Free Questions for 300-440 by go4braindumps

Shared by Abbott on 22-07-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which architecture model establishes internet-based connectivity between on-premises networks and AWS cloud resources?

Options:

A- That establishes an iPsec VPN tunnel with Internet Key Exchange (IKE) for secure key negotiation and encrypted data transmission

B- That relies on AWS Elastic Load Balancing (ELB) for traffic distribution and uses SSL/TLS encryption for secure data transmission.

C- That employs AWS Direct Connect for a dedicated network connection and uses private IP addresses tor secure communication.

D- That uses Amazon CloudFront for caching and distributing content globally and uses HTTPS for secure data transfer.

Answer:

A

Explanation:

The architecture model that establishes internet-based connectivity between on-premises networks and AWS cloud resources is the one that establishes an iPsec VPN tunnel with Internet Key Exchange (IKE) for secure key negotiation and encrypted data transmission.This model is also known as theVPN CloudHubmodel12.It allows multiple remote sites to connect to the same virtual private gateway in AWS, creating a hub-and-spoke topology1.The VPN CloudHub model provides the following benefits12:

It enables secure communication between remote sites and AWS over the public internet, using encryption and authentication protocols such as IPsec and IKE.

It supports dynamic routing protocols such as BGP, which can automatically adjust the routing tables based on the availability and performance of the VPN tunnels.

It allows for redundancy and load balancing across multiple VPN tunnels, increasing the reliability and throughput of the connectivity.

It simplifies the management and configuration of the VPN connections, as each remote site only needs to establish one VPN tunnel to the virtual private gateway in AWS, rather than multiple tunnels to different VPCs or regions.

The other options are not correct because they do not establish internet-based connectivity between on-premises networks and AWS cloud resources. Option B relies on AWS Elastic Load Balancing (ELB) for traffic distribution and uses SSL/TLS encryption for secure data transmission.However, ELB is a service that distributes incoming traffic across multiple targets within a VPC, not across different networks3. Option C employs AWS Direct Connect for a dedicated network connection and uses private IP addresses for secure communication.However, AWS Direct Connect is a service that establishes a private connection between on-premises networks and AWS, bypassing the public internet4. Option D uses Amazon CloudFront for caching and distributing content globally and uses HTTPS for secure data transfer.However, Amazon CloudFront is a service that delivers static and dynamic web content to end users, not to on-premises networks5.

1: Designing and Implementing Cloud Connectivity (ENCC, Track 1 of 5)

2: Cisco ASA Site-to-Site VPN

3: What Is Elastic Load Balancing?

4: What is AWS Direct Connect?

Question 2

Question Type: DragDrop

Refer to the exhibits.

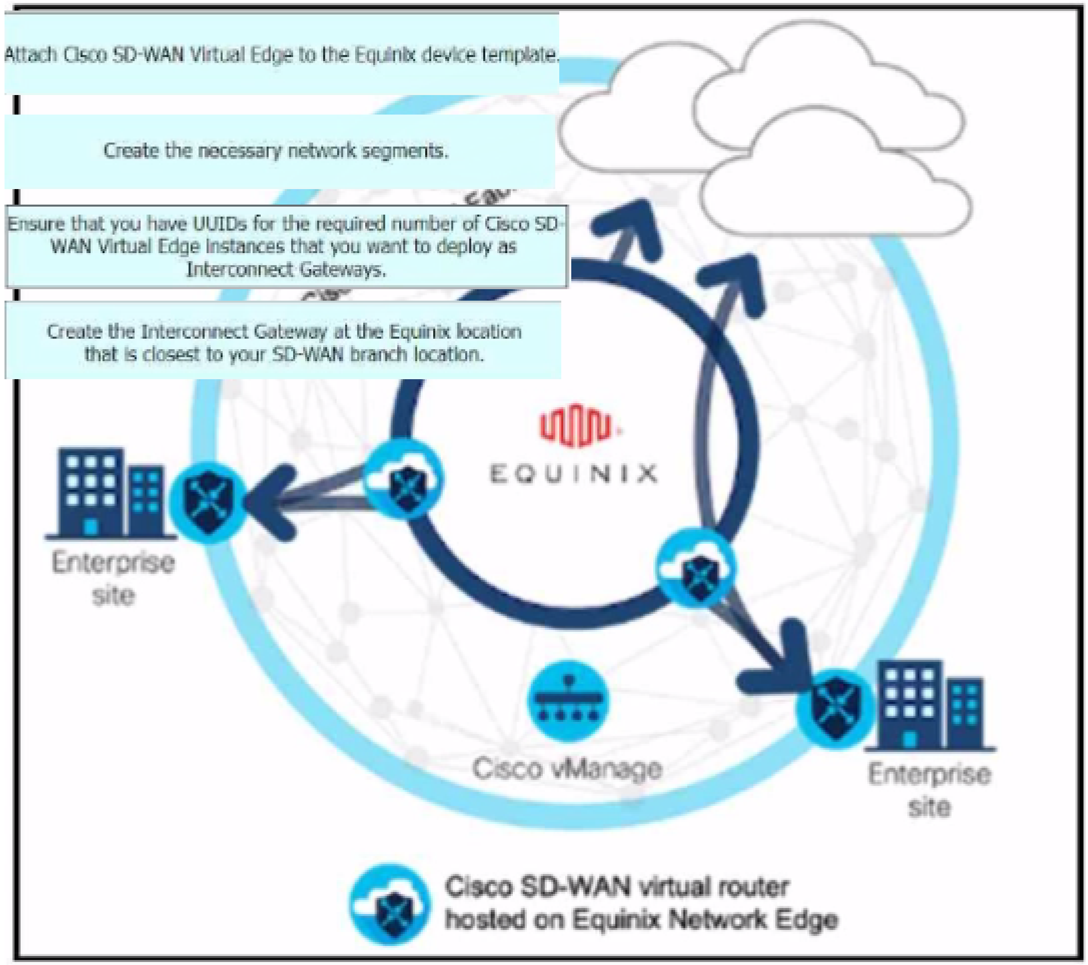

Attach Cisco SD-WAN Virtual Edge to the Equinix device template.

