Cisco

Exam Questions 300-440

Designing and Implementing Cloud Connectivity (ENCC)

NEW QUESTION 1
DRAG DROP

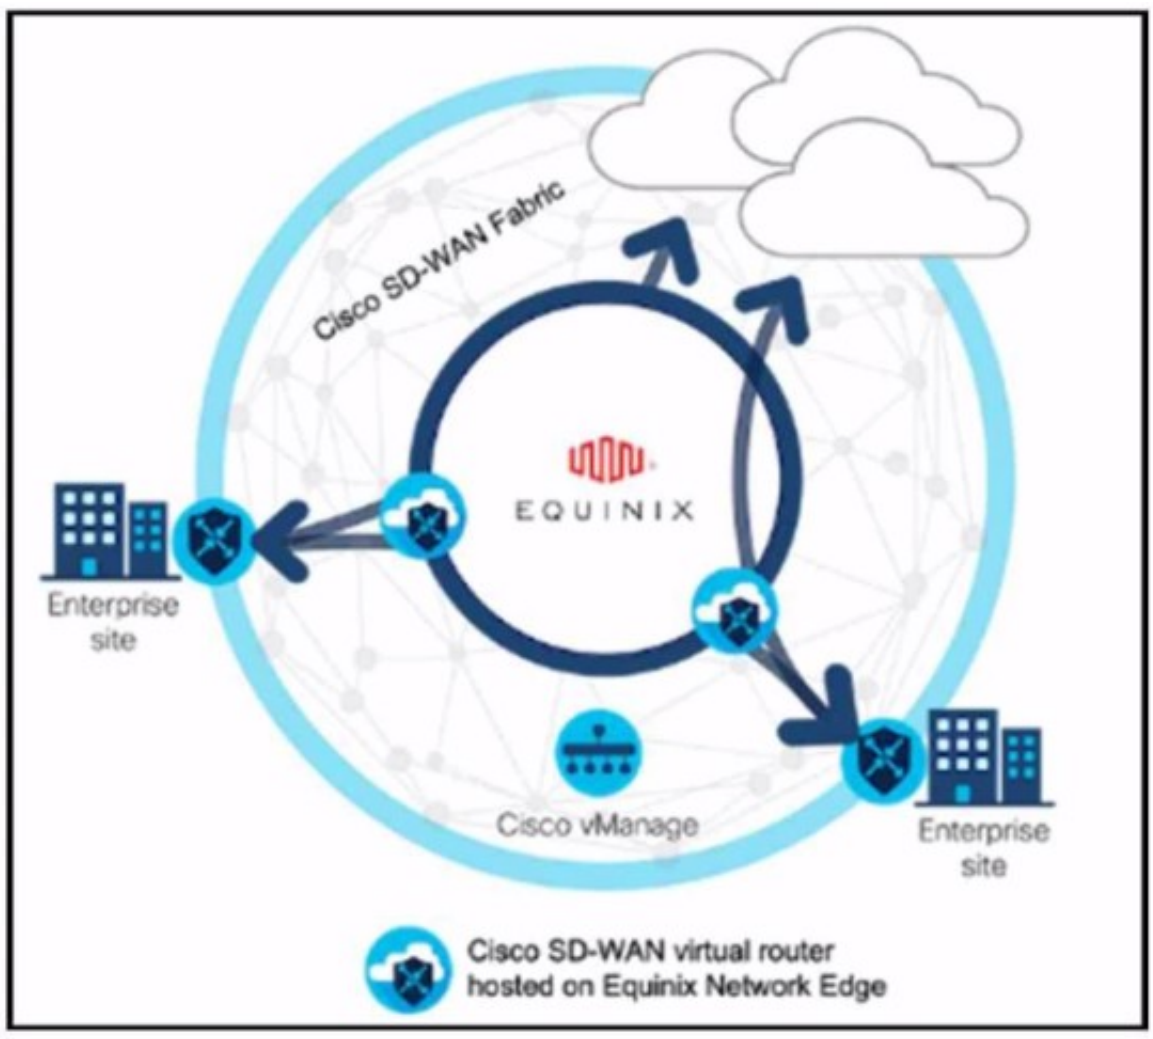

Refer to the exhibit. These configurations are complete:
• Create an account in the Equinix portal.
• Associate the Equinix account with Cisco vManage.
• Configure the global settings for Interconnect Gateways.
Drag the prerequisite steps from the left onto the order on the right to configure a Cisco SD-WAN Cloud Interconnect with Equinix

Attach Cisco SD-WAN Virtual Edge to the Equinix device template.	Step 1
Create the necessary network segments.	Step 2
Ensure that you have UUIDs for the required number of Cisco SD-WAN Virtual Edge instances that you want to deploy as Interconnect Gateways.	Step 3
Create the Interconnect Gateway at the Equinix location that is closest to your SD-WAN branch location.	Step 4

A. Mastered
B. Not Mastered

Answer: A

Explanation:
The process of configuring a Cisco SD-WAN Cloud Interconnect with Equinix involves several steps.
? Ensure that you have UUIDs for the required number of Cisco SD WAN Virtual Edge instances that you want to deploy as Interconnect Gateways: This is the first step where you ensure that you have the necessary UUIDs for the Cisco SD-WAN Virtual Edge instances that you want to deploy.
? Create the necessary network segments: After ensuring the availability of UUIDs, you create the necessary network segments.
? Attach Cisco SD-WAN Virtual Edge to the Equinix device template: After setting up the network segments, you attach the Cisco SD-WAN Virtual Edge to the Equinix device template.
? Create the Interconnect Gateway at the Equinix location that is closest to your SD-
WAN branch location: Finally, you create the Interconnect Gateway at the Equinix location that is closest to your SD-WAN branch location.
References :=
? [Cisco SD-WAN Cloud Interconnect with Equinix]
? [Cisco SD-WAN Cloud OnRamp for CoLocation Deployment Guide]

NEW QUESTION 2
DRAG DROP
An engineer must edit the settings of a site-to-site IPsec VPN connection between an on- premises Cisco IOS XE router and Amazon Web Services (AWS). IPsec must be configured to support multiple peers and failover after 120 seconds of idle time on the first entry of the crypto map named Cisco. Drag and drop the

commands from the left onto the order on the right.

