



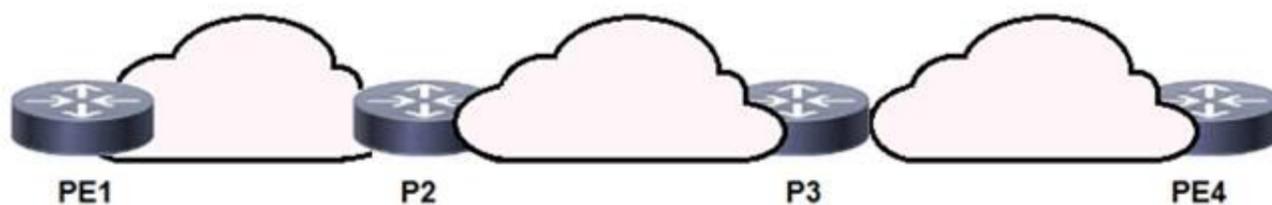
Cisco

Exam Questions 350-501

Implementing and Operating Cisco Service Provider Network Core Technologies

NEW QUESTION 1

Refer to the exhibit:



P3 and PE4 are at the edge of the service provider core and serve as ABR routers. Aggregation areas are on either side of the core. Which statement about the architecture is true?

- A. If each area is running its own IGP
- B. the ABR routers must redistribute the IGP routing table into BGP
- C. To support seamless MPLS
- D. TDP must be used as the label protocol
- E. If each area is running its own IGP
- F. BGP must provide an end-to-end MPLS LSP
- G. To support seamless MPLS, the BGP route reflector feature must be disabled

Answer: C

NEW QUESTION 2

Refer to the exhibit:

```

R2#show mpls ldp neighbor

R2#show mpls ldp discovery detail
Local LDP Identifier:
 2.2.2.2:0
Discovery Sources:
Targeted Hellos:
 2.2.2.2 -> 1.1.1.1 (ldp): active/passive, xmit
Hello interval: 5000 ms; Transport IP addr: 0.0.0.0
    
```

When implementing an LDP protocol, an engineer experienced an issue between two directly connected routers and noticed that no LDP neighbor exists for 1.1.1.1.

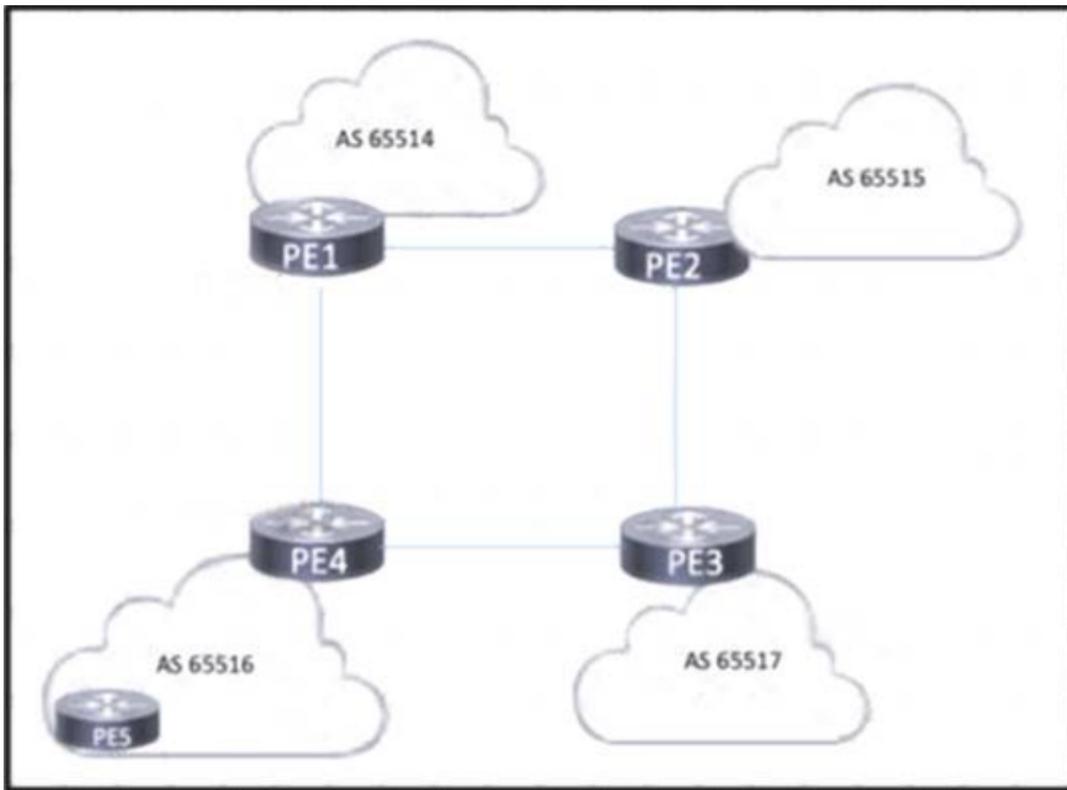
Which factor should be the reason for this situation?

- A. LDP needs to be enabled on the R2 physical interface
- B. R2 does not see any hellos from R1
- C. LDP needs to be enabled on the R2 loopback interface
- D. R2 sees the wrong type of hellos from R1

Answer: B

NEW QUESTION 3

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?

- A. Configure a peering between all five PEs.
- B. Disable BGP synchronization on PE4.
- C. Enable BGP IPv4 unicast on PE4 and PE5
- D. Advertise the route targets for PE5 to the other PEs

Answer: A

NEW QUESTION 4

Refer to the exhibit.

```
snmp-server community ciscotest ro 2
```

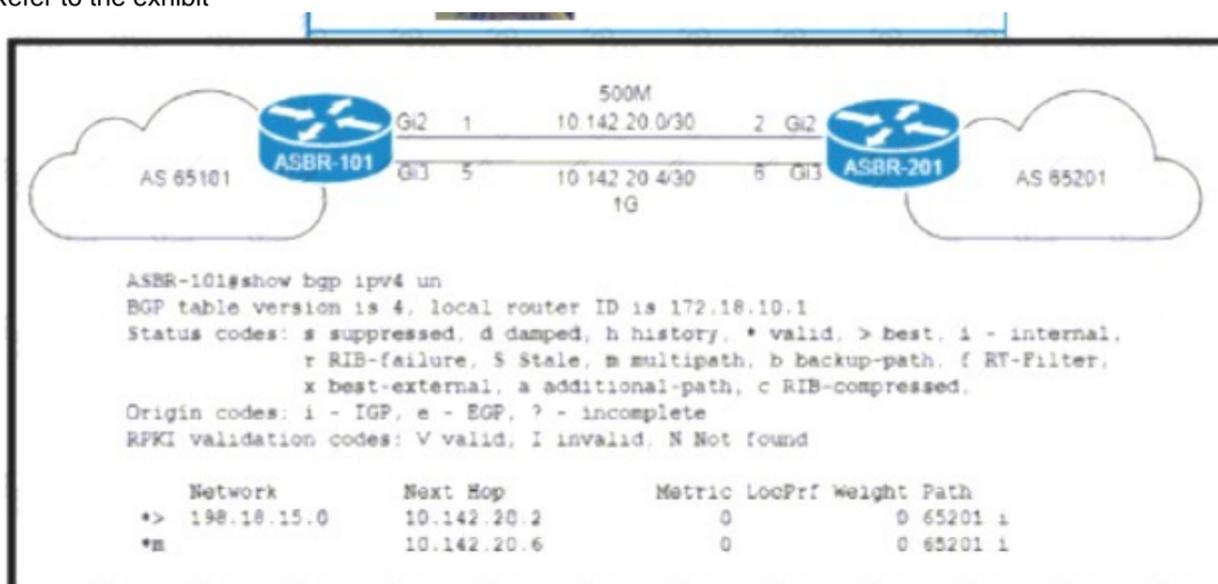
What does the number 2 mean in the configuration?

- A. It dictates the number of sessions that will be open with the SNMP manager
- B. It represents the version of SNMP running.
- C. It indicates two SNMP managers are able to read and write with the agent using community string ciscotest.
- D. It is the numeric name of the ACL that contains the list of SNMP managers with access to the agent.

Answer: D

NEW QUESTION 5

Refer to the exhibit



an engineer working for a private telecommunication company with an employee Id: 4065:96:080 upgrades the WAN link between routers ASBR-101 and ASBR-201 to 1Gb by Installing a new physical connection between the Gi3 Interfaces. Which BGP attribute must the engineer configure on ASBR-201 so that the existing WAN link on Gi2 Is maintained as a backup?

❌ configure terminal
 ip prefix-list ALLOWED_PREFIXES seq 5 permit 198.18.15.0/24

route-map AS65101-OUT permit 10
 match ip address prefix-list ALLOWED_PREFIXES
 set as-path prepend 65101 65101

router bgp 65201
 address-family ipv4
 neighbor 10.142.20.1 route-map AS65101-OUT out
 end

❌ configure terminal
 ip prefix-list ALLOWED_PREFIXES seq 5 permit 198.18.15.0/24

route-map AS65101-OUT permit 10
 match ip address prefix-list ALLOWED_PREFIXES
 set as-path prepend 65101 65101

✅ configure terminal
 ip prefix-list ALLOWED_PREFIXES seq 5 permit 198.18.15.0/24

route-map AS65101-OUT permit 10
 match ip address prefix-list ALLOWED_PREFIXES
 set metric 100

router bgp 65201
 address-family ipv4
 neighbor 10.142.20.1 route-map AS65101-OUT out
 end

❌ configure terminal
 ip prefix-list ALLOWED_PREFIXES seq 5 permit 198.18.15.0/24

route-map AS65101-OUT permit 10
 match ip address prefix-list ALLOWED_PREFIXES
 set metric 100

router bgp 65201
 address-family ipv4
 neighbor 10.142.20.5 route-map AS65101-OUT out
 end

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

Answer: C

NEW QUESTION 6

When configuring traffic engineering tunnels in Cisco MPLS core network, you see the traffic is not taking the expected path in the core. Which command do you use to quickly check path of a TE tunnel?

- A. Traceroute mpls ipv4 -tunnel destination
- B. Ping <tunnel destination IP>
- C. show mpls traffic-engineering tunnels
- D. traceroute <tunnel destination IP>

Answer: A

NEW QUESTION 7

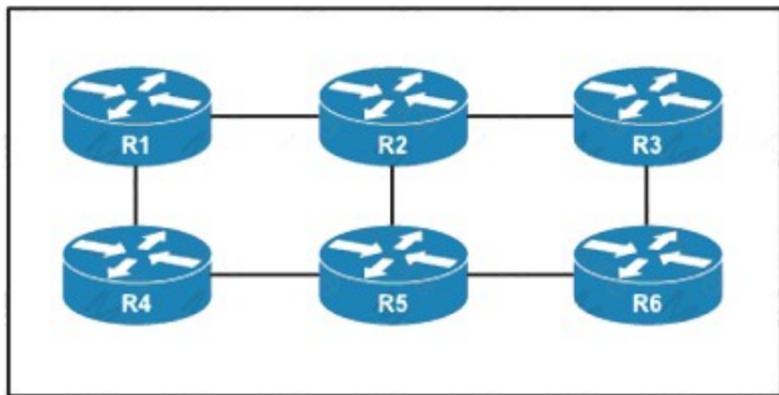
How is a telemetry session established for data analytics?

- A. A router initiates a session using the dial-out to a destination.
- B. A destination initiate a session to a router.
- C. The destination initiate a session using the dial-out more to the router.
- D. A router requests the data using Teinet.

Answer: A

NEW QUESTION 8

Refer to the exhibit.



An engineer is configuring an administrative domain in the given multi-vendor environment with PIM-SM. Which feature must the engineer implement so that devices will dynamically learn the RP?

- A. Auto-RP
- B. BIDIR-PIM
- C. SSM
- D. BSR

Answer: D

NEW QUESTION 10

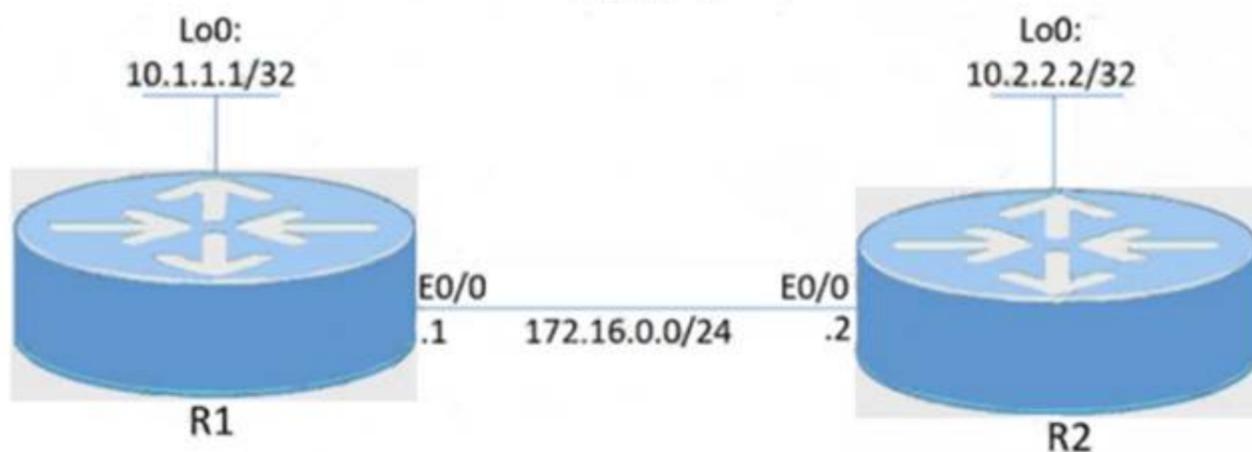
Guidelines



This is a lab item in which tasks will be performed on virtual devices.

- Refer to the Tasks tab to view the tasks for this lab item.
- Refer to the Topology tab to access the device console(s) and perform the tasks.
- Console access is available for all required devices by clicking the device icon or using the tab(s) above the console window.
- All necessary preconfigurations have been applied.
- Do not change the enable password or hostname for any device.
- Save your configurations to NVRAM before moving to the next item.
- Click Next at the bottom of the screen to submit this lab and move to the next question.
- When Next is clicked, the lab closes and cannot be reopened. Topology

**OSPF Process ID 10
Area 0**



Tasks

Configure and verify the OSPF neighbor adjacency between R1 and R2 in OSPF area 0 according to the topology to achieve these goals:

- * 1. Establish R1 and R2 OSPF adjacency. All interfaces must be advertised in OSPF by using the OSPF interface command method. Use Loopback0 as the OSPF ID.
- * 2. There must be no DR/BDR elections in OSPF Area 0 when establishing the neighbor relationship between R1 and R2. OSPF must not generate the host entries /32 for the adjacent interfaces.
- * 3. Enable OSPF MD5 Authentication between both routers at the interface level with password C1sc0!

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Here is the solution:

Graphical user interface, text Description automatically generated

```
R1:
Conf t
Router ospf 10
Router-id 10.1.1.1
```

```
interface e0/0
ip ospf 10 area 0
ip ospf network point-to-point
ip ospf message-digest-key 1 md5 C1sc0!
```

```
int lo0
ip ospf 10 area 0
```

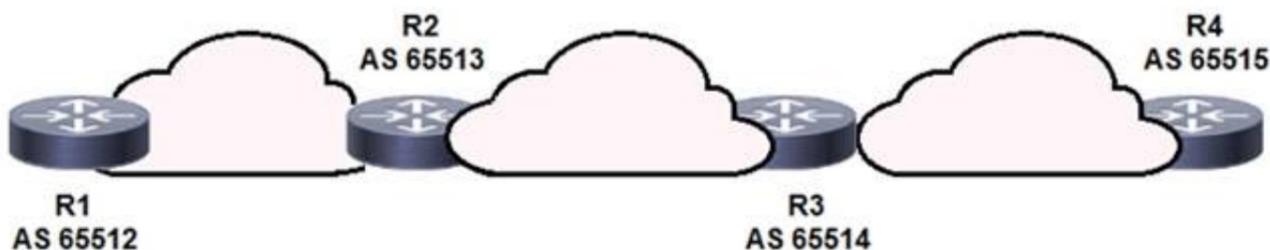
```
R2:
Conf t
Router ospf 10
Router-id 10.2.2.2
```

```
interface e0/0
ip ospf 10 area 0
ip ospf network point-to-point
ip ospf message-digest-key 1 md5 C1sc0!
```

```
int lo0
ip ospf 10 area 0
```

NEW QUESTION 12

Refer to the exhibit:



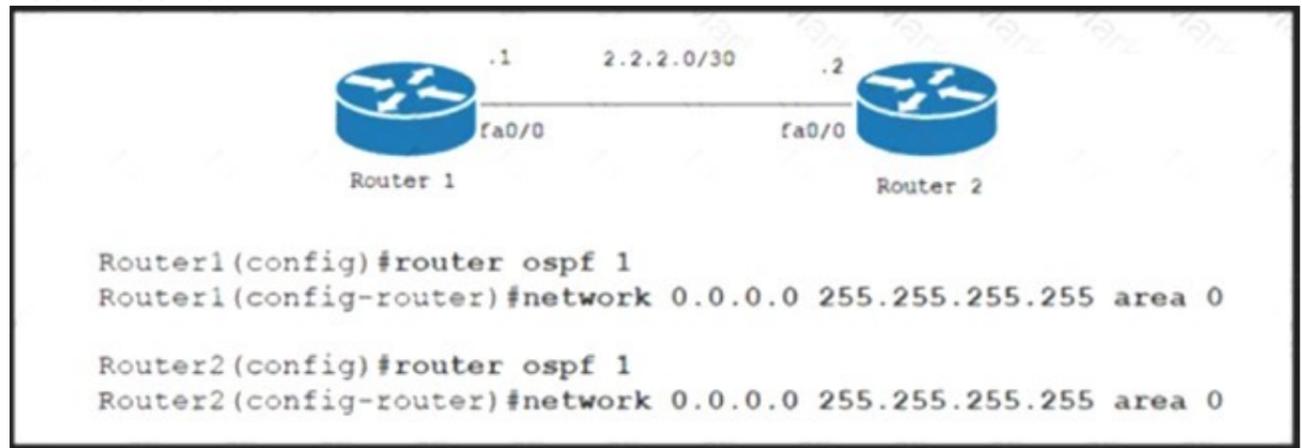
BGPsec is implemented on R1. R2, R3, and R4 BGP peering is established between neighboring autonomous systems Which statement about implementation is true?

- A. BGP updates from the eBGP peers are appended with an additional AS path value that is statically set by the domain administrator
- B. BGP updates from the iBGP peers are appended with a community of local-as
- C. BGP updates from the all BGP peers are appended with a community of no export
- D. BGP updates from the eBGP peers are appended with a BGPsec attribute sequence that includes a public key hash and digital signature

Answer: D

NEW QUESTION 17

Refer to the exhibit.



A network engineer must configure an LDP neighborhood between two newly installed routers that are located in two different offices. Router 1 is the core router in the network and it has already established OSPF adjacency with router 2. On router 1 and router 2, interface fa0/0 is configured for BFD. Which additional configuration must the engineer apply to the two devices to meet the requirement?

- A. Router1(config)#int fa0/0 - Router1(config-if)#mpls ldp autoconfig Router2(config)#router ospf 1 - Router2(config-router)#mpls ip
- B. Router1(config)#int fa0/0 - Router1(config-if)#mpls ip Router1(config-if)#mpls ldp discovery transport-address interface Router2(config)#int fa0/0 Router2(config-

```
if)#mpls ip Router2(config-if)#mpls ldp discovery transport-address interface
C. Router1(config)#int fa0/0 - Router1(config-if)#mpls ldp autoconfig Router1(config-if)#mpls ldp discovery interface Router2(config)#router ospf 1 Router2(config-router)#mpls ldp autoconfig Router2(config-if)#mpls ldp discovery interface
D. Router1(config)#int fa0/0 - Router1(config-if)#mpls ip - Router2(config)#router ospf 1 Router2(config-router)#mpls ldp autoconfig
```

Answer: D

NEW QUESTION 22

Drag and drop the NAT64 descriptions from the left onto the correct NAT64 types on the right.

It is limited on the number of endpoints.	Stateful
It uses address overloading.	
It conserves IPv4 addresses.	
It mandates IPv4-translatable IPv6 address allocation.	Stateless
It has 1:N translation.	