Create the necessary network segments.

Ensure that you have UUIDs for the required number of Cisco SD-WAN Virtual Edge instances that you want to deploy as Interconnect Gateways.

Create the Interconnect Gateway at the Equinix location that is closest to your SD-WAN branch location.

Refer to the exhibit. These configurations are complete:

* Create an account in the Equinix portal.

* Associate the Equinix account with Cisco vManage.

* Configure the global settings for Interconnect Gateways.

Drag the prerequisite steps from the left onto the order on the right to configure a Cisco SD-WAN Cloud Interconnect with Equinix

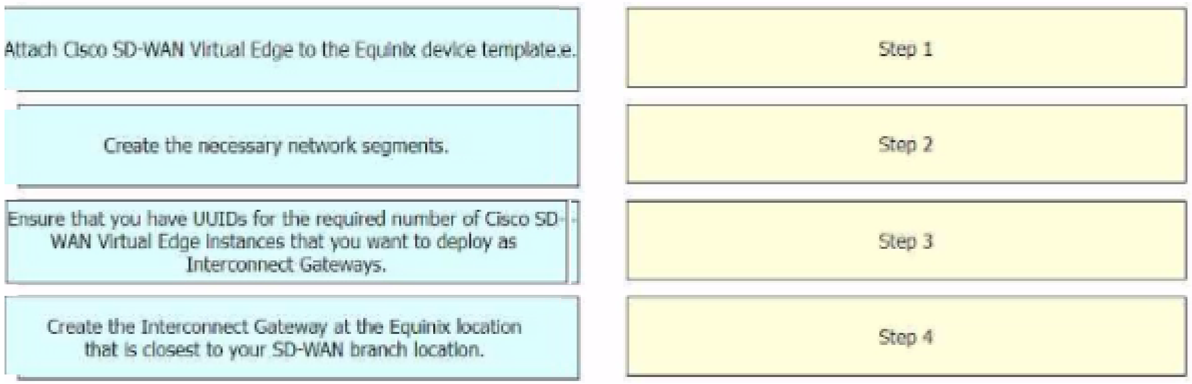

Attach Cisco SD-WAN Virtual Edge to the Equinix device template.e.

Create the necessary network segments.

Ensure that you have UUIDs for the required number of Cisco SD-WAN Virtual Edge instances that you want to deploy as Interconnect Gateways.

Create the Interconnect Gateway at the Equinix location that is closest to your SD-WAN branch location.

Step 1

Step 2

Step 3

Step 4

Explanation:

[Cisco SD-WAN Cloud Interconnect with Equinix]

[Cisco SD-WAN Cloud OnRamp for CoLocation Deployment Guide]

Question 3

Question Type: DragDrop

An engineer must use Cisco vManage to configure an SLA class to specify the maximum packet loss, packet latency, and jitter allowed on a connection. Drag and drop the steps from the left onto the order on the right to complete the configuration.

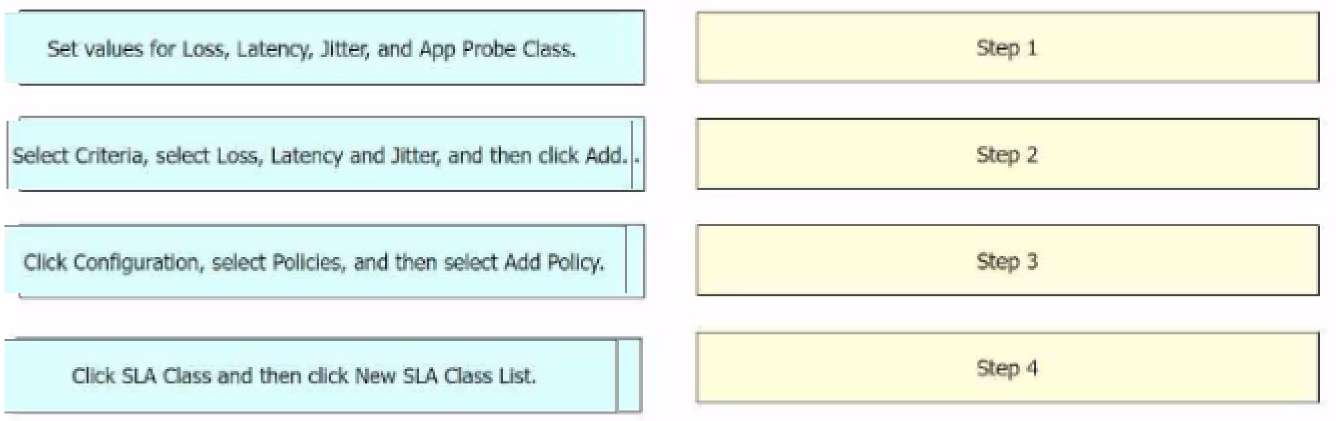

Set values for Loss, Latency, Jitter, and App Probe Class.

