Free Questions for 350-501 by actualtestdumps

Shared by Buckley on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Refer to the exihibit.

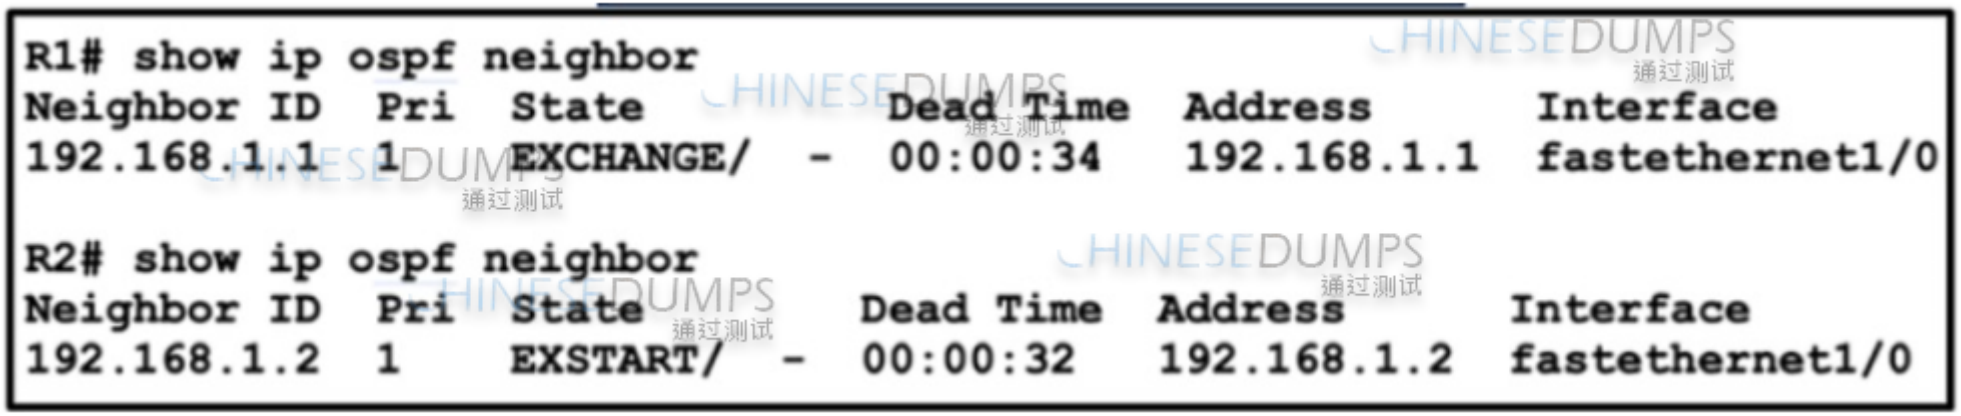

```
R1# show ip ospf neighbor
Neighbor ID   Pri   State           Dead Time   Address         Interface
192.168.1.1   1     EXCHANGE/   -   00:00:34    192.168.1.1     fastethernet1/0

R2# show ip ospf neighbor
Neighbor ID   Pri   State           Dead Time   Address         Interface
192.168.1.2   1     EXSTART/    -   00:00:32    192.168.1.2     fastethernet1/0
```

Refer to the exhibit. A company recently deployed a new network using OSPF in the core to share routes. The network administrator selected OSPF as the routing protocol because of its ability to maintain a route database. When the new network was started up, all routers booted normally, but the link between routers R1 and R2 failed to come up. The two routers are located in the same rack at the data center. Which task should an engineer perform to correct the problem?

Options:

A- Synchronize the dead timers.

B- Change one of the OSPF router IDs so that the router IDs are in different subnets

C- Change the OSPF process ID on one of the devices so that the two IDs match

D- Configure the MTUs on the interface to match.

Answer:

D

Explanation:

When OSPF routers establish adjacency, they must agree on the MTU size on the connecting interface. If the MTUs do not match, the routers will not form an OSPF adjacency, and the link between them will not come up. Therefore, configuring the MTUs on the interfaces to match is essential for the OSPF network to function correctly.

Question 2

Question Type: **MultipleChoice**

Refer to the exihibit.

