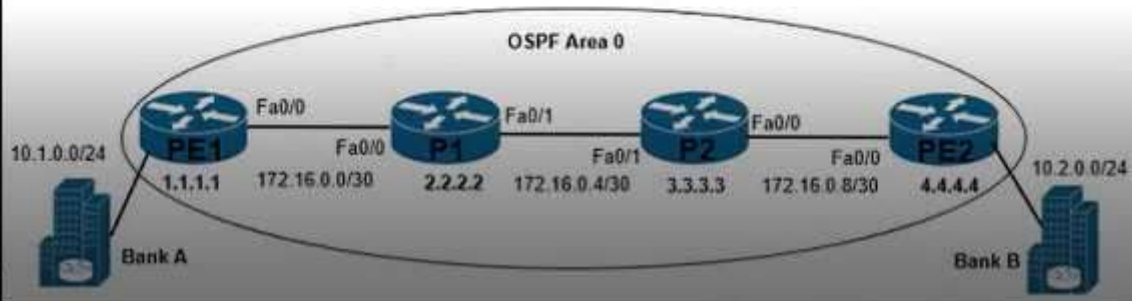
Refer to the exhibit. Network connectivity between bank A and bank B has been lost. Users at bank A and bank B are able to successfully reach their directly connected PE routers. All routers in OSPF area 0 are correctly advertising and learning routing updates. Which action resolves the issue?

```

PE1#sh run | sec router bgp
router bgp 65000
no synchronization
bgp log-neighbor-changes
network 10.1.0.0 mask 255.255.255.0
neighbor 4.4.4.4 remote-as 65000
neighbor 4.4.4.4 update-source Loopback0
no auto-summary

PE2#sh run | sec router bgp
router bgp 65000
no synchronization
bgp log-neighbor-changes
network 10.2.0.0 mask 255.255.255.0
neighbor 1.1.1.1 remote-as 65000
neighbor 1.1.1.1 update-source Loopback0
no auto-summary

PE1#sh ip cef exact-route 10.1.0.1 10.2.0.1
10.1.0.1 -> 10.2.0.1 : FastEthernet0/0 (next hop 172.16.0.2)
  
```

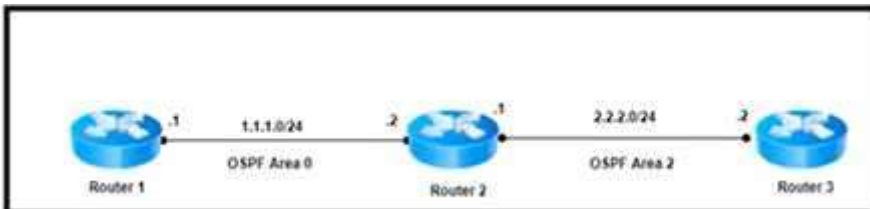


- A. Enable next-hop-self under the iBGP peering configuration on routers PE1 and PE2
- B. Configure the P routers to redistribute BGP routes within OSPF area 0.
- C. Configure router P1 to advertise the IP prefix of PE1.
- D. Configure MPLS with an end-to-end label-switched path on each router.

Answer: D

QUESTION 161

Refer to the exhibit. A network engineer installed a new router (router 3) at the regional hub running MPLS services for scalability. Router 3 is connected to the 10.44.4.0/24, 10.44.5.0/24, 10.44.6.0/24 and 10.44.7.0/24 subnets. The new router has been configured for OSPF area 2, and it is advertising the four connected networks. The engineer noticed that the same networks are listed as interarea summary routes, and they are being flooded into each area on the area borders. Which action resolves the issue?



- A. On router 3, configure an access list to filter the networks.
- B. On router 2, configure a route map to filter the networks.
- C. Under the OSPF configuration on router 3, add area 2 range 10.44.4.0 255.255.252.0.
- D. Under the OSPF configuration on router 2, add area 2 range 10.44.4.0 255.255.252.0.