- A. Mastered
- B. Not Mastered

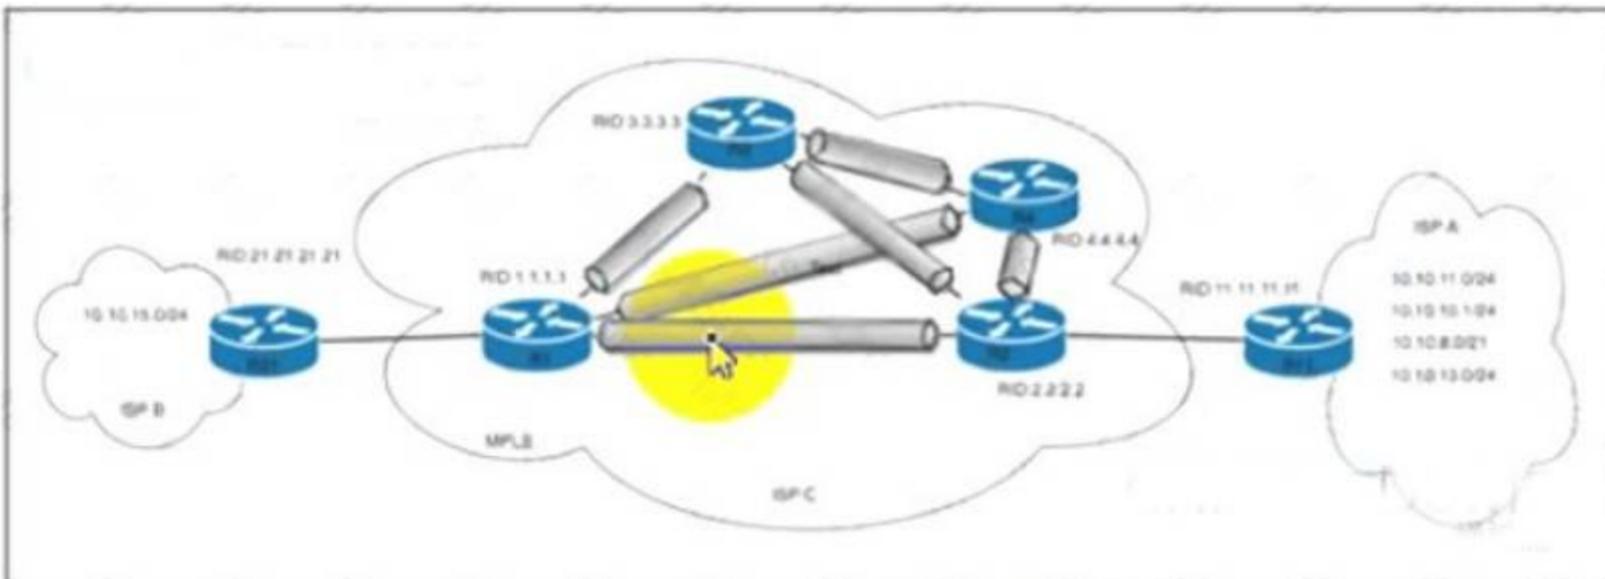
Answer: A

Explanation:

Stateful (It has 1: N translation, It uses address overloading, It conserves IPv4 addresses)
 Stateless (It is limited on the number of endpoints, It mandates IPv4-translatable IPv6 address allocation)

NEW QUESTION 25

Refer to the exhibit



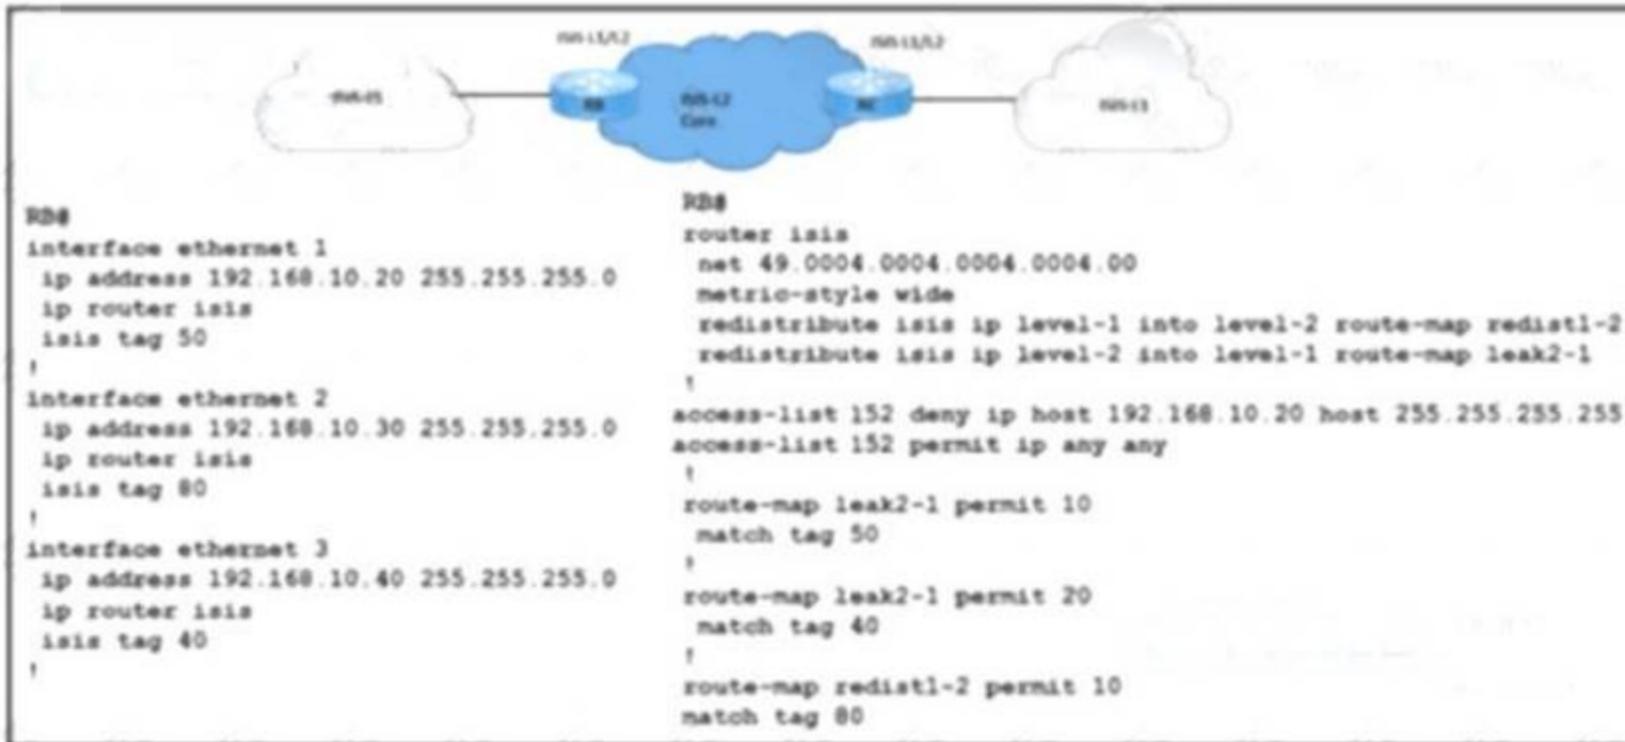
An engineer at ISP C is configuring a new interconnection with ISPs A and B using the BGP protocol. After the initial configuration, the engineer noticed high memory usage and an abnormally large LIB table on router R2. Which two actions must the engineer take on R2 to minimize memory usage? (Choose two.)

- A. Configure Extended ACL 101 with accepted prefixes.
- B. Configure the mpls ldp neighbor 11.11.11.11 labels accept1 command.
- C. Configure Standard ACL 1 with accepted prefixes.
- D. Configure the mpls ldp neighbor 1.1.1.1 labels accept 101 command.
- E. Configure the mpls ldp neighbor 21.21.21.21 labels accept 101 command.

Answer: BC

NEW QUESTION 29

Refer to the exhibit.



A network engineer with an employee ID 4379:43:595 is setting up an IS-IS network with these requirements:

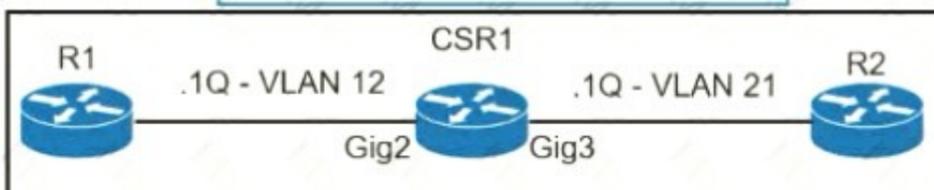
- > Routes with a tag of 80 and IP prefixes other than 192.168.10.20/24 must be redistributed from Level 1 into Level 2.
 - > Route leaking must be configured from Level 2 into the Level 1 domain for routes that are tagged with only 50 or 40.
- Which configuration must be implemented on RB to meet the requirements?

- A. Add match tag 80 in route-map leak2-1
- B. DUMPS Add match ip address 152 in route-map redistrib1-2
- C. Remove match tag 40 from route-map leak2-1
- D. Change match tag 80 to match tag 50 in route-map redistrib1-2.

Answer: D

NEW QUESTION 33

Refer to the exhibit.



A network operator must configure CSR1 interlaces GigabitEthernet2 and GigabitEthernet to rewrite VLAN tags 12 and 21 for traffic between R1 and R2 respectively. Which configurator accomplishes this task?

A)

```

#CSR1
interface GigabitEthernet2
no ip address
service instance 21 ethernet
encapsulation dot1q 21
rewrite ingress tag translate 1-to-1 dot1q 12
rewrite egress tag translate 1-to-1 dot1q 21
bridge-domain 10
!
interface GigabitEthernet3
no ip address
service instance 12 ethernet
encapsulation dot1q 12
rewrite ingress tag translate 1-to-1 dot1q 21
rewrite egress tag translate 1-to-1 dot1q 12
bridge-domain 10
    
```

B)

```
#CSR1
interface GigabitEthernet2
no ip address
service instance 12 ethernet
encapsulation dot1q 12
rewrite ingress tag translate 1-to-1 dot1q 21
rewrite egress tag translate 1-to-1 dot1q 12
bridge-domain 10
!
interface GigabitEthernet3
no ip address
service instance 21 ethernet
encapsulation dot1q 21
rewrite ingress tag translate 1-to-1 dot1q 12
rewrite egress tag translate 1-to-1 dot1q 21
bridge-domain 10
```

C)

```
#CSR1
interface GigabitEthernet2
!
interface GigabitEthernet3
no ip address
service instance 21 ethernet
encapsulation dot1q 21
rewrite ingress tag translate 1-to-1 dot1q 12
rewrite egress tag translate 1-to-1 dot1q 21
bridge-domain 21
```

D)

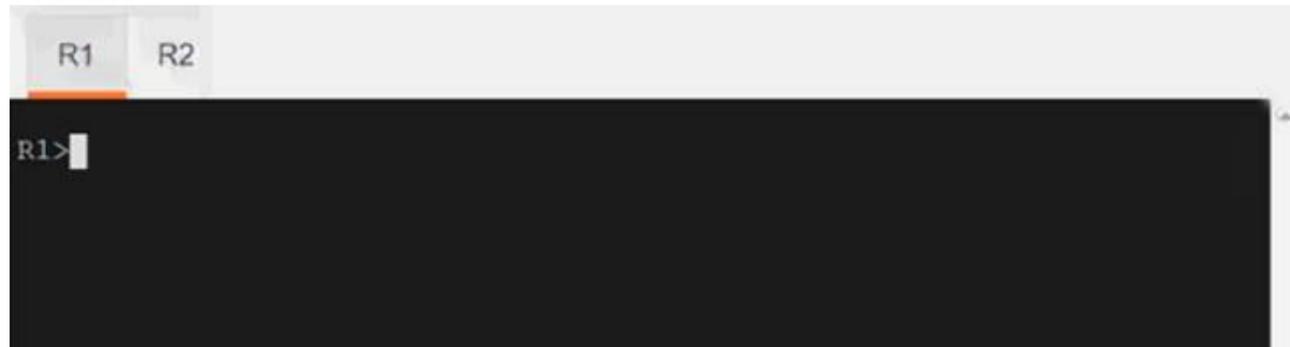
```
#CSR1
interface GigabitEthernet2
no ip address
service instance 12 ethernet
encapsulation dot1q 12
rewrite ingress tag translate 1-to-1 dot1q 21
rewrite egress tag translate 1-to-1 dot1q 12
!
interface GigabitEthernet3
no ip address
service instance 21 ethernet
encapsulation dot1q 21
rewrite ingress tag translate 1-to-1 dot1q 12
rewrite egress tag translate 1-to-1 dot1q 21
```

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

Answer: B

NEW QUESTION 36

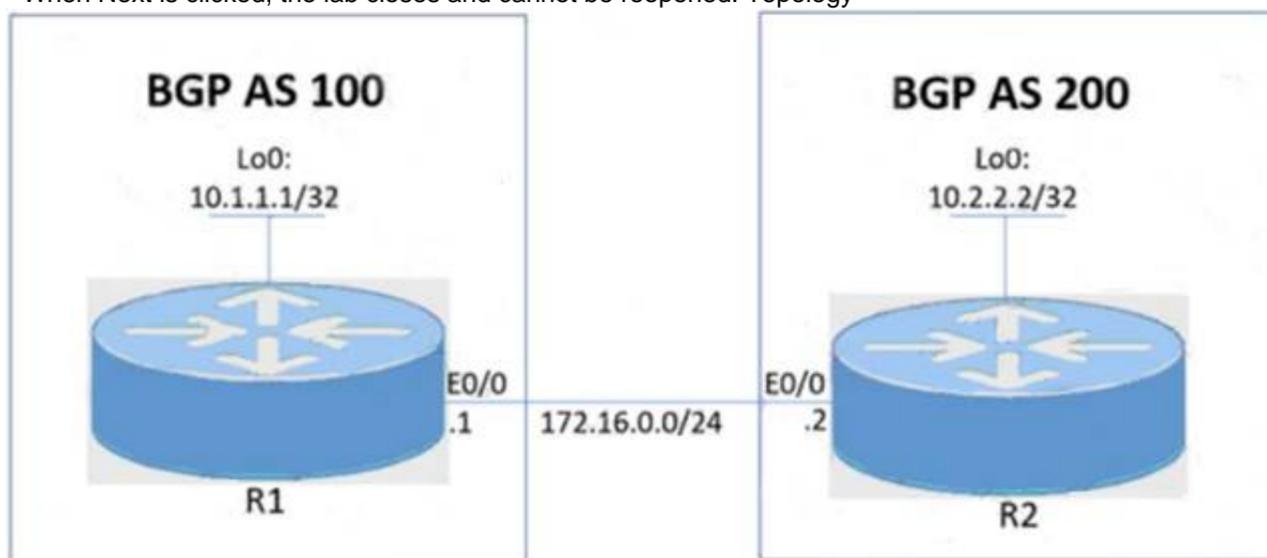
Guidelines



This is a lab item in which tasks will be performed on virtual devices.

- Refer to the Tasks tab to view the tasks for this lab item.
- Refer to the Topology tab to access the device console(s) and perform the tasks.
- Console access is available for all required devices by clicking the device icon or using the tab(s) above the console window.
- All necessary preconfigurations have been applied.
- Do not change the enable password or hostname for any device.
- Save your configurations to NVRAM before moving to the next item.

- Click Next at the bottom of the screen to submit this lab and move to the next question.
- When Next is clicked, the lab closes and cannot be reopened. Topology



Tasks

R1 and R2 are having issues forming an eBGP neighbor relationship. Troubleshoot and resolve the issue to achieve these goals:

- * 1. Configure R1 and R2 to form a BGP neighborship using their Loopback interfaces.
- * 2. Form the neighbor relationship using a BGP multihop mechanism. Use minimal values to solve the issue.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Here is the solution:

Text Description automatically generated

R1:

```
conf t
```

```
ip route 10.2.2.2 255.255.255.255 172.16.0.2
```

```
router bgp 100
```

```
neighbor 10.2.2.2 remote-as 200
```

```
neighbor 10.2.2.2 update-source lo0
```

```
neighbor 10.2.2.2 disable-connected-check
```

```
neighbor 10.2.2.2 ebgp-multihop 2
```

```
address-family ipv4 unicast
```

```
neighbor 10.2.2.2 activate
```

```
do copy running-config startup-config
```

R2:

```
conf t
```

```
ip route 10.1.1.1 255.255.255.255 172.16.0.1
```

```
router bgp 200
```

```
neighbor 10.1.1.1 remote-as 100
```

```
neighbor 10.1.1.1 update-source lo0
```

```
neighbor 10.1.1.1 disable-connected-check
```

```
neighbor 10.1.1.1 ebgp-multihop 2
```

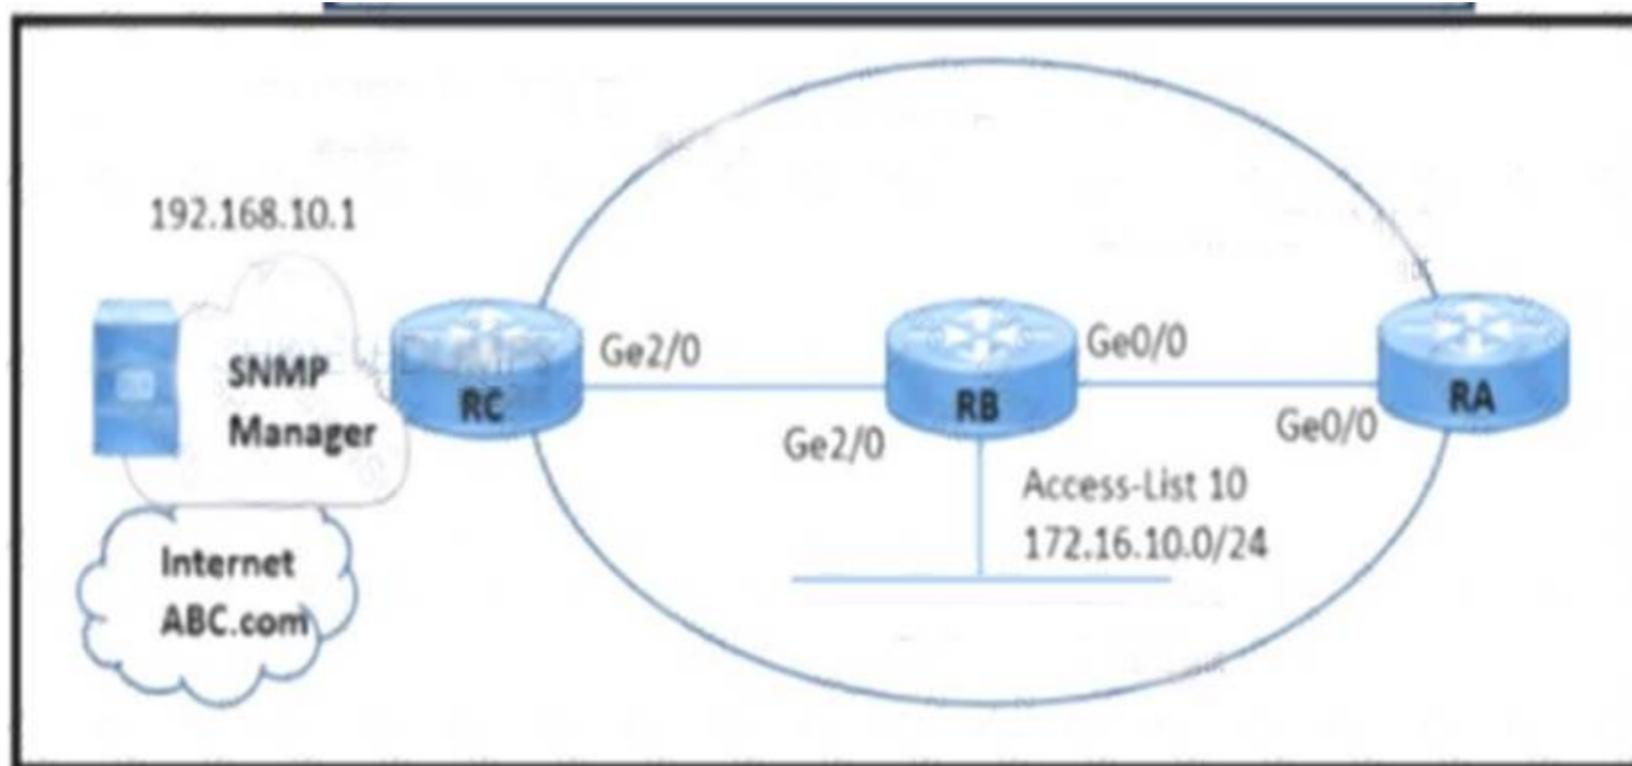
```
address-family ipv4 unicast
```

```
neighbor 10.1.1.1 activate
```

```
do copy running-config startup-config
```

NEW QUESTION 38

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

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

Answer: A

NEW QUESTION 40

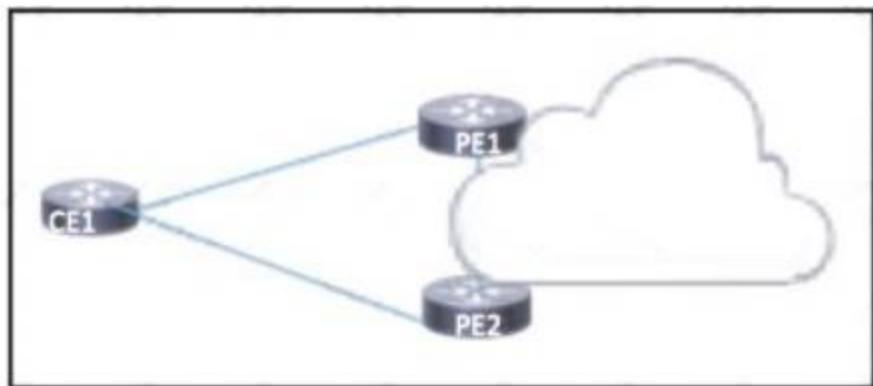
How is RSVP used with MPLS traffic engineering tunnels?

- A. It assigns a tag to a packet as it travels through the tunnel.
- B. It removes and reassigns an MPLS label when the packet enters the tunnel.
- C. It reduces the CPU burden when a packet travels through the tunnel.
- D. It reserves bandwidth along the path of the tunnel.

Answer: C

NEW QUESTION 45

Refer To the exhibit.



Which BGP attribute should be manipulated to have CE1 use PE1 as the primary path to the Internet?

- A. The weight attribute should be manipulated on PE1 on outbound routes advertised to CE1.
- B. The MED should be manipulated on CE1 on inbound routes from PE1.
- C. The local preference attribute should be manipulated on PE2 on inbound routes advertised to CE1.
- D. The origin of all routes should be modified on each router on inbound and outbound routes advertised to CE1.

Answer: B

NEW QUESTION 49

Which OS uses a distributed subsystem architecture?

- A. IOS XE
- B. IOS
- C. IOS XR
- D. CatOS

Answer: C

NEW QUESTION 52

How does Cisco MPLS TE use OSPF extensions to allow for optimized transit between a headend router and a destination router?

- A. Router LSAs share router link advertisements to each router within the MPLS environment so that tunnels can be built bidirectionally.
- B. ASBR Summary LSAs share OSPF domain information so that the two routers know how to reach each other during tunnel setup.
- C. Network LSAs share RSVP information to build the tunnel between the two routers.
- D. Opaque LSAs calculate and establish unidirectional tunnels that are set according to the network constraint.

Answer: D

Explanation:

Cisco MPLS TE uses OSPF extensions to allow for optimized transit between a headend router and a destination router by utilizing Opaque LSAs. Opaque LSAs allow for the calculation and establishment of unidirectional tunnels that are set according to the network constraint. The tunnels are built bidirectionally by utilizing Router LSAs, which share router link advertisements to each router within the MPLS environment. ASBR Summary LSAs are also used to share OSPF domain information so that the two routers know how to reach each other during tunnel setup. Furthermore, Network LSAs are used to share RSVP information which is necessary for setting up the tunnel between the two routers.

NEW QUESTION 53

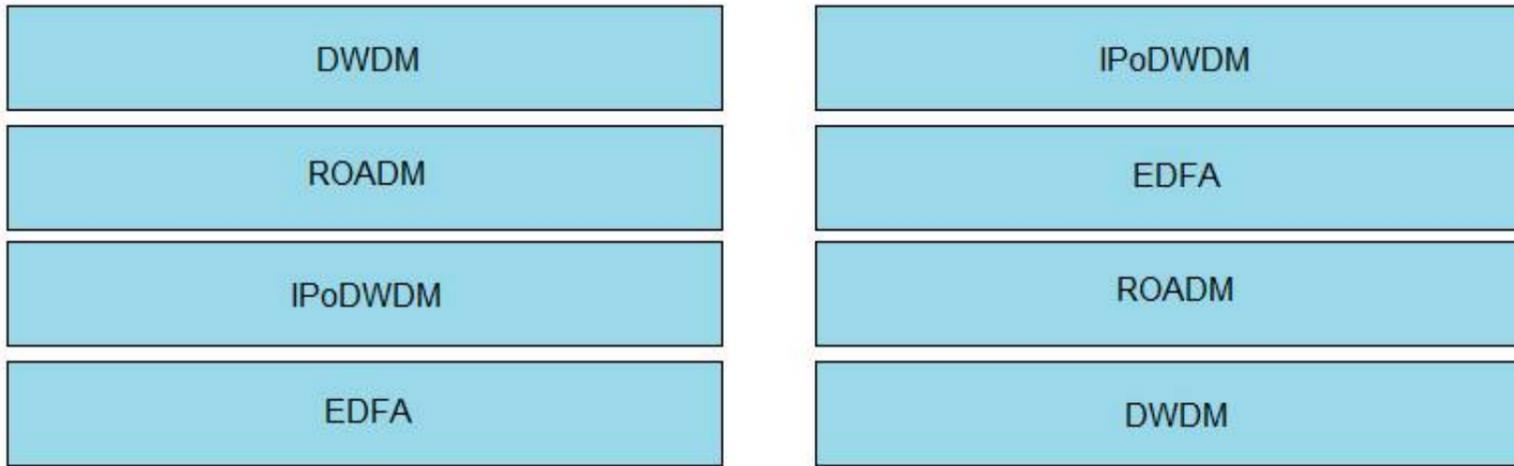
Drag and drop the technologies from the left onto the correct definitions on the right.

DWDM	required for routes and switches to have DWDM and ITU-T G.709 implemented
ROADM	used to amplify an optical signal
IPoDWDM	used to drop certain lambdas within a DWDM ring at a specific location
EDFA	increases bandwidth over a single fiber by using different wavelengths

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 57

A network engineer must implement SNMPv2 with these parameters

- Enable SNMP community string C1sc0 with read-only permissions.
- Enable interface index persistence.
- Restrict the SNMP community to only the monitoring server with IP address 198.18.19.100/32.
- Provide view-only access to ospflEntry and ospfNbrEntry.

Which configuration must the engineer apply?

