Atmic Networks

CERTIFICATION EXAM SUPPORT

EXAM PREPARATION MATERIAL

Practice files designed to help you understand exam patterns, question styles and the overall assessment format.

IMPORTANT NOTE

Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns.

Stay connected with our support team
for the latest file updates and exam-related information.

www.atmicnetworks.com

Knowledge that moves beyond memorization.

Question: 1

Refer to the exhibit.

```
vEdgef show crypto isakap sa

IPv4 Crypto ISAKMP SA dst     src                    state                conn-id                status
203.0.113.1                   203.0.113.2            MN 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. Reference :=
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

Question: 2

Refer to the exhibit.

```
vedgel# show policy from-vsmart
apply-policy
 site-list sitel
  control-policy prefer_local out

policy
 lists
  site-list sitel site-id 100
  tloc-list prefersitel
     tloc 10.1.1.1 color mpls encap ipsec preference 100 control-policy
  prefer_local
```

sequence 10 match route site-list sitel

- action accept set

tloc-list prefer sitel

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. Reference := 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

Question: 3

A company with multiple branch offices wants a connectivity model to meet its network architecture requirements. The company focuses on ensuring low latency and efficient routing for its critical business applications. Which connectivity model meets these requirements?

A. hub-and-spoke topology with SD-WAN technology, using dynamic routing and OSPF as the routing protocol
B. fully meshed topology with SD-WAN technology, using dynamic routing and BGP as the routing protocol
C. point-to-point topology using dedicated leased lines and static routing
D. star topology with internet-based VPN connections and static routing

Answer: B

Explanation:

A fully meshed topology with SD-WAN technology, using dynamic routing and BGP as the routing protocol, meets the requirements of the company because it provides the following benefits:
It allows direct and secure connectivity between any two branch offices, without the need for a central hub or intermediary devices12. This reduces the latency and improves the performance of the critical business applications.
It leverages SD-WAN technology to optimize the traffic flow and application quality of service (QoS) across the WAN13. SD-WAN can dynamically select the best path for each application based on the network conditions and policies13. SD-WAN can also provide redundancy, security, and visibility for the WAN13.
It uses dynamic routing and BGP as the routing protocol to exchange routing information and establish connectivity between the branch offices14. BGP is a scalable and flexible protocol that can support multiple address families, such as IPv4 and IPv6, and multiple routing policies, such as local preference and route filtering14. BGP can also enable

seamless integration with the cloud service providers (CSPs) and internet service providers (ISPs)14.

Reference :=

1: Designing and Implementing Cloud Connectivity (ENCC, Track 1 of 5) (Cisco U. login required)

2: Cisco SD-WAN Design Guide

Question: 4

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.

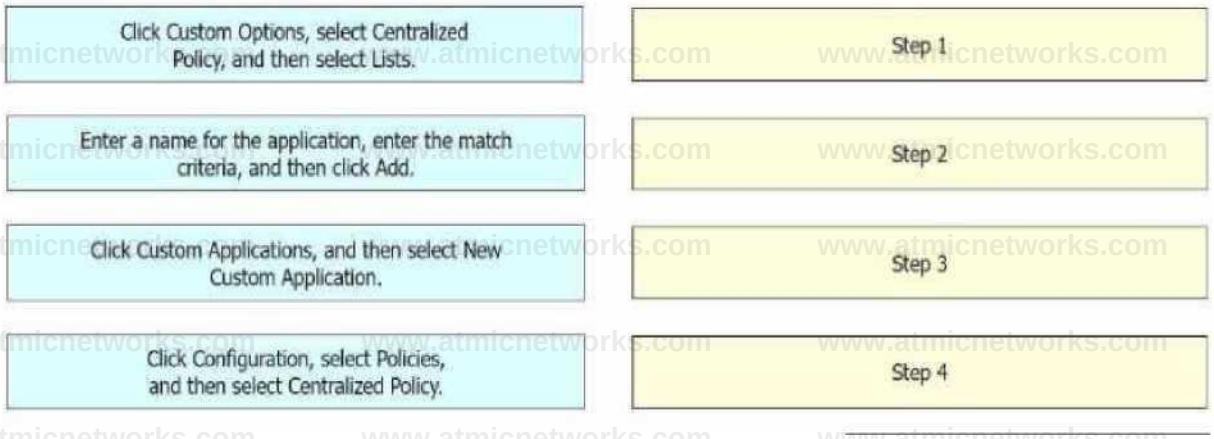

Answer:

Explanation:

To configure a custom application with Cisco SD-WAN centralized policy, you need to follow these steps25:

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

Click Custom Options, select Centralized Policy, and then select Lists.

Click Custom Applications, and then select New Custom Application.

Enter a name for the application, enter the match criteria, and then click Add.

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.

Reference :=

Cisco Catalyst SD-WAN Policies Configuration Guide, Cisco IOS XE

Question: 5

Which Microsoft Azure service enables a dedicated and secure connection between an on-premises infrastructure and Azure data centers through a colocation provider?

A. Azure Private Link
B. Azure ExpressRoute
C. Azure Virtual Network
D. Azure Site-to-Site VPN

Answer: B

Explanation:
Azure ExpressRoute is a service that enables a dedicated and secure connection between an onpremises infrastructure and Azure data centers through a colocation provider. A colocation provider is a third-party data center that offers network connectivity services to multiple customers. Azure ExpressRoute allows customers to bypass the public internet and connect directly to Azure services, such as virtual machines, storage, databases, and more. This provides benefits such as lower latency, higher bandwidth, more reliability, and enhanced security. Azure ExpressRoute also supports hybrid scenarios, such as connecting to Office 365, Dynamics 365, and other SaaS applications hosted on Azure. Azure ExpressRoute requires a physical connection between the customer's network and the colocation provider's network, as well as a logical connection between the customer's network and the Azure virtual network. The logical connection is established using a Border Gateway Protocol (BGP) session, which exchanges routing information between the two networks. Azure ExpressRoute supports two models: standard and premium. The standard model offers connectivity to all Azure regions within the same geopolitical region, while the premium model offers connectivity to all Azure regions globally, as well as additional features such as increased route limits, global reach, and Microsoft peering. Reference: Designing and Implementing Cloud Connectivity (ENCC) v1.0, Learning Plan: Designing and Implementing Cloud Connectivity v1.0 (ENCC 300-440) Exam Prep, ENCC | Designing and Implementing Cloud Connectivity | Netec

Question: 6

An engineer must enable the OMP advertisement of BGP routes for a specific VRF instance on a Cisco IOS XE SD-WAN device. What should be configured after the global address-family ipv4 is configured?

A. Set the VRF-specific route advertisements.
B. Enable bgp advertisement.
C. Enter sdwan mode.
D. Disable bgp advertisement.

Answer: B

Explanation:

To enable the OMP advertisement of BGP routes for a specific VRF instance on a Cisco IOS XE SD- WAN device, the engineer must first configure the global address-family ipv4 and then enable bgp advertisement under the vrf definition. This will allow the device to advertise the BGP routes learned from the cloud provider to the OMP control plane, which will then distribute them to the other SD- WAN devices in the overlay network1
Reference := 1: Designing and Implementing Cloud Connectivity (ENCC) v1.0, Module 3: Implementing Cloud Connectivity, Lesson 3: Configuring IPsec VPN from Cisco IOS XE to AWS, Topic: Configuring BGP on the Cisco IOS XE Device, Page 3-24.

Question: 7

Refer to the exhibit.