Answer: D

[300-510 Exam Dumps](#) [300-510 Exam Questions](#) [300-510 PDF Dumps](#) [300-510 VCE Dumps](#)

<https://www.braindump2go.com/300-510.html>

QUESTION 162

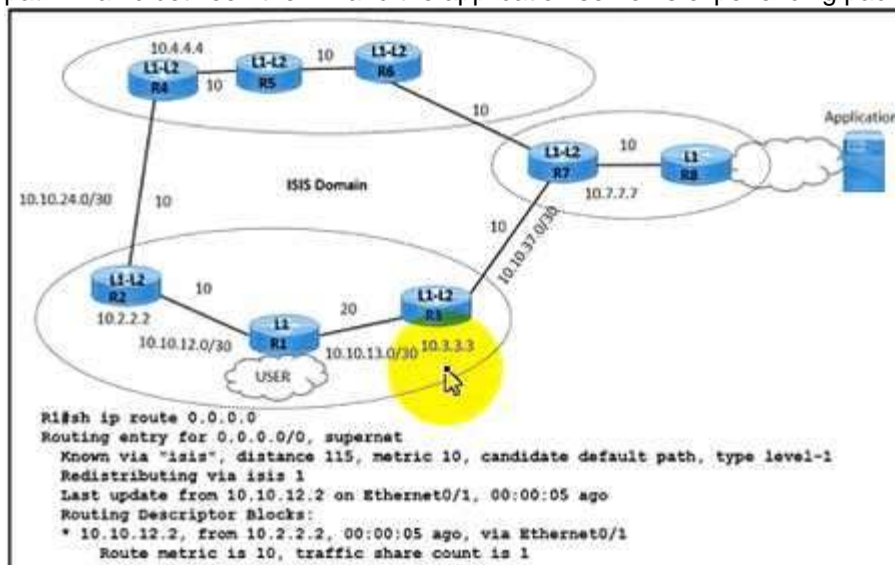
How does SRv6 function on the control plane?

- A. It enables SRH-capable nodes to terminate IPv6 packets at the network egress to carry the SRv6 locator.
- B. The Ingress node of the SR domain adds a uSID format IPv6 header to carry the SRv6 locator.
- C. The ingress node of the SR domain swaps the SRv6 header for the IPv6 header
- D. The egress node of SR domain imposes a new outer header

Answer: B

QUESTION 163

Refer to the exhibit. An engineer is troubleshooting a networking issue with several symptoms. The shortest path from router R1 to R8 is underused and the longest path is overused. Traffic from 10.1.1.1 to 10.8.8.8 is routed on the longest path. Traffic between the R1 and the application server is experiencing packet drops and latency problems.



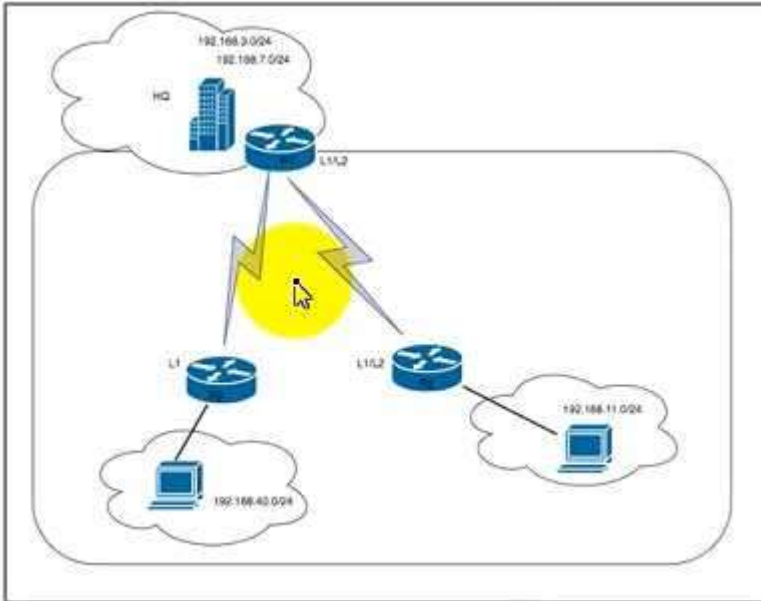
Which action resolves the issue?

- A. Configure route leaking for the IP address Of the application server on router R1.
- B. Increase the R1 to R2 link metric to 20.
- C. Configure a Level 2 IS-IS domain on router R1.
- D. Block the advertisement of the application server IP address to router R6.