- configure terminal**
access-list 5 permit 198.18.19.100 0.0.0.0
snmp-server view BLOCKED_VIEW internet excluded
snmp-server view BLOCKED_VIEW ospflEntry included
snmp-server view BLOCKED_VIEW ospfNbrEntry included
snmp-server community c1sc0 view BLOCKED_VIEW RO 5
snmp ifmib ifindex persist
end
- configure terminal**
access-list 5 permit 198.18.19.100 0.0.0.0
snmp-server view BLOCKED_VIEW internet excluded
snmp-server view BLOCKED_VIEW ospflEntry included
snmp-server view BLOCKED_VIEW ospfNbrEntry included
snmp-server community c1sc0 view BLOCKED_VIEW RW 5
snmp ifmib ifindex persist
end
- configure terminal**
access-list 5 permit 198.18.19.100 0.0.0.0
snmp-server view BLOCKED_VIEW internet included
snmp-server view BLOCKED_VIEW ospflEntry included
snmp-server view BLOCKED_VIEW ospfNbrEntry included
snmp-server community c1sc0 view BLOCKED_VIEW RO
snmp ifmib ifindex persist
end
- configure terminal**
access-list 5 permit 198.18.19.100 0.0.0.0
snmp-server view BLOCKED_VIEW internet excluded
snmp-server view BLOCKED_VIEW ospflEntry included
snmp-server view BLOCKED_VIEW ospfNbrEntry included
snmp-server community c1sc0 view BLOCKED_VIEW RO
snmp ifmib ifindex persist
end

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

Answer: A

NEW QUESTION 60

Which programmable API allows the service provider to plan and optimize the automation of network operations and achieve closed-loop operations?

- A. Network Services Orchestrator
- B. WAN Automation Engine
- C. Evolved Programmable Network Manager
- D. Crosswork Network Automation

Answer: D

NEW QUESTION 62

How much must the MTU be increased when configuring the 802.1q VLAN tag?

- A. 2 bytes
- B. 4 bytes
- C. 8 bytes
- D. 12 bytes

Answer: B

NEW QUESTION 66

Refer to the exhibit

```

Sep 30 03:12:33: ISIS-Adj: Rec serial IIH from *HDLC* (Serial1/1), cir type L1L2
Sep 30 03:12:33: ISIS-Adj: rcvd state DOWN, old state UP, new state INIT
Sep 30 03:12:33: ISIS-Adj: Action = GOING DOWN
Sep 30 03:12:33: %CLNS-5-ADJCHANGE: ISIS: Adjacency to R1 (Serial1/1) Down, nes
Sep 30 03:12:33: ISIS-Adj: L2 adj count 0
Sep 30 03:12:33: ISIS-Adj: Sending serial IIH on Serial1/1, length 1699
Sep 30 03:12:41: ISIS-Adj: Rec serial IIH from *HDLC* (Serial1/1), cir type L1L2
Sep 30 03:12:41: ISIS-Adj: rcvd state DOWN, old state DOWN, new state INIT
Sep 30 03:12:41: ISIS-Adj: Action = GOING UP, new type = L2
Sep 30 03:12:41: ISIS-Adj: New serial adjacency
Sep 30 03:12:41: ISIS-Adj: Sending serial IIH on Serial1/1, length 1699
Sep 30 03:12:47: ISIS-Adj: Rec serial IIH from *HDLC* (Serial1/1), cir type L1L2
Sep 30 03:12:47: ISIS-Adj: rcvd state DOWN, old state INIT, new state INIT
Sep 30 03:12:47: ISIS-Adj: Action = GOING UP, new type = L2
Sep 30 03:12:47: ISIS-Adj: Sending serial IIH on Serial1/1, length 1699
Sep 30 03:12:47: ISIS-Adj: Sending serial IIH on Serial1/1, length 1699
    
```

Routers R1 and R2 are connected via a serial link and use the IS-IS routing protocol for route exchange. After a configuration change on R2, IS-IS connectivity is interrupted. A network engineer confirmed that the interfaces are in the UP state and connectivity exists between the two routers. Which two actions must the engineer perform to resolve the problem? (Choose two.)

- A. Disable padding for hello packets under the serial interface on R2 DUMPS
- B. Change the hello interface timer to 10 seconds on R1.
- C. Change the MTU to 1500 bytes on R2.
- D. Enable hello packet padding globally on R1.
- E. Change R2 to an IS-IS Level 1 router.

Answer: CE

NEW QUESTION 68

Refer to the exhibit.

```

R1
interface gigabitethernet1/0/0
  ipv6 enable ipv6 ospf 1 area 1
interface gigabitethernet2/0/0
  ipv6 enable ipv6 ospf 1 area 2
    
```

An engineer implemented OSPF neighbor relationship on an IOS device. Which configuration must be applied to get the OR/BOR election removed from interfaces running OSPF?

- A. ip ospf network broadcast on interfaces running OSPF
- B. ip ospf network point-to-point on interfaces running OSPF
- C. ip ospf network multipoint-point on interfaces running OSPF
- D. ip ospf network non-broadcast on interfaces running OSPF

Answer: B

NEW QUESTION 69

Refer to the exhibit:

```

ip flow-export source loopback 0
ip flow-export destination 192.168.1.1
ip flow-export version 5 origin-as
    
```

If the NetFlow configuration is updated to version 9, which additional piece of information can be reported"?

- A. IPv6 flow information
- B. flow sequence numbers
- C. BGP AS information
- D. IPv4 flow information

Answer: A

NEW QUESTION 71

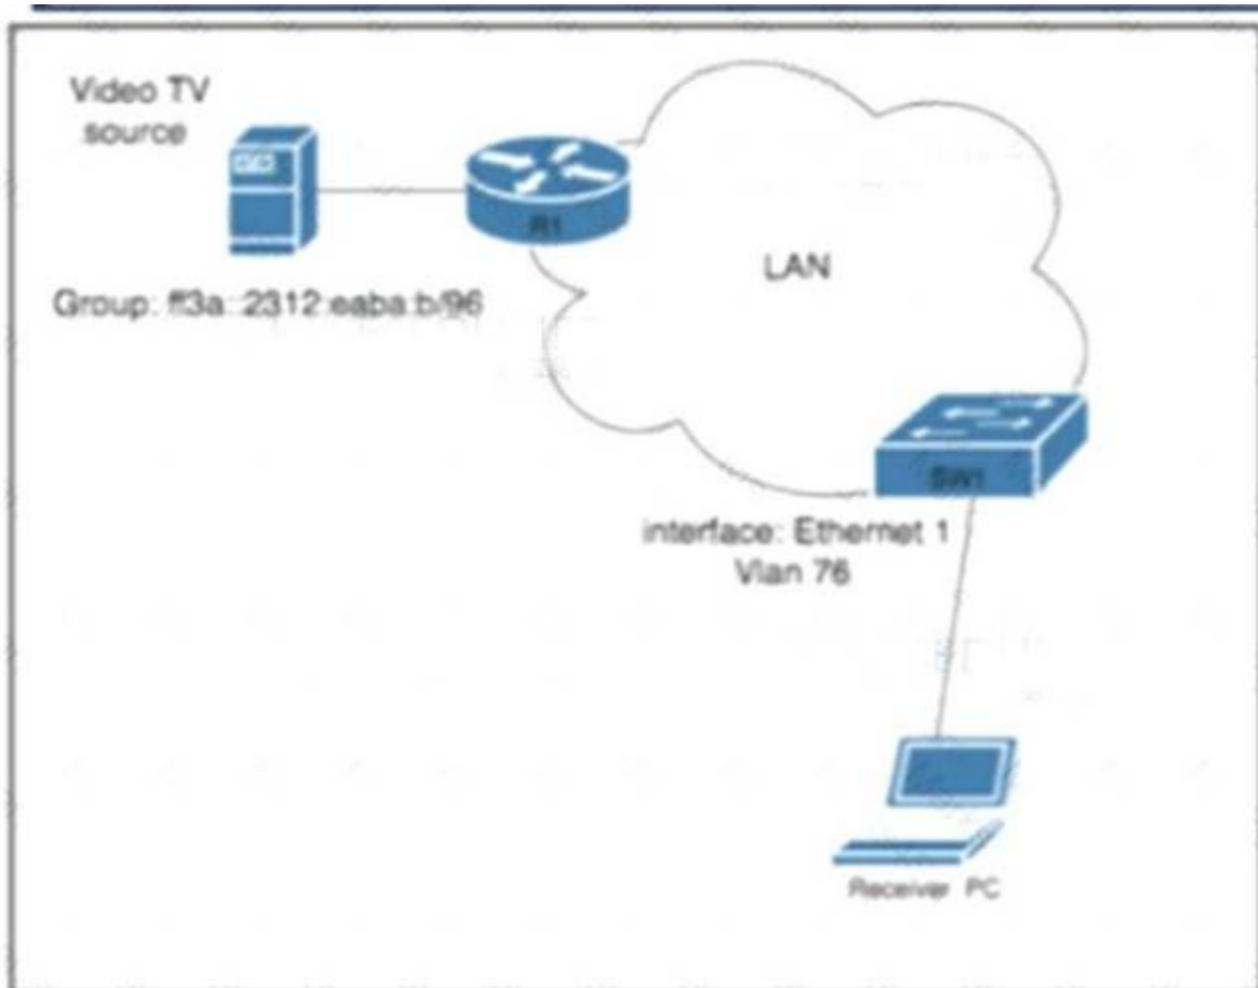
An engineer working for a private service provider with employee id: 3994 37 650 is configuring a Cisco device to redistribute OSPF into BGP. Which task enables the device to filter routes?

- A. Configure a distribute list and associate it to the BGP peer interface
- B. Configure a prefix list and associate it to the BGP peer interface
- C. Configure a route map and reference it with the redistribute command
- D. Configure an access list and reference it with the redistribute command

Answer: C

NEW QUESTION 72

Refer to the exhibit.



A network engineer working for a telecommunication company with an employee ID: 4602:62:646 is configuring security controls for the IPv6 multicast group, which is used for video TV. The solution from the engineer should reduce network usage and minimize the leave latency for the user that is connected to VLAN 76. Which two configurations meet this goal? (Choose two.)

A)

Apply the following commands globally on SW1:

```
ipv6 mld vlan 76 fast-leave vlan 76
ipv6 mld security join vlan 76
```

B)

Configure an ACL to limit the IPv6 multicast group with the entry `permit ipv6 any ff3a::2312:eaba:b/96`.

C)

Configure an ACL to limit the IPv6 multicast group with the entries `ipv6 access-list security_access_list` and `permit ipv6 ff3a::2312:eaba:b/96 any`.

D)

Apply the following commands globally on SW1:

```
ipv6 mld vlan 76 immediate-leave
ipv6 mld snooping
```

E)

Apply the following commands globally on SW1:

```
ipv6 mld snooping multicast optimise-multicast-flood
ipv6 mld snooping fast-leave group security_access_list
```

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

- D. Option D
- E. Option E

Answer: DE

NEW QUESTION 76

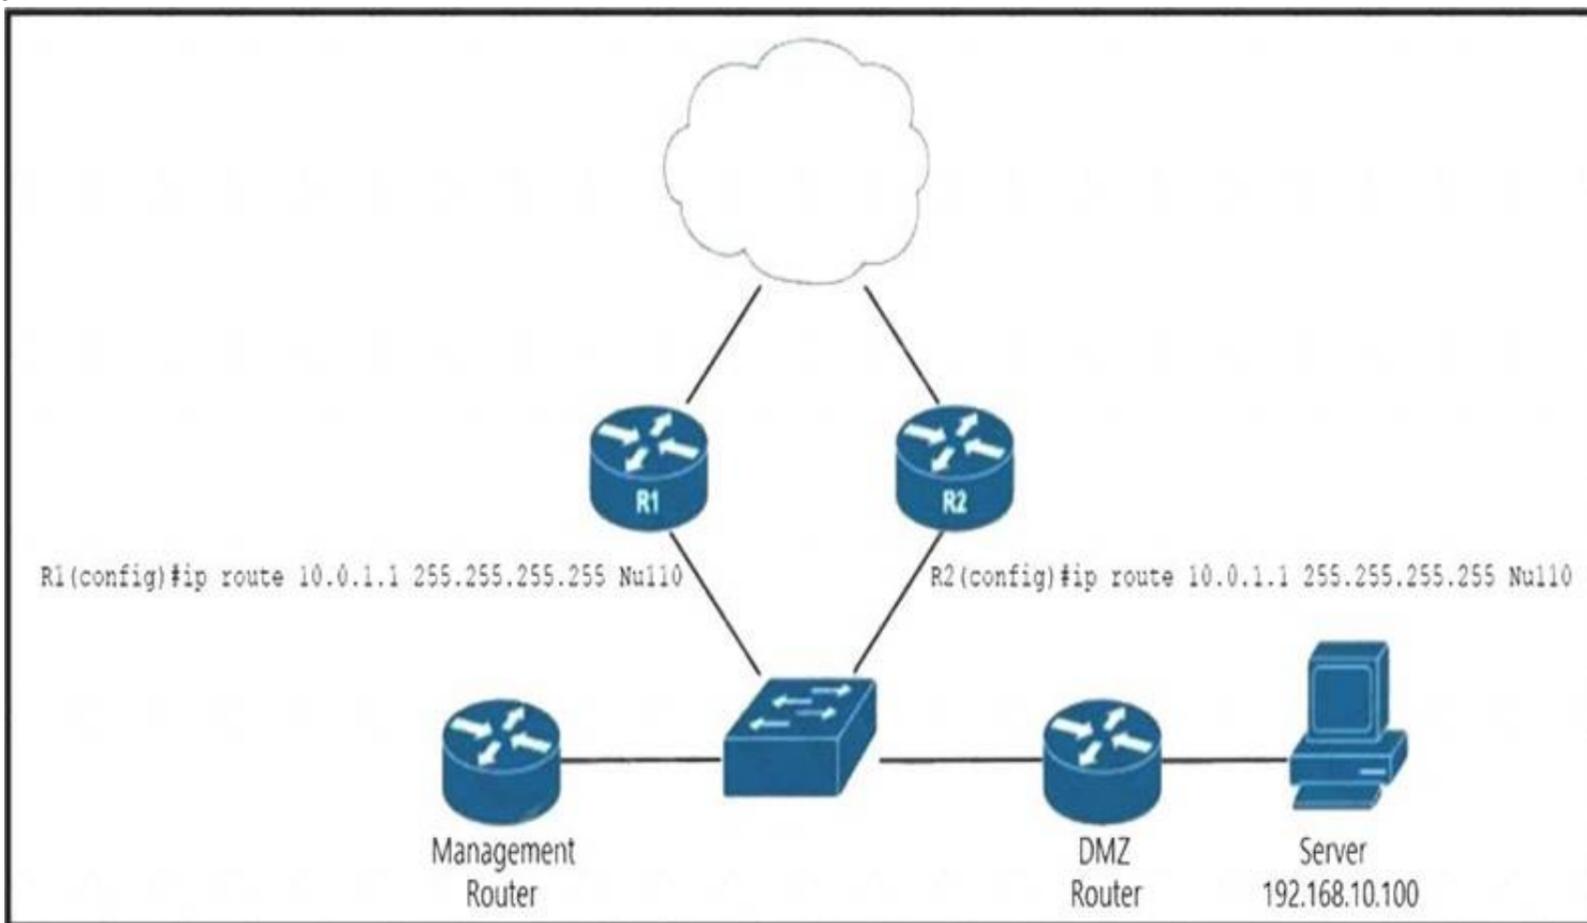
Which protocol is used for communication between the PCE and PCC?

- A. ICMP
- B. PCEP
- C. CEF
- D. POP

Answer: B

NEW QUESTION 81

Refer to the exhibit.



router(config)# route-map blackhole-trigger router(config-route-map)# match tag 777 router(config-route-map)# set ip next-hop 10.0.1.1 router(config-route-map)# set origin igp router(config-route-map)# set community no-export
 EIGRP is running across the core to exchange internal routes, and each router maintains iBGP adjacency with the other routers on the network. An operator has configured static routes on the edge routers R1 and R2 for IP address 10.0.1.1, which is used as a black hole route as shown. Which configuration should the operator implement to the management router to create a route map that will redistribute tagged static routes into BGP and create a static route to blackhole traffic with tag 777 that is destined to the server at 192.168.10.100?

- A. router(config)# router bgp 55100router(config-router)# redistribute static route-map blackhole-trigger router(config)# ip route 10.0.1.1 255.255.255.255 Null0 tag 777
- B. router(config)# router bgp 55100router(config-router)# redistribute static route-map blackhole-trigger router(config)# ip route 192.168.10.100 255.255.255.255 Null0 tag 777
- C. router(config)# router bgp 55100 router(config-router)# redistribute connectedrouter(config)# ip route 192.168.10.100 255.255.255.255 tag 777
- D. router(config)# router bgp 55100router(config-router)# redistribute connected route-map blackhole-trigger router(config)# ip route 192.168.10.100 255.255.255.255 Null0 tag 777

Answer: B

NEW QUESTION 85

What is a primary benefit of IPoATM or MPLS over ATM backbone service provider networks?

- A. dedicated circuits
- B. variable-length packets
- C. isochronous system
- D. fixed-length cells

Answer: A

NEW QUESTION 88

Refer to the exhibit.

```

R1
ip multicast-routing
ip pim rp-candidate GigabitEthernet1/0/0

interface g1/0/0
  ip pim sparse-mode

R2
ip multicast-routing
ip pim bsr-candidate GigabitEthernet1/0/0

interface g1/0/0
  ip pim sparse-mode
  
```

An engineer configured multicast routing on client's network. What is the effect of this multicast implementation?

- A. R2 floods information about R1 throughout the multicast domain.
- B. R2 is unable to share information because the ip pim autorp listener command is missing.
- C. R1 floods information about R2 throughout the multicast domain.
- D. R2 is elected as the RP for this domain.

Answer: B

NEW QUESTION 90

A network engineer is deploying VPLS configuration between multiple PE routers so that customer's remote offices have end-to-end LAN connectivity. Which additional configuration should the engineer perform on the PE routers to enable the virtual switch instance?

A)

```

interface Vlan 5
xconnect vfi ciscotest
  
```

B)

```

I2 vfi ciscotest manual
vpn id 100
neighbor 192.168.2.2 encapsulation mpls
neighbor 192.168.3.3 encapsulation mpls
  
```

C)

```

interface GigEthernet1/1
switchport mode trunk
switchport trunk encap dot1q
switchport trunk allow vlan 2-10
  
```

D)

```

interface Vlan 100
xconnect vfi ciscotest
ip address 192.168.1.1 255.255.255
  
```

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

Answer: B

Explanation:

https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l2_vpns/configuration/xe-3s/mp-l2-vpns-xe-3s-book/mp

NEW QUESTION 92

Which service is a VNF role?

- A. Compute
- B. Network
- C. Firewall
- D. Storage

Answer: B

NEW QUESTION 94

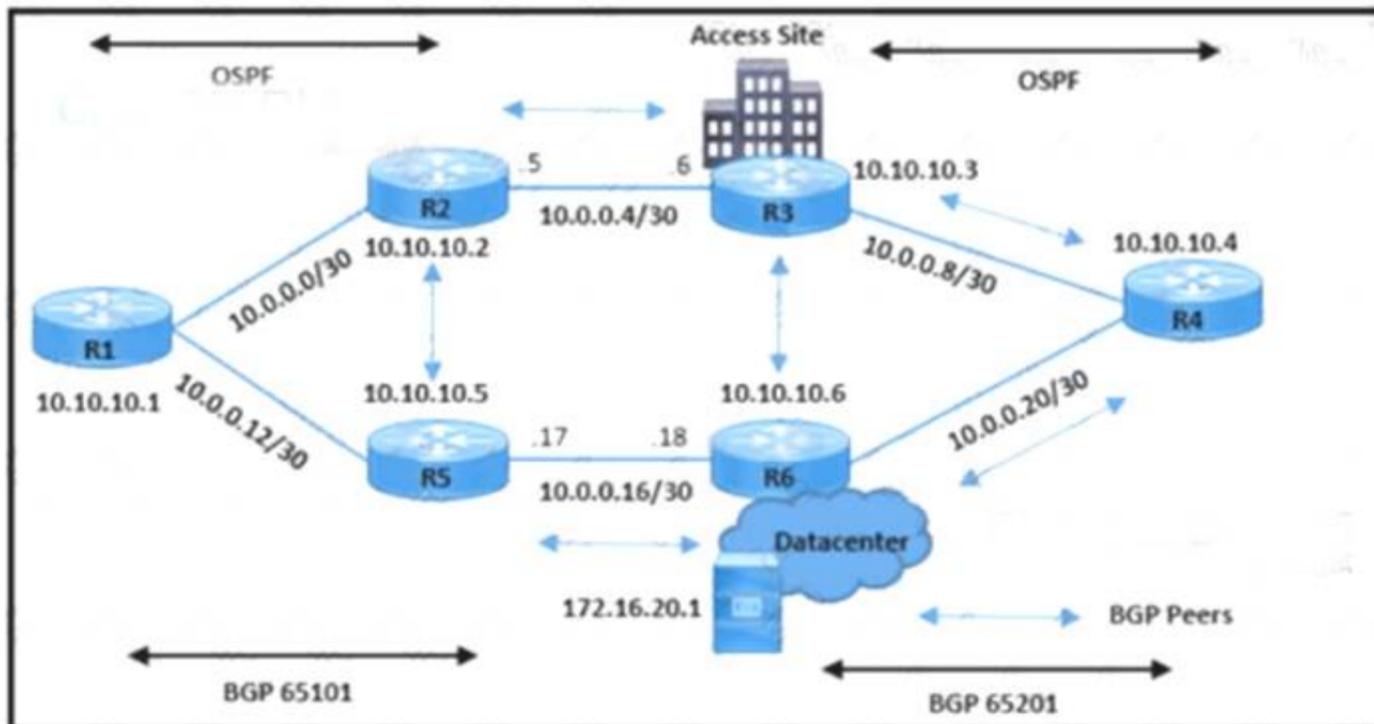
What is a feature of model-driven telemetry?

- A. It occasionally streams to multiple servers in the network.
- B. It is less secure because it uses community strings.
- C. It uses the pull model to send requested data to a client when polled.
- D. It uses the push model to stream data to desired destinations.