```
local7.debug: Mar 11 11:31:11 VEDGE-1 VDAEMON[1136]: vdaemon_disable_my_tloc[1308]:
%VDAEMON_DBG_EVENT3-1: 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
```

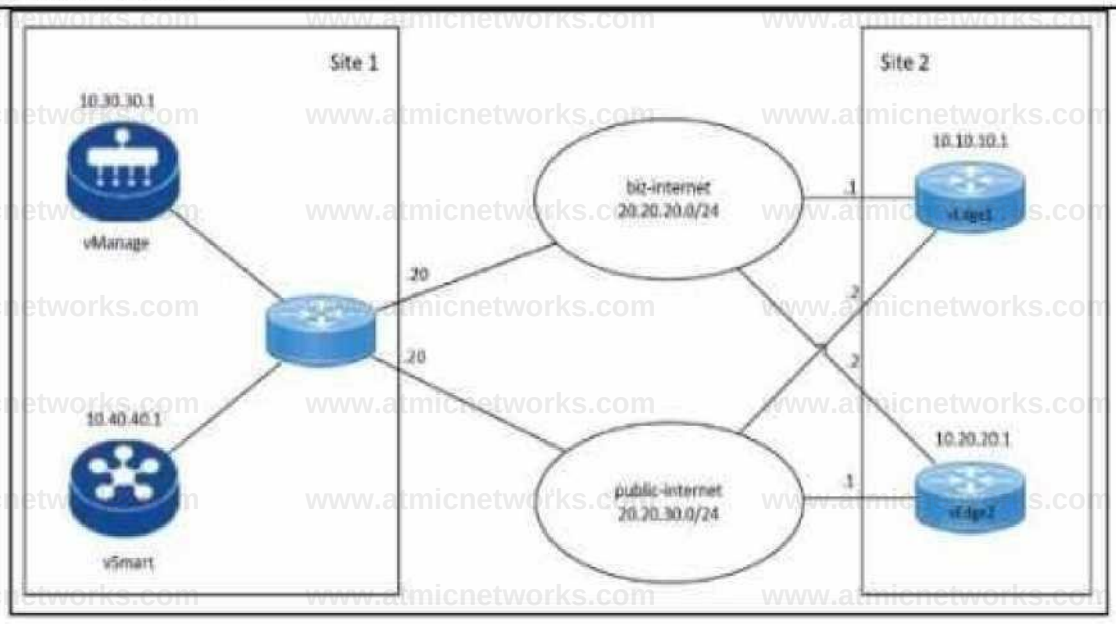

Refer to the exhibits. An engineer troubleshoots a Cisco SD-WAN connectivity issue between an onpremises data center WAN Edge and a public cloud provider WAN Edge. The engineer discovers that BFD is Dapping on vEdge1. What is the problem?

A. The remote Edge device BFD is down.
B. The remote Edge device 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, it considers 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 porthop event on the device. Reference: Bidirectional Forwarding Detection Flap-Reason Definitions on Cisco vEdge Routers, Cisco Catalyst SD-WAN BFD, Cisco SD WAN: BFD (Bidirectional Forwarding Detection)

Question: 8

An engineer is implementing a highly secure multitier application in AWS that includes S3. RDS, and some additional private links. What is critical to keep the traffic safe?

A. VPC peering and bucket policies
B. specific routing and bucket policies
C. EC2 super policies and specific routing policies
D. gateway load balancers and specific routing policies

Answer: B

Explanation:

A highly secure multitier application in AWS that includes S3, RDS, and some additional private links requires specific routing and bucket policies to keep the traffic safe. The reasons are as follows: Specific routing policies are needed to ensure that the traffic between the tiers is routed through the private links, which provide secure and low-latency connectivity between AWS services and onpremises resources12. The private links can also prevent the exposure of the data and the application logic to the public internet12.