Select Criteria, select Loss, Latency and Jitter, and then click Add.

Click Configuration, select Policies, and then select Add Policy.

Click SLA Class and then click New SLA Class List.

Step 1

Step 2

Step 3

Step 4

Explanation:

Information About Application-Aware Routing - Cisco

Policies Configuration Guide for vEdge Routers, Cisco SD-WAN Release 20

Question 4

Question Type: MultipleChoice

Refer to the exhibit.

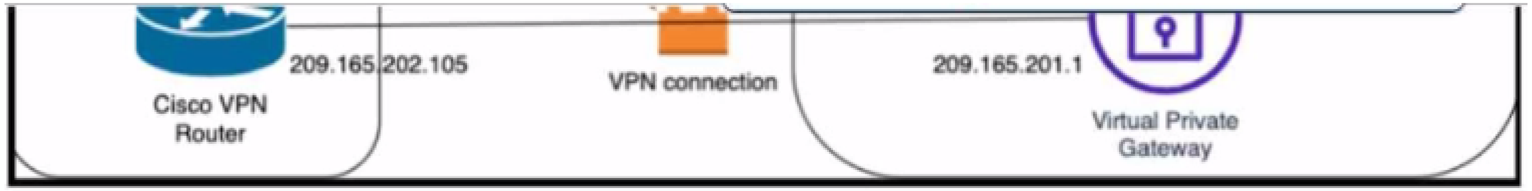

Which Cisco IKEv2 configuration brings up the IPsec tunnel between the remote office router and the AWS virtual private gateway?

A)

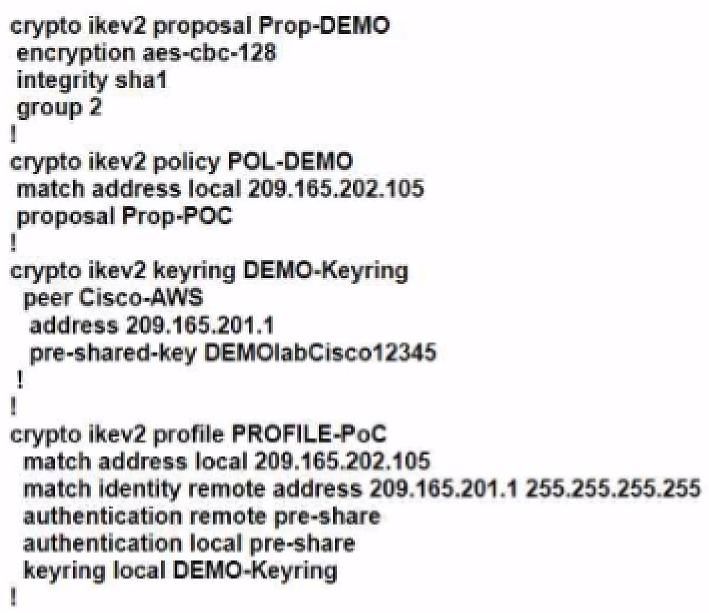

```
crypto ikev2 proposal Prop-DEMO
 encryption aes-cbc-128
 integrity sha1
 group 2
!
crypto ikev2 policy POL-DEMO
 match address local 209.165.202.105
 proposal Prop-POC
!
crypto ikev2 keyring DEMO-Keyring
  peer Cisco-AWS
   address 209.165.201.1
   pre-shared-key DEMOlabCisco12345
 !
!
crypto ikev2 profile PROFILE-PoC
  match address local 209.165.202.105
  match identity remote address 209.165.201.1 255.255.255.255
  authentication remote pre-share
  authentication local pre-share
  keyring local DEMO-Keyring
!
```

B)

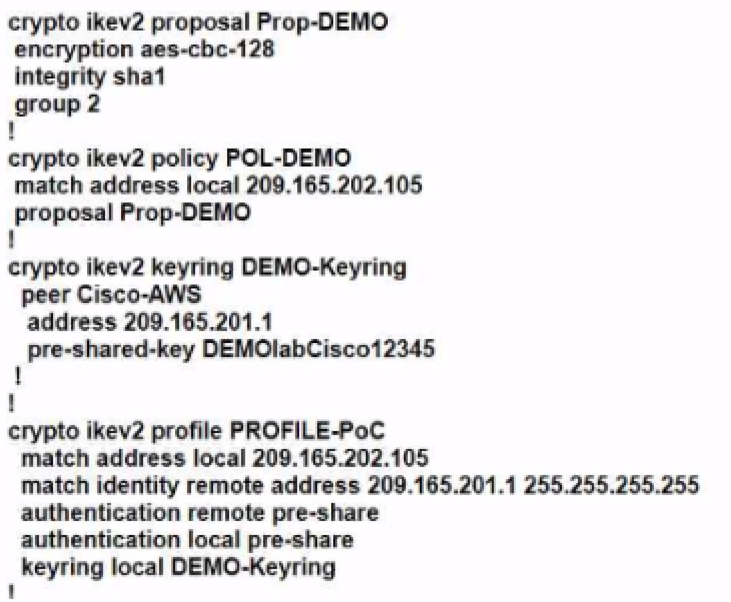

```
crypto ikev2 proposal Prop-DEMO
 encryption aes-cbc-128
 integrity sha1
 group 2
!
crypto ikev2 policy POL-DEMO
 match address local 209.165.202.105
 proposal Prop-DEMO
!
crypto ikev2 keyring DEMO-Keyring
  peer Cisco-AWS
   address 209.165.201.1
   pre-shared-key DEMOlabCisco12345
 !
!
crypto ikev2 profile PROFILE-PoC
  match address local 209.165.202.105
  match identity remote address 209.165.201.1 255.255.255.255
  authentication remote pre-share
  authentication local pre-share
  keyring local DEMO-Keyring
!
```