Answer: D

NEW QUESTION 97

Refer to the exhibit.



```
R3#show ip route
 192.168.30.0/32 is subnetted, 1 subnets
 B   192.168.30.1 [200/0] via 10.10.10.4, 00:39:23
 172.16.0.0/32 is subnetted, 2 subnets
 O   172.16.20.1 [110/3] via 10.0.0.10, 00:05:39, GigabitEthernet2/0
 B   172.16.10.10 [200/0] via 10.10.10.1, 00:39:23
 10.0.0.0/8 is variably subnetted, 15 subnets, 3 masks
 C   10.0.0.8/30 is directly connected, GigabitEthernet2/0
 O   10.0.0.12/30 [110/3] via 10.0.0.5, 00:41:16, FastEthernet0/0
 S   10.10.10.2/32 [1/0] via 10.0.0.5
 C   10.10.10.3/32 is directly connected, Loopback0
 O   10.0.0.0/30 [110/2] via 10.0.0.5, 00:41:16, FastEthernet0/0

 O   10.10.10.1/32 [110/3] via 10.0.0.5, 00:41:16, FastEthernet0/0
 O   10.10.10.6/32 [110/2] via 10.0.0.29, 00:41:16, FastEthernet1/0
 O   10.10.10.4/32 [110/2] via 10.0.0.10, 00:41:16, GigabitEthernet2/0
 C   10.0.0.4/30 is directly connected, FastEthernet0/0
```

```
 O   10.10.10.1/32 [110/3] via 10.0.0.5, 00:41:16, FastEthernet0/0
 O   10.10.10.6/32 [110/2] via 10.0.0.29, 00:41:16, FastEthernet1/0
 O   10.10.10.4/32 [110/2] via 10.0.0.10, 00:41:16, GigabitEthernet2/0
 C   10.0.0.4/30 is directly connected, FastEthernet0/0
 O   10.10.10.5/32 [110/12] via 10.0.0.5, 00:41:16, FastEthernet0/0
 O   10.0.0.24/30 [110/11] via 10.0.0.5, 00:41:16, FastEthernet0/0
 C   10.0.0.28/30 is directly connected, FastEthernet1/0
 B   10.0.0.16/30 [200/0] via 10.10.10.5, 00:39:23
 O   10.0.0.20/30 [110/2] via 10.0.0.10, 00:41:16, GigabitEthernet2/0
 192.168.1.0/32 is subnetted, 1 subnets

 R4#show ip route 172.16.20.1
Routing entry for 172.16.20.1/32
  Known via "ospf 10", distance 110, metric 2, type intra area
  Last update from 10.0.0.21 on FastEthernet1/0, 00:06:51 ago
Routing Descriptor Blocks:
 * 10.0.0.21, from 172.16.20.1, 00:06:51 ago, via FastEthernet1/0
   Route metric is 2, traffic share count is 1
```

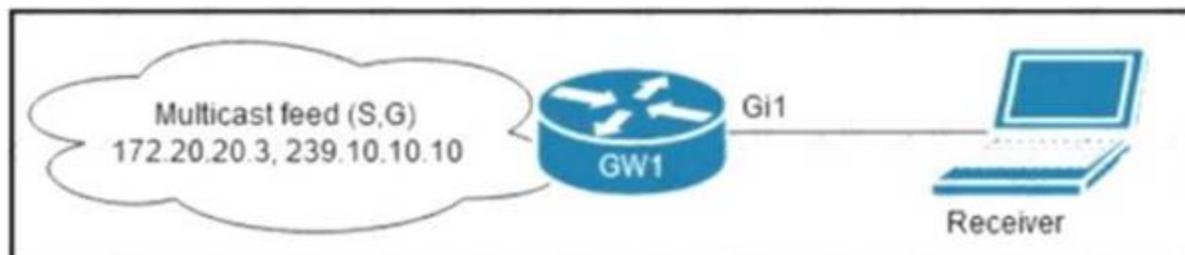
The network operations team reported that the access site that is connected to R3 is not connecting to the application server in the data center and that all packets that are sent from the application server to the access site are dropped. The team verified that OSPF and BGP peerings are up in BGP AS 65101 and BGP AS 65201. R4 is expected to receive traffic from the application server route via OSPF. Which action resolves this issue?

- A. Remove the route-map on R4 when advertising 172.16.20.1 in BGP to R3.
- B. Advertise application server 172.16.20.1 in the OSPF routing table on R6.
- C. Allow 172.16.20.1 in the BGP advertisement on R3 in the route-map.
- D. Add the next-hop-self command on R6 to enable R3 iBGP peering.

Answer: D

NEW QUESTION 102

Refer to the exhibit.



A network administrator is implementing IGMP to enable multicast feed transmission to the receiver. Which configuration must the administrator deploy on GW1 to permit IGMP Joins only to the assigned (S, G) feed?

- A)


```

config t
access-list 100 permit igmp host 0.0.0.0 host 239.10.10.10
access-list 100 deny igmp any any
interface GigabitEthernet1
ip igmp access-group 100
ip igmp version 3
end
      
```
- B)


```

config t
access-list 100 permit igmp host 0.0.0.0 host 239.10.10.10
access-list 100 permit igmp host 172.20.20.3 host 239.10.10.10
access-list 100 deny igmp any any
interface GigabitEthernet1
ip igmp access-group 100
ip igmp version 3
end
      
```
- C)


```

config t
access-list 100 permit igmp host 0.0.0.0 host 239.10.10.10
access-list 100 deny igmp any any
interface GigabitEthernet1
ip igmp access-group 100
ip igmp version 2
end
      
```
- D)


```

config t
access-list 100 permit igmp host 0.0.0.0 host 239.10.10.10
access-list 100 permit igmp host 172.20.20.3 host 239.10.10.10
access-list 100 deny igmp any any
interface GigabitEthernet1
ip igmp access-group 100
ip igmp version 2
end
      
```

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

Answer: B

Explanation:

How IGMP Checks an Extended Access List

When an IGMP extended access list is referenced in the `ip igmp access-group` command on an interface, the (S, G) pairs in the `permit` and `deny` statements of the extended access list are matched against the (S, G) pair of the IGMP reports received on the interface. For example, if an IGMP report with (S1, S2...Sn, G) is received, first the group (0.0.0.0, G) is checked against the access list statements. The convention (0.0.0.0, G) means (*, G), which is a wildcard source with a multicast group number. If the group is denied, the entire IGMP report is denied. If the group is permitted, each individual (S, G) pair is checked against the access list. Denied sources are taken out of the IGMP report, thereby denying the sources access to the multicast traffic.

NEW QUESTION 105

A customer site is being connected to a Frame Relay network via a T1 link. The customer has a contract for 512 kbps service with a Tc value of 125 ms. Under peak line conditions, customer traffic can reach four times the contracted speed. Which QoS configuration must the service provider implement to limit the customer to the contracted values?

- policy-map policy_map
class class_map
police cir 512000 bc 64000 pir 20480000 be 192000
conform-action transmit
exceed-action drop
- policy-map policy_map
class class_map
police cir 512kbps bc 256kbps pir 2Mbps be 9600 kbps
conform-action transmit
exceed-action set-de-bit transmit
violate-action drop
- policy-map policy_map
class class_map
police cir 512000 bc 128000 pir 256000 be 32000
conform-action transmit
exceed-action set-be-bit transmit
exceed-action drop
- policy-map policy_map
class class_map
police cir 512000 bc 32000 pir 64000 be 6400
conform-action transmit
violate-action set-dscp-transmit default
exceed-action drop

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

Answer: A

NEW QUESTION 110

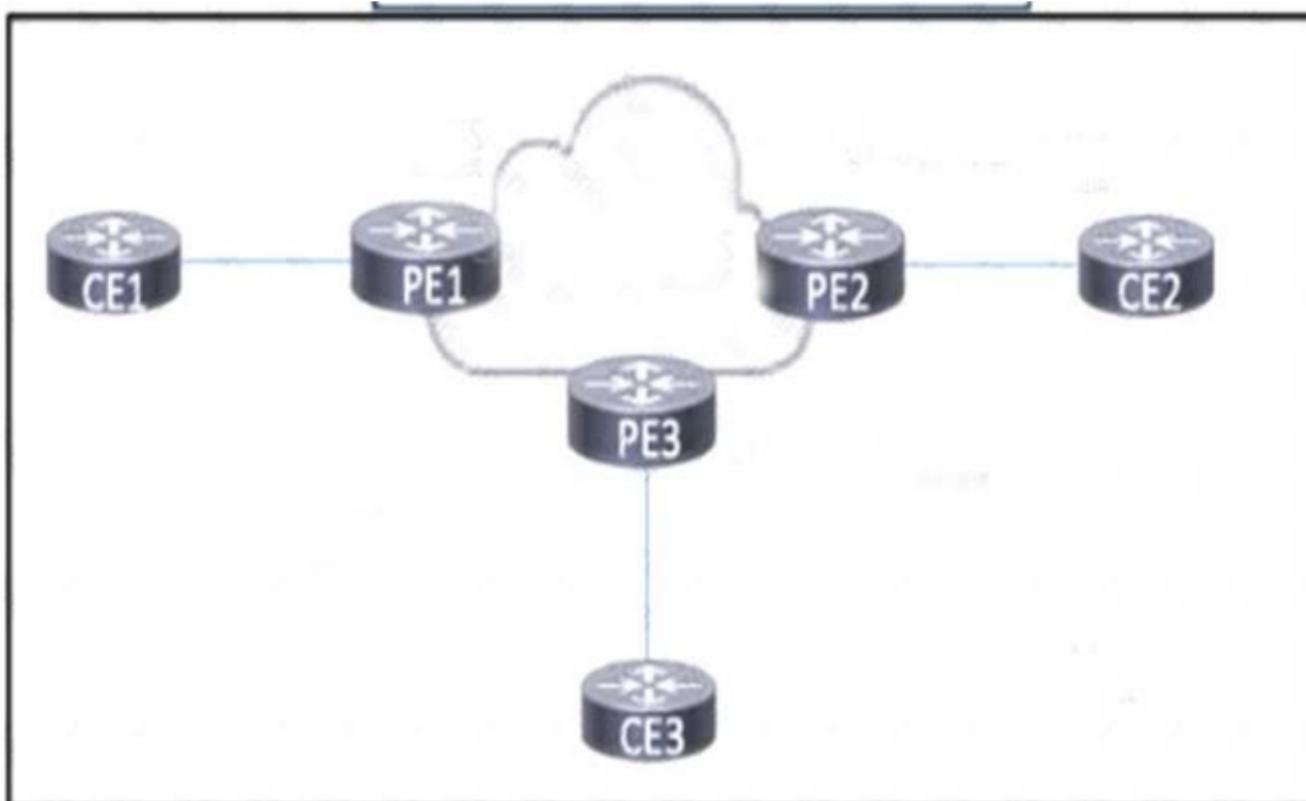
What must a network engineer consider when designing a Cisco MPLS TE solution with OSPF?

- A. The OSPF extensions and RSVP-TE must be enabled on all routers in the network.
- B. OSPF extensions for RSVP-TE are supported in Area 1.
- C. The OSPF extensions and RSVP-TE must be enabled on the egress routers.
- D. OSPF extensions for RSVP-TE are implemented in Type 6, 7, and 8 LSAs.

Answer: A

NEW QUESTION 111

Refer to the exhibit.



A large enterprise has multiple branch offices that span several geographic regions. The enterprise runs MPLS within the core to propagate VPNv4 routes using BGP. After a recent series of DDoS attacks disrupted the network, a network engineer has been asked to reconfigure BGP to help mitigate future attacks. Which configuration must the engineer apply?

A)

```
router bgp 100
address-family ipv4 flowspec
neighbor 192.168.1.1 activate
```

B)

```
router bgp 100
address-family ipv4 mdt
neighbor 192.168.1.1 activate
```

C)

```
router bgp 100
address-family ipv4
neighbor 192.168.1.1 activate
```

D)

```
router bgp 100
address-family vpnv4
neighbor 192.168.1.1 activate
```

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

Answer: A

NEW QUESTION 116

An engineer must implement QoS to prioritize traffic that requires better service throughout the network. The engineer started by configuring a class map to identify the high-priority traffic. Which additional tasks must the engineer perform to implement the new QoS policy?

- A. Attach the class map to a policy map that sets the minimum bandwidth allocated to the classified traffic and designates the action to be taken on the traffic.
- B. Attach the class map to a policy map that designates the action to be taken on the classified traffic and then attach the policy map to an interface using a service policy.
- C. Attach the class map to a policy map within a VRF to segregate the high-priority traffic and then attach the policy map to an interface in another VRF.
- D. Create a route map to manipulate the routes that are entered into the routing table and then attach the route map to an interface using a service policy.

Answer: B

NEW QUESTION 119

Refer to the exhibit:

```
telemetry model-driven
sensor-group cisco
sensor-path Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface/latest/generic-counters
commit
```

This configuration is being applied on an IOS XR router. Which statement about this configuration is true?

- A. It is used to create a subscription to specify the streaming interval
- B. It is used to identify traps for SNMP polling
- C. It is used to identify MIB entries and has a list of YANG models
- D. It is used to create a sensor-group and has a list of YANG models for streaming

Answer: D

NEW QUESTION 122

Refer to the exhibit:

<pre>PE-A ! interface FastEthernet0/0 ip address 10.10.10.1 255.255.255.252 ip ospf authentication null ip ospf 1 area 0 duplex full end ! router ospf 1 log-adjacency-changes passive-interface Loopback0 network 10.10.10.0 0.0.0.3 area 0 default-metric 200 !</pre>	<pre>PE-B ! interface FastEthernet0/0 ip address 10.10.10.2 255.255.255.252 ip ospf authentication null ip mtu 1400 ip ospf 1 area 0 duplex half end ! R1#sho run b router ospf router ospf 1 log-adjacency-changes passive-interface Loopback10 network 10.10.10.0 0.0.0.255 area 0 default-metric 100</pre>
--	---

Which configuration prevents the OSPF neighbor from establishing?

- A. mtu
- B. duplex
- C. network statement
- D. default-metric

Answer: A

NEW QUESTION 123

An network engineer is deploying VRF on ASBR router R1. The interface must have connectivity over an MPLS VPN inter-AS Option AB network. Which configuration must the engineer apply on the router to accomplish this task?

- A)


```
R1(config)# interface ethernet 1/0
R1(config-if)# ip vrf forwarding CISCO
R1(config-if)# mpls ip
```
- B)


```
R1(config)# interface ethernet 1/0
R1(config-if)# ip address 192.168.1.254 255.255.255.0
R1(config-if)# ip vrf forwarding CISCO
R1(config-if)# shutdown
```
- C)


```
R1(config)# interface ethernet 1/0
R1(config-if)# ip vrf forwarding CISCO
R1 (config-if)# ip ospf 1 area 0
```
- D)

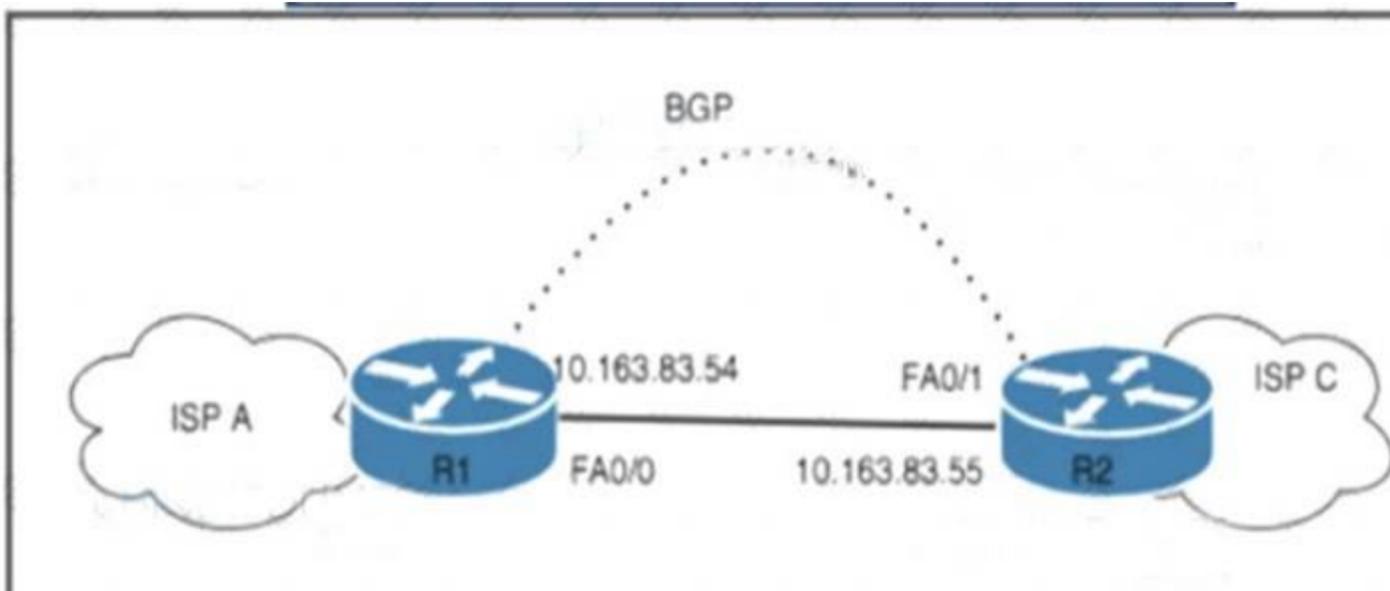

```
R1(config)# interface ethernet 1/0
R1(config-if)# ip vrf forwarding CISCO
R1(config-if)# mpls bgp forwarding
```

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

Answer: D

NEW QUESTION 127

Refer to the exhibit.



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?

- A. Configure bgp maxas-limit 1 in the IPv4 address family under the 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-security hops 2 under the global BGP configuration.

Answer: D

NEW QUESTION 129

Refer to the exhibit:

```
ip flow-export source loopback 0
ip flow-export destination 192.168.1.1
ip flow-export version 9 origin-as
```

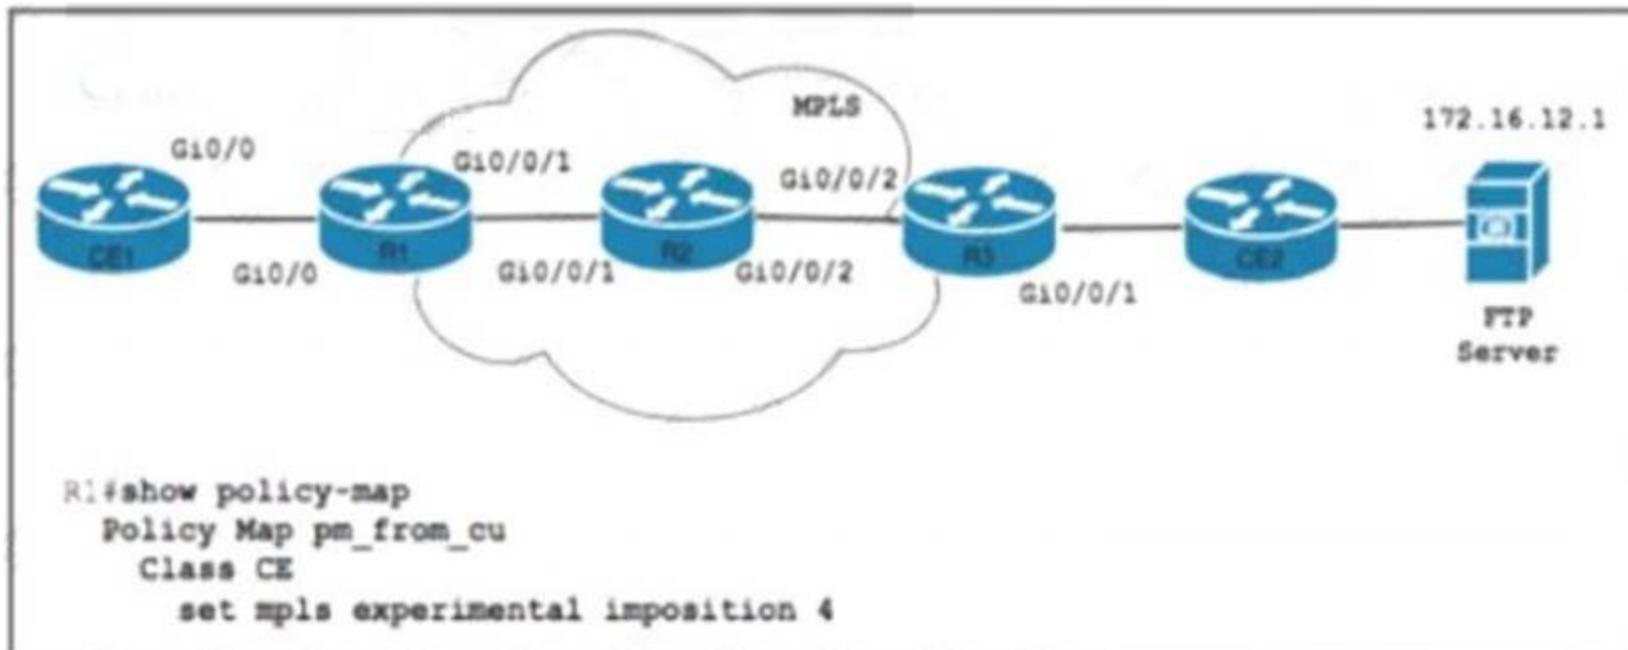
Export statistics received do not include the BGP next hop. Which statement about the NetFlow export statistics is true?

- A. Only the origin AS of the source router will be included in the export statistics.
- B. Loopback 0 must be participating in BGP for it to be included in the export statistics.
- C. The origin AS and the peer-as will be included in the export statistics.
- D. To include the BGP next hop in the export statistics, those keywords must be included with the version 9 entry.

Answer: D

NEW QUESTION 134

Refer to the exhibit.



Router R1 is configured with class map CE with match Ip precedence critical to align with customer contract SLAs. The customer is sending all traffic from CE1 toward the FTP server with IP precedence 5. A network engineer must allow 10% of interface capacity on router R3. Which two actions must the engineer take to accomplish the task? (Choose two)

- A. Implement a class map on R1 to match all packets with QoS IP precedence value 100.
- B. Implement a class map on R3 to match all packets with QoS IP precedence value 101.
- C. Apply a policy map to R1 to reserve the remaining 10% of interface bandwidth.
- D. Apply a policy map to R3 to reserve 10% of interface bandwidth.
- E. Implement a class map on R3 to match all packets with QoS IP precedence.

Answer: BD

NEW QUESTION 139

What is the primary role of a BR router in a 6rd environment?

- A. It provides connectivity between end devices and the IPv4 network.
- B. It embeds the IPv4 address in the 2002::/16 prefix.
- C. It connects the CE routers with the IPv6 network.
- D. It provides IPv4-in-IPv6 encapsulation.

Answer: C

NEW QUESTION 140

Refer to the exhibit.



An engineer is scripting ACLs to handle traffic on the given network. The engineer must block users on the network between R1 and R2 from leaving the network through R5. but these users must still be able to access all resources within the administrative domain. How must the engineer implement the ACL configuration?

- A. Configure an ACL that permits traffic to any internal address, and apply it to the R5 interfaces to R3 and R4 in the egress direction
- B. Configure a permit any ACL on the R1 interface to R2 in the egress direction, and a deny any ACL on the interface in the ingress direction
- C. Configure an ACL that permits traffic to all internal networks and denies traffic to any external address, and apply it to the R2 interface to R1 in the ingress direction.
- D. Configure an ACL that denies traffic to any internal address and denies traffic to any external address, and apply it to the R5 interfaces to R3 and R4 in the ingress direction

Answer: C

NEW QUESTION 142

Which two IS-IS parameters must match before two Level 2 peers can form an adjacency? (Choose two)

- A. authentication settings
- B. area ID
- C. system ID
- D. MTU
- E. hello timer setting

Answer: AD

NEW QUESTION 143

Refer to the exhibit.

```
R1(config)# router isis area1
R1(config-router)# net 49.0001.0000.0000.000b.00

R1(config-router)# interface loopback 0
R1(config-if)# ipv6 address 2001:0000:1001:1000::1/128
R1(config-if)# exit

R1(config)# interface Ethernet 1/2
R1(config-if)# ipv6 address 2001:0000:1001:100A::1/64
R1(config-if)# ipv6 router isis area1
R1(config-if)# exit
```

A network engineer with an employee id: 3812:12:993 has started to configure router R1 for IS-IS as shown. Which additional configuration must be applied to configure the IS-IS instance to advertise only network prefixes associated to passive interfaces?

- R1(config)# router isis area1
 - R1(config-router)# passive-interface loopback 0
 - R1(config-router)# address-family ipv6
 - R1(config-router-af)# advertise passive-only
- R1(config-router)# address-family ipv6
 - R1(config-router-af)# advertise passive-only
- R1(config)# router isis area1
 - R1(config-router)# loopback 0 passive-interface
 - R1(config-router)# address-family ipv6
 - R1(config-router-af)# prc-interval 20
- R1(config)# router isis area1
 - R1(config-router)# passive-interface loopback 0

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

Answer: A

NEW QUESTION 144

A network operator needs to implement PIM-SSM multicast configuration on customer's network so that users in different domains are able to access and stream live traffic. Which two actions must the engineer perform on the network to make the streaming work? (Choose two.)