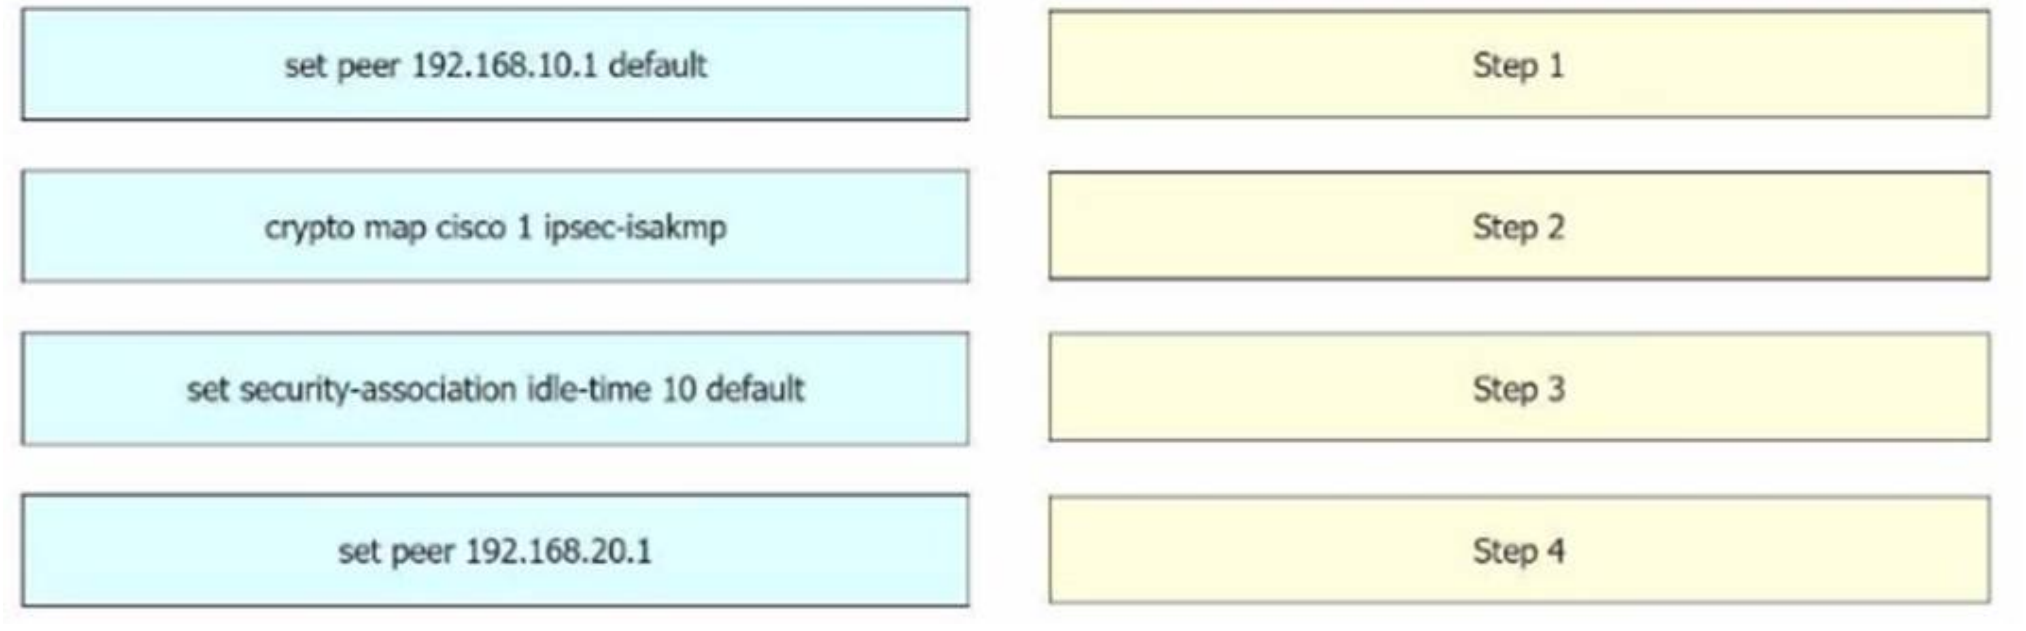

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1 = crypto map cisco 1 ipsec-isakmp Step 2 = set peer 192.168.10.1 default Step 3 = set peer 192.168.20.1 Step 4 = set security-association idle-time 120 default
The process of editing the settings of a site-to-site IPsec VPN connection between an on- premises Cisco IOS XE router and Amazon Web Services (AWS), and configuring IPsec to support multiple peers and failover after 120 seconds of idle time on the first entry of the crypto map named Cisco involves several steps123456.
? crypto map cisco 1 ipsec-isakmp: This command is used to create a new entry in the crypto map named ??cisco??. The ??1?? is the sequence number of the entry, and ??ipsec-isakmp?? specifies that the IPSec security associations (SAs) should be established using the Internet Key Exchange (IKE) protocol13.
? set peer 192.168.10.1 default: This command is used to specify the IP address of the default peer for the crypto map entry. In this case, the default peer is at IP address 192.168.10.115.
? set peer 192.168.20.1: This command is used to add an additional peer to the crypto map entry. In this case, the additional peer is at IP address 192.168.20.1. This allows the IPsec VPN to support multiple peers56.
? set security-association idle-time 120 default: This command is used to set the idle time for the security association. If no traffic is detected over the VPN for the specified idle time (in this case, 120 seconds), the security association is deleted, and the VPN connection fails over to the next peer46.
References :=
? Configure a Site-to-Site IPSec IKEv1 Tunnel Between an ASA and a Cisco IOS Router - Cisco
? Configure IOS-XE Site-to-Site VPN Connection to Amazon Web Services - Cisco Community
? Configuring Site to Site IPSec VPN Tunnel Between Cisco Routers
? Configure Failover for IPSec Site-to-Site Tunnels with Backup ISP Links on FTD Managed by FMC - Cisco
? Does Setting Multiple Peers in a Crypto Map Also Support Parallel IPSec Connections - Cisco Community
? Multiple WAN Connections — IPsec in Multi-WAN Environments | pfSense Documentation
? Multiple Set Peer for VPN Failover - Server Fault

NEW QUESTION 3
Refer to the exhibit.

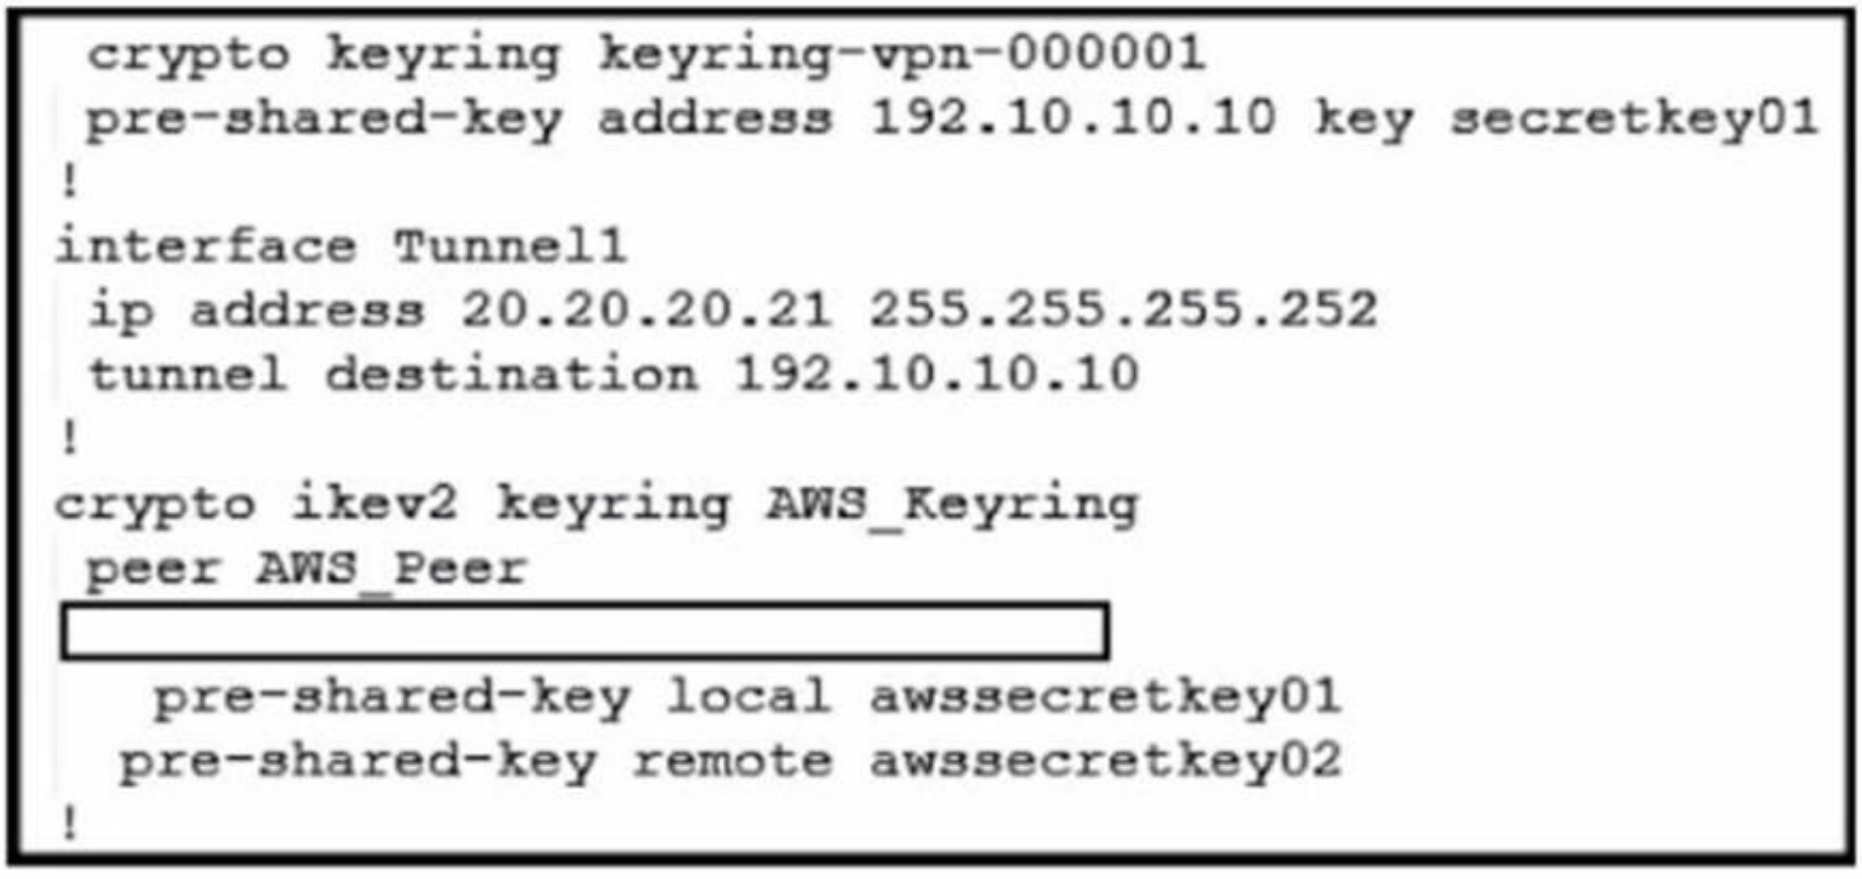

An engineer needs to configure a site-to-site IPsec VPN connection between an on-premises Cisco IOS XE router and Amazon Web Services (AWS). Which configuration command must be placed in the blank in the code to complete the tunnel configuration?

A. address 20.20.20.21

B. address 192.10.10.10
C. tunnel source 20.20.20.21
D. tunnel source 192.10.10.10

Answer: C

Explanation:
In the given scenario, an engineer is configuring a site-to-site IPsec VPN connection between an on-premises Cisco IOS XE router and AWS. The correct command to complete the tunnel configuration is ??tunnel source 20.20.20.21??. This command specifies the source IP address for the tunnel, which is essential for establishing a secure connection between two endpoints over the internet or another network1. References:
? Configure IOS-XE Site-to-Site VPN Connection to Amazon Web Services - Cisco
Community
? [Security for VPNs with IPsec Configuration Guide, Cisco IOS XE Release 3S - Config

NEW QUESTION 4
DRAG DROP
An engineer must configure cloud connectivity with Cisco Umbrella Secure Internet Gateway (SIG) in active/backup mode. The engineer already configured the SIG Credentials and SIG Feature Templates. Drag and drop the steps from the left onto the order on the right to complete the configuration.