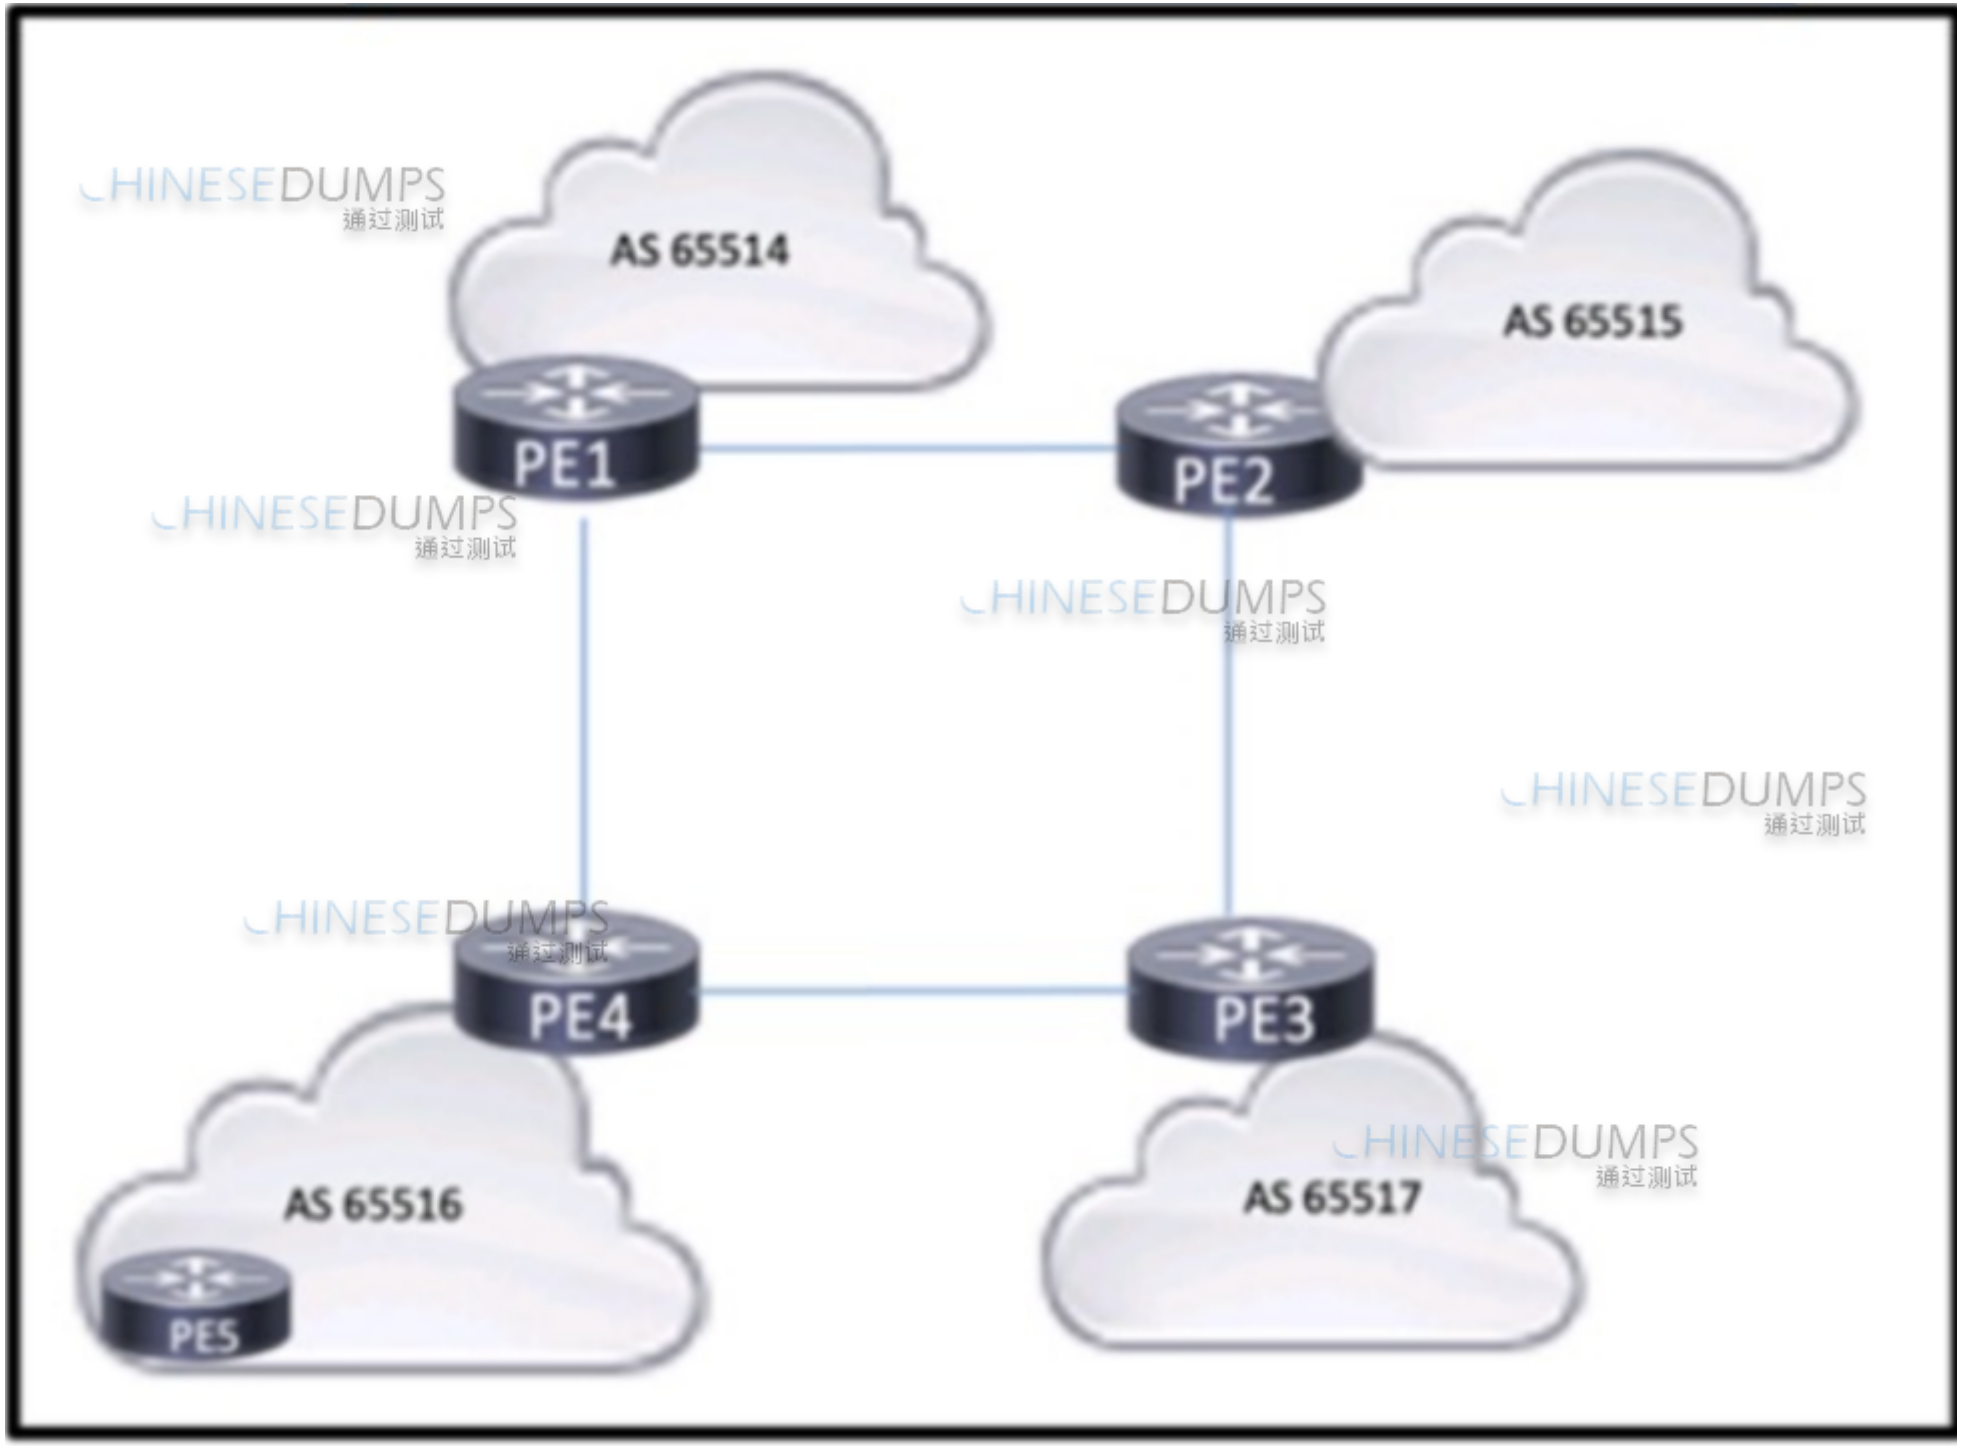
AS 65514
AS 65515
PE1
PE2
PE4
PE3
AS 65516
PE5
AS 65517