C)

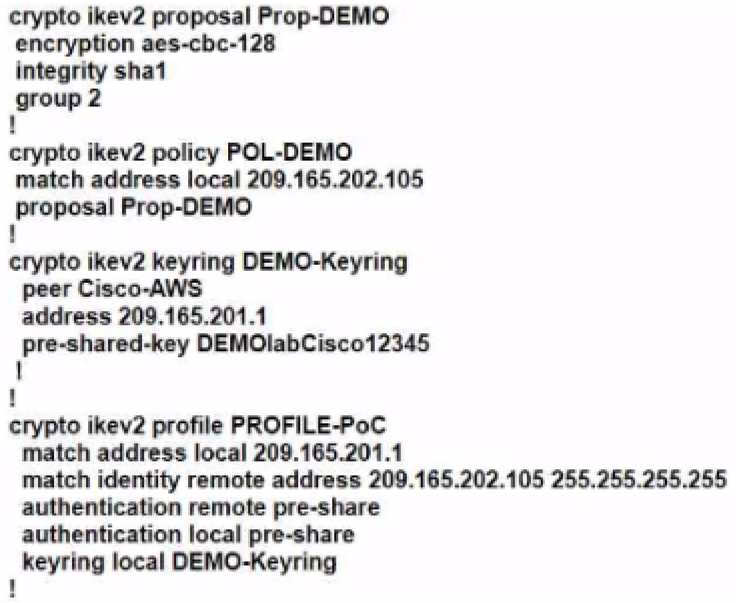

```
crypto ikev2 proposal Prop-DEMO
 encryption aes-cbc-128
 integrity sha1
 group 2
!
crypto ikev2 policy POL-DEMO
 match address local 209.165.202.105
 proposal Prop-DEMO
!
crypto ikev2 keyring DEMO-Keyring
  peer Cisco-AWS
  address 209.165.201.1
  pre-shared-key DEMOlabCisco12345
 !
!
crypto ikev2 profile PROFILE-PoC
  match address local 209.165.201.1
  match identity remote address 209.165.202.105 255.255.255.255
  authentication remote pre-share
  authentication local pre-share
  keyring local DEMO-Keyring
!
```

Options:

A- Option A

B- Option B

C- Option C

Answer:

C

Explanation:

Option C is the correct answer because it configures the IKEv2 profile with the correct match identity, authentication, and keyring parameters. It also configures the IPsec profile with the correct transform set and lifetime parameters. Option A is incorrect because it does not specify the match identity remote address in the IKEv2 profile, which is required to match the AWS virtual private gateway IP address. Option B is incorrect because it does not specify the authentication pre-share in the IKEv2 profile, which is required to authenticate the IKEv2 peers using a pre-shared key.Option C also matches the configuration example provided by AWS1and Cisco2for setting up an IKEv2 IPsec site-to-site VPN between a Cisco IOS-XE router and an AWS virtual private gateway.Reference:=

1: AWS VPN Configuration Guide for Cisco IOS-XE

2: Configure IOS-XE Site-to-Site VPN Connection to Amazon Web Services

Question 5

Question Type: **MultipleChoice**

An engineer must configure an IPsec tunnel to the cloud VPN gateway. Which Two actions send traffic into the tunnel? (Choose two.)

Options:

A- Configure access lists that match the interesting user traffic.

B- Configure a static route.

C- Configure a local policy in Cisco vManage.

D- Configure an IPsec profile and match the remote peer IP address.

E- Configure policy-based routing.

Answer:

A, E

Explanation:

To send traffic into an IPsec tunnel to the cloud VPN gateway, the engineer must configure two actions:

Configure access lists that match the interesting user traffic. This is the traffic that needs to be encrypted and sent over the IPsec tunnel. The access lists are applied to the crypto map that defines the IPsec parameters for the tunnel.

Configure policy-based routing (PBR). This is a technique that allows the engineer to override the routing table and forward packets based on a defined policy. PBR can be used to send specific traffic to the IPsec tunnel interface, regardless of the destination IP address. This is useful when the cloud VPN gateway has a dynamic IP address or when multiple cloud VPN gateways are available for load balancing or redundancy.Reference:

Designing and Implementing Cloud Connectivity (ENCC) v1.0, Module 3: Implementing Cloud Connectivity, Lesson 3: Implementing IPsec VPNs to the Cloud, Topic: Configuring IPsec VPNs on Cisco IOS XE Routers

Security for VPNs with IPsec Configuration Guide, Cisco IOS XE, Chapter: Configuring IPsec VPNs, Topic: Configuring Crypto Maps

[Cisco IOS XE Gibraltar 16.12.x Feature Guide], Chapter: Policy-Based Routing, Topic: Policy-Based Routing Overview

Question 6

Question Type: MultipleChoice

Refer to the exhibits.