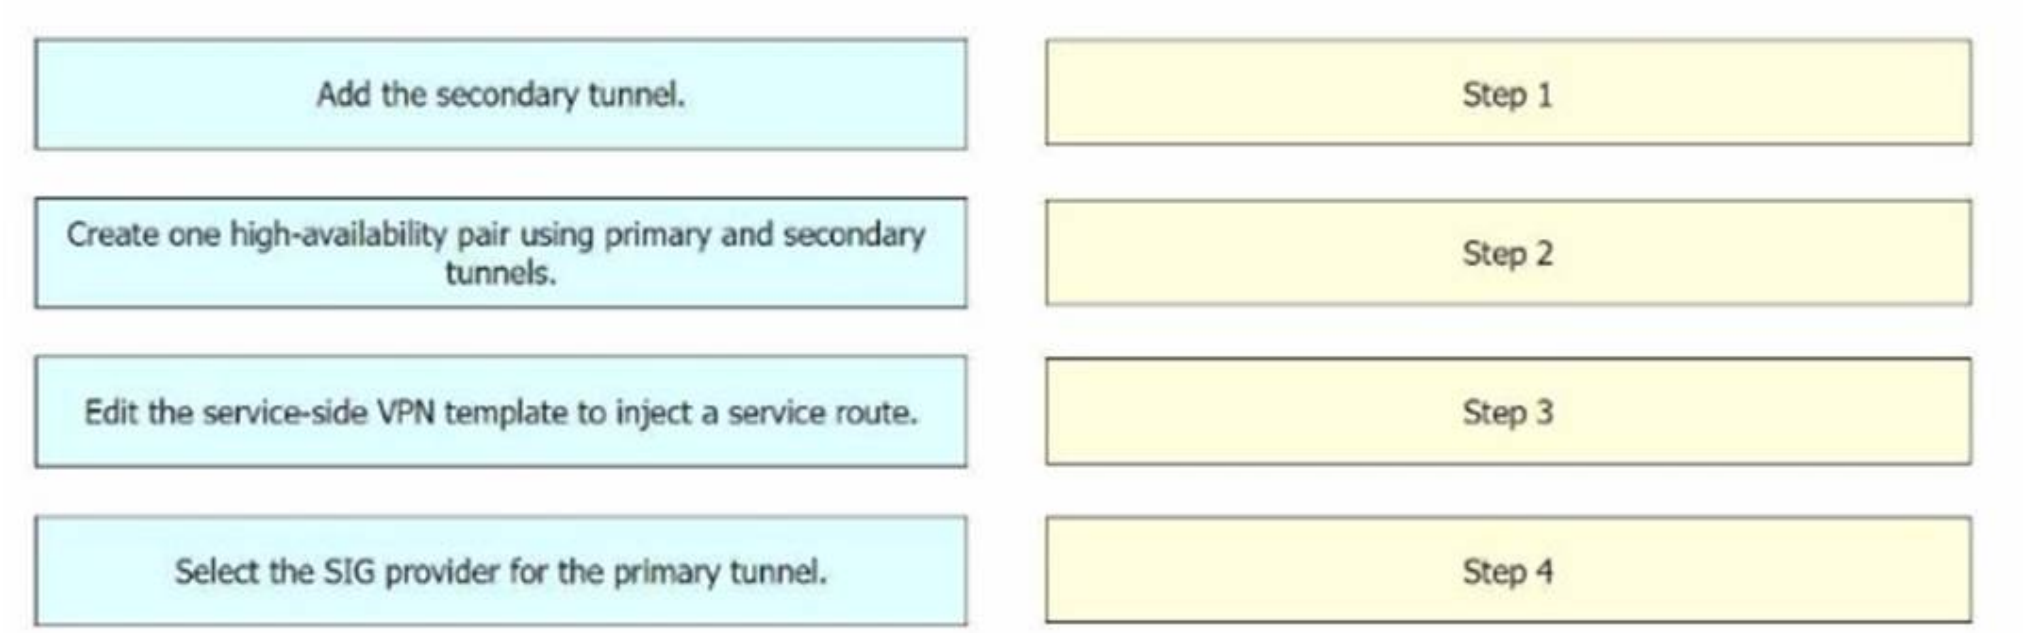

A. Mastered
B. Not Mastered

Answer: A

Explanation:
The configuration of cloud connectivity with Cisco Umbrella Secure Internet Gateway (SIG) in active/backup mode involves several steps. After configuring the SIG Credentials and SIG Feature Templates, the engineer must:
? Select the SIG provider for the primary tunnel: This is the first step in setting up the active/backup mode. The primary tunnel is the main connection path for the cloud connectivity.
? Add the secondary tunnel: The secondary tunnel serves as a backup in case the primary tunnel fails. It ensures that the cloud connectivity remains uninterrupted even if there are issues with the primary tunnel.
? Create one high-availability pair using primary and secondary tunnels: This step involves pairing the primary and secondary tunnels to create a high-availability pair. Thisensures that the cloud connectivity will switch over to the secondary tunnel seamlessly if the primary tunnel fails.
? Edit the service-side VPN template to inject a service route: The final step involves modifying the VPN template on the service side to include a service route. This ensures that the traffic is correctly routed through the primary or secondary tunnel as needed.
References :=
? Designing and Implementing Cloud Connectivity (ENCC) v1.01
? Learning Plan: Designing and Implementing Cloud Connectivity v1.0 (ENCC 300-440) Exam Prep2
? Configure Umbrella SIG Tunnels for Active/Backup or Active/Active Scenarios - Cisco3

NEW QUESTION 5
DRAG DROP
An engineer must configure a site-to-site IPsec VPN connection between an on-premises Cisco IOS XE router In Controller mode and AWS. The IKE version must be changed from IKEv1to IKEv2 in Cisco vManage. Drag and drop the steps from the left onto the order on the right to complete the configuration.