Refer to the exhibit. Four midsize service providers provide access to different customers that use Layer 3 VPN services to enable communication across geographic regions. The service providers are connected as shown in the exhibit, and the PEs have established eBGP relationships. PE4 has an IBGP relationship with PE5. The routes that PE4 learns from PE5 must reach the other PE routers, but they are absent from the routing tables on the other PES. Which action should the engineers take to correct the problem?

Options:

A- Configure a peering between all five Pes.

B- Disable BGP synchronization on PE4.

C- Enable BOP IPv4 unicast on PE4 and PE5

D- Advertise the route targets for PE5 to the other PEs

Answer:

D

Explanation:

The absence of routes learned by PE4 from PE5 in other PE routers' routing tables is likely due to missing route target configurations. Route targets are crucial for importing and exporting VPN routes between VRFs in MPLS Layer 3 VPNs.

Option D is correct because advertising the route targets for PE5 to the other PEs will allow them to import routes from PE5 into their respective VRFs, facilitating communication across the Layer 3 VPN.

Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR) source book and study guide.

Question 3

Question Type: **MultipleChoice**

Refer to the exihibit.

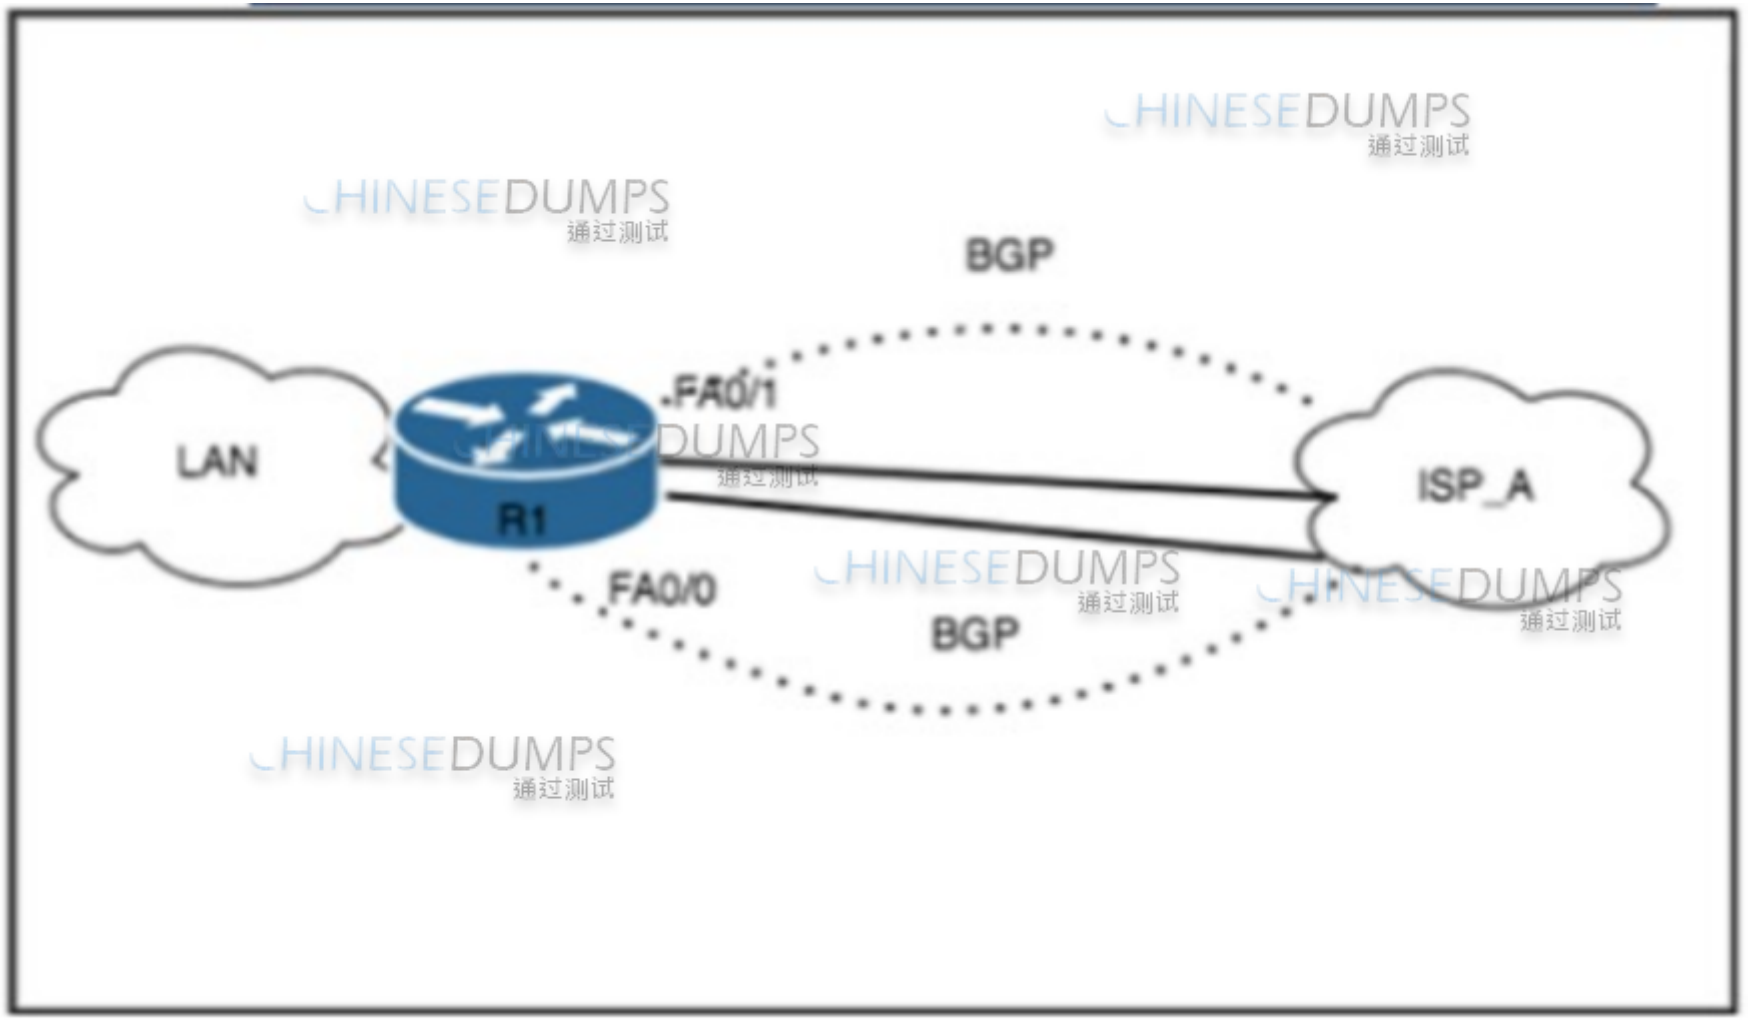

Refer to the exhibit. A network engineer must deny access from spoofed addresses to the LAN. The edge router currently has two active BGP sessions established with Tier 1 ISP_

Options:

A- Due to asymmetric routing, no ACL is configured on either interface. Which two configurations must the engineer perform on the edge router to complete the task? (Choose two.)

A- ip verify unicast source reachable-via tx under FA0/0

B- ip verify unicast source reachable-via under FA0/1

C- ip verify unicast source reachable-via any under FA0/1

D- ip verify unicast source reachable-via both under FA0/0

E- ip verify unicast source reachable-via any under FA0/0

Answer:

A, A, C