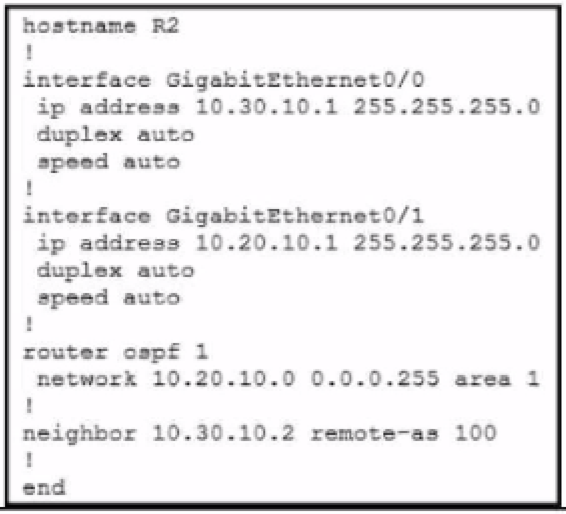

```
hostname R2
!
interface GigabitEthernet0/0
 ip address 10.30.10.1 255.255.255.0
 duplex auto
 speed auto
!
interface GigabitEthernet0/1
 ip address 10.20.10.1 255.255.255.0
 duplex auto
 speed auto
!
router ospf 1
 network 10.20.10.0 0.0.0.255 area 1
!
neighbor 10.30.10.2 remote-as 100
!
end
```

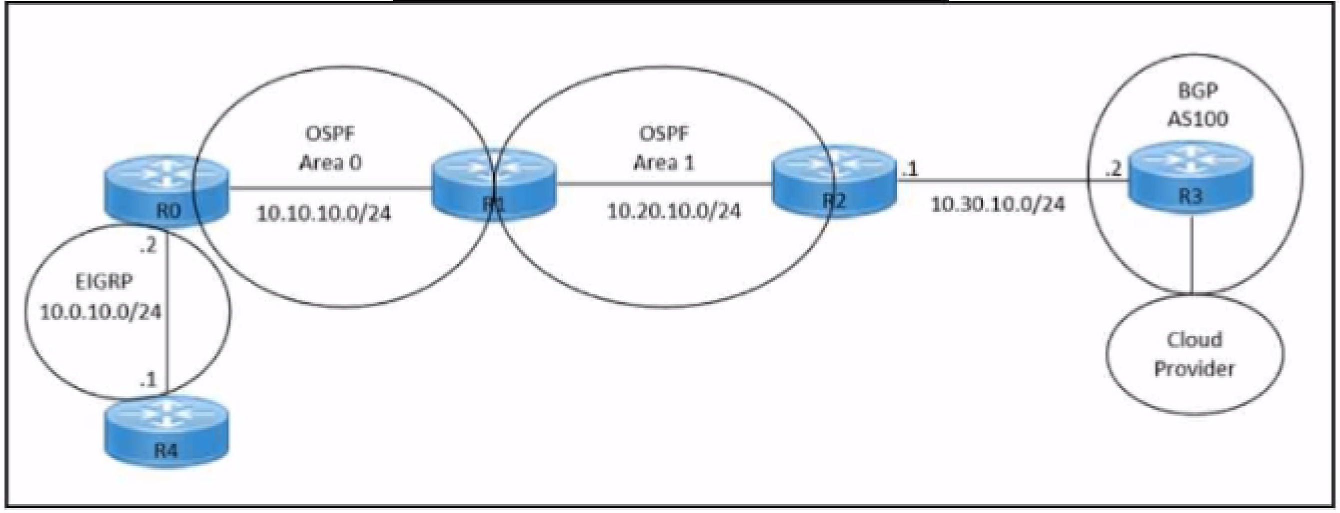

Refer to the exhibits. An engineer must redistribute OSPF internal routes into BGP to connect an on-premises network to a cloud provider without introducing extra routes. Which two commands must be configured on router R2? (Choose two.)

Options:

A- router ospf 1

B- router bgp 100

C- redistribute ospf 1

D- redistribute bgp 100

E- redistribute ospf 1 match internal external

Answer:

B, E

Explanation:

To redistribute OSPF internal routes into BGP, the engineer needs to configure two commands on router R2. The first command isrouter bgp 100, which enables BGP routing process and specifies the autonomous system number of 100. The second command isredistribute ospf 1 match internal external, which redistributes the routes from OSPF process 1 into BGP, and matches both internal and external OSPF routes. This way, the engineer can avoid introducing extra routes that are not part of OSPF process 1, such as the default route or the connected routes.Reference: =Designing and Implementing Cloud Connectivity (ENCC) v1.0, [ENCC: Configuring IPsec VPN from Cisco IOS XE to AWS], [Deploying Cisco IOS VTI-Based Point-to-Point IPsec VPNs]

Question 7

Question Type: **MultipleChoice**

Refer to the exhibits.

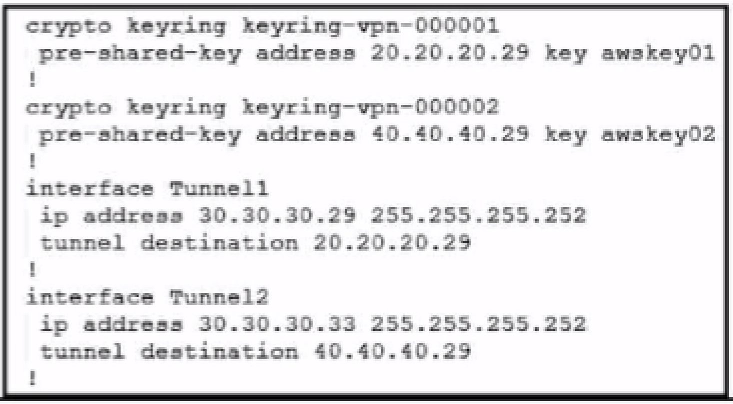

```
crypto keyring keyring-vpn-000001
 pre-shared-key address 20.20.20.29 key awskey01
!
crypto keyring keyring-vpn-000002
 pre-shared-key address 40.40.40.29 key awskey02
!
interface Tunnel1
 ip address 30.30.30.29 255.255.255.252
 tunnel destination 20.20.20.29
!
interface Tunnel2
 ip address 30.30.30.33 255.255.255.252
 tunnel destination 40.40.40.29
!
```