Answer: C

QUESTION 164

Refer to the exhibit. The branch office in area 10 is connected to HQ via Frame Relay uplinks with bandwidth constraints. After a recent implementation of QoS on the R2 and R3 networks the system has been logging %SYS-2-MALLOCFAIL: Memory allocation of 65536 bytes failed from 0x224E321, alignment 0 messages. To reduce traffic load and memory utilization on R2 and R3 the network engineer configured R1 to announce only one user subnet per location by issuing the summary address 192.168.0.0 255.255.248.0 command on R1. However, the engineer noticed that router R2 still has two routes and a summary address from HQ and R3 also has two routes from HQ. Which two actions must the engineer take on R1 to fix the issue so that only one route is announced? (Choose two.)

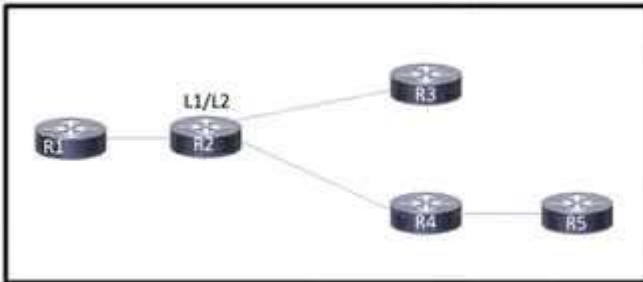


- A. Configure the subnet range with the network command under the IS-IS process.
- B. Configure R1 as a Level 1 device.
- C. Redistribute both routes into the ISIS process.
- D. Configure a summary route for Level 1-Level 2 devices.
- E. Configure a summary route for Level 1 devices.

Answer: DE

QUESTION 165

Refer to the exhibit. Routers R2, R3, R4 and R5 all reside in the same area, with R1 in a different area R3 is overutilized and the engineer wants to reduce its CPU load. The engineer configured router R4 to summarize routes that it receives from R5, but R3 is still receiving all of the R5 routes. Which action resolves the issue?