Bucket policies are needed to control the access to the S3 buckets that store the application data34. Bucket policies can specify the conditions under which the requests are allowed or denied, such as the source IP address, the encryption status, the request time, etc.34. Bucket policies can also enforce encryption in transit and at rest for the data in S334.

Reference :=

1: AWS PrivateLink
2: AWS PrivateLink FAQs
3: Using Bucket Policies and User Policies
4: Bucket Policy Examples

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.

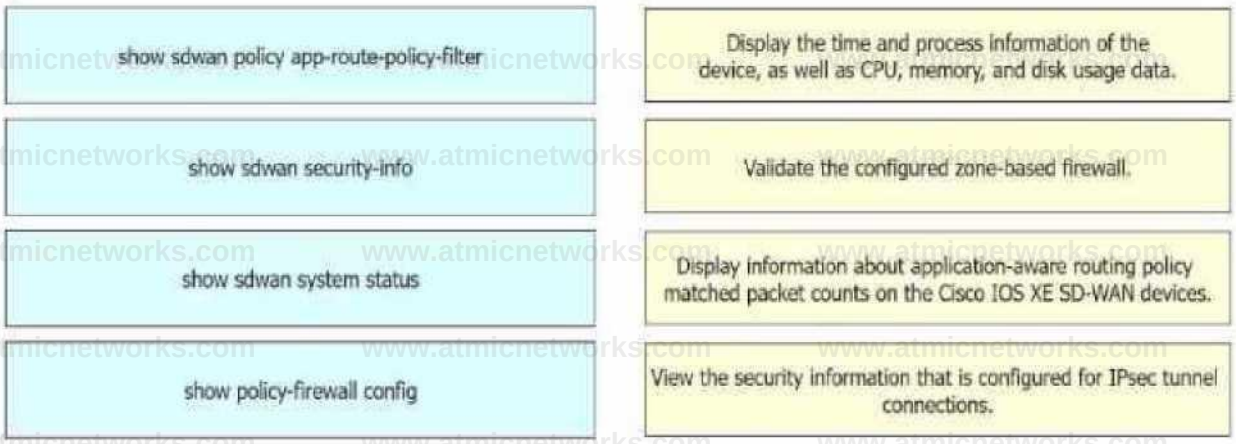

Answer:

Explanation:

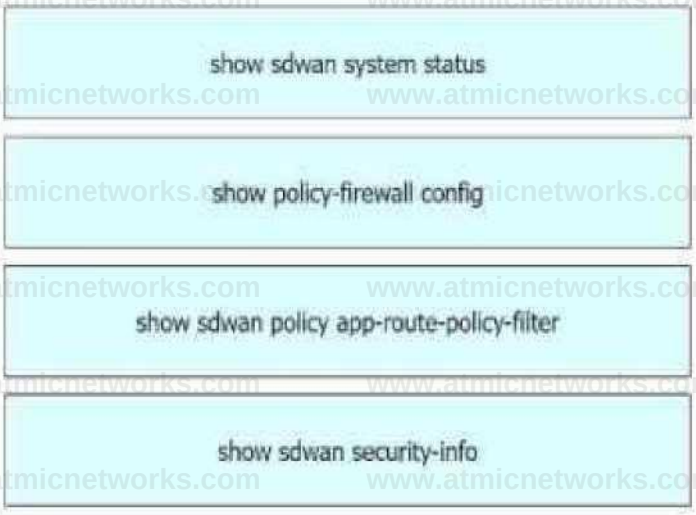

Display the time and process information of the device, as well as CPU, memory, and disk usage dat a. = show sdwan system status1

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 securityinfo

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.

Reference :=

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

Question: 10

Refer to the exhibits.

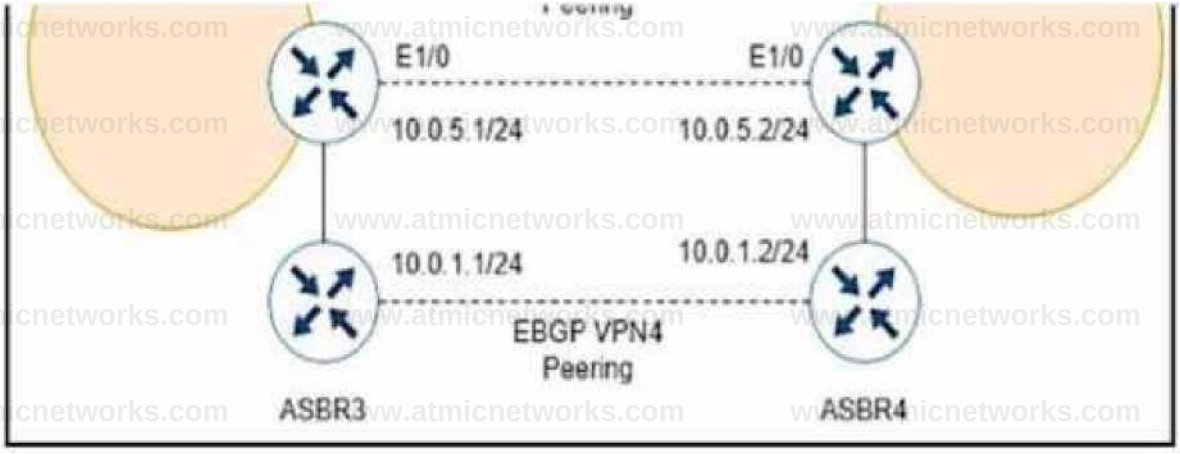

While troubleshooting, a network engineer discovers that the backup path fails between ASBR3 and ASBR4 for traffic between BGP AS6000 and BGP AS6500 when the connection between ASBR1 and ASBR2 goes down. The following configurations were performed on ASBR1:

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

Which command is missing?

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

Question: 11

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 onpremises 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: 12

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
```