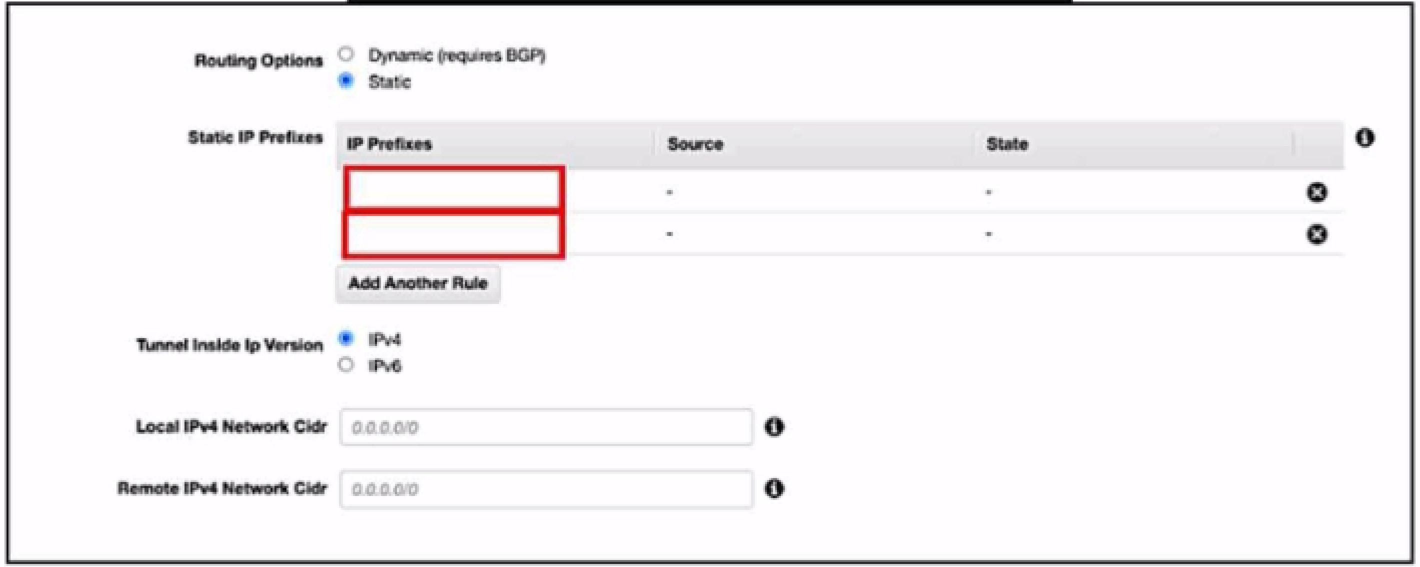

Routing Options ○ Dynamic (requires BGP) ● Static

Static IP Prefixes

IP Prefixes	Source	State	
	-	-	⊗
	-	-	⊗

Tunnel Inside Ip Version ● IPv4 ○ IPv6

Local IPv4 Network Cidr 0.0.0.0/0

Remote IPv4 Network Cidr 0.0.0.0/0

Refer to the exhibits. An engineer needs to configure a site-to-site IPsec VPN connection between an on premises Cisco IOS XE router and Amazon Web Services (AWS). Which two IP prefixes should be used to configure the AWS routing options? (Choose two.)

Options:

A- 30.30.30.0/30

B- 20.20.20.0/24

C- 30.30.30.0/24

D- 50.50.50.0/30

E- 40.40.40.0/24

Answer:

A, E

Explanation:

The correct answer is A and E because they are the IP prefixes that match the tunnel interfaces on the Cisco IOS XE router. The AWS routing options should include the local and remote IP prefixes that are used for the IPsec tunnel endpoints. The other options are either the public IP addresses of the routers or the LAN subnets that are not relevant for the IPsec tunnel configuration.Reference:=Designing and Implementing Cloud Connectivity (ENCC) v1.0,Configure IOS-XE Site-to-Site VPN Connection to Amazon Web Services,Site-to-Site VPN with Amazon Web Services

Question 8

Question Type: **MultipleChoice**

Refer to the exhibit.