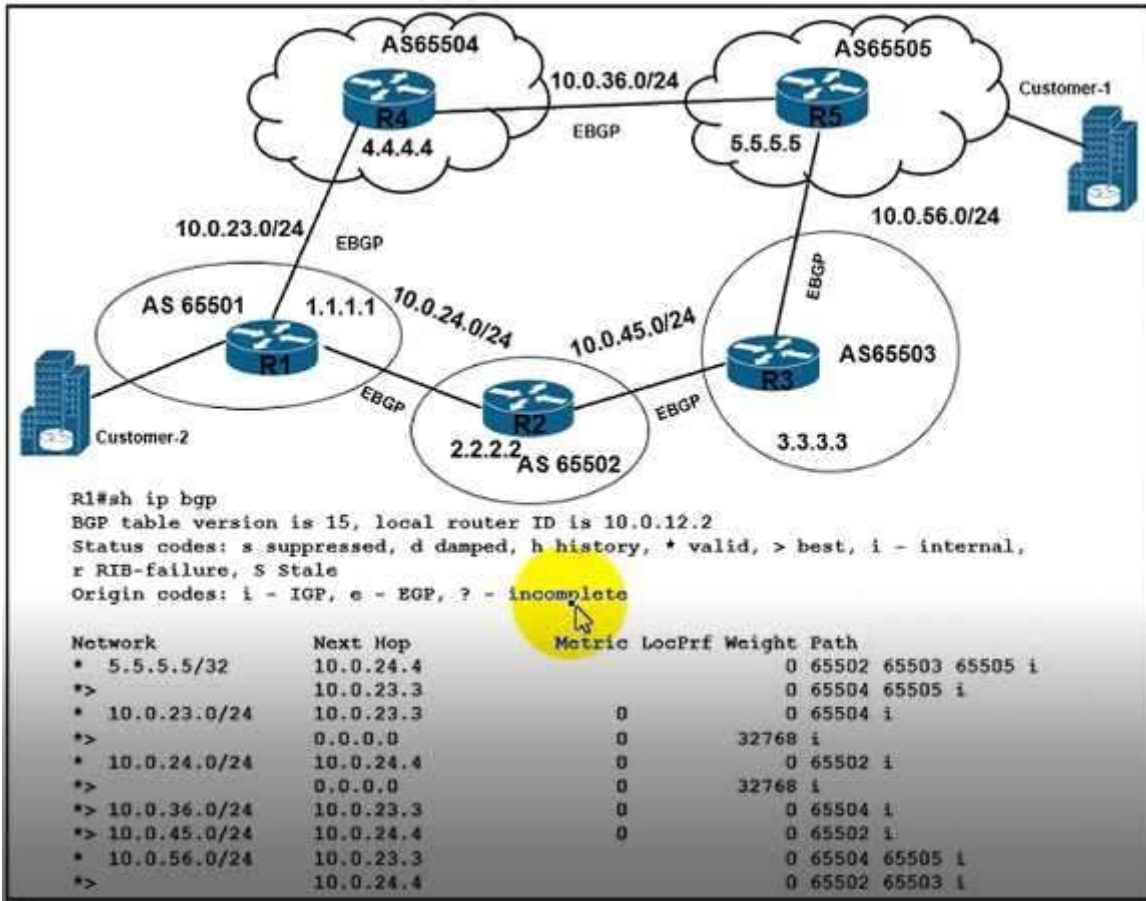


- A. Configure R3 in a new area,
- B. Configure R2 as a Level 1 router
- C. Configure the summary routes on R5.
- D. Configure R4 as a Level I-Level 2 router

Answer: A

QUESTION 166

Refer to the exhibit. There is a BGP traffic path issue between Customer-1 and Customer-2. Users from Customer-2 have reported file transfer issues High utilization on the path between both customers causes many packet drops. Which configuration resolves the issue?



- R1#neighbor 10.0.24.4 route-map LOCAL-PREF-150 in
 route-map LOCAL-PREF-150
 set local-preference 150
 ip prefix-list 5-5-5-5 seq 5 permit 5.5.5.5/32
 route-map LOCAL-PREF-150 permit 10
 match ip address prefix-list 5-5-5-5
 set local-preference 150
- R4#router bgp 65504
 neighbor 10.0.23.3 remote-as 65501
 neighbor 10.0.23.3 filter-list 1 out
 ip as-path access-list 1 deny ^65505\$
 ip as-path access-list 1 permit .^
- R4#router bgp 65504
 address-family ipv4 unicast
 neighbor 10.0.23.3 remote-as 65501
 neighbor 10.0.23.3 activate
 neighbor 10.0.23.3 route-map PREPEND in
 exit-address-family
 exit
 route-map PREPEND permit 10
 set as-path prepend 65506 65507
- R1#neighbor 10.0.23.3 route-map LOCAL-PREF-150 out
 route-map LOCAL-PREF-150
 set local-preference 150

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

Answer: B

QUESTION 167

What is a characteristic of a segment routing mapping server?

- A. It must be placed in the core of the network.
- B. It serves multiple VRFs.
- C. It must have an IGP adjacency.
- D. It applies SID mappings from one IGP instance to another IGP instance.

Answer: C

QUESTION 168

Refer to the exhibit. A network engineer is investigating a report of packet drops between the branch site and the central site.

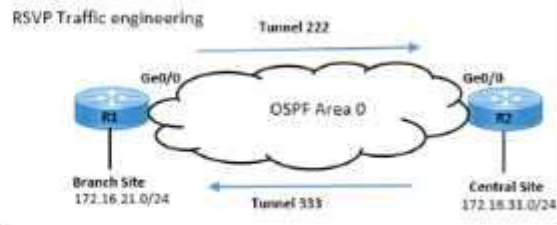
The two sites are connected via OSPF and RSVP-TE tunnels. Traffic from the central site to the branch site is passing normally. Technicians at both sites successfully ping the loopback IP addresses on routers R1 and R2.

Which configuration corrects the packet-drop problem?

```

R1#sh run inter tu222
interface Tunnel222
description R1>msR2
bandwidth 33000
ip unnumbered Loopback0
load-interval 30
tunnel destination 10.10.11.1
tunnel mode mpls traffic-eng
tunnel mpls traffic-eng autoroute announce
tunnel mpls traffic-eng priority 1 1 2
tunnel mpls traffic-eng path-option 10 dynamic
tunnel mpls traffic-eng record-route
no routing dynamic
End
R1# show ip rsvp reservation