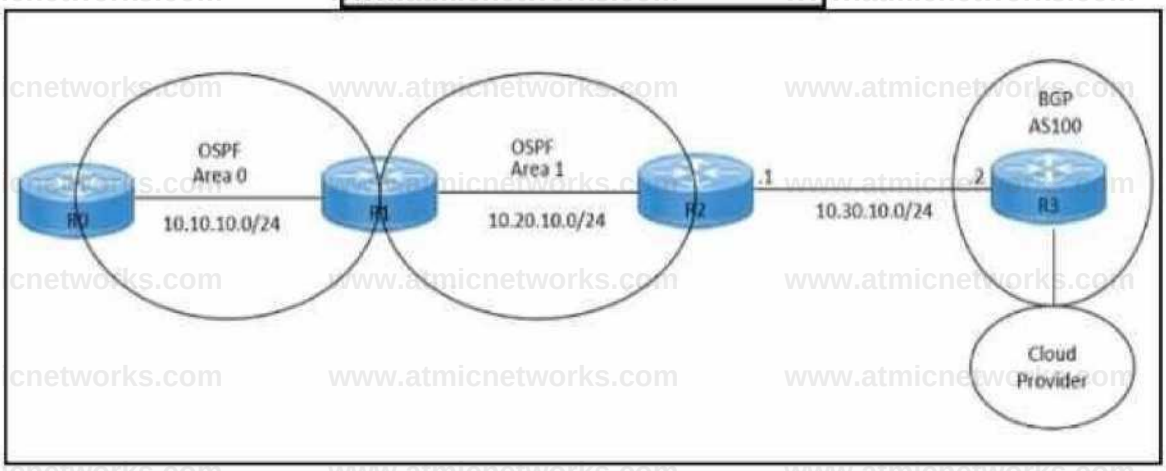

Refer to the exhibits. An engineer must redistribute IBGP routes into OSPF to connect an on-premises network to a cloud provider. Which command must be configured on router R2?

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.

Reference: 1: Learning Plan: Designing and Implementing Cloud Connectivity v1.0 (ENCC 300-440) Exam Prep 2: Designing and Implementing Cloud Connectivity (ENCC) v1.0 3: Cisco Multiprotocol Label Switching 4: Exploring Cisco Cloud OnRamp for Colocation 5: ENCC: Configuring IPsec VPN from Cisco IOS XE to AWS : [Deploying Cisco IOS VTI-Based Point-to-Point IPsec VPNs]

Question: 13

```
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
!
```

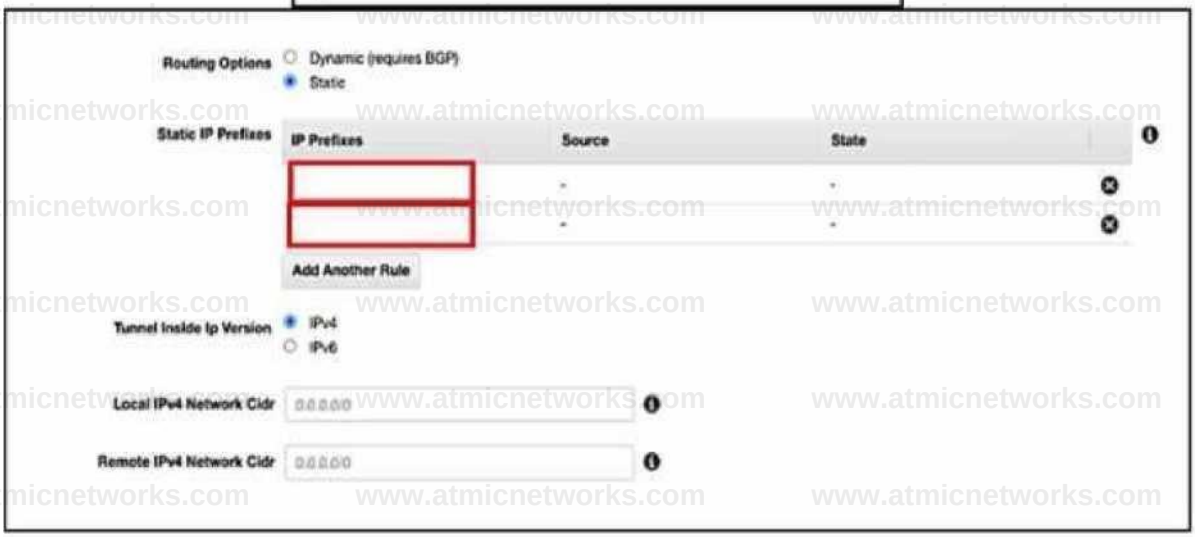

Refer to the exhibits. An engineer needs to configure a site-to-site IPsec VPN connection between an on premises Cisco IOS XE router and Amazon Web Services (AWS). Which two IP prefixes should be used to configure the AWS routing options? (Choose two.)