Explanation:

To prevent access from spoofed addresses to the LAN, the engineer should enable Unicast Reverse Path Forwarding (uRPF). uRPF checks the reachability of the source address in packets being forwarded, which helps mitigate issues caused by malformed or forged IP source addresses.

Option A is correct because it enables uRPF on interface FA0/0 and verifies that the source address is reachable via the outgoing interface, which is essential when the routing is symmetric.

Option C is also correct as it enables uRPF on interface FA0/1 and allows the source address to be reachable via any interface, accommodating environments with asymmetric routing.

Implementing and Operating Cisco Service Provider Network Core Technologies (SPCOR) source book and study guide.

Question 4

Question Type: **MultipleChoice**

Refer to the exihibit.

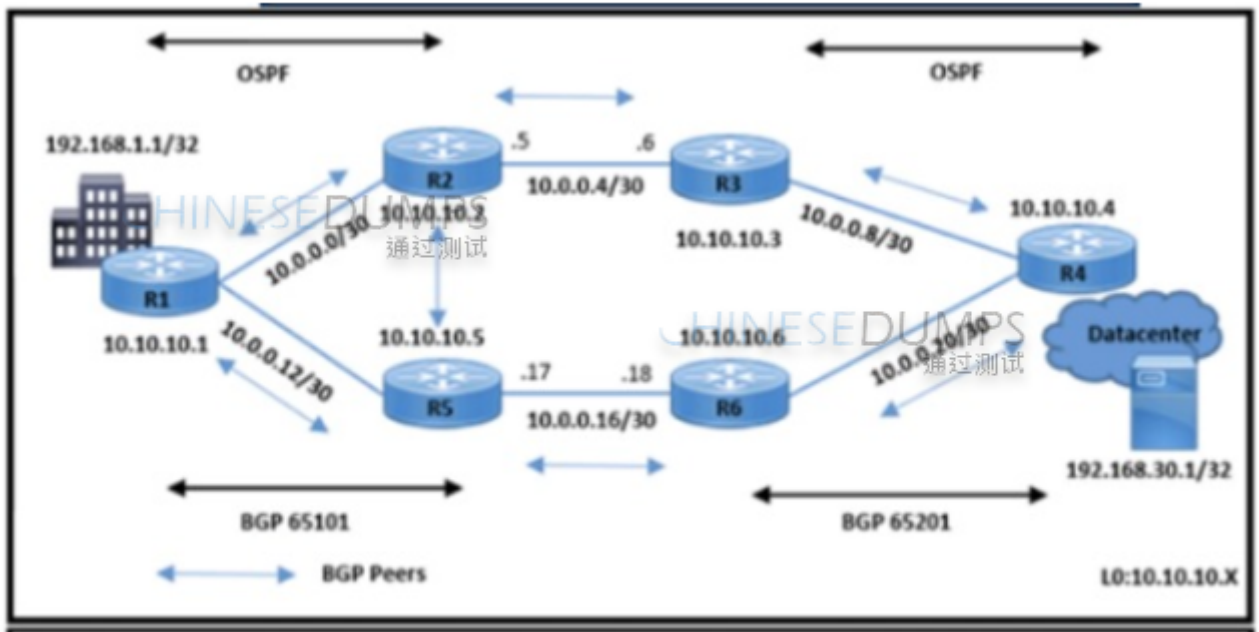


```
R5#show ip bgp 192.168.1.1/32
BGP routing table entry for 192.168.1.1/32, version 25
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        3
  Local
    10.10.10.1 (metric 2) from 10.10.10.1 (192.168.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best

R2#show ip bgp 192.168.1.1/32
BGP routing table entry for 192.168.1.1/32, version 13
Paths: (1 available, no best path)
  Not advertised to any peer
  Local
    10.10.10.1 (metric 2) from 10.10.10.1 (192.168.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal, not synchronized

R1#show ip bgp 192.168.1.1/32
BGP routing table entry for 192.168.1.1/32, version 15
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        1
  Local
    0.0.0.0 from 0.0.0.0 (192.168.1.1)
      Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced, local, best
```

Refer to the exhibit. All BGP peering in AS 65101 and 65201 is enabled. The operations team is told that traffic destined to 192.168.1.1/32 from R4 does not use the path R3-R2-R1 as expected. An engineer debugs the issue and determines that 192.168.1.1/32 is advertised in the BGP routing table on R1. Which action resolves the issue?