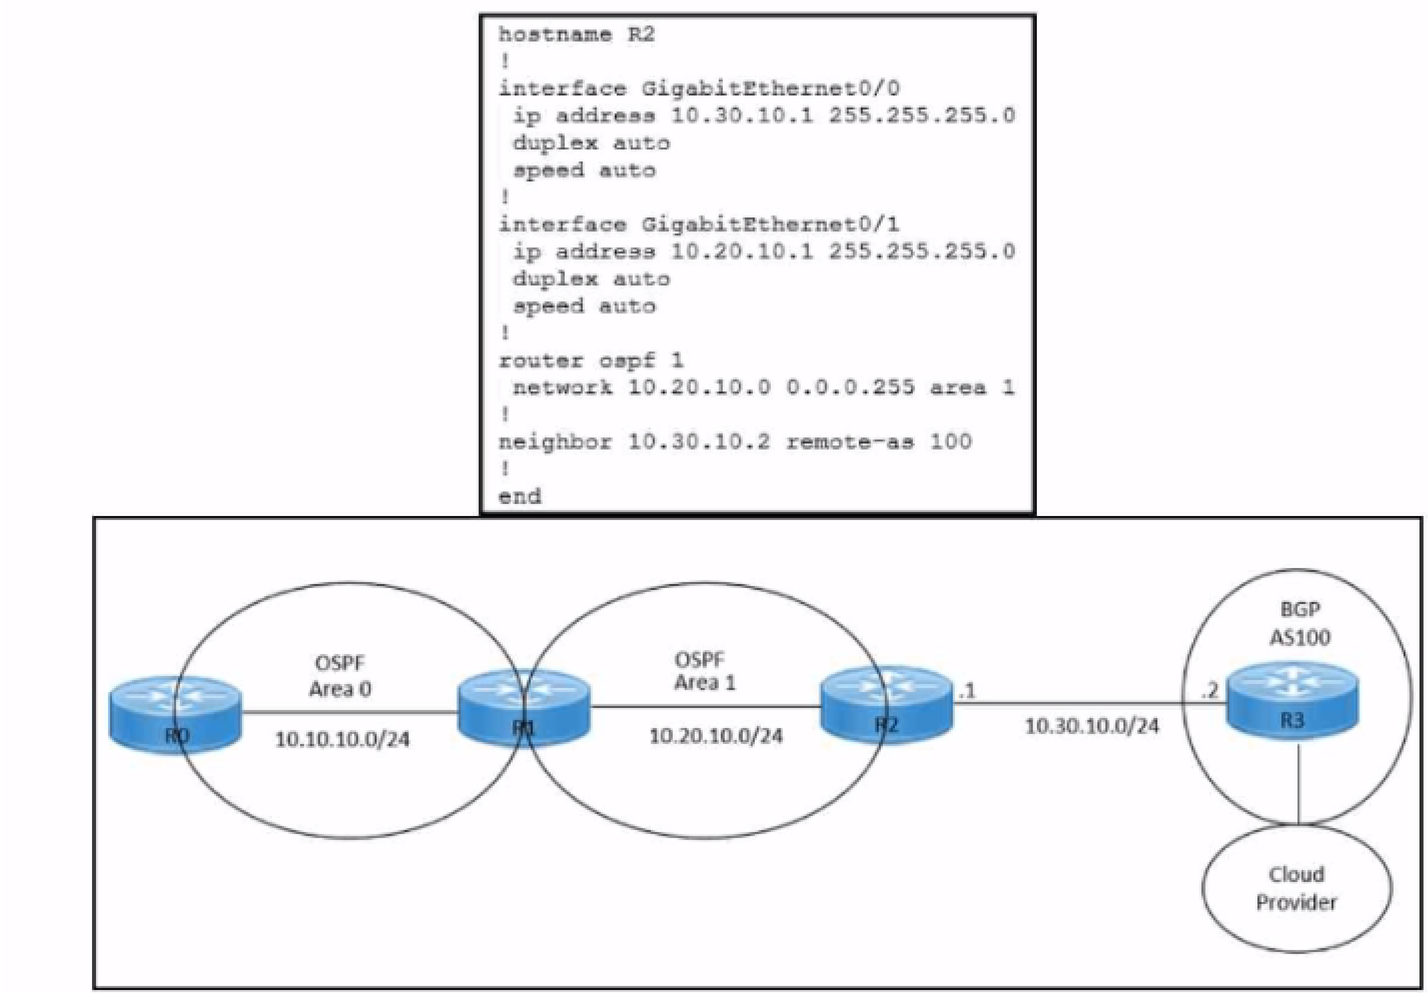

```
hostname R2
!
interface GigabitEthernet0/0
 ip address 10.30.10.1 255.255.255.0
 duplex auto
 speed auto
!
interface GigabitEthernet0/1
 ip address 10.20.10.1 255.255.255.0
 duplex auto
 speed auto
!
router ospf 1
 network 10.20.10.0 0.0.0.255 area 1
!
neighbor 10.30.10.2 remote-as 100
!
end
```

Refer to the exhibits. An engineer must redistribute IBGP routes into OSPF to connect an on-premises network to a cloud provider. Which command must be configured on router R2?

Options:

A- redistribute ospf 1

B- redistribute bgp 100 ospf 1

C- redistribute bgp 100 subnets

D- bgp redistrlbute-lnternal

Answer:

B

Explanation:

This command redistributes the routes learned from BGP AS100 into OSPF Area 1, which allows router R2 to advertise those routes to router R1 and connect the on-premises network to the cloud provider.The other options are incorrect because they either redistribute the wrong routes or use the wrong syntax5.

I hope this helps you understand the question and the answer. If you have any other questions or requests, please let me know. I am always happy to help.

Question 9

Question Type: **MultipleChoice**

What is the role of service providers to establish private connectivity between on-premises networks and Google Cloud resources?