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: 14

```
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
```

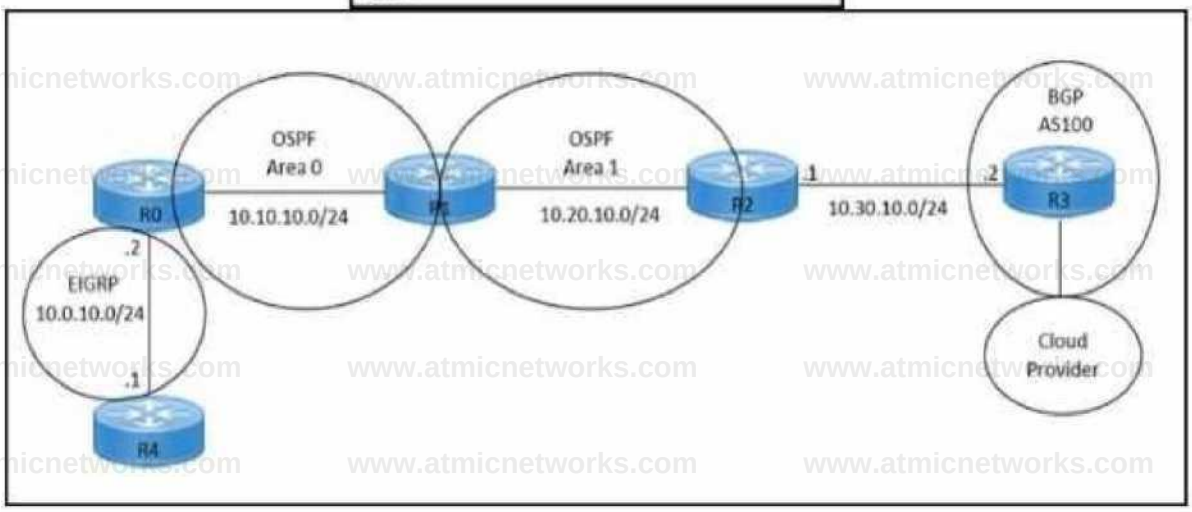

Refer to the exhibits. An engineer must redistribute OSPF internal routes into BGP to connect an onpremises network to a cloud provider without introducing extra routes. Which two commands must be configured on router R2? (Choose two.)

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

Answer: B, E

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.
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: 15

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

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: 16

Refer to the exhibit.

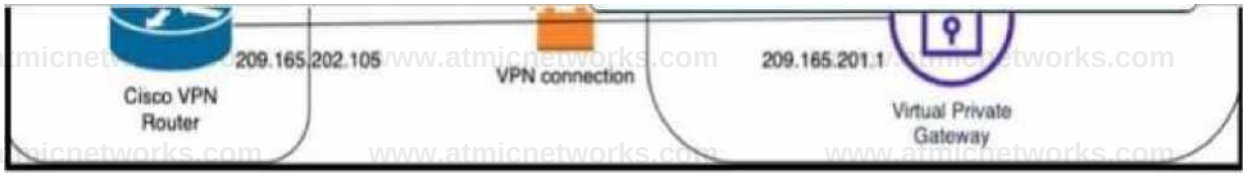

Which Cisco lKEv2 configuration brings up the IPsec tunnel between the remote office router and the AWS virtual private gateway?