Options:

A- Enable no synchronization on R2 m AS65101.

B- Apply route-map High-LP out for prefix 192.168.1.1/32 on R1 with R2 BGP peering.

C- Apply redistribute ospf 10 on R1 in BGP AS 65101.

D- Configure network 192.168.1.1 mask 255.255.255.255 in BGP AS 65101 on R2

Answer:

B

Explanation:

The issue described indicates that the preferred BGP path for traffic destined to 192.168.1.1/32 is not being used. By applying a route-map with a higher Local Preference (LP) on R1 for the prefix 192.168.1.1/32, the BGP path attributes can be manipulated to prefer the path through R2. Local Preference is a BGP attribute that determines the preferred exit path out of an AS. A higher Local Preference value is preferred over a lower one. Therefore, configuring the route-mapHigh-LPwith a higher Local Preference on R1's BGP peering with R2 will influence the BGP path selection process, ensuring that traffic for 192.168.1.1/32 prefers the path R3-R2-R1 as expected.

Question 5

Question Type: MultipleChoice

Refer to the exihibit.

Refer to the exhbit. ISP A has a BGP peering with ISP C with the maximum-prefix 150 configuration on R1. After a recent security breach on the ISP A network, a network engineer has been asked to enable a lightweight security mechanism to protect the R1 CPU and BGP membership from spoofing attacks. Which solution must ISP A implement?

Options:

A- Configure bgp maxas-limit 1 in the IPv4 address family urateUhe global BGP configuration.

B- Configure neighbor 10.163.83.54 enable-connected-check under the BGP IPv4 address family.

C- Configure neighbor 10.163.83.55 password Cisco under the global BGP IPv4 address family.

D- Configure neighbor 10.163.83.55 ttl-stcurity hops 2 under the global BGP configuration.

Answer:

D

Explanation:

The Time to Live (TTL) security mechanism is a lightweight security feature that can protect against CPU and memory resource exhaustion due to a denial-of-service (DoS) attack. It works by setting a limit on the number of hops (TTL value) that BGP packets can traverse. By configuringneighbor 10.163.83.55 ttl-security hops 2, ISP A ensures that BGP packets received from ISP C must have a TTL value of at least 254 when they reach R1, as BGP packets decrement the TTL by 1 with each hop. This effectively prevents spoofing attacks from outside the directly connected network because any packets spoofed from further away would have a TTL that drops below the threshold before reaching R1.

Question 6

Question Type: MultipleChoice

Refer to the exihibit.

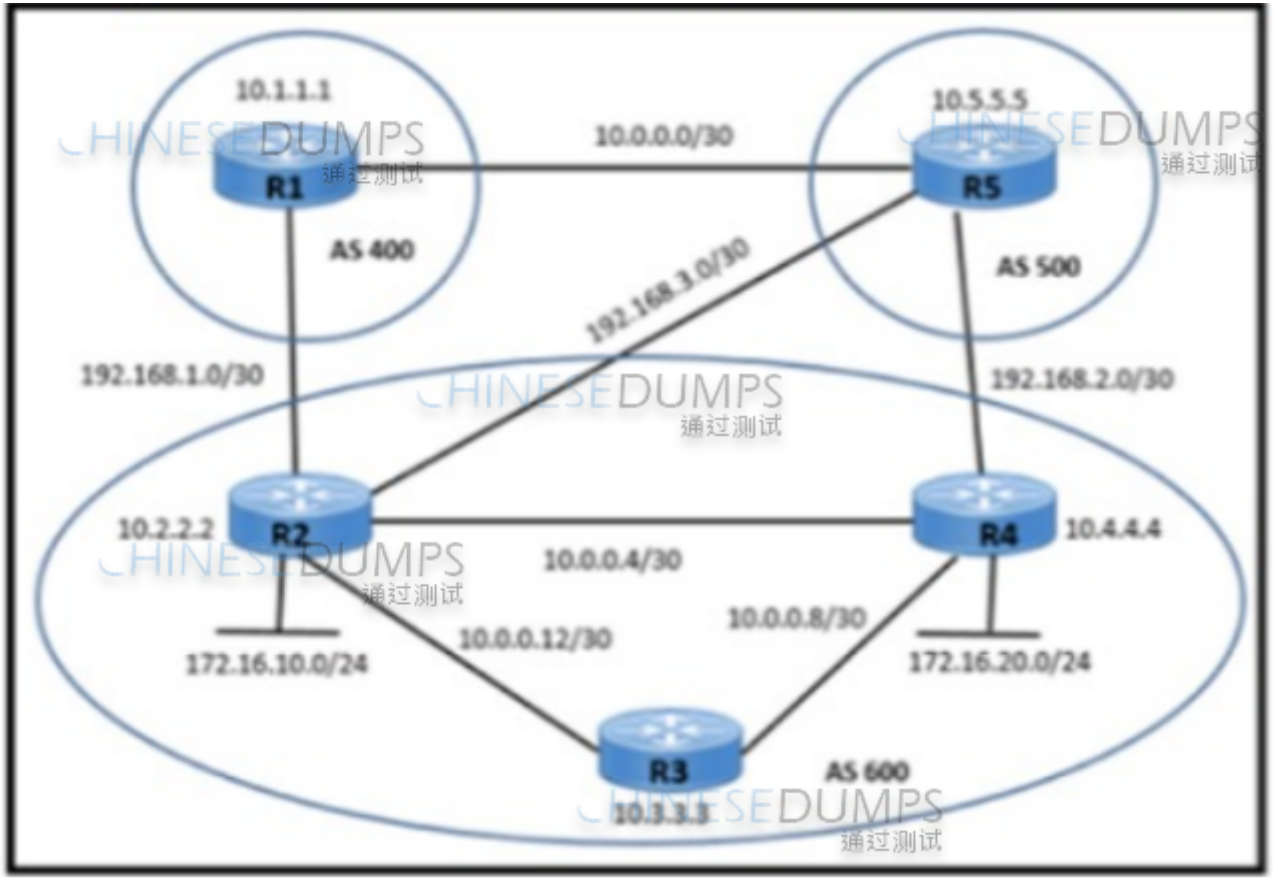

Refer to the exhibit. A network engineer is implementing iBGP and eBGP between AS 600 and AS 500 with these requirements:

R2 must establish eBGP peering on 192.168.3.0/30 with R5 for sending unicast and multicast traffic

R2 must wait for 30 seconds before sending BGP updates to R5 for multicast traffic.