- A. Configure at least one MSDP peer on the network
- B. Enable IGMP version 2 at the interface lever.
- C. Enable PIM sparse mode on the device.
- D. Enable IGMP version 3 at the interface level.
- E. Enable PM dense mode on the device.

Answer: AD

NEW QUESTION 148

Refer to the exhibit:

```
interface gigabitethernet1/0/1
switchport mode access
switchport access vlan 5
channel-group 1 mode desirable
```

An engineer is preparing to implement link aggregation configuration. Which statement al about this configuration is true?

- A. The switch port actively sends packets to negotiate an EtherChannel using PAgP
- B. The switch port accepts LACP and PAgP packets from a connected peer and negotiate an EtherChannel using the common EtherChannel mode.
- C. The switch port passively negotiates an EtherChannel if it receives PAgP packets from a connected peer
- D. The switch port negotiates an EtherChannel if it receives LACP packets from a connected peer

Answer: A

NEW QUESTION 152

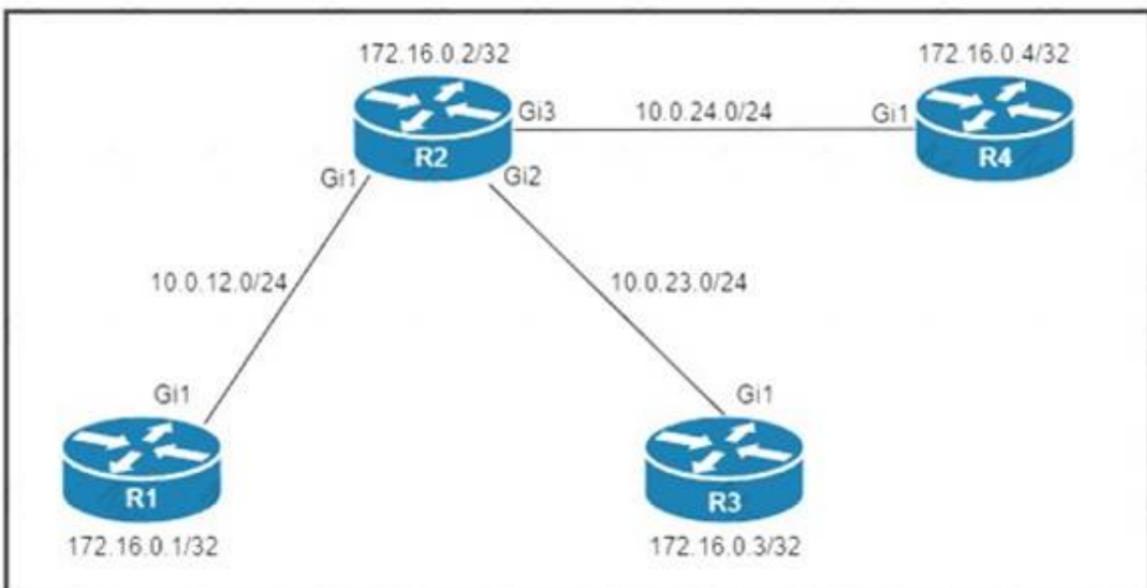
How does Cisco DNA Center enhance network automation?

- A. It allows network administrators to quickly deploy Cisco Layer 2 devices without requiring STP and broadcast transport.
- B. It allows network administrators to reduce inconsistencies when they deploy and validate network configurations.
- C. It allows network administrators to reduce the number of VRFs in a multi customer environment by automatically implementing a single VLAN per customer.
- D. It allows network administrators to combine voice and data networks into a single topology without manual configuration.

Answer: B

NEW QUESTION 153

Refer to the exhibit.



Which configuration must be applied to each of the four routers on the network to reduce LDP LIB size and advertise label bindings for the /32 loopback IP space only?

```

config t
ip prefix-list LOOPBACKS seq 5 permit 0.0.0.0/0 le 32
mpls ldp label
allocate global prefix-list LOOPBACKS
end
    
```

```

config t
access-list 10 permit 172.16.0.0 0.0.0.7
access-list 20 permit 10.0.0.0 0.0.31.255
no mpls ldp advertise-labels
mpls ldp advertise-labels for 10 to 20
end
    
```

```

config t
access-list 10 permit 172.16.0.0 0.0.0.7
access-list 20 permit 172.16.0.0 0.0.0.7
no mpls ldp advertise-labels
mpls ldp advertise-labels for 10 to 20
end
    
```

```

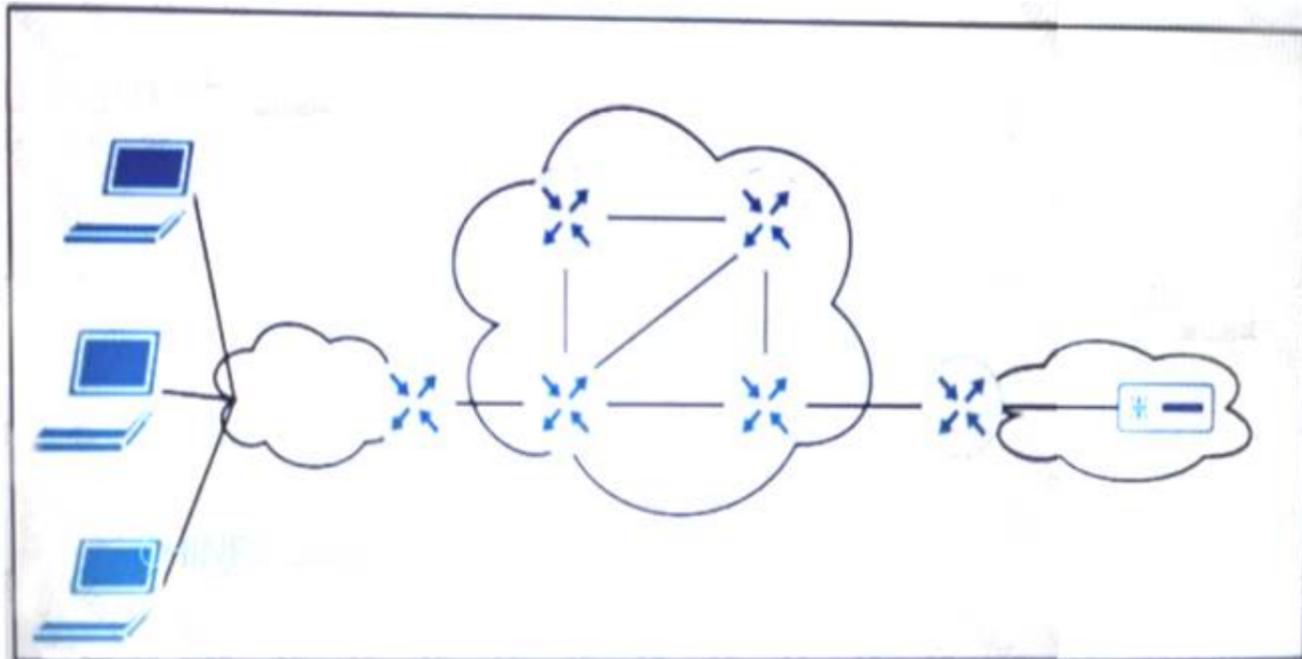
config t
mpls ldp label
allocate global host-routes
end
    
```

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

Answer: A

NEW QUESTION 157

Refer to the exhibit.



ISP A provides VPLS services and DDoS protection to Company XYZ to connect their branches across the North America and Europe regions. The uplink from the data center to the ISP is Mbps. The company XYZ security team asked the ISP to redirect ICMP requests which are currently going to the web server to a new local security appliance which configuration must an ISPP engineer apply to router R2 to redirect the ICMP traffic?

A)

```

class-map type traffic match-all B_210.10.65.1
match destination-address ipv4 210.10.65.1
match protocol 7
match ipv4 icmp-type 3
    
```

B)

```

class-map type traffic match-all B_210.10.65.1
match destination-address ipv4 210.10.65.1
match protocol 3
match ipv4 icmp-type 5
    
```

C)

```
class-map type traffic match-all B_210.10.65.1
match destination-address ipv4 210.10.65.1
match protocol 6
match ipv4 icmp-type 9
```

D)

```
class-map type traffic match-all B_210.10.65.1
match destination-address ipv4 210.10.65.1
match protocol 1
match ipv4 icmp-type 8
```

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

Answer: D

NEW QUESTION 160

Refer to the exhibit.

```
RZ#
*Dec 8 06:25:39.147: OSPF: Rcv hello from 10.10.10.2 area 0 from GigabitEthernet2/0 10.0.0.25
*Dec 8 06:25:39.151: OSPF: End of hello processing
*Dec 8 06:25:39.747: OSPF: Send hello to 224.0.0.5 area 100 on FastEthernet0/0 from 10.0.0.14
*Dec 8 06:25:40.015: OSPF: Rcv hello from 192.168.10.1 area 100 from FastEthernet0/0 10.0.0.13
*Dec 8 06:25:40.019: OSPF: Hello from 10.0.0.13 with mismatched Stub/Transit area option bit
RZ#
*Dec 8 06:25:47.287: OSPF: Send hello to 224.0.0.5 area 0 on GigabitEthernet2/0 from 10.0.0.26
*Dec 8 06:25:48.187: OSPF: Send hello to 224.0.0.5 area 0 on FastEthernet1/0 from 10.0.0.17
RZ#
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.10.10.5	1	FULL/BDR	00:00:39	10.0.0.26	Ethernet3/0

A network engineer received a complaint about these problems in OSPF stub area 100:

- > The Ethernet link is down between routers RX and RY because the fiber was cut.
- > CE site A traffic to the hub site is being dropped. Which action resolves these issues?

- A. Set the OSPF authentication type to MD5 between RX and RY DUMPS
- B. Change the OSPF area 100 type to stub on RZ.
- C. Change the OSPF priority to 100 on the interfaces that connect RX and RY.
- D. DUMPS Set the OSPF MTU to 1500 on the link between RX and RZ.

Answer: B

NEW QUESTION 162

Refer to the exhibit.

```
Router 1:
snmp-server group group1 v3 noauth
snmp-server user testuser group1 remote 192.168.0.254
snmp-server host 192.168.0.254 informs version 3 noauth testuser config
```

A network engineer is deploying SNMP configuration on client's routers. Encrypted authentication must be included on router 1 to provide security and protect message confidentially. Which action should the engineer perform on the routers to accomplish this task?

- A. snmp-server host 192.168.0.254 informs version 3 auth testuser config.
- B. snmp-server user testuser group 1 remote 192.168.0.254 v3 auth md5 testpassword
- C. snmp-server group group 1 v3 auth.
- D. snmp-server community public

Answer: B

NEW QUESTION 167

What do Chef and Puppet have in common?

- A. use Ruby
- B. use a master server
- C. require modules to be created from scratch
- D. manage agents referred to as minions

Answer: B

NEW QUESTION 172

Drag and drop the OSs from the left onto the correct deceptions on the right.

IOS XR	It is a monolithic architecture that runs all modules on one memory space.
IOS	It runs over a Linux platform and pulls the system functions out of the main kernel and into separate processes.
IOS XE	It segments ancillary processes into separate memory spaces to prevent system crashes from errant bugs.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

IOS XR	IOS
IOS	IOS XE
IOS XE	IOS XR

NEW QUESTION 176

Refer to the exhibit.

```

R1(config)# ipv6 unicast-routing
R1(config)# ipv6 router ospf 100
R1(config-rtr)# router-id 1.1.1.1
```

An engineer is configuring router R1 for OSPFv3 as shown. Which additional configuration must be performed so that the three active interfaces on the router will advertise routes and participate in OSPF IPv6 processes?

- A)

```
R1(config)# interface Ethernet1/1
R1(config-if)# ipv6 ospf 100 area 0
```

```
R1(config)# interface Ethernet1/2
R1(config-if)# ipv6 ospf 100 area 10
```

```
R1(config)# interface Ethernet1/3
R1(config-if)# ipv6 ospf 100 area 20
```

B)

```
R1(config)# interface Ethernet1/1
R1(config-if)# ip ospf hello-interval 1
R1(config-if)# ip ospf 1 area 0
```

```
R1(config)# interface Ethernet1/2
R1(config-if)# ip ospf hello-interval 1
R1(config-if)# ip ospf 1 area 10
```

```
R1(config)# interface Ethernet1/3
R1(config-if)# ip ospf hello-interval 1
R1(config-if)# ip ospf 1 area 20
```

C)

```
R1(config)# interface Ethernet1/1
R1(config-if)# ip ospf 1 area 0
```

```
R1(config)# interface Ethernet1/2
R1(config-if)# ip ospf 1 area 10
```

```
R1(config)# interface Ethernet1/3
R1(config-if)# ip ospf 1 area 20
```

A.

Answer: A

NEW QUESTION 179

Which action does the ingress VTEP perform on traffic between EVPN VXLAN overlays?

- A. routing and tunneling when doing symmetric IRB
- B. routing when doing asymmetric IRB
- C. routing and bridging when doing asymmetric IRB
- D. bridging when doing symmetric IRB

Answer: C

Explanation:

Asymmetric IRB

With asymmetric IRB, the ingress VTEP performs both Layer-2 bridging and Layer-3 routing lookup, whereas the egress VTEP performs only Layer-2 bridging lookup.

<https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/guide-c07-734107.html>

NEW QUESTION 183

Refer to the exhibit.

```

R1# show ip bgp summary
Neighbor      V  AS   MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
11.11.11.11   4  5400  0         0         0       0    0     never    Active

R1
interface Loopback0
 ip address 2.2.2.2 255.255.255.255
interface Ethernet1/0
 ip address 11.11.11.11 255.255.255.0
router bgp 5400
 neighbor 11.11.11.12 remote-as 5400
 neighbor 11.11.11.12 update-source Loopback0
 ip route 1.1.1.1 255.255.255.255 11.11.11.12

R2
interface Loopback0
 ip address 1.1.1.1 255.255.255.255
interface Ethernet1/0
 ip address 11.11.11.12 255.255.255.0
router bgp 5400
 neighbor 11.11.11.11 remote-as 5400
 neighbor 11.11.11.11 update-source Loopback0
 ip route 2.2.2.2 255.255.255.255 11.11.11.11

```

Router R1 is reporting that its BGP neighbor adjacency to router R2 is down, but its state is Active as shown. Which configuration must be applied to routers R1 and R2 to fix the problem?

A)

```

R1
router bgp 5400
neighbor 2.2.2.2 remote-as 5400

```

```

R2
router bgp 5400
neighbor 1.1.1.1 remote-as 5400

```

B)

```

R1
router bgp 5400
 neighbor 11.11.11.11 remote-as 5400
 neighbor 11.11.11.11 update-source Loopback0

```

```

R2
router bgp 5400
 neighbor 11.11.11.12 remote-as 5400
 neighbor 11.11.11.12 update-source Loopback0

```

C)

```

R1
router bgp 5400
 neighbor 1.1.1.1 remote-as 5400
 neighbor 1.1.1.1 update-source Loopback0

```

```

R2
router bgp 5400
 neighbor 2.2.2.2 remote-as 5400
 neighbor 2.2.2.2 update-source Loopback0

```

D)

```

R1
router bgp 5400
 neighbor 2.2.2.2 remote-as 5400
 neighbor 2.2.2.2 update-source Loopback0

```

```

R2
router bgp 5400
 neighbor 1.1.1.1 remote-as 5400
 neighbor 1.1.1.1 update-source Loopback0

```

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

Answer: C

NEW QUESTION 187

Refer to the exhibit:

```

Router 1:

ip route 192.0.2.0 255.255.255.0 null 0
ip route 192.168.1.0 255.255.255.0 null 0 tag 1

route-map ddos
 match tag 1
 set ip next-hop 192.0.2.1
 set local-preference 150
 set community no export

route-map ddos permit 20

router bgp 65513
 redistribute static route-map ddos

Router 2:

ip route 192.0.2.0 255.255.255.0 null 0
    
```

An engineer is preparing to implement data plane security configuration. Which statement about this configuration is true?

- A. Router 1 drops all traffic with a local-preference set to 150
- B. All traffic is dropped
- C. All traffic to 192.168.1.0/24 is dropped
- D. Router 1 and Router 2 advertise the route to 192.0.2.0/24 to all BGPFD peers.

Answer: C

NEW QUESTION 190

Refer to the exhibit:

<https://192.168.1.100/api/mo/uni/tn-ciscotest.xml>

What is the URL used for with REST API?

- A. It is used to contact a URL filter to determine the efficacy of a web address
- B. It is used to send a TACACS+ authentication request to a server
- C. It is used to send a message to the APIC to perform an operation on a managed object or class operator
- D. It is used to initiate an FTP session to save a running configuration of a device.

Answer: C

NEW QUESTION 191

Refer to the exhibits:

“*Apr 30 14:33:43.619: %CLNS-4-AUTH_FAIL: ISIS: LAN IIH authentication failed”.

```

R1#show isis neighbors

Tag TEST:
System Id  Type Interface  IP Address  State Holdtime Circuit Id
R2         L2    Fa0/0      UP    9         R2.01

R2#show isis neighbors

Tag TEST:
System Id  Type Interface  IP Address  State Holdtime Circuit Id
R2         L1    Fa0/0      INIT  22         R2.01
R2         L2    Fa0/0      UP    24         R2.01
    
```

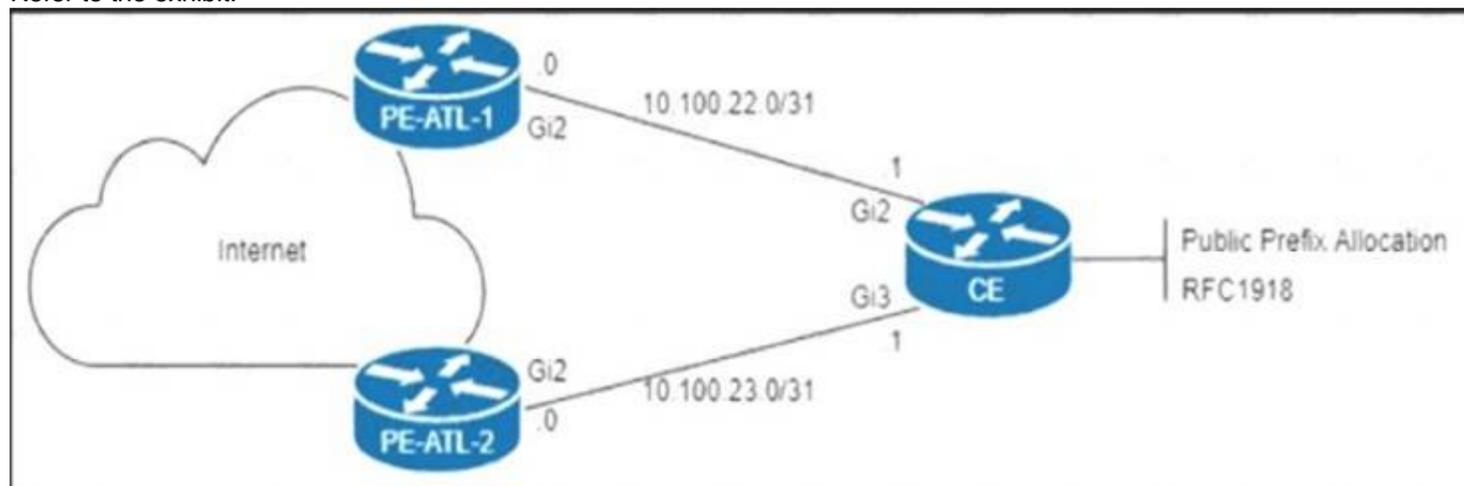
R1 and R2 are directly connected and IS-IS routing has been enabled between R1 and R2 R1 message periodically Based on this output, which statement is true?

- A. IS-IS neighbor authentication is failing for Level 2 first and then for Level 1 PDUs
- B. IS-IS neighbor authentication is failing for Level 1 and Level 2 PDUs .
- C. IS-IS neighbor authentication is failing for Level 1 PDUs only
- D. IS-IS neighbor authentication is failing for Level 2 PDUs only.

Answer: C

NEW QUESTION 193

Refer to the exhibit.



The CE router is peering with both PE routers and advertising a public prefix to the internet. Routing to and from this prefix will be asymmetric under certain network conditions, but packets must not be discarded. Which configuration must an engineer apply to the two PE routers so that they validate reverse packet forwarding for packets entering their Gi2 interfaces and drop traffic from the RFC1918 space?

- A. ip verify unicast source reachable-via rx allow-default
- B. interface GigabitEthernet 2 ip verify unicast source reachable-via rx
- C. ip verify unicast source reachable-via any allow-default interface GigabitEthernet 2
- D. ip verify unicast source reachable-via any

Answer: D

NEW QUESTION 196
 SIMULATION 4

Guidelines **Topology** Tasks

EBGP Neighbor Adjacency

ASN: 65515

lo0:
1.1.1.1/32
2000:cc13:cc13:1::1/128

R1

ASN: 65516

lo0:
2.2.2.2/32
2000:cc13:cc13:2::1/128

R2

EO/O .1 --- 192.168.1.0/24 --- EO/O .2
 2000:cc13:cc13:cc13::/48

R1 R2

R1>

Guidelines **Topology** Tasks

Configure the BGP routing protocol for R1 and R2 according to the topology to achieve these goals:

1. Configure EBGP neighbor adjacency for the IPv4 and IPv6 address family between R1 and R2 using Loopback0 IPv4 and IPv6 addresses. All BGP updates must come from the Loopback0 interface as the source. Do not use IGP routing protocols to complete this task.
2. Configure MD5 Authentication for the EBGP adjacency between R1 and R2. The password is clear text **C1sc0!**.

R1 R2

R1>

[Submit feedback about this item](#)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```

R1
Router bgp 65515
No bgp default ipv4-unicast Neig 2.2.2.2 remote-as 65516
Nei 2.2.2.2 update-soc loopback0 Nei 2.2.2.2 ebgp-multihop 2 Neig 2.2.2.2 pass C1sc0!
Nei 2000:cc13:cc13:2::1 remote-as 65516 Nei 2000:cc13:cc13:2::1 update-so loopback0 Nei 2000:cc13:cc13:2::1 pass C1sc0!
Nei 2000:cc13:cc13:2::1 ebgp-multihop 2 Address-family ipv4
Neig 2.2.2.2 activate Address-family ipv6
Nei 2000:cc13:cc13:2::1 activate
Ip route 2.2.2.2 255.255.255.255 192.168.1.2
Ipv6 route 2000:cc13:cc13:2::1/128 2000:cc13:cc13:cc13:2 R2
Router bgp 65516
No bgp default ipv4-unicast Neig 1.1.1.1 remote-as 65515
Nei 1.1.1.1 update-soc loopback0 Nei 1.1.1.1 pass C1sc0!
Nei 1.1.1.1 ebgp-multihop 2
Nei 2000:cc13:cc13:1::1 remote-as 65515 Nei 2000:cc13:cc13:1::1 update-so loopback0 Nei 2000:cc13:cc13:1::1 pass C1sc0!
Nei 2000:cc13:cc13:1::1 ebgp-multihop 2 Address-family ipv4
Neig 1.1.1.1 activate Address-family ipv6
Nei 2000:cc13:cc13:1::1 activate
Ip route 1.1.1.1 255.255.255.255 192.168.1.1
Ipv6 route 2000:cc13:cc13:1::1/128 2000:cc13:cc13:cc13:1
    
```

NEW QUESTION 197

Refer to the exhibit.

```
restconf_headers["Content-Type"] = "application/ yang-data+json"

loopback = {"name": "Loopback101",
            "description": "Router-1",
            "ip": "192.168.11.11",
            "netmask": "255.255.255.0"}
data = {
    "ietf-interfaces:interface": {
        "name": loopback["name"],
        "description": loopback["description"],
        "type": "iana-if-type:softwareLoopback",
        "enabled": True,
        "ietf-ip:ipv4": {
            "address": {
                ("ip": loopback["ip"],
                "netmask": loopback["netmask"])
            }
        }
    }
}

url = interface_url.format(ip= corel_ip, int_name= loopback["name"])
r = requests.put(url,
                 headers = restconf_headers,
                 auth=(username, password),
                 json= data,
                 verify=False)
print("Request Status Code: {}".format(r.status_code))
```

An engineer at a new ISP must configure many Cisco devices in the data center. To make the process more efficient, the engineer decides to automate the task with a REST API. Which action does this JSON script automate?