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 i
crypto ikev2 keyring DEMO-Keyring peer Cisco-AWS
  address 209.165.201.1
  pre-shared-key DEMOlabCisco12345
I !
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 209.165.202.10 5
    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
!
```

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 AWS1 and Cisco2 for 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: 17

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.

Answer:

Explanation:

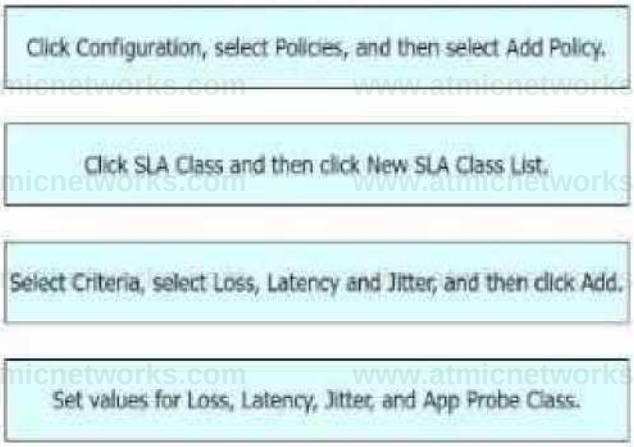

The process of configuring an SLA class to specify the maximum packet loss, packet latency, and jitter allowed on a connection using Cisco vManage involves several steps12.

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.

Reference :=

Information About Application-Aware Routing - Cisco

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

Question: 18

DRAG DROP

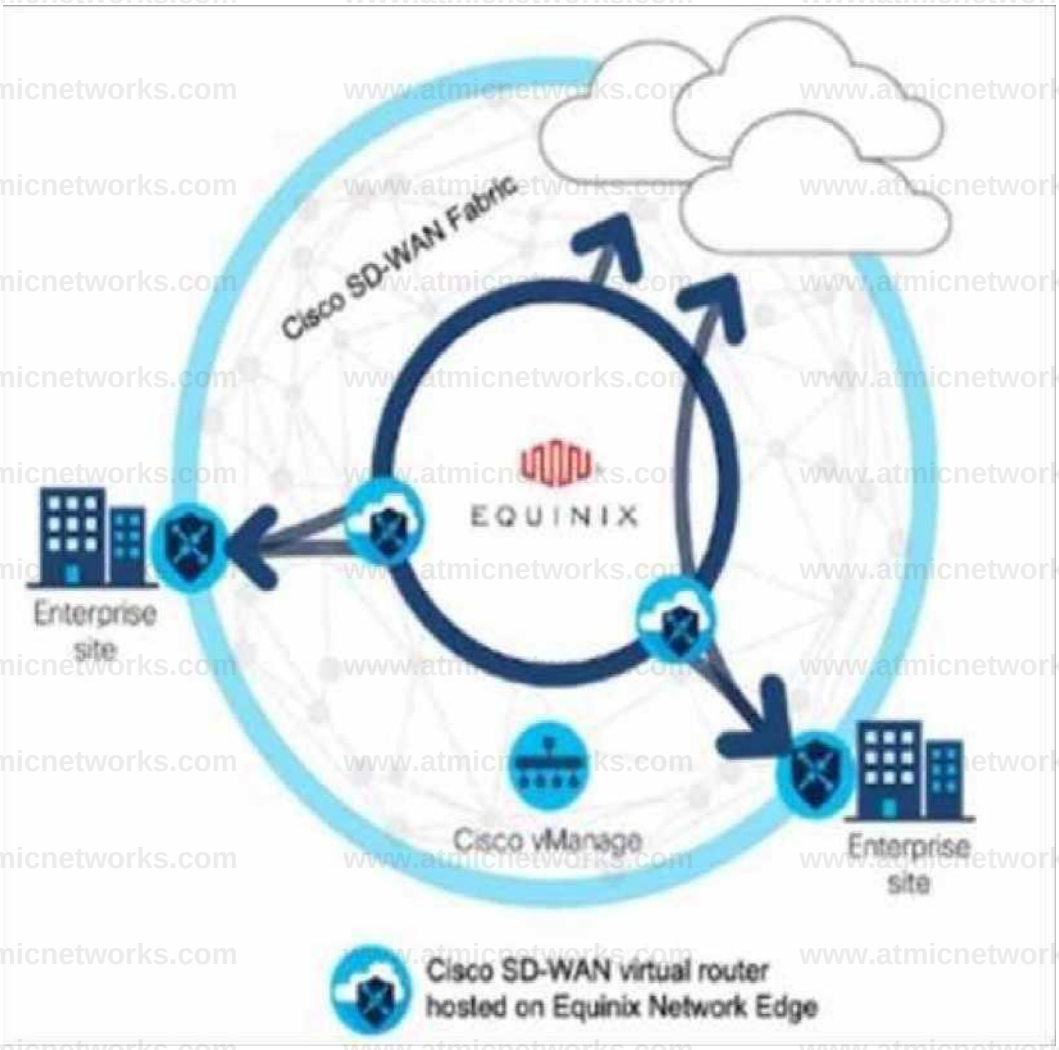

Refer to the exhibit. These configurations are complete:

- 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

Answer:

Explanation:

Insure that you have IRJIDs tor the rwjrirwi number of Chien 50* WAN Virtual Edge instances tlvMyau want to deploy as Interconnect Gateway*

C ro»tf the necessary notwvk segment;

Mtoch Cbco SO WAN Virtual Edge to the Egum* device template

Create the Interconnect Gateway at the Equinix location tM b ik^t In «xB SO WAN branch lucutbiL

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.
Reference :=
[Cisco SD-WAN Cloud Interconnect with Equinix]
[Cisco SD-WAN Cloud OnRamp for CoLocation Deployment Guide]

Question: 19

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 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 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.
Reference:
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: 20

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.

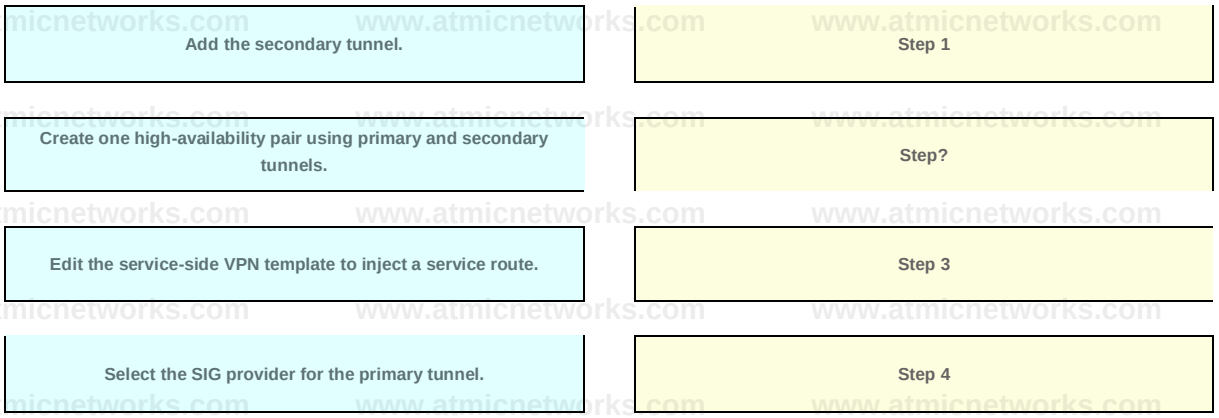

Answer:

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. This ensures 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.
Reference :=
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

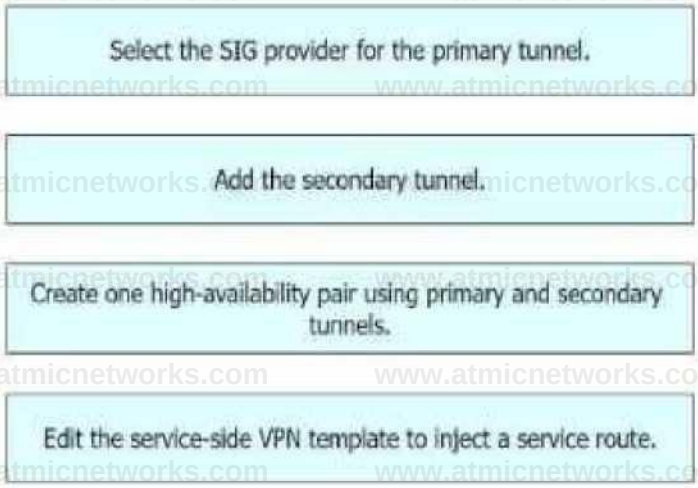

Question: 21