Which action must be taken on R2 to meet the requirements?

Options:

A- Configure advertisement-interval 30 In address-family ipv4 unicast

B- Configure advertisement-Interval 30 in address-family Ipv4 multicast

C- Apply timers bgp 30 in address-family ipv4 unicast

D- Apply timers bgp 30 in address-family ipv4 multicast.

Answer:

B

Explanation:

The requirement is for R2 to delay sending BGP updates to R5 for multicast traffic by 30 seconds. This can be achieved by configuring the advertisement interval for IPv4 multicast traffic on R2, which controls the frequency of sending BGP updates for multicast routes.Reference:=Cisco SPCOR

Question 7

Question Type: MultipleChoice

Refer to the exihibit.

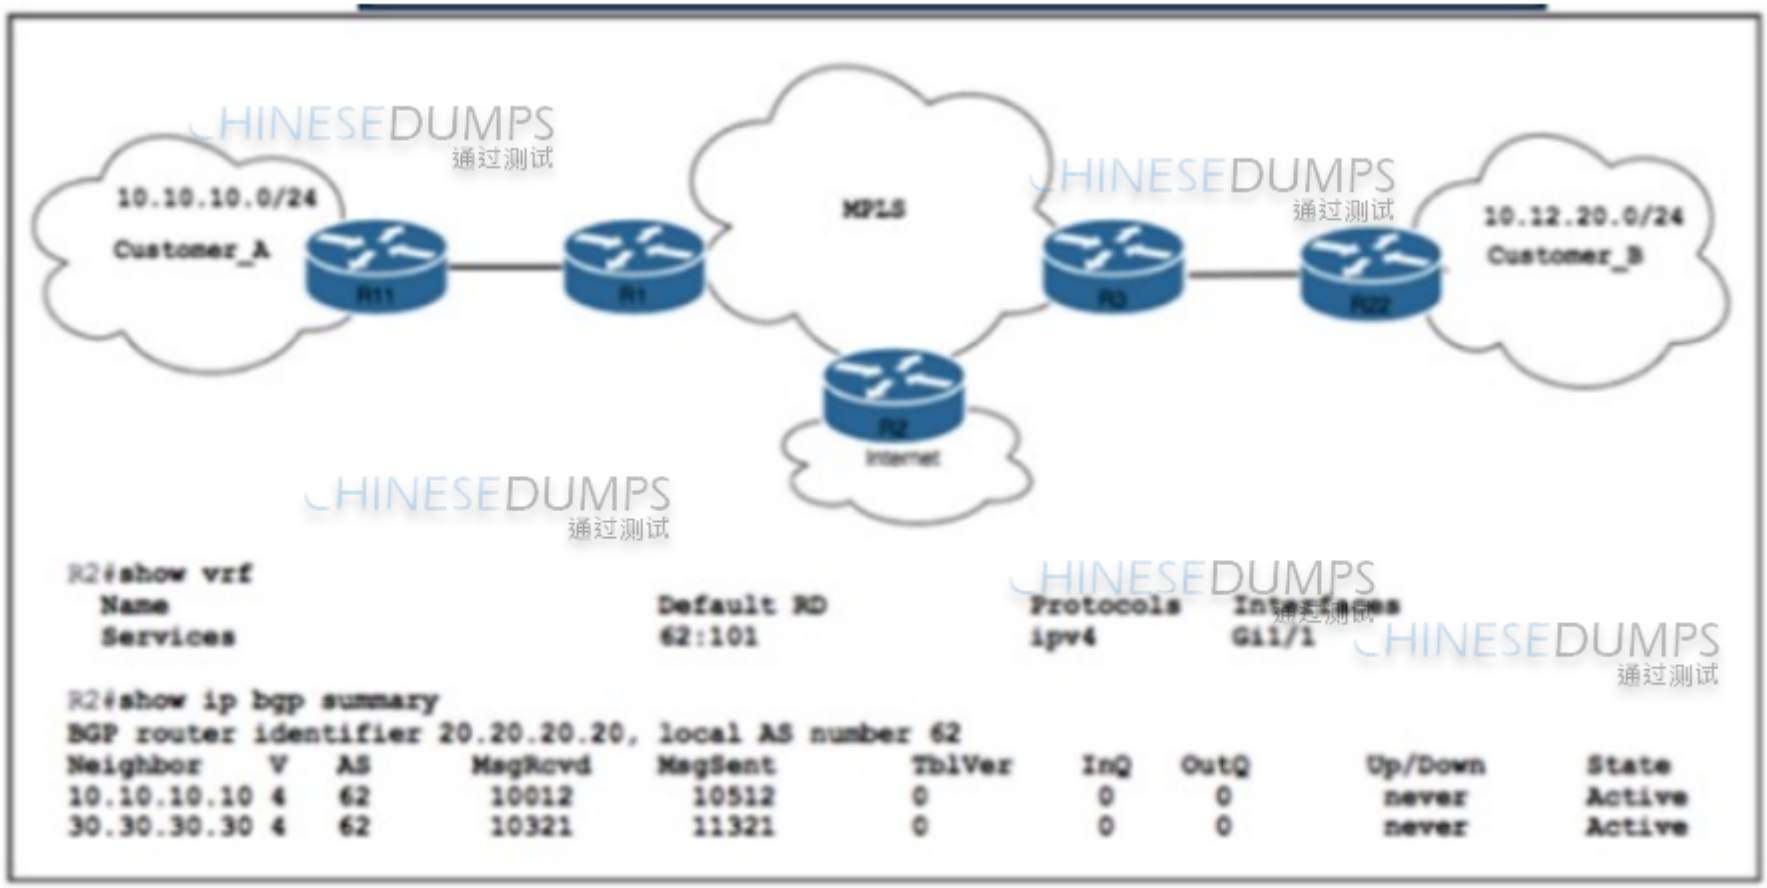

Refer to the exhibit. ISP_A is about to launch a new internet service. ISP_A is already providing MPLS VPN Layer 3 services to Customer_A and Customer_B, which are connected to ISP_A via OSPF. A network engineer completed the BGP and VRF configurations on R2 to support the new internet service. Which additional action completed the launch?