- A. Configure the IP address for the existing loopback interface.
- B. Configure a physical interface on the router with an IP address and then create a loopback interface.
- C. Configure a physical interface on the router with an IP address.
- D. Delete the existing loopback Interface and replace it with a new loopback interface.

Answer: A

NEW QUESTION 198

When Cisco IOS XE REST API uses HTTP request methods what is the purpose of a PUT request?

- A. retrieves the specified resource or representation
- B. submits data to be processed to the specified resource
- C. updates the specified resource with new information
- D. creates a new resource

Answer: C

Explanation:

PUT	<p>Updates the specified resource with new information. The data that is included in the PUT operation replaces the previous data.</p> <ul style="list-style-type: none"> • The PUT operation is used to replace or modify an existing resource. The PUT operation cannot be used to create a new resource. • The request body of a PUT operation must contain the complete representation of the mandatory attributes of the resource.
-----	--

NEW QUESTION 202

Refer to the exhibit.

```
RP/0/RP0/CPU0:XR1#do sh bundle

Bundle-Ether11
  Status: Up
  Local links <active/standby/configured>: 1 / 2 / 3
  Local bandwidth <effective/available>: 1000000 (1000000) kbps
  MAC address (source): 0007.ec14.cc2b (Chassis pool)
  Inter-chassis link: No
  Minimum active links / bandwidth: 1 / 1 kbps
  Maximum active links: 1
  Wait while timer: 2000 ms
  Load balancing:
    Link order signaling: Not configured
    Hash type: Default
    Locality threshold: None
  LACP: Operational
    Flap suppression timer: Off
    Cisco extensions: Disabled
    Non-revertive: Disabled
  mLACP: Not configured
  IPv4 BFD: Not configured
  IPv6 BFD: Not configured

Port          Device      State      Port ID          B/W, kbps
-----
Gi0/0/0/0    Local      Standby    0x8000, 0x0003   1000000
  Link is Standby due to maximum-active links configuration
Gi0/0/0/1    Local      Standby    0x8000, 0x0002   1000000
  Link is Standby due to maximum-active links configuration
Gi0/0/0/2    Local      Active     0x8000, 0x0001   1000000
  Link is Active
```

A network operator needs to shut down interface Gi0/0/0/2 for maintenance. What occurs to the interface states of Gi0/0/0/0 and Gi0/0/0/1?

- A. Gi0/0/0/1 and Gi0/0/0/0 become active
- B. Gi0/0/0/1 and Gi0/0/0/0 remains standby
- C. Gi0/0/0/0 becomes active
- D. Gi0/0/0/1 remains standby
- E. Gi0/0/0/1 becomes active Gi0/0/0/0 remains standby

Answer: D

NEW QUESTION 203

Refer to the exhibit.

```
Control Plane Interface
Service policy CoPP-normal
Hardware Counters:
class-map: CoPP-normal (match-all)
Match: access-group 100
police :
6000 bps 1000 limit 1000 extended limit
Earl in slot 3 :
0 bytes
5 minute offered rate 0 bps
aggregate-forwarded 0 bytes action: transmit
exceeded 0 bytes action: drop
aggregate-forward 0 bps exceed 0 bps
Earl in slot 5 :
0 bytes
5 minute offered rate 0 bps
aggregate-forwarded 0 bytes action: transmit
exceeded 0 bytes action: drop
aggregate-forward 0 bps exceed 0 bps
```

Which show command shows statistics for the control plane policy and is used to troubleshoot?

- A. show control-plane CoPP
- B. show control-plane
- C. show policy-map control-plane
- D. show policy control-plane

Answer: C

Explanation:

```
Router# show policy-map control-plane
```

```
Control Plane
```

```
Service-policy input:TEST
```

```
Class-map:TEST (match-all)
```

```
 20 packets, 11280 bytes
```

```
 5 minute offered rate 0 bps, drop rate 0 bps
```

```
Match:access-group 101
```

```
police:
```

```
 8000 bps, 1500 limit, 1500 extended limit
```

```
 conformed 15 packets, 6210 bytes; action:transmit
```

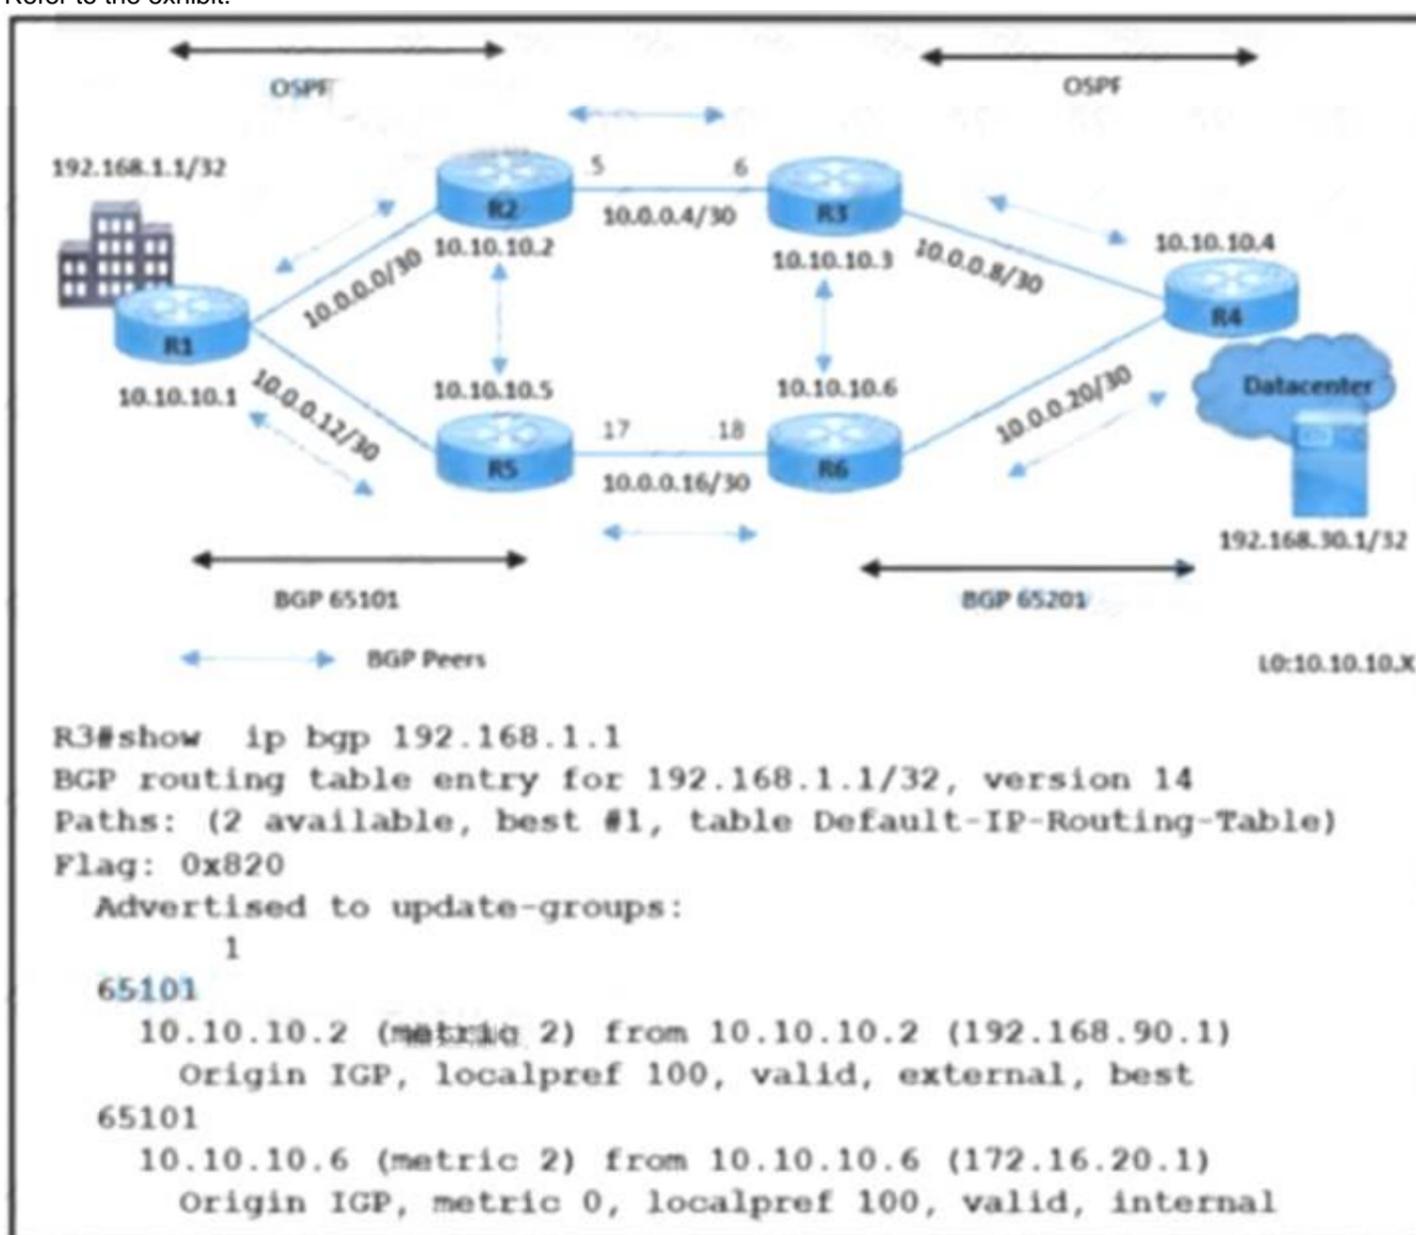
```
 exceeded 5 packets, 5070 bytes; action:drop
```

```
 violated 0 packets, 0 bytes; action:drop
```

```
 conformed 0 bps, exceed 0 bps, violate 0 bps
```

NEW QUESTION 208

Refer to the exhibit.



A network engineer is implementing BGP in AS 65101 and AS 65201. R3 sends data traffic to 192.168.1.1 /32 via the path R3-R2-R1. The traffic must travel via alternate path R6-R5 for prefix 192.168.1.1/32. Which action must be taken to meet the requirement?

- A. Apply route-map HIGH-MED out on R2 for neighbor R3.
- B. Apply route-map HIGH-LP in on R3 for neighbor R6
- C. Apply route-map LOW-LP out on R2 for neighbor R3.
- D. Apply route-map LOW-MED in on R5 for neighbor R2

Answer: A

NEW QUESTION 212

A router is advertising multiple networks to its BGP neighbor in AS 5200 with peer IP address 1.1.1.1. Which configuration must be applied so that the router permits updates only for networks with a prefix mask length less than or equal to 21?

- router bgp 5100
neighbor 1.1.1.1 remote-as 5200
neighbor 1.1.1.1 prefix-list SELECTED in

ip prefix-list SELECTED seq 10 permit 0.0.0.0/0 le 21
- router bgp 5100
neighbor 1.1.1.1 remote-as 5200
neighbor 1.1.1.1 prefix-list SELECTED

ip prefix-list SELECTED seq 10 permit 0.0.0.0/0 ge 8 le 24
- router bgp 5100
neighbor 1.1.1.1 remote-as 5200
neighbor 1.1.1.1 prefix-list SELECTED

ip prefix-list SELECTED seq 10 permit 0.0.0.0/0 ge 21
- router bgp 5100
neighbor 1.1.1.1 remote-as 5200
neighbor 1.1.1.1 prefix-list SELECTED in

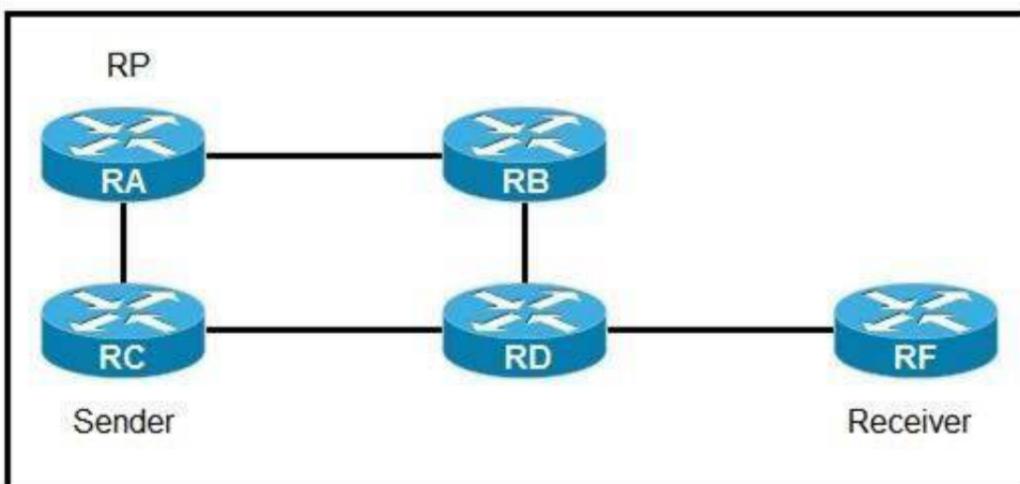
ip prefix-list SELECTED seq 10 permit 0.0.0.0/0 ge 21 le 24

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

Answer: A

NEW QUESTION 217

Refer to the exhibit:



If router A is the RP, which PIM mode can you configure so that devices will send multicast traffic toward the RP?

- A. PIM-SM
- B. PIM-DM
- C. BIDIR-PIM
- D. PIM-SSM

Answer: C

NEW QUESTION 219

Refer to the exhibit:

```

R1
interface fastethernet1/0
 ip address 192.168.1.3 255.255.255.0
router bgp 65000
 router-id 192.168.1.1
 neighbor 192.168.1.2 remote-as 65012

R2
interface fastethernet1/0
 ip address 192.168.1.2 255.255.255.0
router bgp 65012
 router-id 192.168.1.1
 neighbor 192.168.1.3 remote-as 65000
 neighbor 192.168.1.3 local-as 65112

```

Assume all other configurations are correct and the network is otherwise operating normally. Which conclusion can you draw about the neighbor relationship between routers R1 and R2?

- A. The neighbor relationship will be up only if the two devices have activated the correct neighbor relationships under the IPv4 address family
- B. The neighbor relationship is down because R1 believes R2 is in AS 65012.
- C. The neighbor relationship is up
- D. The neighbor relationship is down because the local-as value for R2 is missing in the R1 neighbor statement

Answer: B

NEW QUESTION 221

A service provider requires continuous real-time network monitoring to provide reliable SLAs to its customers. To satisfy this requirement, a network administrator is implementing gRPC dial out on an ASR with TLS. Receiver 192.168.10.2 will be assigned one of the subscriptions, and it will manage the ASR. Which configuration must the engineer apply to the router as part of the configuration process?

- A. snmp-server community public snmp-server enable trapssnmp-server host 192.168.10.2 version 2c public.
- B. telemetry model-driven destination-group DGroup1address family ipv4 192.168.10.2 1 port 10 encoding self-describing-gpb
- C. snmp-server community public snmp-server enable trapssnmp-server enable traps snmp authentication snmp-server managersnmp-server manager session-timeout 1000
- D. telemetry model-driven destination-group ciscotestaddress family ipv4 192.168.10.2 port 10 encoding self-describing-gpbprotocol grpc tis-hostname ciscotest.com

Answer: D

NEW QUESTION 226

An engineer must extend Layer 2 Between two campus sites connected through an MPLS backbone that encapsulates Layer 2 and Layer 3 data Which action must the engineer perform on the routers to accomplish this task?

- A. Configure a EtherChannel for E-LAN.
- B. Configure a pseudowire for E-LINE.
- C. Configure Cisco MPLS TE for use with E-TREE.
- D. Configure QoS for MPLS and E-ACCESS

Answer: B

NEW QUESTION 229

Refer to the exhibit.

```
EDGE-GW-1#show bgp ipv4 unicast summary
BGP router identifier 198.19.45.6, local AS number 65502
BGP table version is 19, main routing table version 19

Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.26.2  4      65503    0      0        1    0    0 00:0956  Idle

EDGE-GW-1#show log
Log Buffer (4096 bytes):
BGP Notification sent
Dec 7 08:02:29.619: %BGP-5-ADJCHANGE: neighbor 192.168.26.2 passive Down BGP Notification sent
Dec 7 08:02:32.695: %BGP-3-NOTIFICATION: sent to neighbor 192.168.26.2 active 2/2 (peer in wrong AS) 2 bytes FE63
Dec 7 08:02:32.695: %BGP-4-MSGDUMP: unsupported or mal-formatted message received from 192.168.26.2:
FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF 0039 0104 FE63 00B4 0AFF FF02 1C02 0601
0400 0100 0102 0280 0002 0202 0002 0246 0002 0641 0400 00FE 63
Dec 7 08:02:36.558: %BGP-3-NOTIFICATION: sent to neighbor 192.168.26.2 passive 2/2 (peer in wrong AS) 2 bytes FE63
Dec 7 08:02:36.558: %BGP-4-MSGDUMP: unsupported or mal-formatted message received from 192.168.26.2:
FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF 0039 0104 FE63 00B4 0AFF FF02 1C02 0601
0400 0100 0102 0280 0002 0202 0002 0246 0002 0641 0400 00FE 63
Dec 7 08:02:37.812: %BGP-5-NBR_RESET: Neighbor 192.168.26.2 active reset (BGP Notification sent)
Dec 7 08:02:37.812: %BGP-5-ADJCHANG: neighbor 192.168.26.2 active Down BGP Notification sent
Dec 7 08:02:37.812: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.26.2 IPv4 Unicast topology base removed from session
BGP Notification sent
Dec 7 08:02:40.883: %BGP-5-NBR_RESET: Neighbor 192.168.26.2 passive reset (BGP Notification sent)
Dec 7 08:02:40.884: %BGP-5-ADJCHANGE: neighbor 192.168.26.2 passive Down BGP Notification sent
Dec 7 08:02:47.822: %BGP-3-NOTIFICATION: sent to neighbor 192.168.26.2 passive 2/2 (peer in wrong AS) 2 bytes FE63
Dec 7 08:02:77.822: %BGP-4-MSGDUMP: unsupported or mal-formatted message received from 192.168.26.2:
FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF 0039 0104 FE63 00B4 0AFF FF02 1C02 0601
0400 0100 0102 0280 0002 0202 0002 0246 0002 0641 0400 00FE 63
```

A network support engineer for ASN 65502 receives a technical support ticket from a customer in ASN 65503 who reports that an eBGP session is down. The engineer determines that the peering failed after a recent change to the device at 192.168.26.2. EDGE-GW-1 must establish an eBGP session with the peering router 192.168.26.2. Which configuration establishes this session?

- A. configure terminal no router bgp 65502 router bgp 65503 neighbor 192.168.26.2 remote-as 65503 address-family ipv4 neighbor 192.168.26.2 activate end
- B. configure terminal router bgp 65502 address-family ipv4 neighbor 192.168.26.2 activate end
- C. configure terminal no router bgp 65502 router bgp 65503 neighbor 192.168.26.2 remote-as 65123 address-family ipv4 neighbor 192.168.26.2 activate end
- D. configure terminal router bgp 65502 no neighbor 192.168.26.2 remote-as 65503 neighbor 192.168.26.2 remote-as 65123 address-family ipv4 neighbor 192.168.26.2 activate end

Answer: B

NEW QUESTION 231

Which control plane protocol is used between Cisco SD-WAN routers and vSmart controllers?

- A. OTCP
- B. OMP
- C. UDP
- D. BGP

Answer: B

NEW QUESTION 234

Refer to the exhibit.

```
interface GigabitEthernet 1/0/1
 ip address 192.168.1.1 255.255.255.0
 ip router isis
 isis tag 15
 route-map match-tag permit 10
 match tag 15
```

A large organization is merging the network assets of a recently acquired competitor with one of its own satellite offices in the same geographic area. The newly acquired network is running different routing protocol than the company's primary network. As part of the merger a network engineer implemented this route map. Which task must the engineer perform to complete the implementation?

- A. Attach the route map to an IS-IS network statement to advertise the routes learned on this interface to IS-IS
- B. Enable metric style wide to allow the use of extended metrics from the protocols
- C. Attach the route map to the redistribution command to manipulate the routes as they are shared

D. Configure an additional route map sequence to override the implicit deny at the end of the route map

Answer: C

NEW QUESTION 237

Refer to the exhibit.

```
Router 1:
router isis
 net 49.0011.0000.0000.0001.00

Router 2:
router isis
 net 49.0001.0000.0000.0001.00

Router 3:
router isis
 net 49.0011.0000.0000.0002.00
```

Router 4 is added to the network and must be in the same area as router 1. Which NET should the engineer assign?

- A. 49.0001.0000.0000.0004.00
- B. 49.0111.0000.0000.0001.00
- C. 49.0011.0000.0000.0003.00
- D. 49.0011.0000.0000.0002.00

Answer: C

NEW QUESTION 239

In an MPLS network, which protocol can be used to distribute a Segment Prefix?

- A. OSPF
- B. LDP
- C. RSVP-TE
- D. EIGRP

Answer: A

NEW QUESTION 244

Why is the keyword none needed when implementing management plane security using TACACS?

- A. It allows the local database to query a RADIUS server when the TACACS+ server is unreachable.
- B. It allows the local database to authenticate when the TACACS+ server is unreachable.
- C. It allows authentication to succeed when the TACACS+ server is unreachable.
- D. It prevents all users from accessing router 1 unless the TACACS+ server is reachable,

Answer: C

NEW QUESTION 249

Refer to the exhibit.

```
!
telemetry model-driven
 destination-group DGroup2
 address family ipv4
 172.10.10.10 port 57500
 encoding self-describing-gpb
 protocol grpc
 commit
!
```

A network engineer at a large ISP is configuring telemetry streams to monitor the health status of PE routers on the network using gRPC dial-out. The PE routers are located at several data centers in different physical locations, and they are using IS-IS and BGP for routing. Which additional configuration must the engineer implement on the PE routers to meet the goal?