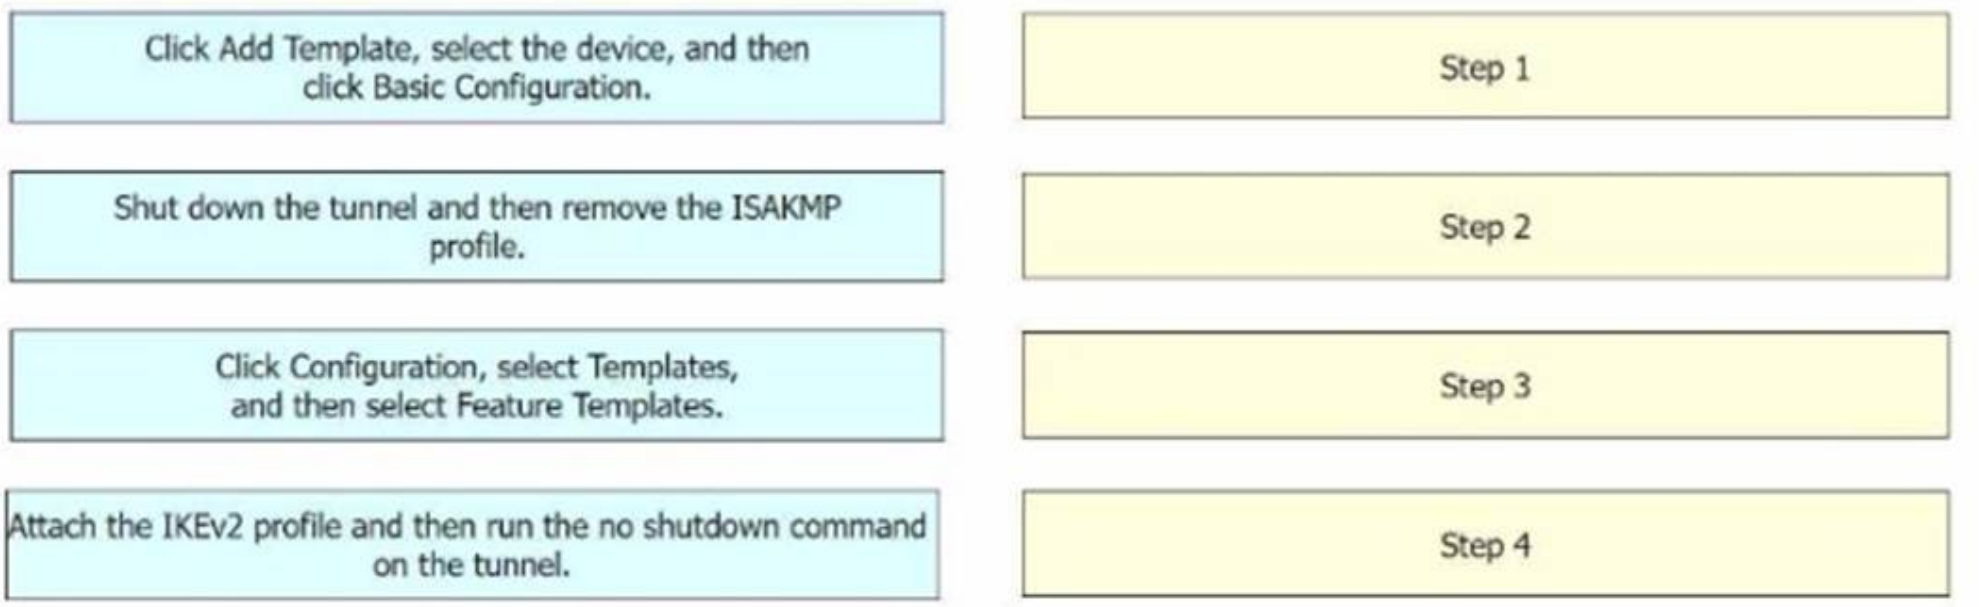

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1 = Click Configuration, select Templates, and then select Feature Templates. Step 2 = Click Add Template, select the device, and then click Basic Configuration. Step 3 = Shut down the tunnel and then remove the ISAKMP profile. Step 4 = Attach the IKEv2 profile and then run the no shutdown command on the tunnel.
The process of configuring a site-to-site IPsec VPN connection between an on-premises Cisco IOS XE router in Controller mode and AWS, and changing the IKE version from IKEv1 to IKEv2 in Cisco vManage involves several steps123.
? Click Configuration, select Templates, and then select Feature Templates: This is the first step where you navigate to the Templates section in the Configuration menu of Cisco vManage1.
? Click Add Template, select the device, and then click Basic Configuration: In this step, you add a new template for the device and proceed with the basic configuration1.
? Shut down the tunnel and then remove the ISAKMP profile: Before changing the IKE version, you need to shut down the existing tunnel and remove the ISAKMP profile that is configured for IKEv12.
? Attach the IKEv2 profile and then run the no shutdown command on the tunnel:
Finally, you attach the newly created IKEv2 profile to the tunnel and bring the tunnel back up2.
References :=
? Configuring Internet Key Exchange Version 2 (IKEv2) - Cisco
? Switch from IKEv1 to IKEv2 on Cisco Routers - Cisco Community
? Configure IOS-XE Site-to-Site VPN Connection to Amazon Web Services - Cisco Community

NEW QUESTION 6
Which method is used to create authorization boundary diagrams (ABDs)?

A. identify only interconnected systems that are FedRAMP-authorized
B. show all networks in CIDR notation only
C. identify all tools as either external or internal to the boundary
D. show only minor or small upgrade level software components

Answer: C

Explanation:
According to the FedRAMP Authorization Boundary Guidance document1, the method used to create authorization boundary diagrams (ABDs) is to identify all tools as either external orinternal to the boundary. The ABD is a visual representation of the components that make up the authorization boundary, which includes all technologies, external and internal services, and leveraged systems and accounts for all federal information, data, and metadata that a Cloud Service Offering (CSO) is responsible for. The ABD should illustrate a CSP??s scope of control over the system and show components or services that are leveraged from external services or controlled by the customer1. The other options are incorrect because they do not capture the full scope and details of the authorization boundary as required by FedRAMP. References := FedRAMP Authorization Boundary Guidance document1

NEW QUESTION 7
Which architecture model establishes internet-based connectivity between on-premises networks and AWS cloud resources?

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 CloudFrontfor 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 the VPN CloudHub model12. 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.
References:
? 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?

NEW QUESTION 8
Refer to the exhibit.

```
vedge1# show policy from-vsmart
apply-policy
 site-list site1
  control-policy prefer_local out
!
policy
 lists
  site-list site1
   site-id 100
  tloc-list prefer_site1
    tloc 10.1.1.1 color mpls encap ipsec preference 100
  control-policy prefer_local
   sequence 10
    match route
     site-list site1
   !
   action accept
    set
      tloc-list prefer_site1
```

A network engineer discovers that the policy that is configured on an on-premises Cisco WAN edge router affects only the route tables of the specific devices that are listed in the site list. What is the problem?

A. An inbound policy must be applied.
B. The action must be set to deny
C. A localized data policy must be configured.
D. A centralized data policy must be configured

Answer: D

Explanation:
A centralized data policy is a policy that is applied to all devices in the overlay network, regardless of the site list. A localized data policy is a policy that is applied only to the devices that are listed in the site list. In this case, the network engineer wants to apply the policy to all devices in the overlay network, not just the specific devices in the site list. Therefore, a centralized data policy must be configured on the on-premises Cisco WAN edge router. References :=
? Designing and Implementing Cloud Connectivity (ENCC) v1.0, Module 3:
Implementing Cloud Connectivity, Lesson 3: Implementing Cisco SD-WAN Cloud OnRamp for Colocation, Topic: Centralized Data Policy
? [Cisco SD-WAN Cloud OnRamp for Colocation Deployment Guide], Chapter:
Configuring Centralized Data Policy

NEW QUESTION 9
DRAG DROP
Drag and drop the commands from the left onto the purposes on the right to identify issues on a Cisco IOS XE SD-WAN device.

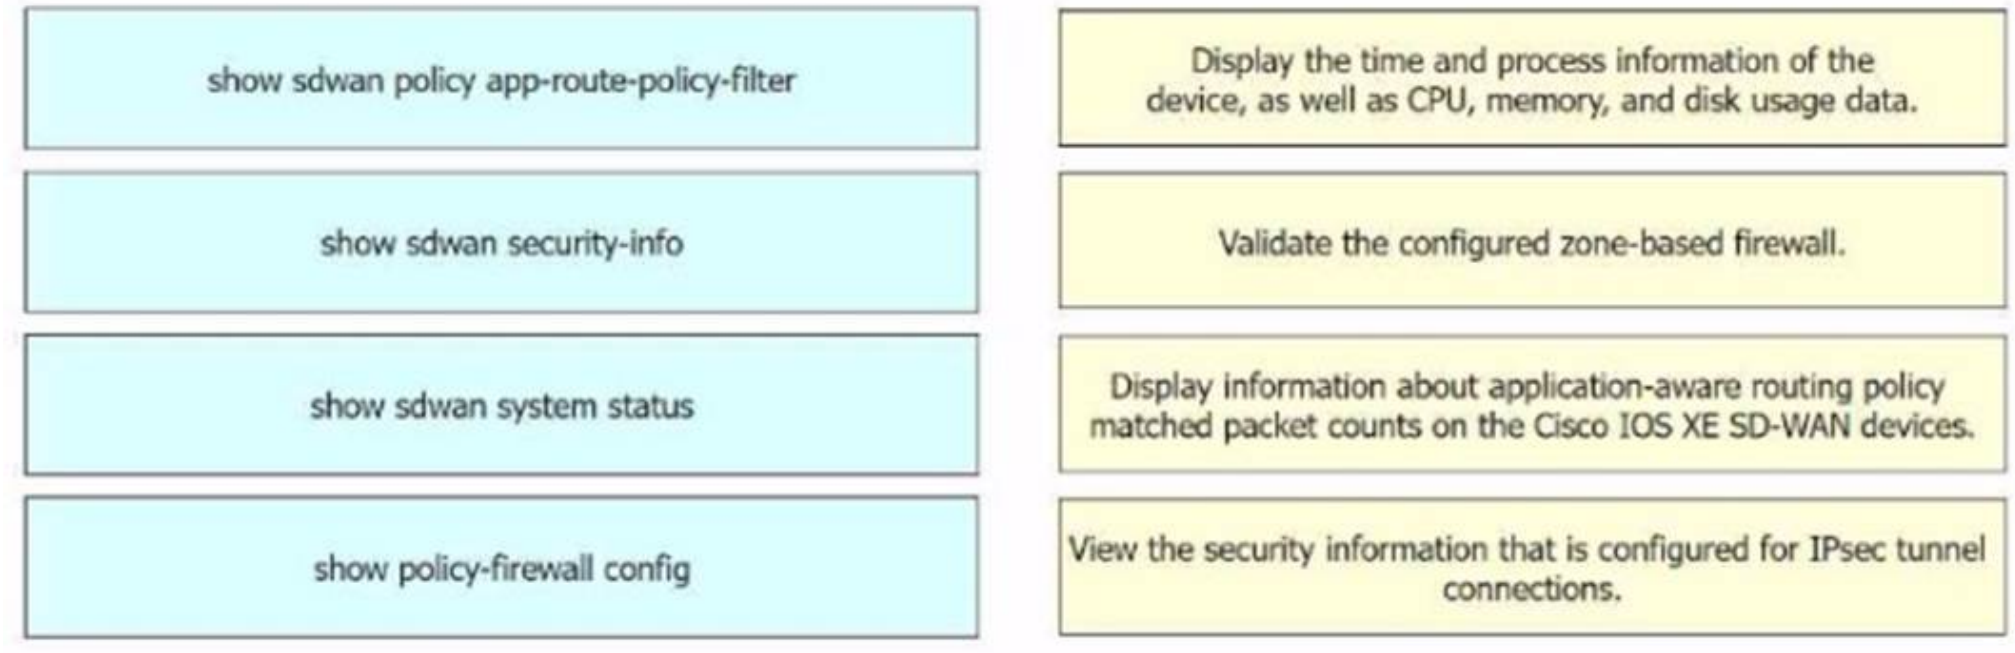

A. Mastered
B. Not Mastered

Answer: A

Explanation:
? Display the time and process information of the device, as well as CPU, memory, and disk usage data. = show sdwan system status1
? Validate the configured zone-based firewall. = show policy-firewall config1
? Display information about application-aware routing policy matched packet counts on the Cisco IOS XE SD-WAN devices. = show sdwan policy app-route-policy-filter1
? View the security information that is configured for IPsec tunnel connections. = show sdwan security-info
The commands used to identify issues on a Cisco IOS XE SD-WAN device are as follows1:
? show sdwan system status: This command is used to display the time and process information of the device, as well as CPU, memory, and disk usage data1.
? show policy-firewall config: This command is used to validate the configured zone- based firewall1.
? show sdwan policy app-route-policy-filter: This command is used to display information about application-aware routing policy matched packet counts on the Cisco IOS XE SD-WAN devices1.
? show sdwan security-info: This command is used to view the security information that is configured for IPsec tunnel connections1.
References :=
? Cisco IOS XE Catalyst SD-WAN Qualified Command Reference
? Cisco Catalyst SD-WAN Command Reference
? Cisco Catalyst SD-WAN Systems and Interfaces Configuration Guide, Cisco IOS XE
? SD-WAN Tunnel Interface Commands - Cisco

NEW QUESTION 10
What is the role of service providers to establish private connectivity between on-premises networks and Google Cloud resources?