Options:

A- Implement the BGP routing protocol in the customer VRFs on R1 and R2

B- Import route-target 62:101 into the customer VRFs on R1 and R3.

C- Enable the route-replicate command under the customer VRFs on R1 and R2

D- Activate NAT CE in the customer VRFs on R1. R2. and R3.

Answer:

B

Explanation:

To complete the launch of the new internet service, ISP_A needs to ensure that routes are shared between different VRFs for Customer_A and Customer_B. Importing route-target 62:101 into the customer VRFs on R1 and R3 allows for the exchange of routes between these VRFs, facilitating the necessary connectivity for the new service.Reference:=Cisco SPCOR

Question 8

Question Type: MultipleChoice

Refer to the exihibit.

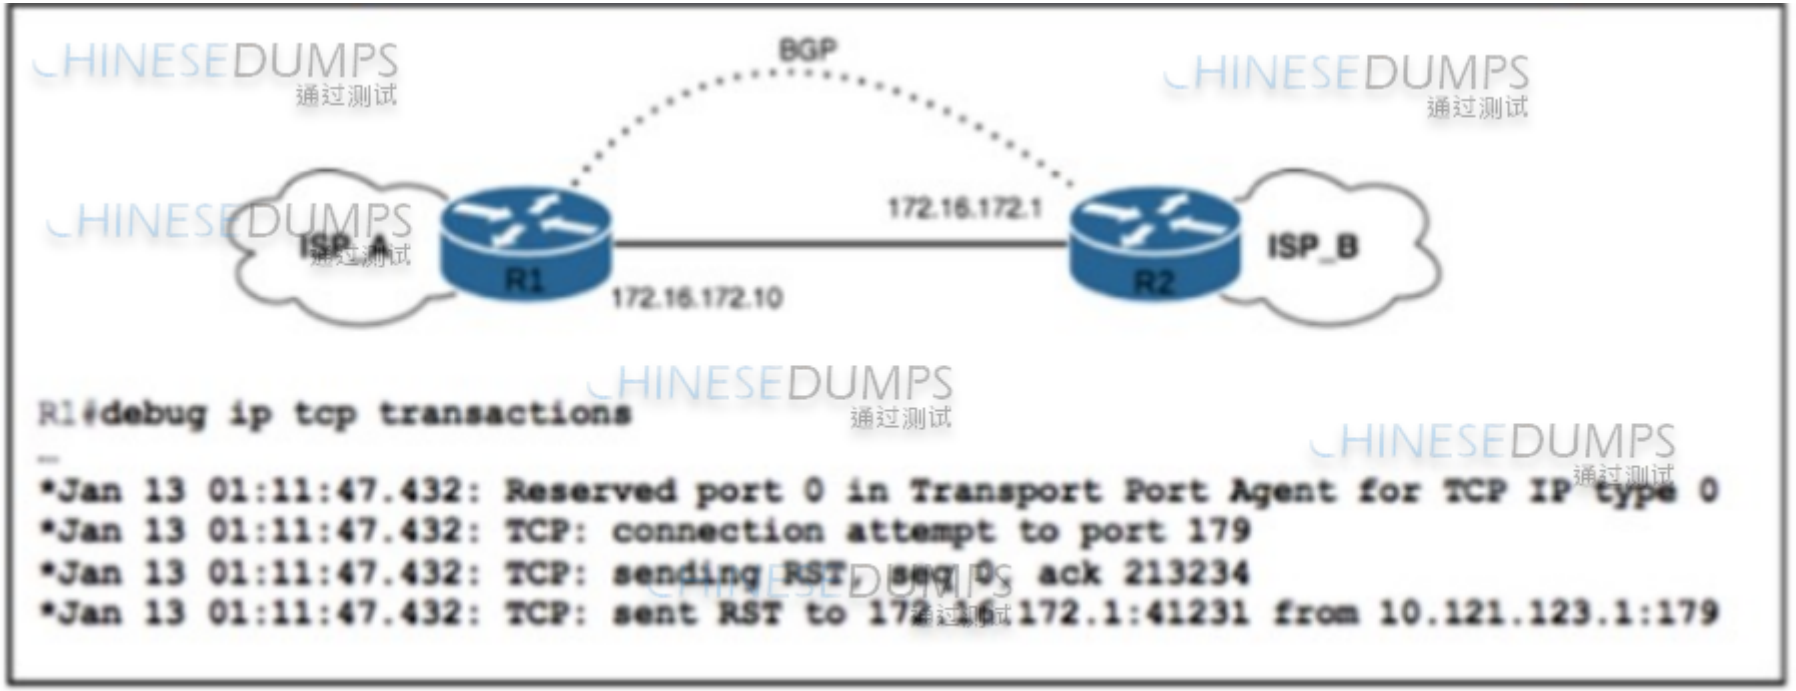

Refer to the exhibit. ISP_A and ISP_B use AS numbers 38321 and 16213 respectively. After a network engineer reloaded router R1, the BGP session with R2 failed to establish. The engineer confirmed BGP next-hop availability with a connectivity test between the router loopback addresses 10.121.123.2 and 10.121.123.1, as well as between interfaces Gi1/1 and Gi1/2. EBGP multihop has been configured on both routers. Which action must the engineer take to resolve the issue?

Options:

A- Configure transport connection-mod passive on R2.

B- Configure neighbor 172.16.172.1 authentication on R1

C- Configure neighbor update-source lo0 on R2

D- Configure remote-as 16213 on R1.