- A. Text, letter Description automatically generated

```

sensor-group SGroup2
  sensor-path openconfig-interfaces:interfaces/interface
!
subscription Sub3
  sensor-group-id SGroup3 sample-interval 30000
  
```

B. Text Description automatically generated

```

sensor-group SGroup2
  sensor-path Cisco-IOS-XR-plat-chas-invmgr-oper:platform-inventory/racks/rack
!
subscription Sub1
  sensor-group-id SGroup1 sample-interval 30000
  destination-id DGroup1
  
```

C. Graphical user interface, text Description automatically generated

```

sensor-group SGroup2
  sensor-path Cisco-IOS-XR-infra-statsd-oper:infra-statistics/interfaces/interface/latest/generic-cou
!
subscription Sub1
  sensor-group-id SGroup1 sample-interval 30000
  destination-id DGroup1
  
```

D. Text, letter Description automatically generated

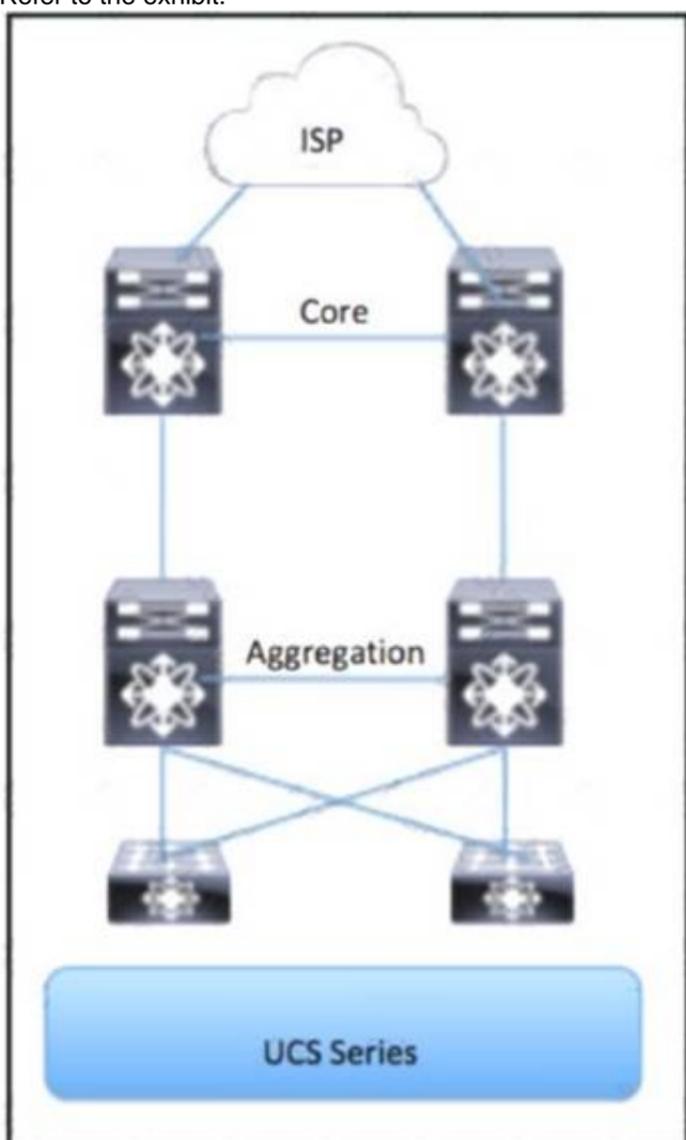
```

sensor-group SGroup2
  sensor-path Cisco-IOS-XR-nto-misc-oper:memory-summary/nodes/node/summ
!
subscription Sub2
  sensor-group-id SGroup2 sample-interval 30000
  destination-id DGroup2
  
```

Answer: D

NEW QUESTION 253

Refer to the exhibit.



Which part of the diagram will host OpenStack components?

- A. Aggregation
- B. UCS Series
- C. Access
- D. Core

Answer: C

NEW QUESTION 257

Refer to the exhibit.

```
router(config)# router ospf 11
router(config-if)# passive-interface default
```

An engineer started to configure a router for OSPF. Which configuration must the engineer perform on the router without changing any interface configuration so that the router establishes an OSPF neighbor relationship with its peer?

- A. router(config)# router ospf 11router(config-if)# no passive-interface ethernet 1/1
- B. router(config)# interface ethernet 1/1router(config-if)# no shutdown
- C. router(config)# interface ethernet 1/1router(config-if)# ip ospf hello-interval
- D. router(config)# interface ethernet 1/1router(config-if)# ip ospf priority 0

Answer: A

NEW QUESTION 262

Refer to the exhibit.

```
ASBR-1#show bgp ipv4 unic | begin Network
Network      Next Hop      Metric LocPrf Weight Path
* > i 198.18.15.0 172.31.255.1 0      100    0 65001 ?
* i          172.31.255.2 0      100    0 65001 ?

EDGE-1#show bgp ipv4 un | begin Netowrk
Network      Next Hop      Metric LocPrf Weight Path
* > 198.18.15.0/25 100.65.0.2 0      0      0 65001 ?
* > 198.18.15.0 100.65.0.2 0      0      0 65001 ?
* i          172.31.255.2 0      100    0 65001 ?

EDGE-1#show bgp ipv4 un 198.18.15.0
BGP routing table entry for 198.18.15.0/25, version 9
Paths: (1 available, best #1, table default, not advertised to any peer)
Not advertised to any peer
Refresh Epoch 1
65001
 100.65.0.2 from 100.65.0.2 (198.18.100.1)
  Origin incomplete, metric 0, localpref 100, valid, external, best
  Community: 64611:65001 no-advertise

RP/0/0/CPU0:INT-R1#show bgp ipv4 unicast | begin Network
Network      Next Hop      Metric LocPrf Weight Path
* > 198.18.15.0/24 0.0.0.0 0      32768 ?
* > 198.18.15.0/25 0.0.0.0 0      32768 ?
```

The network engineer who manages ASN 65001 is troubleshooting suboptimal routing to the 198.18.15.0/24 prefix. According to the network requirements: Routing to IP destinations in the 198.18.15.0/25 block must be preferred via the EDGE-1 PE. Routing to IP destinations in the 198.18.15.128/25 block must be preferred via the EDGE-2 PE. More specific prefixes of the 198.18.15.0/24 block must not be advertised beyond the boundaries of ASN 64611. Routing to 198.18.15.0/24 must be redundant in case one of the uplinks on INT-R1 fails. Which configuration must the network engineer implement on INT-R1 to correct the suboptimal routing and fix the issue?

- A. configure terminalroute-policy ASN65001-SPECIFIC-OUT if destination in (198.18.15.0/25) then set community (no-export, peer-as:65001) done end if destination in (198.18.15.0/24) then prepend as-path 65001 3 done end if drop end policy!router bgp 65001 neighbor 100.65.0.1 address-family ipv4 unicast route-policy ASN65001-SPECIFIC-OUT out end
- B. configure terminalroute-policy ASN65001-SPECIFIC-OUT if destination in (198.18.15.0/25) then set community (internal, peer-as:65001) done end if destination in (198.18.15.0/24) then done end if drop end policy!router bgp 65001 neighbor 100.65.0.1 address-family ipv4 unicast route-policy ASN65001-SPECIFIC-OUT out end
- C. configure terminalroute-policy ASN65001-SPECIFIC-OUT if destination in (198.18.15.0/25) then set community (no-advertise, peer-as:65001) done end if destination in (198.18.15.128/25) then prepend as-path 65001 3 done end if drop end policy!router bgp 65001 neighbor 100.65.0.1 address-family ipv4 unicast route-policy ASN65001-SPECIFIC-OUT out end
- D. configure terminalroute-policy ASN65001-SPECIFIC-OUT if destination in (198.18.15.0/25) then set community (no-export, peer-as:65001) done end if destination in (198.18.15.128/25) then prepend as-path 65001 3 done end if drop end policy!router bgp 65001 neighbor 100.65.0.1 address-family ipv4 unicast route-policy ASN65001-SPECIFIC-OUT in end

Answer: B

NEW QUESTION 267

Drag and drop the characteristics from the left onto the automation tool on the right.

Answer Area

It is the standard transport protocol for communicating with network devices.	NETCONF <div style="background-color: #FFFF00; height: 20px; width: 100%;"></div> <div style="background-color: #FFFF00; height: 20px; width: 100%;"></div> <div style="background-color: #FFFF00; height: 20px; width: 100%;"></div>
It is a standard data modeling language.	
It retrieves operational data.	
It develops data models.	
It shapes state data.	
It sets and reads configuration data.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

It is the standard transport protocol for communicating with network devices.	NETCONF <div style="background-color: #ADD8E6; height: 20px; width: 100%;"></div> <div style="background-color: #ADD8E6; height: 20px; width: 100%;"></div> <div style="background-color: #ADD8E6; height: 20px; width: 100%;"></div>
It is a standard data modeling language.	
It retrieves operational data.	
It develops data models.	
It shapes state data.	
It sets and reads configuration data.	

NEW QUESTION 269

Refer to the exhibit.

```
interface gigabitethernet 0/2
no ip directed-broadcast
```

Which type of DDoS attack will be mitigated by this configuration?

- A. SYN flood
- B. smurf attack
- C. SIP INVITE flood attacks
- D. teardrop attack

Answer: B

NEW QUESTION 272

How does an untrusted interface at the boundary of an administrative domain handle incoming packets?

- A. It remarks all values to a CoS of 0.
- B. It forwards only traffic with a DSCP value of 48.

- C. It translates the IP precedence value to the corresponding DSCP value.
- D. It drops all traffic ingressing the network.

Answer: A

NEW QUESTION 276

An engineering team must implement Unified MPLS to scale an MPLS network. Devices in the core layer use different IGPs, so the team decided to split the network into different areas. The team plans to keep the MPLS services as they are and introduce greater scalability. Which additional action must the engineers take to implement the Unified MPLS?

- A. Redistribute the IGP prefixes from one IGP into the other routers to ensure end-to-end LSPs.
- B. Configure the ABR routers as route reflectors that redistribute IGP into BGP.
- C. Redistribute the IGP prefixes into another IGP to ensure end-to-end LSPs.
- D. Move the IGP prefixes into IS-IS as the loopback prefixes of the PE routers to distribute the prefixes to other routers to create end-to-end LSPs.

Answer: B

NEW QUESTION 280

An engineer is implementing a router redistribution within BGP. The route map must be configured to permit all unmatched routes. Which action must the engineer perform to complete this task?

- Include a **permit** statement as the first entry
- Include at least one explicit **deny** statement
- Remove the implicit **deny** entry
- Include a **permit** statement as the last entry

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

Answer: D

NEW QUESTION 284

An engineer implemented LDP protocol on the ISP network. The engineer must ensure that there are no packet loss issues when IGP and LDP protocols are not synchronized. Which configuring must the engineer implement so that the IGP routing protocol will wait until LDP convergence is completed?

- A. Disable IP CEF routers running LDP and enable LDP protocol.
- B. Configure MPLS LDP IGP synchronization on the network.
- C. Configure LDP sessions protection on the network.
- D. Disable MPLS LDP IGP synchronization on the network.

Answer: B

NEW QUESTION 285

Which statement about Network Services Orchestrator (NSO) is true?

- A. It is used only in service provider environments
- B. It can be used only with XML coding
- C. It uses YANG modeling language to automate devices
- D. It must use SDN as an overlay for addressing

Answer: C

NEW QUESTION 289

How can a network administrator secure rest APIs?

- A. They can allow read and write privileges to all users
- B. They can ensure that user sessions are authenticated using TACACS+ only
- C. They can have a general administrator login for multiple users to access that has command entries logged
- D. They can authenticate user sessions and provide the appropriate privilege level

Answer: D

NEW QUESTION 294

Which module refers to the network automation using Ansible?

- A. the iosxr_system module to collect facts from remote devices
- B. the iosxr_user module to manage banners for users in the local database
- C. the iosxr_logging module to run debugging for severity levels 2 to 5
- D. the iosxr_command module to issue run commands on remote devices

Answer: D

Explanation:

https://docs.ansible.com/ansible/latest/collections/cisco/iosxr/iosxr_command_module.html#ansible-collections-

NEW QUESTION 296

Refer to the exhibit:

```

mpls label protocol ldp
mpls ldp router-id loopback 0
mpls ip
ip cef
  
```

A network operator working for service provider with an employee id 3715 15:021 applied this configuration to a router. Which additional step should the engineer use to enable LDP?

- A. Disable Cisco Express Forwarding globally
- B. Delete the static router ID
- C. Enable MPLS LDP on the interface
- D. Configure the both keyword to enable LDP globally

Answer: C

NEW QUESTION 300

Refer to the exhibit.

```

mpls label range 16 100000 static 100002 1048570
mpls label protocol ldp

mpls ldp graceful-restart
!
interface Loopback0
!
ip address 10.20.20.20 255.255.255.255
no ip directed-broadcast
no ip mroute-cache
!
interface Gi1/1/0
ip address 10.12.0.2 255.255.0.0
no ip directed-broadcast
mpls label protocol ldp
mpls ip
!
router ospf 100
log-adjacency-changes
nsf cisco enforce global
redistribute connected subnets
network 10.20.20.20 0.0.0.0 area 0
network 10.12.0.0 0.0.255.255 area 0
!
mpls ldp router-id Loopback0 force
  
```

A network administrator implemented MPLS LDP changes on PE-A LSR device. The engineer must ensure there are no LDP peer are fully operational. Which LDP feature must the engineer apply to the existing configuration to eliminate the problem?

- A. Configure MPLS LDP IGP synchronization on the network.
- B. Configure MPLS LDP NSR for all LDP sessions.
- C. Enable LDP session protection under the routing protocol.
- D. Disable IP CEF on routers running LDP and enable LDP.

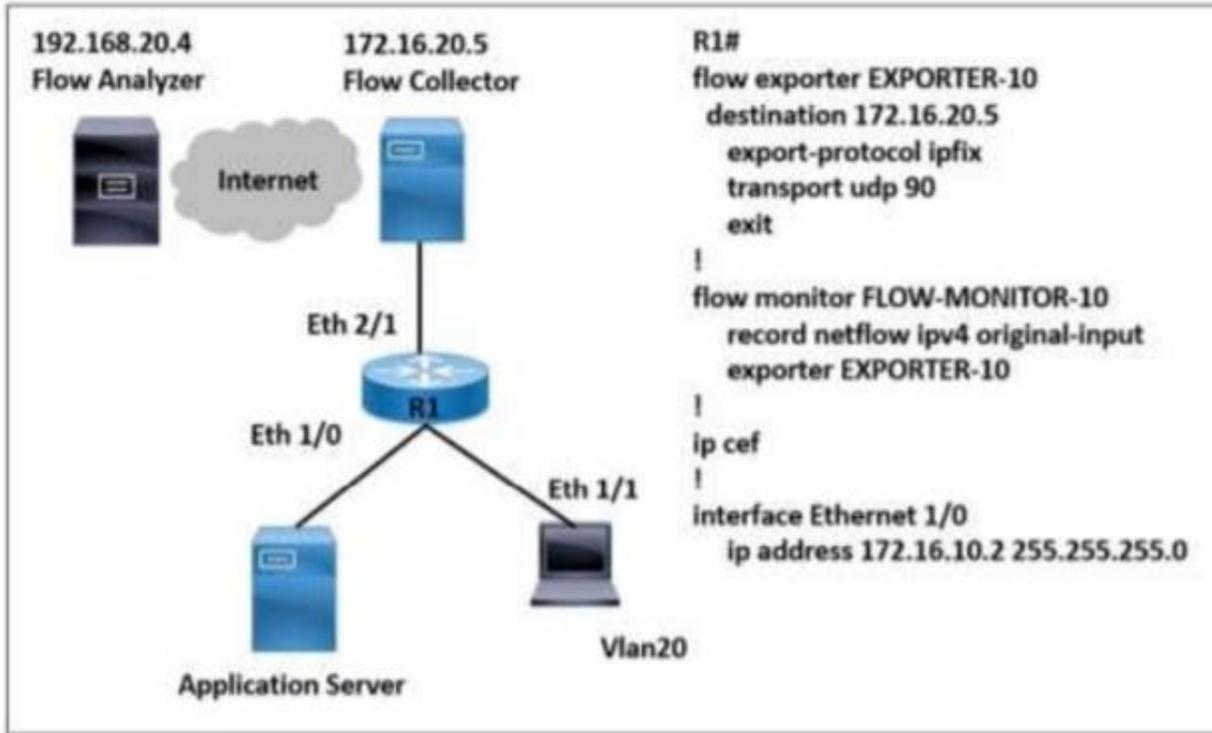
Answer: B

Explanation:

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/misp/configuration/xr-3s/mp-ha-xr-3s-book/mp-nsr-ldp-supp>

NEW QUESTION 301

Refer to the exhibit.



A network engineer wants to monitor traffic from the application server and send the output to the external monitoring device at 172.16.20.5. Application server traffic should pass through the R1 Eth2/1 interface for further analysis after it is monitored. Which configuration must be applied on the R1 router?

- A. Configure the FLOW-MONITOR-20 command.
- B. Configure the flow exporter EXPORTER-10 destination 192.168.20.4 command.
- C. Configure the ip flow monitor FLOW-MONITOR-10 input command on the Ethernet1/0 interface.
- D. Configure the ip flow monitor FLOW-MONITOR-10 output command on the Ethernet 2/1 interface.

Answer: C

NEW QUESTION 303

Refer to the exhibit:

```

R1:
interface FastEthernet0/0
ip address 10.1.12.1 255.255.255.0
duplex full
end
!
!
!
R1(config)#interface FastEthernet0/0
R1(config-if)#ospfv3 1 area 1 ipv4
% IPv6 routing not enabled
    
```

A network engineer is implementing an OSPF configuration Based on the output, which statement is true?

- A. In the ospfv3 1 area 1 ipv4 command, area 0 must be configured instead of area 1.
- B. OSPFv3 does not run for IPv4 on FastEthernet0/0 until IPv6 routing is enabled on the router and IPv6 is enabled on interface FastEthernet0/0
- C. OSPFv3 cannot be configured for IPv4; OSPFv3 works only for IPv6.
- D. "IPv6 routing not enabled" is just an informational message and OSPFv3 runs for IPv4 on interface FastEthernet0/0 anyway

Answer: B

NEW QUESTION 306

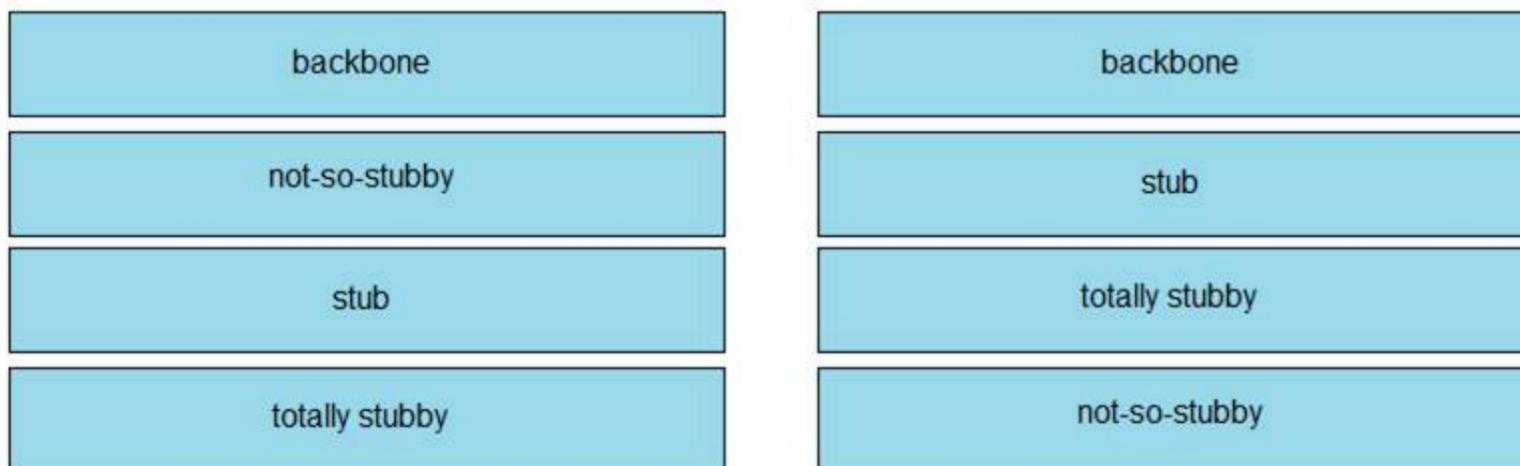
Drag and drop the OSPF area types from the left onto the correct statements on the right

backbone	required area that allows interarea communication
not-so-stubby	area that can learn interarea routes and the default route
stub	area that can learn only the default route and routes within its own area
totally stubby	area that can serve as a redistribution point for external routes to enter the OSPF domain

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 309

An engineer is trying to implement BGP in a multihomed architecture. What must the engineer configure to influence inbound path selection?

- A. A route map with WEIGHT attribute to control the inbound traffic.
- B. An offset list to set the metric for routes received from neighboring autonomous systems.
- C. An access list to identify traffic and enable it on both of the provider-facing interfaces.
- D. A route map with AS_PATH attribute to control the inbound traffic.

Answer: D

NEW QUESTION 312

What is the function of Cisco NFV infrastructure platform?

- A. It does not have a security audit feature.
- B. It does not offer high availability.
- C. It offers consistent performance.
- D. It offers decentralized logging.

Answer: C

NEW QUESTION 314

Refer to the exhibit.

```
Router(config)# ip access-list standard Suppressed
Router(config-std-nacl)# permit 10.16.6.0 0.0.0.255
Router(config)# route-map SuppressMap
Router(config-route-map)# match ip address Suppressed
```

An engineer is implementing BGP selective prefix suppression. The router must advertise only 10.16.4.0/24, 10.16.5.0/24, and summarized route 10.16.0.0/21, and suppress 10.16.6.0/24. Which configuration must the engineer apply to the router?

- A)


```
Router (config)# router bgp 300
Router(config-router)# aggregate-address 10.16.6.0 255.255.252.0 as-set suppress-map SuppressMap
```
- B)


```
Router (config)# router bgp 300
Router(config-router)# aggregate-address 10.16.0.0 255.255.248.0 as-set suppress-map SuppressMap
```
- C)


```
Router (config)# router bgp 300
Router(config-router)# aggregate-address 10.16.6.0 255.255.255.0 as-set suppress-map SuppressMap
```
- D)

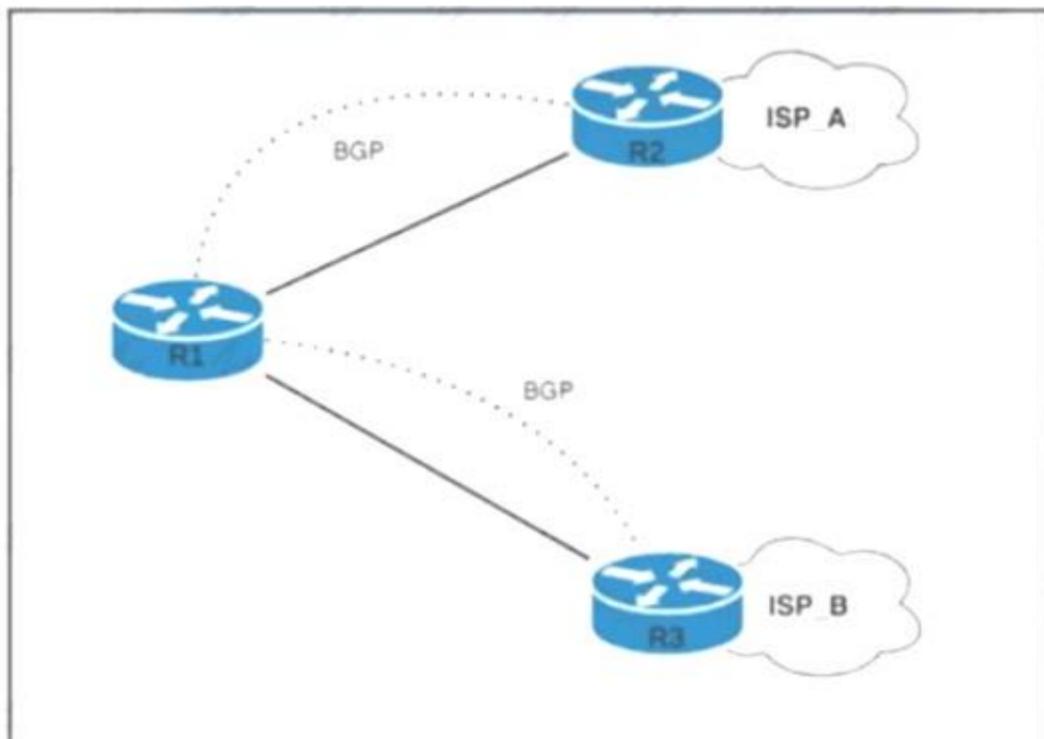

```
Router (config)# router bgp 300
Router(config-router)# aggregate-address 10.16.0.0 255.255.255.0 as-set suppress-map unSuppressMap
```

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

Answer: B

NEW QUESTION 315

Refer to the exhibit.