Options:

A- facilitate direct, dedicated network connections through Google Cloud Interconnect

B- enable intelligent routing and dynamic path selection using software-defined networking

C- provide end-to-end encryption for data transmission using native IPsec

D- accelerate content delivery through integration with Google Cloud CDN

Answer:

A

Explanation:

The role of service providers to establish private connectivity between on-premises networks and Google Cloud resources is to facilitate direct, dedicated network connections through Google Cloud Interconnect. Google Cloud Interconnect is a service that allows customers to connect their on-premises networks to Google Cloud through a service provider partner. This provides low latency, high bandwidth, and secure connectivity to Google Cloud services, such as Google Compute Engine, Google Cloud Storage, and Google BigQuery. Google Cloud Interconnect also supports hybrid cloud scenarios, such as extending on-premises networks to Google Cloud regions, or connecting multiple Google Cloud regions together. Google Cloud Interconnect offers two types of connections: Dedicated Interconnect and Partner Interconnect. Dedicated Interconnect provides physical connections between the customer's network and Google's network at a Google Cloud Interconnect location. Partner Interconnect provides virtual connections between the customer's network and Google's network through a supported service provider partner. Both types of connections use VLAN attachments to establish private connectivity to Google Cloud Virtual Private Cloud (VPC) networks.Reference:

Designing and Implementing Cloud Connectivity (ENCC) v1.0

[Google Cloud Interconnect Overview]

[Google Cloud Interconnect Documentation]

Question 10

Question Type: **MultipleChoice**

Refer to the exhibits.

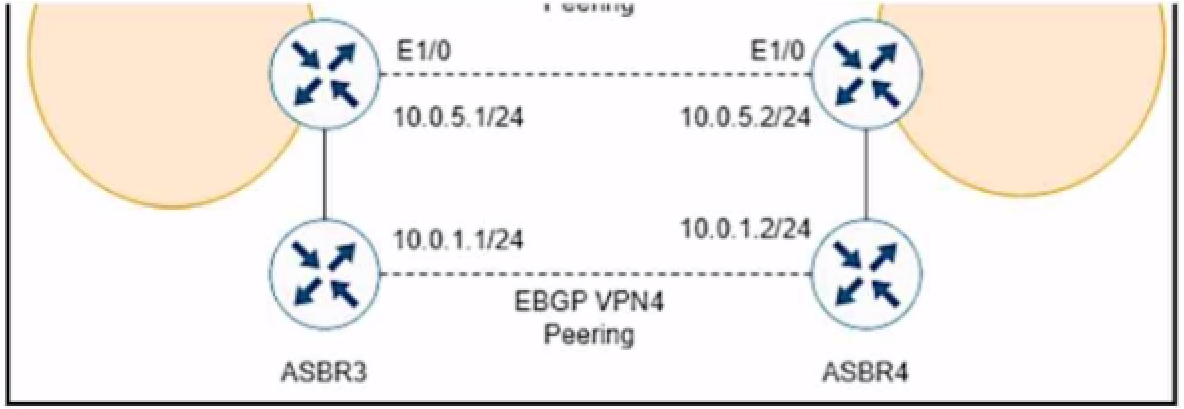

While troubleshooting, a network engineer discovers that the backup path fails between ASBR3 and ASBR4 for traffic between BGP AS6000 and BGP AS6500 when the connection between ASBR1 and ASBR2 goes down. The following configurations were performed on ASBR1:

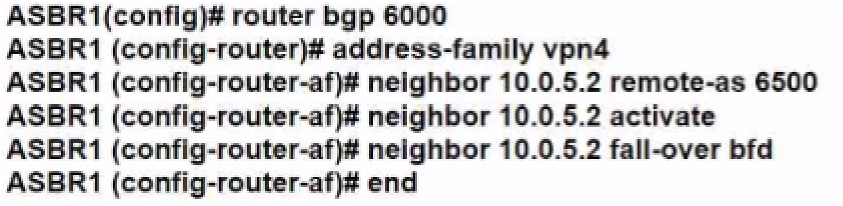

```
ASBR1(config)# router bgp 6000
ASBR1 (config-router)# address-family vpn4
ASBR1 (config-router-af)# neighbor 10.0.5.2 remote-as 6500
ASBR1 (config-router-af)# neighbor 10.0.5.2 activate
ASBR1 (config-router-af)# neighbor 10.0.5.2 fall-over bfd
ASBR1 (config-router-af)# end
```

Which command is missing?

Options:

A- bgp additional-paths Install

B- bgp additional-paths select

C- redistribute static

D- bgp advertise-best-external

Answer:

D

Explanation:

The bgp advertise-best-external command is used to enable the advertisement of the best external path to internal BGP peers. This command is useful when there are multiple exit points from the local AS to other ASes, and the local AS wants to use the closest exit point for each destination. By default, BGP only advertises the best path to its peers, and the best path is usually the one with the lowest IGP metric to the next hop. However, this may not be the optimal path for traffic leaving the local AS, as it may result in suboptimal hot-potato routing or MED oscillations. The bgp advertise-best-external command allows BGP to advertise the best external path, which is the path with the lowest MED among the paths from different neighboring ASes, in addition to the best path. This way, the internal BGP peers can choose the best exit point based on the MED value, rather than the IGP metric. In this scenario, ASBR1 is configured to receive additional paths from ASBR2, which is a route reflector. ASBR2 receives two paths for the same prefix from AS6500, one from ASBR3 and one from ASBR4. ASBR2 selects the best path based on the IGP metric to the next hop, and advertises it to ASBR1. However, this path may not be the best external path, as it may have a higher MED value than the other path. If the connection between ASBR1 and ASBR2 goes down, ASBR1 will not have any backup path to reach AS6500, as it does not know the other path from ASBR4. To prevent this situation, ASBR1 should be configured with the bgp advertise-best-external command, so that it can receive the best external path from ASBR2, along with the best path. This way, ASBR1 will have a backup path to reach AS6500, in case the primary path fails.Reference:=IP Routing: BGP Configuration Guide - BGP Additional Paths ... - Cisco,BGP Additional Paths