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. References:
? Designing and Implementing Cloud Connectivity (ENCC) v1.0
? [Google Cloud Interconnect Overview]
? [Google Cloud Interconnect Documentation]

NEW QUESTION 10
Refer to the exhibit.

```
1-Aug-2021 20:12:11 EDT] Failed to apply policy - Failed to
process device request -
Error type : application
Error tag : operation-failed
Error Message : /apply-policy/site-list[name='All-Site']:
Overlapping apply-policy site-list Hub site id 200-299 with
site-list All-Site
Error info : <error-info>
<bad-element>site-list</bad-element>
</error-info>
```

A company uses Cisco SD-WAN in the data center. All devices have the default configuration. An engineer attempts to add a new centralized control policy in Cisco vManage but receives an error message. What is the problem?

A. A centralized control policy is already applied to the specific site ID and direction
B. The policy for "Hub" should be applied in the outbound direction, and the policy for "All- Site" should be applied inbound.
C. Apply an additional outbound control policy to override the site ID overlaps.
D. Site-list "All-Site" should be configured with a new match sequence that is lower than the sequence for site-list "Hub*.

Answer: D

Explanation:
The problem is that the site-list ??All-Site?? has a higher match sequence than the site-list ??Hub??, which means that the policy for ??All-Site?? will take precedence over the policy for ??Hub?? for any site that belongs to both lists. This creates a conflict and prevents the engineer from adding a new centralized control policy in Cisco vManage. To resolve this issue, the site-list ??All-Site?? should be configured with a new match sequence that is lower than the sequence for site-list ??Hub??, so that the policy for ??Hub?? will be applied first and then the policy for ??All-Site?? will be applied only to the remaining sites that are not in the ??Hub?? list. References :=
? Designing and Implementing Cloud Connectivity (ENCC, Track 1 of 5), Module 3:
Cisco SD-WAN Cloud OnRamp for Colocation, Lesson 3: Cisco SD-WAN Cloud OnRamp for Colocation - Centralized Control Policies
? Cisco SD-WAN Cloud OnRamp for Colocation Deployment Guide, Chapter 4:
Configuring Centralized Control Policies
? Cisco SD-WAN Configuration Guide, Release 20.3, Chapter: Centralized Policy Framework, Section: Policy Configuration Overview

NEW QUESTION 15
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
```

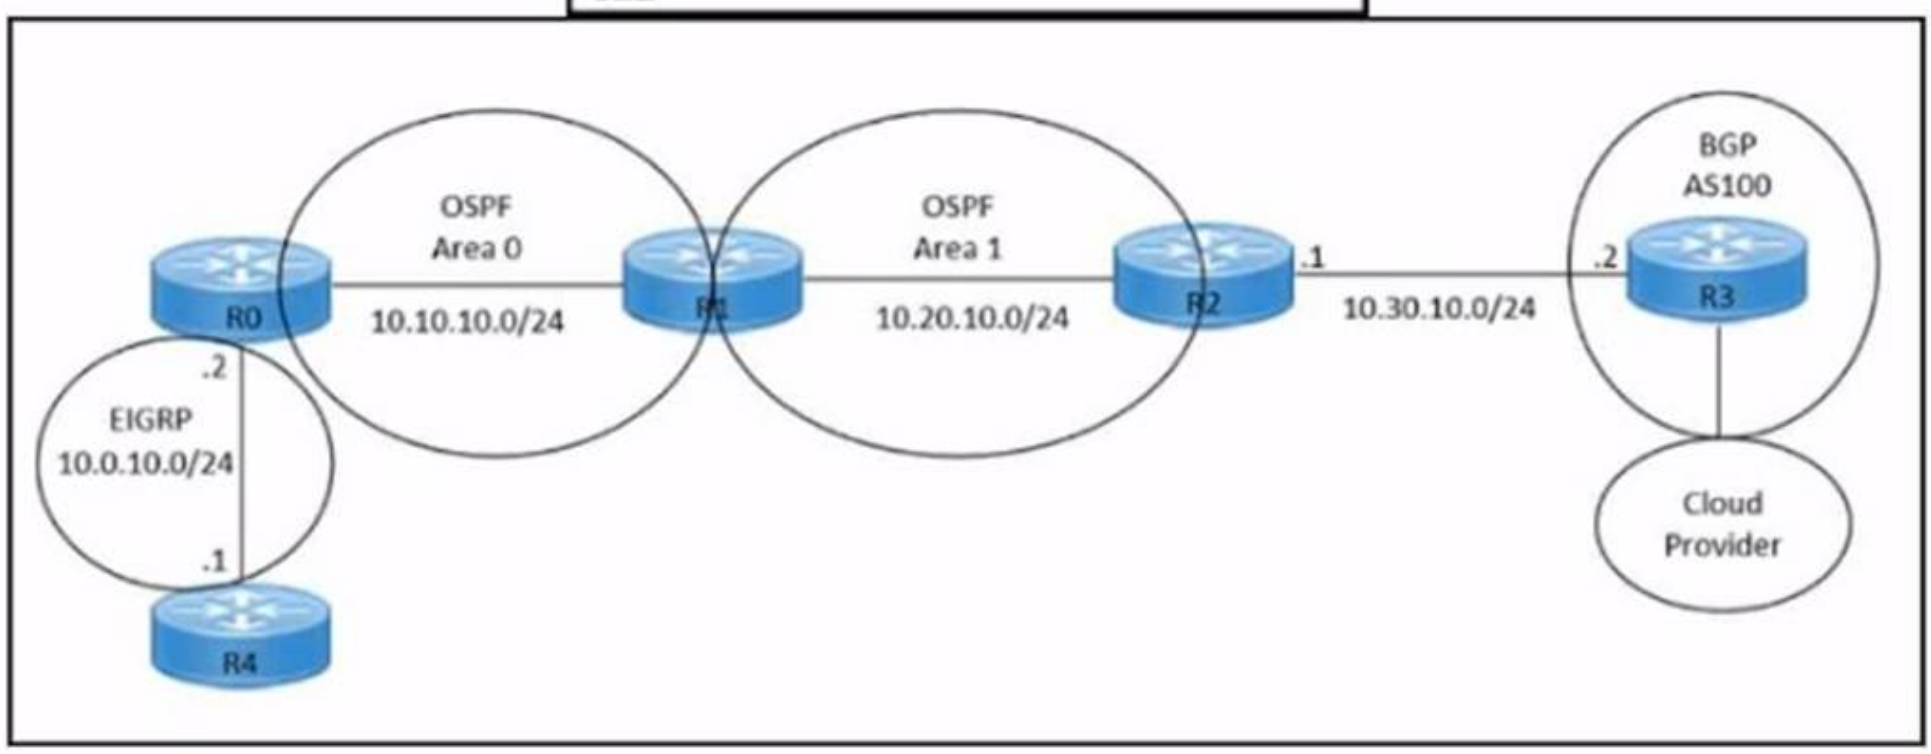

An engineer must redistribute OSPF internal routes into BGP to connect an on-premises network to a cloud provider without introducing extra routes. Which two commands must be configured on router R2? (Choose two.)

A. router ospf 1
B. router bgp 100
C. redistribute ospf 1
D. redistribute bgp 100
E. redistribute ospf 1 match internal external

Answer: BE

Explanation:
To redistribute OSPF internal routes into BGP, the engineer needs to configure two commands on router R2. The first command is router bgp 100, which enables BGP routing process and specifies the autonomous system number of 100. The second command is redistribute 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. References: = 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]

NEW QUESTION 20
Refer to the exhibit.