R1 has two upstream Tier 1 service providers. BGP is in use as the exterior routing protocol, and ISP_A and ISP_B are sending the full BGP table. A network engineer must assign local-preference 70 to all routes with multiple exit discriminator 30. Which configuration must the network engineer apply?

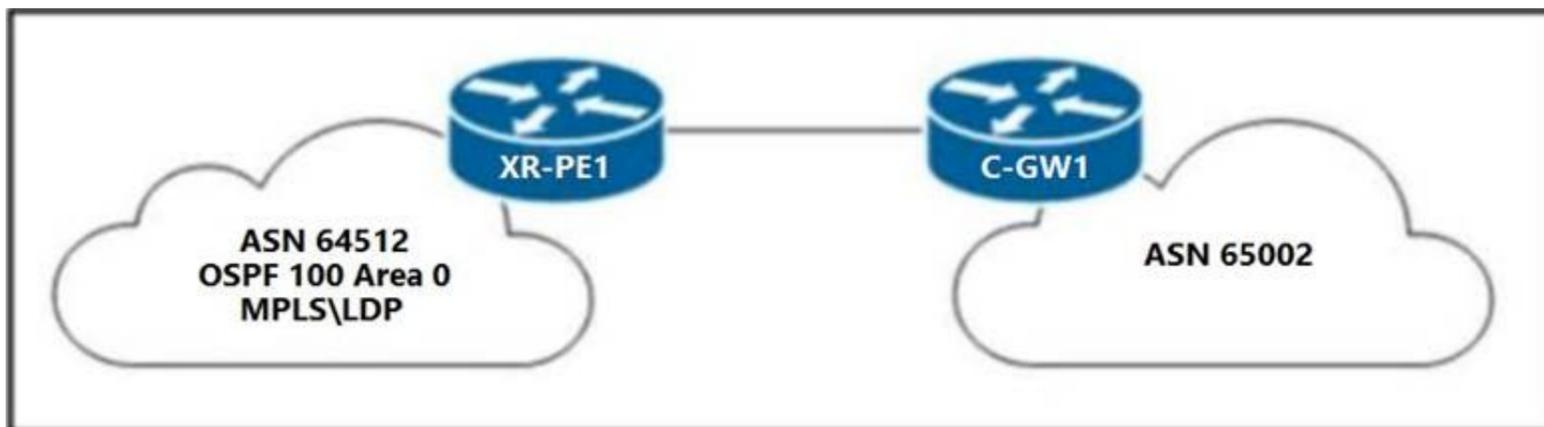
- route-policy routepolicy
 if destination in (0.0.0.0/0) and (med = 30) then
 set local-preference 170
 else
 set local-preference 70
 drop
 endif
 end-policy
- route-policy routepolicy
 if destination 0.0.0.0/0 and med 30 then
 set local-preference 70
 else
 drop
 endif
 end-policy
- route-policy routepolicy
 if med eq 30 then
 set local-preference 70
 else pass
 endif
 end-policy
- route-policy routepolicy
 if destination in (.*) and med eq 70 then
 set local-preference 30
 else
 drop
 endif
 end-policy

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

Answer: C

NEW QUESTION 318

Refer to the exhibit.



A network engineer must configure XR-PE1 for uninterruptible failover from active RP to the standby RP. Neither peer devices CGW1 nor the network of ASN 64512 support restart extensions. Which configuration must the engineer apply to XR-PE1 to complete tasks?

- A)


```
router bgp 64512 nsr
router ospf 100 nsr
mpls ldp nsr
```
- B)


```
nsr process-failures switchover
router ospf 100 nsf cisco
```
- C)


```
nsr process-failures switchover
router ospf 100 nsf ietf
```
- D)


```
nsr process-failures switchover
router bgp 64512 nsr
router ospf 100 nsr
mpls ldp nsr
```

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

Answer: D

NEW QUESTION 319

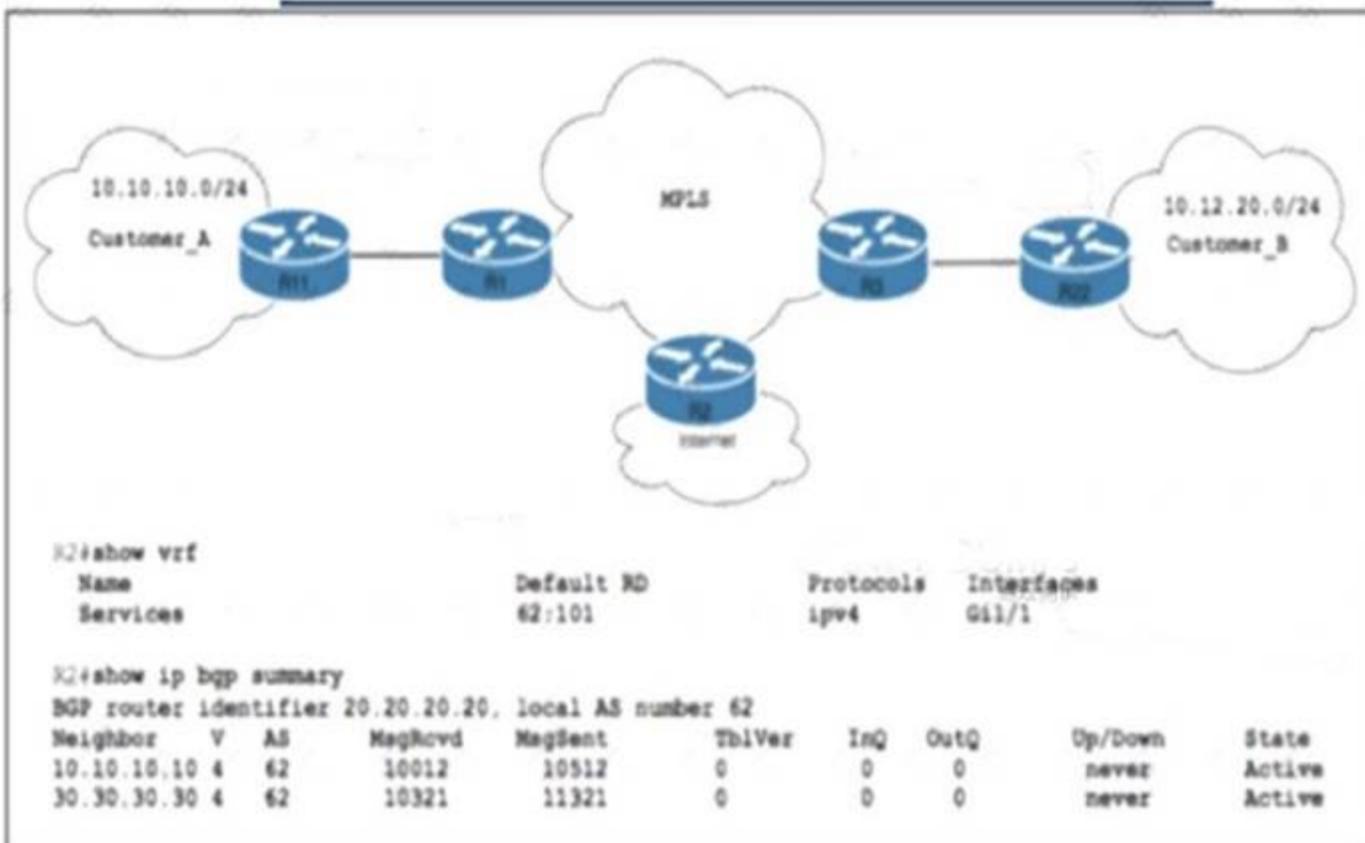
What is the characteristic of the TI-LFA?

- A. It guarantees a loop-free path for all interfaces in the OSPF- super backbone .
- B. It applies on each area and instance and makes all the interfaces inherit the configuration
- C. It guarantees a loop-free path for all areas configured in OSPF
- D. It applies only on the instance and makes at the interfaces inherit the configuration

Answer: A

NEW QUESTION 323

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?

- 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: A

NEW QUESTION 328

Refer to the exhibit.

```

R1#show ip bgp
BGP table version is 3, local router ID is 50.50.50.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 22.22.22.22/32   50.50.50.2         0         100 500 ?
*                   40.40.40.2         0          200   0 400 ?
*                   30.30.30.2         0         300 300 ?
*                   20.20.20.2         0         200 ?

R1#show ip bgp 22.22.22.22
BGP routing table entry for 22.22.22.22/32, version 3
Paths: (4 available, best #1, table Default-IP-Routing-Table)
Flag: 0x820
  Advertised to update-groups:
    1
  500
    50.50.50.2 from 50.50.50.2 (50.50.50.2)
      Origin incomplete, metric 0, localpref 100, weight 100, valid, external, best
  400
    40.40.40.2 from 40.40.40.2 (40.40.40.2)
      Origin incomplete, metric 0, localpref 200, valid, external
  300 300
    30.30.30.2 from 30.30.30.2 (30.30.30.2)
      Origin incomplete, metric 0, localpref 100, valid, external
  200
    20.20.20.2 from 20.20.20.2 (20.20.20.2)
      Origin incomplete, metric 0, localpref 100, valid, external
    
```

An engineer wants to determine which paths are best, second best, third best, and fourth best. Drag and drop the peer addresses on the left to the corresponding BGP best-path selection order on the right.

20.20.20.2	Best Path
30.30.30.2	2nd Best Path
40.40.40.2	3rd Best Path
50.50.50.2	4th Best Path

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Best – 50.50.50.2
 2nd Best – 40.40.40.2
 3rd Best – 20.20.20.2
 4th Best – 30.30.30.2

NEW QUESTION 329

Refer to the exhibit.

```

R1
ip cef distributed
mpls ldp graceful-restart
interface GigabitEthernet 0/0/1
    mpls ip
    mpls label protocol ldp
    
```

What is the effect of this configuration?

- A. R1 supports a graceful restart operation on the peer, even if graceful restart is disabled on the peer.
- B. R1 supports a peer that is configured for LDP SSO/NSF as the peer recovers from an outage.
- C. R1 failovers only to a peer that is configured for LDP SSO/NSF.
- D. R1 failovers to any peer.

Answer: B

NEW QUESTION 333

Refer to the exhibit:

```

PE-A#show ip bgp vpv4 vrf Customer-A neighbors 10.10.10.2 routes
BGP table version is 13148019, local router ID is 10.10.10.10
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
Route Distinguisher: 65000:1111 (default for vrf Customer-A)
*> 192.168.0.0/19   10.10.10.2         0           0 4282 65001 ?
*> 192.168.0.0/17   10.10.10.2         0           0 4282 65001 ?
*> 192.168.0.0/16   10.10.10.2         0           0 4282 65001 ?

Total number of prefixes 5

PE-A#config t
Enter configuration commands, one per line. End with CNTL/Z.
PE-A(config)#ip prefix-list ALLOW permit 192.168.0.0/16 ge 17 le 19
PE-A(config)#router bgp 65000
PE-A(config-router)#address-family ipv4 vrf Customer-A
PE-A(config-router-af)#neighbor 10.10.10.2 prefix-list ALLOW in
    
```

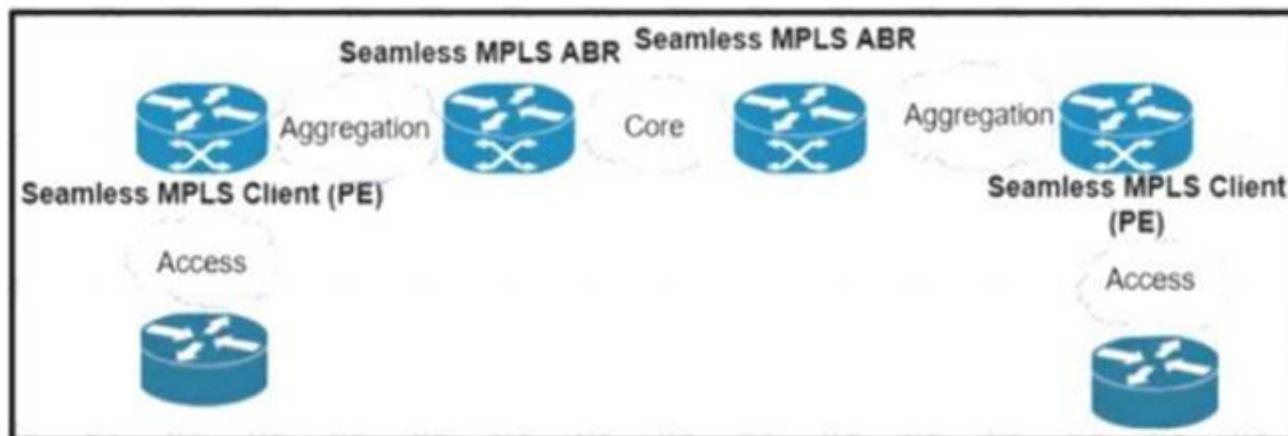
Which three outcomes occur if the prefix list is added to the neighbor? (Choose three)

- A. 192.168 0.0/19 is denie
- B. 192.168 0.0/17 is denied.
- C. 192.168 0.0/17 is permitted
- D. 192.168.0.0/16 is denied
- E. 192.168 0.0/16 is permitted
- F. 192.168 0.0/19 is permitted

Answer: CDF

NEW QUESTION 336

Refer to the exhibit.



A network operator working for a telecommunication company with an employee 3994:37:650 is implementing a cisco Unified MPLS solution. What is the effect of this implementation?

- A. EIGRP is deployed between the PEs and ABRs with RFC 3107.
- B. OSPF is deployed between the PEs and ABRs with RFC 3107.
- C. IS-IS is deployed between the PEs and ABRs with RFC 3107.
- D. BGP is deployed between the PEs and ABRs with RFC 3107.

Answer: D

Explanation:

Carry Label Information in BGP-4 (RFC 3107)

It is a prerequisite to have a scalable method in order to exchange prefixes between network segments. You could simply merge the IGP's (Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), or Enhanced Interior Gateway Routing Protocol (EIGRP)) into a single domain. However an IGP is not designed to carry 100,000s of prefixes. The protocol of choice for that purpose is BGP. It is a

NEW QUESTION 340

Refer to the exhibit.

```

PE-1#show xconnect name ENNI-ID-100150
Legend:  XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State
         UP=Up                DN=Down          AD=Admin Down    IA=Inactive
         SB=Standby           HS=Hot Standby  RV=Recovering   NH=No Hardware

XC ST Segment 1                               S1 Segment 2                               S2
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
UP pri ac Gi2:150 (Eth VLAN)                   UP mpls 172.20.20.2:100150                   UP

PE-2#show xconnect name UNI-ID-100150
Legend:  XC ST=Xconnect State  S1=Segment1 State  S2=Segment2 State
         UP=Up                DN=Down          AD=Admin Down    IA=Inactive
         SB=Standby           HS=Hot Standby  RV=Recovering   NH=No Hardware

XC ST Segment 1                               S1 Segment 2                               S2
-----+-----+-----+-----+-----+-----+-----+-----+-----+
UP pri ac Gi2:10 (Eth VLAN)                    UP mpls 172.20.20.1:100150                   UP

CE-2#show run interface gigabitEthernet 2.10
interface GigabitEthernet2.10
 encapsulation dot1q 10
 ip address 100.65.0.2 255.255.255.252

CE-1#show run interface gigabitEthernet 0/0/0/1.150
interface GigabitEthernet0/0/0/1.150
 ipv4 address 100.65.0.1 255.255.255.252
 encapsulation dot1ad 150 dot1q 10
    
```

An Ethernet access provider is configuring routers PE-1 and PE-2 to provide E-Access EVPL service between UNI and ENNI. ENNI service multiplexing is based on 802.1ad tag 150, and service-multiplexed UNI is based on 802.1q tag 10. Which EFP configurations must the provider implement on PE-1 and PE-2 to establish end-to-end connectivity between CE-1 and CE-2?

- A. On PE-1:interface GigabitEthernet2 service instance 100 ethernet encapsulation dot1ad 150rewrite ingress tag pop 1 symmetric On PE-2:interface GigabitEthernet2 service instance 2 ethernet encapsulation dot1q 10
- B. On PE-1:interface GigabitEthernet2 service instance 100 ethernet encapsulation dot1q 150rewrite ingress tag pop 1 symmetric On PE-2:interface GigabitEthernet2 service instance 2 ethernet encapsulation dot1q 10
- C. On PE-1:interface GigabitEthernet2 service instance 100 ethernet encapsulation dot1ad 150 dot1q 10rewrite ingress tag pop 2 symmetric On PE-2:interface GigabitEthernet2 service instance 2 ethernet encapsulation dot1q 10
- D. On PE-1:interface GigabitEthernet2 service instance 100 ethernet encapsulation dot1ad 150rewrite ingress tag pop 1 symmetric On PE-2:interface GigabitEthernet2 service instance 2 ethernet encapsulation dot1q 10rewrite ingress tag pop 1 symmetric

Answer: C

NEW QUESTION 344

An engineer is implementing MPLS to monitor within the MPLS domain. Which must the engineer perform to prevent packets from being forwarded beyond the service provider domain when the LSP is down?

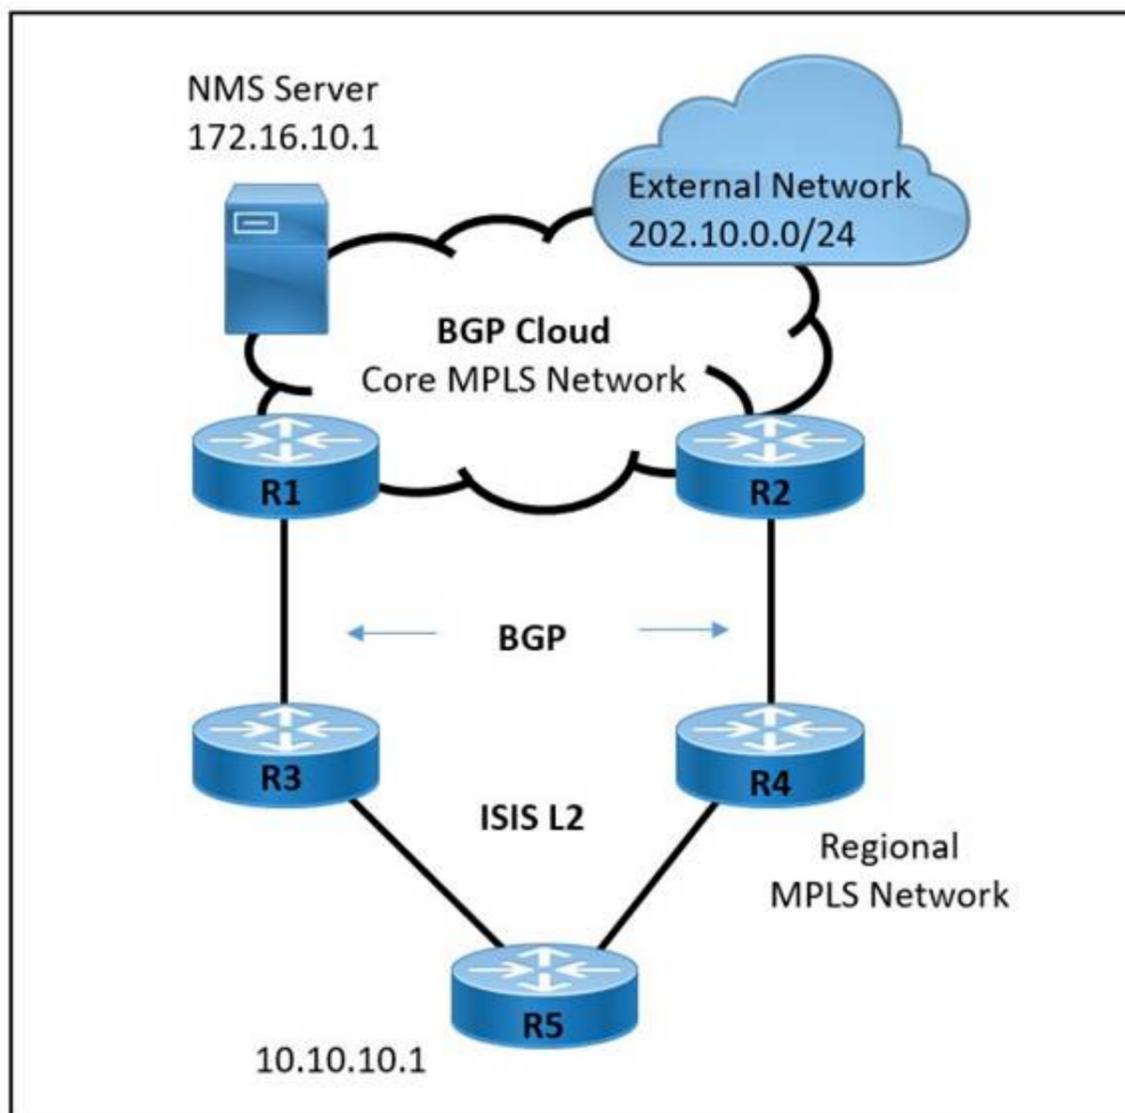
- Disable IP redirects only on outbound interfaces.
- Implement the destination address for the LSP echo request packet in the 127 x y z/8 network
- Disable IP redirects on all ingress interfaces
- Configure a private IP address as the destination address of the headend router of Cisco MPLS TE.

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

Answer: B

NEW QUESTION 349

Refer to the exhibit.



A large service provider is migrating device management from Layer 2 VLAN-based to Layer 3 IP-based solution. An engineer must configure the ISIS solution with these requirements:

Network management server IP 172.16.10.1 must be advertised from the core MPLS network to the regional domain.

The external network 202.10.0.0/24 must not establish ISIS peering with the R5 router.

The regional network must prevent sending unnecessary hello packets and flooding the routing tables of the R5 router.

Which two ISIS parameters must be implemented to meet these requirements? (Choose two.)

- A. LSP lifetime maximum
- B. advertise-passive-only
- C. overload bit passive
- D. attached bit on ISIS instance
- E. passive-interface Loopback0

Answer: AD

NEW QUESTION 354

Which type of attack is a Protocol attack?

- A. HTTP flood
- B. TFTP flood
- C. SYN flood
- D. Slowloris

Answer: C

Explanation:

Protocol Attacks

Includes SYN floods, fragmented packet attacks, Ping of Death, Smurf DDoS and more. This type of attack consumes actual server resources,

NEW QUESTION 357

A router is configured to perform MPLS LDP graceful restart.

Which three steps are included when the RP sends an LDP initialization to a neighbor to establish an LDP session? (Choose three)

- A. Reconnect Timeout field
- B. Learn from Neighbor (N) flag, set to 1
- C. Graceful restart capability in OPEN message
- D. Recovery Time field
- E. Learn from Network (L.) flage, set to 1
- F. Type-9 LSA

Answer: ADE

NEW QUESTION 360

In an EVPN operation, how does the PE determine and advertise Ethernet segment reachability?

- A. The PE discovers the remote PEs in the EVI and builds a flood list linked with the EVI.
- B. The PE discovers and shared routing information for the B-MAC addresses associated with local Ethernet segments.
- C. The PE discovers other PEs in the same Ethernet segment and elects a DF.
- D. The PE discovers remote ESIs and determines their redundancy mode.

Answer: A

NEW QUESTION 363

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

350-501 Practice Exam Features:

- * 350-501 Questions and Answers Updated Frequently
- * 350-501 Practice Questions Verified by Expert Senior Certified Staff
- * 350-501 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * 350-501 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 350-501 Practice Test Here](#)