Answer:

C

Explanation:

After a router reload, BGP sessions may fail to re-establish due to incorrect source interface configurations for BGP sessions. Since EBGP multihop has been configured, it's essential that the BGP session uses the correct source interface for establishing the connection. By configuring the neighbor update-source command with the loopback interface on R2, the BGP session will use the stable loopback address rather than the physical interface, which can change state.This ensures that the BGP session remains stable and is not affected by interface states.Reference: BGP Troubleshooting Cheat Sheet With Examples - Catchpoint2, BGP session stuck in Idle state after NSRP failover - Juniper Networks3, BGP Session is Disconnected - Network Playbook

Question 9

Question Type: MultipleChoice

Refer to the exihibit.

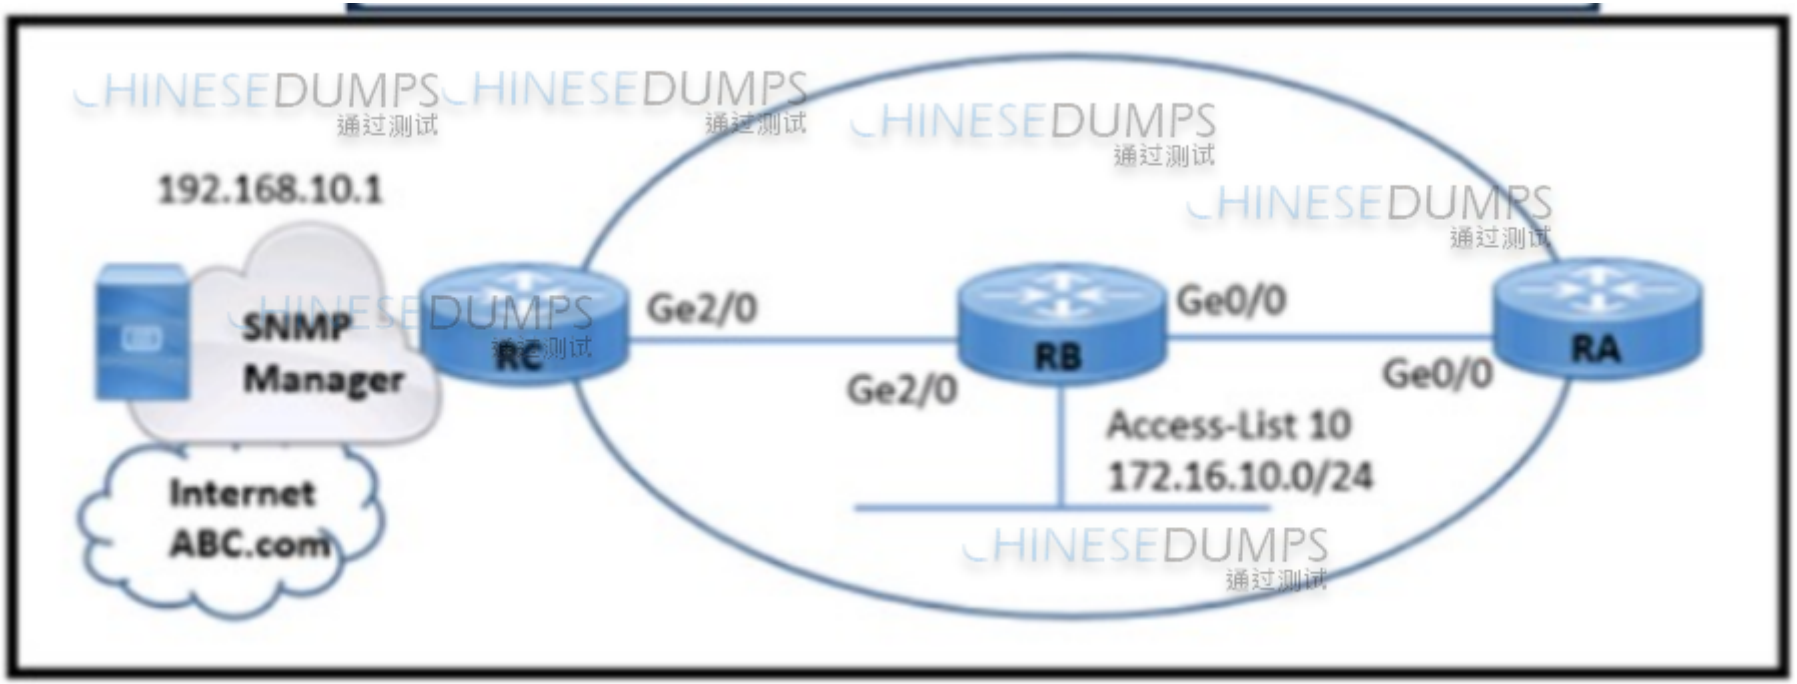

Refer to the exhibit. A network engineer is configuring an SNMP community on router RB with these requirements:

Allow read-only access for all objects to members of Access-List 10 that use the comaccess community string.

Other SNMP managers must not have access to objects.

SNMP authentication failure traps must be sent to SNMPv2c and then to the host using SNMPv2c with the public community string.

Which configuration meets these requirements?

- RB(config)# snmp-server community comaccess ro 10
  RB(config)# snmp-server enable traps snmp authentication
  RB(config)# snmp-server host ABC.com version 2c public

- RB(config)# snmp-server community comaccess ro 10
  RB(config)# snmp-server enable traps snmp authentication
  RB(config)# snmp-server host ABC.com
  RB(config)# snmp-server host informs ABC.com restricted entity

- RB(config)# snmp-server community comaccess ro 10
  RB(config)# snmp-server enable traps snmp authentication
  RB(config)# snmp-server enable traps entity
  RB(config)# snmp-server host informs ABC.com restricted entity

- RB(config)# snmp-server community comaccess ro 10
  RB(config)# snmp-server enable traps
  RB(config)# snmp-server host 192.168.10.1 informs version 2c public
  RB(config)# snmp-server host ABC.com public

Options:

A- Option A

B- Option B

C- Option C

D- Option D

Answer:

B

Explanation:

Option B is the correct answer because it meets all the given requirements. The configuration "snmp-server community comaccess ro 10" allows read-only access for all objects to members of Access-List 10 that use the comaccess community string. The "snmp-server host ABC.com informs version 2c public" command ensures SNMP authentication failure traps are sent to SNMPv2c and then to the host using SNMPv2c with the public community string.