```
vEdge# show crypto isakmp sa

IPv4 Crypto ISAKMP SA
dst                  src                  state          conn-id    status
203.0.113.1          203.0.113.2          MM_KEY_EXCH    14526      Active
```

While troubleshooting an IPsec connection between a Cisco WAN edge router and an Amazon Web Services (AWS) endpoint, a network engineer observes that the security association status is active, but no traffic flows between the devices What is the problem?

A. wrong ISAKMP policy
B. identity mismatch
C. wrong encryption
D. IKE version mismatch

Answer: B

Explanation:
An identity mismatch occurs when the local and remote identities configured on the IPsec peers do not match. This can prevent the establishment of an IPsec

tunnel or cause traffic to be dropped by the IPsec policy. In this case, the network engineer should verify that the local and remote identities configured on the Cisco WAN edge router and the AWS endpoint match the values expected by each peer. The identities can be an IP address, a fully qualified domain name (FQDN), or a distinguished name (DN). The identities are exchanged during the IKE phase 1 negotiation and are used to authenticate the peers. If the identities do not match, the peers will reject the IKE proposal and the IPsec tunnel will not be established or will be torn down. References :=
? Configure IOS-XE Site-to-Site VPN Connection to Amazon Web Services, Topic:
Troubleshooting
? Designing and Implementing Cloud Connectivity (ENCC) v1.0, Module 3: Implementing Cloud Connectivity, Lesson 2: Implementing Cisco SD-WAN Cloud OnRamp for IaaS, Topic: Troubleshooting Cisco SD-WAN Cloud OnRamp for IaaS
? Cisco IOS Security Configuration Guide, Release 15M&T, Chapter: Configuring IPsec Network Security, Topic: Configuring IPsec Identity and Peer Addressing

NEW QUESTION 25
DRAG DROP
An engineer signs in to Cisco vManage and needs to configure a custom application with a Cisco SD-WAN centralized policy. Drag and drop the steps from the left onto the order on the right to complete the configuration.

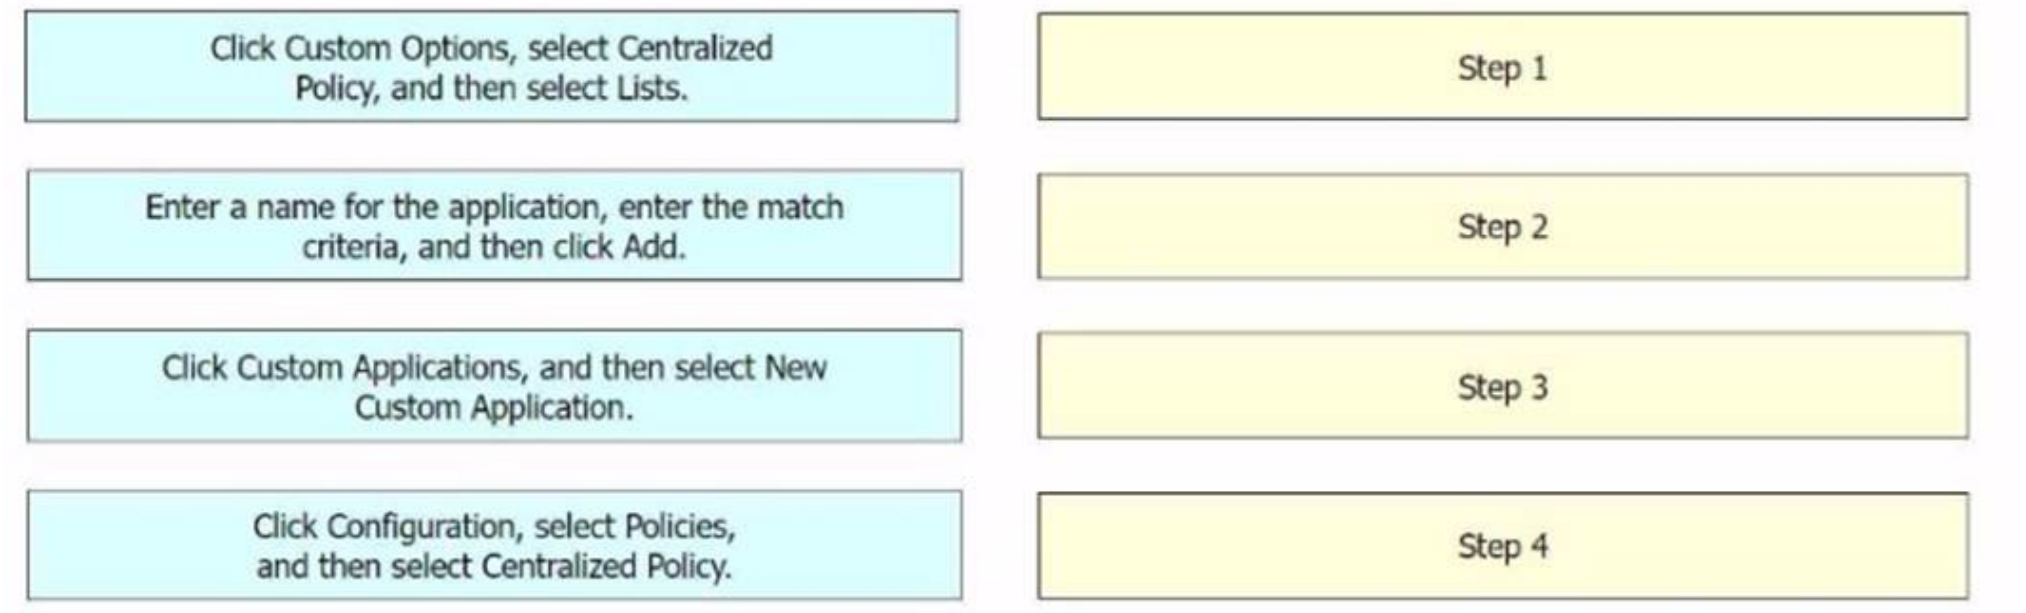

A. Mastered
B. Not Mastered

Answer: A

Explanation:
? To configure a custom application with Cisco SD-WAN centralized policy, you need to follow these steps25:
The process of configuring a custom application with a Cisco SD-WAN centralized policy using Cisco vManage involves several steps1.
? Click Configuration, select Policies, and then select Centralized Policy: This is the first step where you navigate to the Policies section in the Configuration menu of Cisco vManage1.
? Click Custom Options, select Centralized Policy, and then select Lists: In this step, you select the Custom Options, then select Centralized Policy, and finally select Lists1.
? Click Custom Applications, and then select New Custom Application: After setting up the Lists, you click on Custom Applications and then select New Custom Application1.
? Enter a name for the application, enter the match criteria, and then click Add:
Finally, you enter a name for the application, specify the match criteria, and then click Add to complete the configuration1.
References :=
? Cisco Catalyst SD-WAN Policies Configuration Guide, Cisco IOS XE

NEW QUESTION 26
DRAG DROP
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.

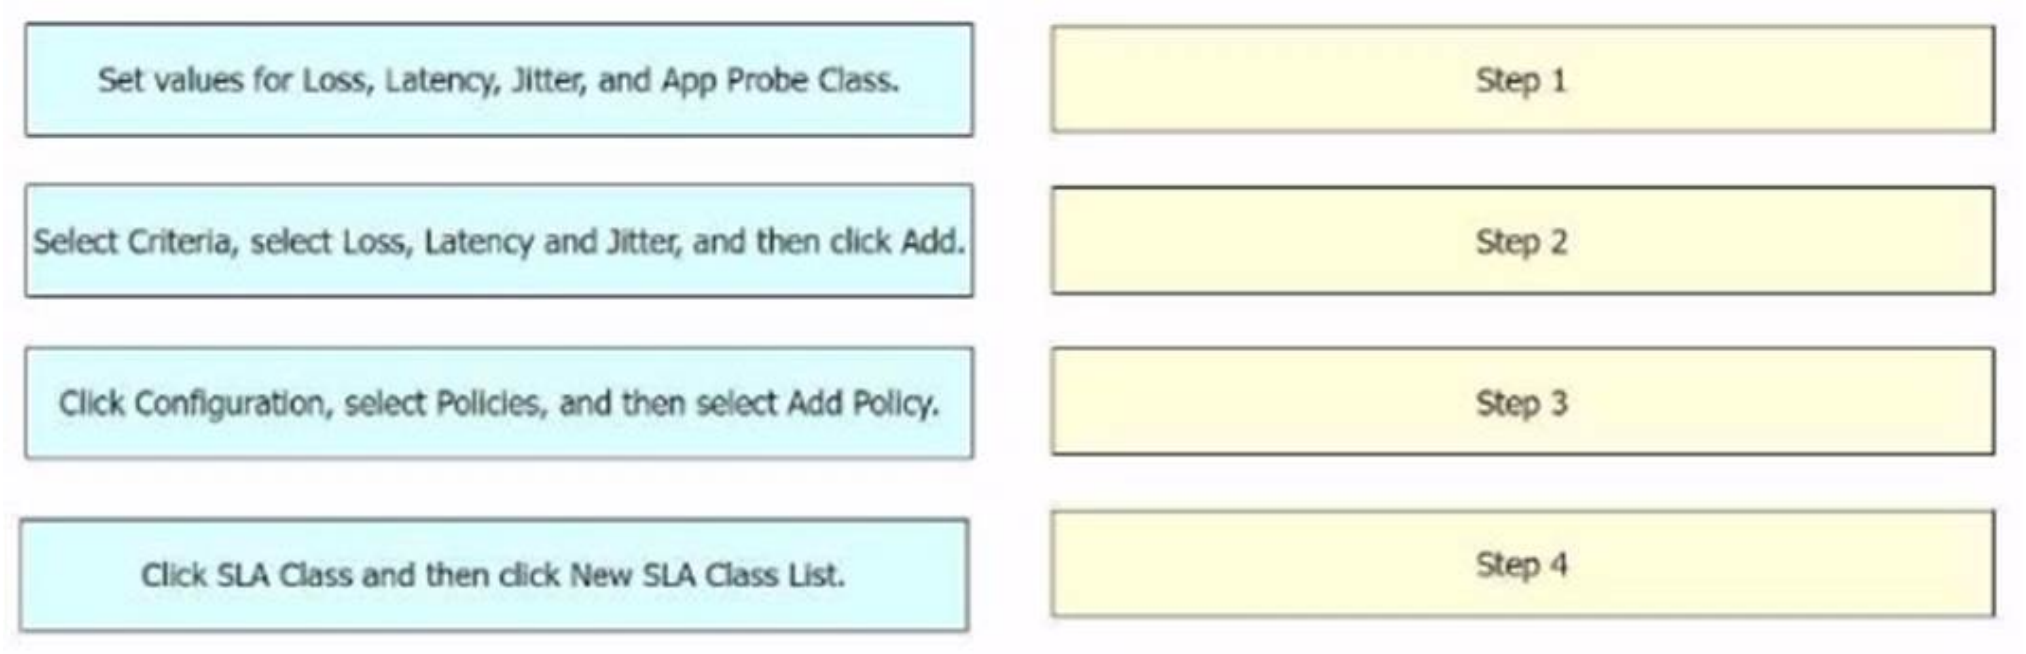

A. Mastered
B. Not Mastered

Answer: A

Explanation:
The process of configuring an SLA class to specify the maximum packet loss, packet latency, and jitter allowed on a connection using Cisco vManage involves several steps12.
? Click Configuration, select Policies, and then select Add Policy: This is the first step where you navigate to the Policies section in the Configuration menu of Cisco vManage1.
? Click SLA Class and then click New SLA Class List: In this step, you create a new SLA Class List1.
? Select Criteria, select Loss, Latency and Jitter, and then click Add: After setting up the SLA Class List, you select the criteria for the SLA class. In this case, the criteria are Loss, Latency, and Jitter1.
? Set values for Loss, Latency, Jitter, and App Probe Class: Finally, you set the values for Loss, Latency, Jitter, and App Probe Class1.
References :=
? Information About Application-Aware Routing - Cisco
? Policies Configuration Guide for vEdge Routers, Cisco SD-WAN Release 20

NEW QUESTION 29
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
!
router bgp 100
 neighbor 10.30.10.2 remote-as 100
 redistribute ospf 1
!
```