```



RSVP Traffic engineering

```

R1#sh mpls traffic-eng tunnels tu222
Name: R1>msR2
(Tunnel222) Destination: 10.10.11.1
Status:
Admin: up Oper: down Path: valid Signalling: RSVP signalling proceeding
path option 10, type dynamic (Basis for Setup, path weight 2)
Config Parameters:
Bandwidth: 33000 kbps (Global) Priority: 1 1 Affinity: 0x0/0xFFFF
Metric Type: TE (default)
AutoRoute: enabled LockDown: disabled Loadshare: 33000 bw-based
auto-bw: disabled
RSVP Signalling Info:
Src 10.0.1.4, Dst 10.10.11.1, Tun_Id 222, Tun_Instance 73
Shortest Unconstrained Path Info:
Path Weight: 2 (TE)
Explicit Route: 10.0.4.254 10.10.11.1
History:
Tunnel:
Time since created: 6 hours, 10 minutes
Time since path change: 1 minutes, 22 seconds
Current LSP:
Setup Time: 3 minutes, 37 seconds remaining
Prior LSP:
ID: path option 10 [72]
Removal Trigger: setup timed out

```

```
R1(Config)# interface Tunnel222
R1(Config-if)# tunnel mpls traffic-eng bandwidth 33000

R2(Config)# interface Tunnel333
R2(Config-if)# tunnel mpls traffic-eng bandwidth 33000

R2(Config)# interface Ge0/0
R2(Config-if)# ip rsvp bandwidth 33000 3300

R1(Config)# interface Ge0/0
R1(Config-if)# ip rsvp bandwidth 33000 3300
```

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

Answer: D

QUESTION 169

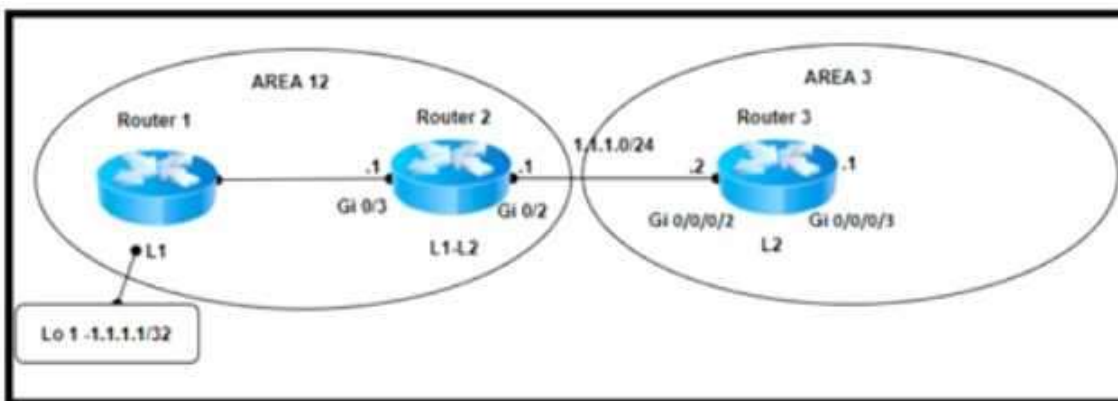
Which difference must an engineer consider when Implementing Inter-domain and Intra-domain multicast routing on the network?

- A. Intra-domain routing allows the service provider to control incoming and outgoing multicast data streams on its network, but inter-domain routing limits the service provider's control.
- B. Intra-domain routing uses the PIM and MBGP protocols for multicast routing, but inter-domain routing must use PIM.SSM or MSDP.
- C. Intra-domain routing is dependent on the RP router within the same SP network, but inter-domain routing reduces the dependency on the other SP network.
- D. Inter-domain routing supports policy routing to connect different multicast domains using PIM.SM, but intra-domain routing supports policy routing using PIM-SM only within a single domain.

Answer: C

QUESTION 170

Refer to the exhibit. A network engineer configured three new PE routers to expand the network. The new routers run in the IS-IS routing protocol and reside in the data center in the same exchange as the existing routers. However, the network is now experiencing suboptimal routing. The Layer 2 configuration and VLANs are configured correctly to provide segregation between networks, but the Level 1 routes are not being converted to Level 2 routes. Which action resolves the issue?



- A. On Router 1, redistribute the routes into IGP.
- B. On Router 1, summarize internal routes between areas.
- C. On Router 2, redistribute the routes into IGP.
- D. On Router 2, summarize internal routes between areas.

[300-510 Exam Dumps](#) [300-510 Exam Questions](#) [300-510 PDF Dumps](#) [300-510 VCE Dumps](#)

<https://www.braindump2go.com/300-510.html>

Answer: D