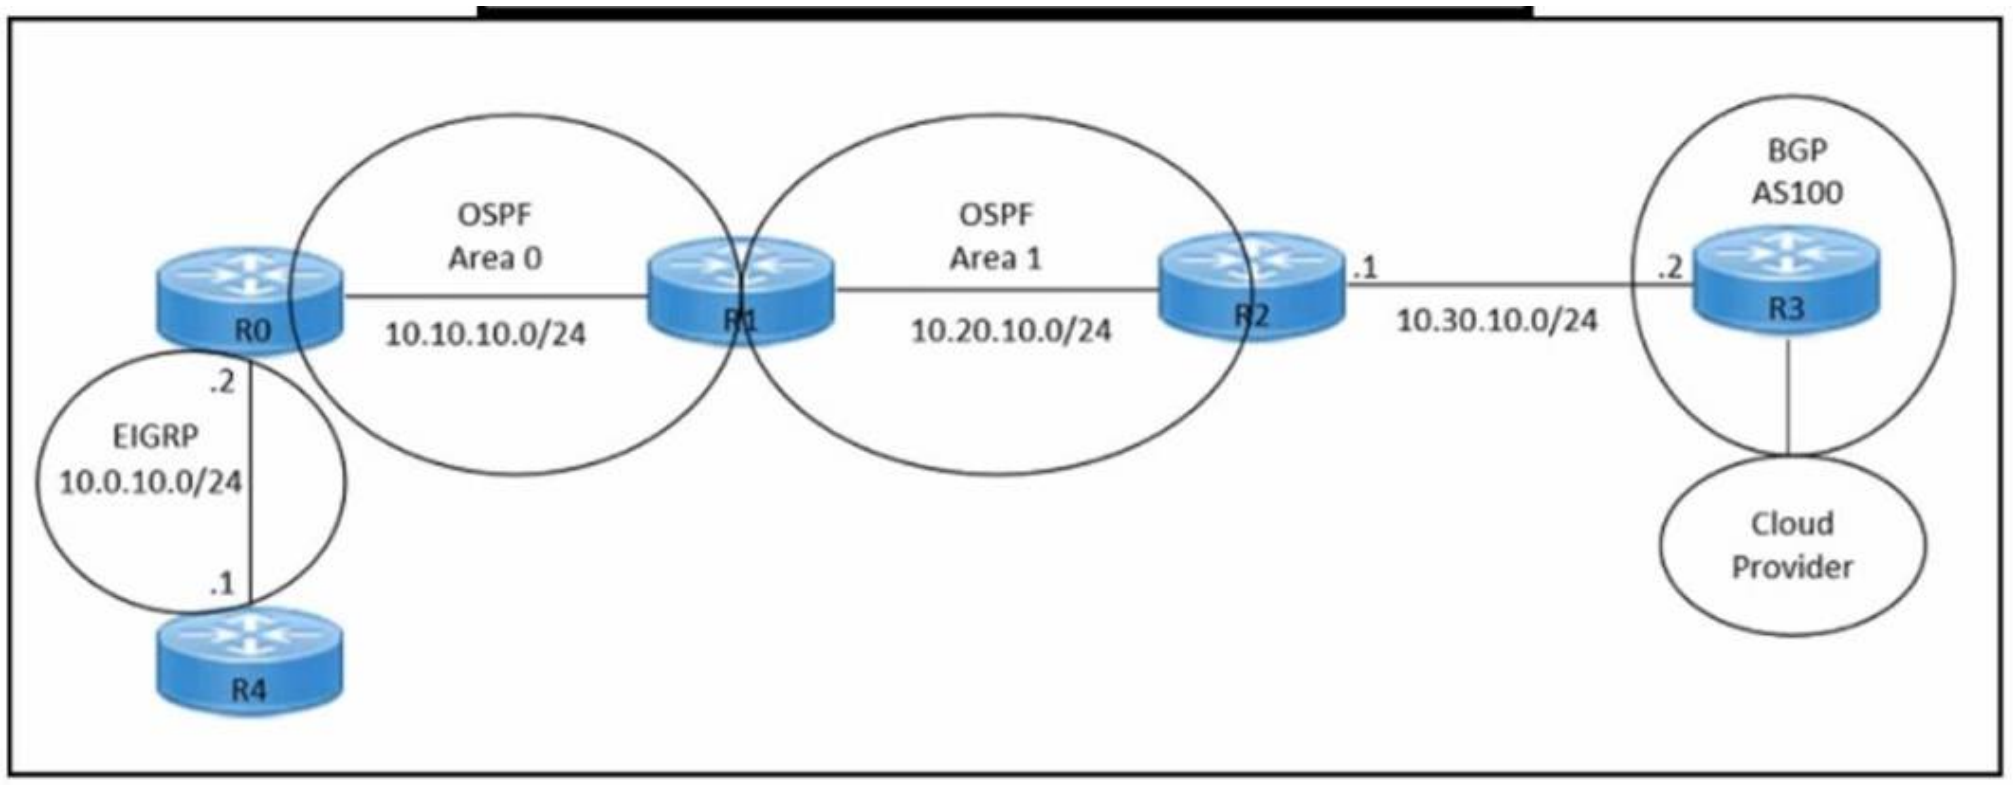

An engineer must redistribute only the 10.0.10.0/24 network into BGP to connect an on-premises network to a public cloud provider. These routes are currently redistributed:

*10.10.10.0/24
*10.20.10.0/24

Which command is missing on router R2?

A. neighbor 10.0.10.2 remote-as 100
B. redistribute ospf 1 match internal
C. redistribute ospf 1 match external
D. neighbor 10.0.10.0/24 remote-as 100

Answer: C

Explanation:
The command redistribute ospf 1 match external is missing on router R2. This command is needed to redistribute only the external OSPF routes into BGP. The external OSPF routes are those that are learned from another routing protocol or redistributed into OSPF. In this case, the 10.0.10.0/24 network is an external OSPF route, as it is redistributed from EIGRP into OSPF on router R1. The other commands are either already present or not relevant for this scenario.
References :=
? Designing and Implementing Cloud Connectivity (ENCC) v1.0, Module 3:
Implementing Cloud Connectivity, Lesson 3.1: Implementing IPsec VPN from Cisco IOS XE to AWS, Topic 3.1.2: Configure BGP on the Cisco IOS XE Router
? Security for VPNs with IPsec Configuration Guide, Cisco IOS XE, Chapter:
Configuring IPsec VPNs with Dynamic Routing Protocols, Section: Configuring BGP over IPsec VPNs

NEW QUESTION 32
Refer to the exhibit.

```
local7.debug: Mar 11 11:31:11 VEDGE-1 VDAEMON[1136]: vdaemon_disable_my_tloc[1308]:
%VDAEMON_DBG_EVENTS-1: Disabling tloc ge0_1.
local7.info: Mar 11 11:31:11 VEDGE-1 VDAEMON[1136]: %Viptela-VEDGE-1-vdaemon-6-INFO-1400002:
Notification:
 3/11/2023 11:31:11 control-connection-state-change severity-level:major host-name:"VEDGE-1"
 system-ip:10.10.10.1
 personality:vEdge peer-type:vmanage peer-system-ip:10.30.30.1 peer-vmanage-system-ip:0.0.0.0
 public-ip:20.20.20.20
 public-port:12947 src-color:biz-internet remote-color:public-internet uptime:"0:01:36:34" new-
 state:down
local7.info: Mar 11 11:31:11 VEDGE-1 FTMD[1126]: %Viptela-VEDGE-1-ftmd-6-INFO-1400002:
Notification:
 3/11/2023 11:31:11 bfd-state-change severity-level:major host-name:"VEDGE-1" system-
 ip:10.10.10.1 src-ip:20.20.30.2
dst-ip:20.20.30.20   proto:ipsec src-port:12406 dst-port:12347 local-system-ip:10.10.10.1 local-
color:"biz-internet"
  emote-system-ip:10.10.10.4 remote-color:"public-internet" new-state:down deleted:false flap-
  reason:bfd-deleted
```

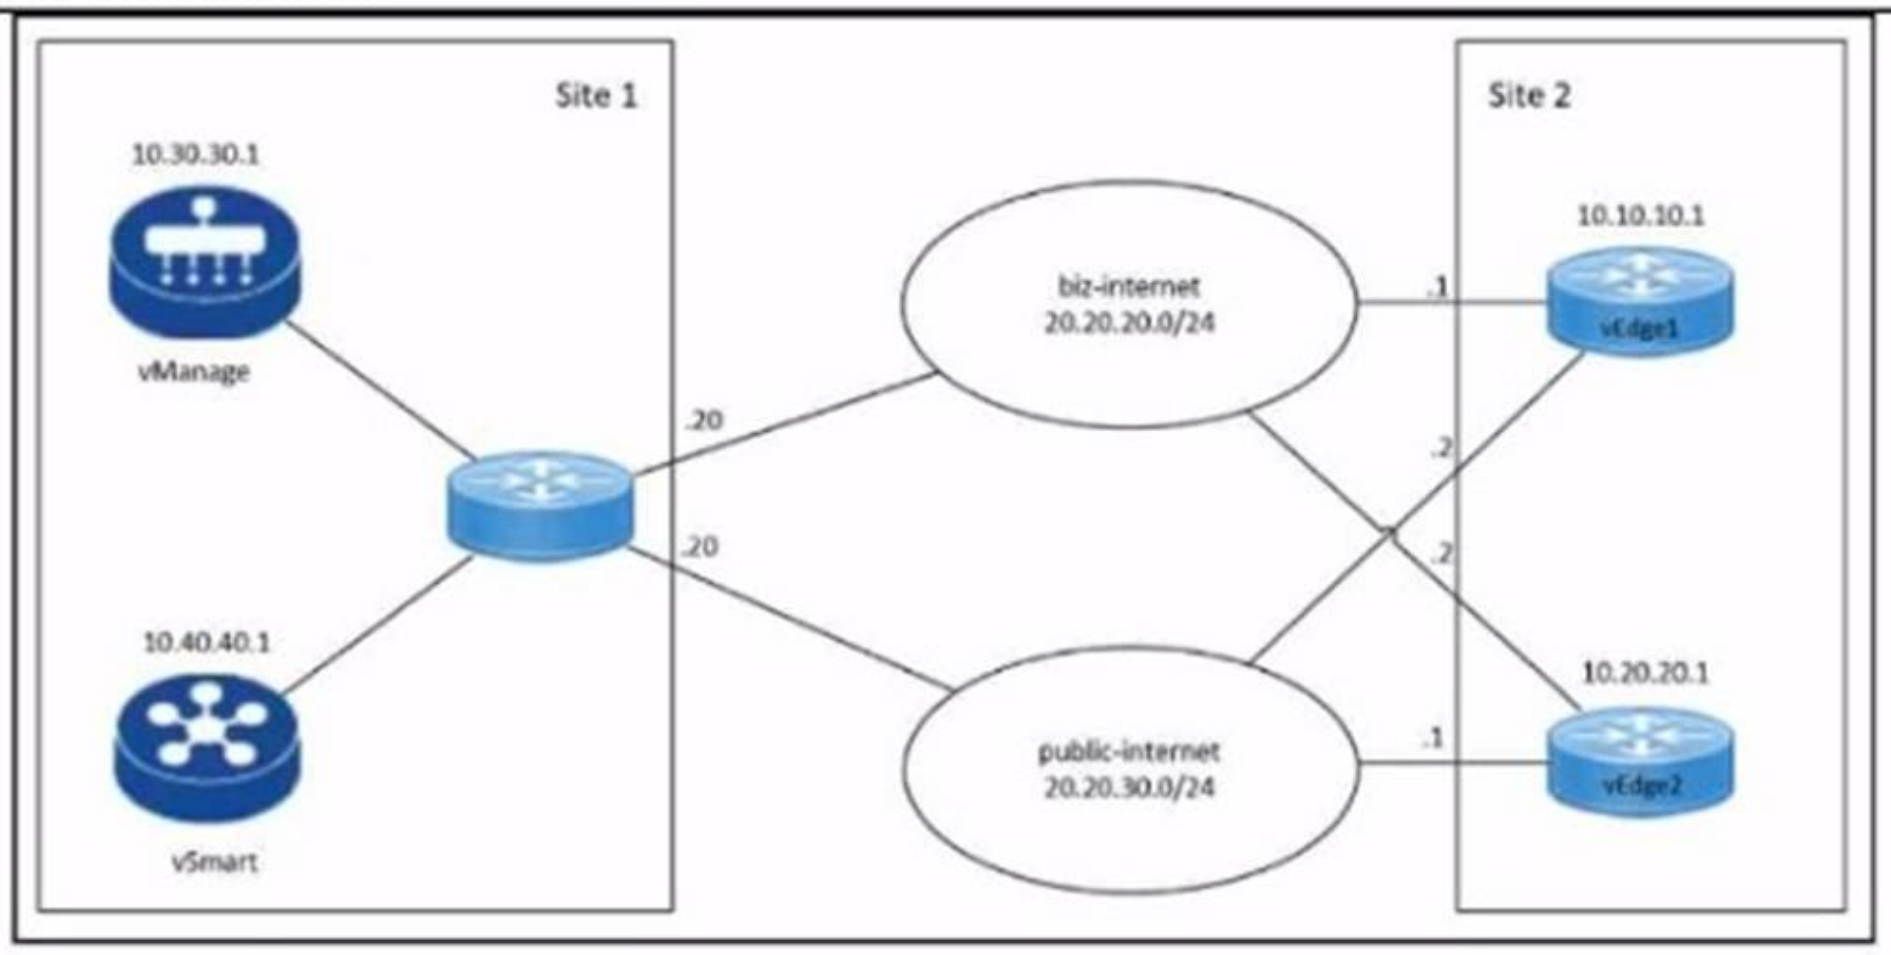

Refer to the exhibits. An engineer troubleshoots a Cisco SD-WAN connectivity issue between an on-premises data center WAN Edge and a public cloud provider WAN Edge. The engineer discovers that BFD is Dapping on vEdge1. What is the problem?

A. The remote Edge device BFD is down.
B. The remote Edgedevice failed to respond BFD keepalives.
C. The remote Edge device has a duplicate IP address.
D. The control plane deleted the BFD session.

Answer: B

Explanation:
BFD (Bidirectional Forwarding Detection) is a protocol that detects failures in the overlay tunnel between Cisco SD-WAN devices. BFD packets are sent and received periodically by each device to check the liveliness and quality of the connection. If a device does not receive a BFD packet from its peer within a specified timeout interval, itconsiders the peer to be unreachable and reports a BFD down event. This event triggers a control connection state change and a possible route change in the SD-WAN fabric.
In this scenario, the engineer discovers that BFD is flapping on vEdge1, which means that the BFD session between vEdge1 and the remote Edge device is going up and down repeatedly. This indicates a connectivity issue between the two devices, such as network congestion, packet loss, or misconfiguration. The most likely cause of the problem is that the remote Edge device failed to respond BFD keepalives within the timeout interval, which resulted in a BFD timeout event on vEdge1. This event caused vEdge1 to mark the remote Edge device as down and notify the control plane. The control plane then tried to establish a new BFD session with the remote Edge device, which may have succeeded or failed depending on the network condition. This cycle of BFD session creation and deletion caused the BFD flapping on vEdge1.
The other options are less likely to be the cause of the problem. Option A is incorrect because if the remote Edge device BFD was down, vEdge1 would not receive any BFD packets from it and would not flap. Option C is incorrect because if the remote Edge device had a duplicate IP address, vEdge1 would not be able to establish a BFD session with it in the first place. Option D is incorrect because the control plane does not delete the BFD session unless there is a configuration change or a port-hop event on the device. References: Bidirectional Forwarding Detection Flap-Reason Definitions on Cisco vEdge Routers, Cisco Catalyst SD-WAN BFD, Cisco SD WAN: BFD (Bidirectional Forwarding Detection)

NEW QUESTION 34
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

300-440 Practice Exam Features:

* 300-440 Questions and Answers Updated Frequently

* 300-440 Practice Questions Verified by Expert Senior Certified Staff

* 300-440 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* 300-440 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year