A cloud engineer is setting up a new set of nodes in the AWS EKS cluster to manage database integration with Mongo Atlas. The engineer set up security to Mongo but now wants to ensure that the nodes are also secure on the network side. Which feature in AWS should the engineer use?

A. EC2 Trust Lock
B. security groups
C. tagging
D. key pairs

Answer: B

Explanation:
Security groups are a feature in AWS that allow you to control the inbound and outbound traffic to your instances. They act as a virtual firewall that can filter the traffic based on the source, destination, protocol, and port. You can assign one or more security groups to your instances, and each security group can have multiple rules. Security groups are stateful, meaning that they automatically allow the response traffic for any allowed inbound traffic, and vice versa. Security groups are essential for securing your nodes in the AWS EKS cluster, as they can prevent unauthorized access to your Mongo Atlas database or other resources. You can also use security groups to isolate your nodes from other instances in the same VPC or subnet, or to allow communication between nodes in different clusters or regions.
Reference :=
AWS Security Groups
Security Groups for Your VPC
Security Groups for Your Amazon EC2 Instances

Security Groups for Your Amazon EKS Cluster

Question: 22

DRAG DROP

An engineer needs to configure enhanced policy-based routing (ePBR) for IPv4 by using Cisco vManage. Drag and drop the steps from the left onto the order on the right to complete the configuration of the ePBR using the CLI add-on template.

Answer:

Explanation:

Enhanced Policy-Based Routing (ePBR) is used to direct packets that arrive at an interface to a specified next-hop. It is very useful in managing a large number of configured access lists more efficiently. In ePBR, the router drops the traffic packets if the next hop configured in the PBR policy is not reachable. To avoid packet loss in such scenarios, you must configure multiple next hops for each access control entry.

Here are the steps to configure ePBR for IPv4 using Cisco vManage:

Configure an extended ACL: This step involves defining the network or the host. For example, you can permit IPv4 traffic from any source to specific hosts.

Configure a class map that matches the ACL: Class maps match the parameters in the ACLs. For instance, you can create a class map of type traffic and match it with the previously created ACL.

Configure the policy map with the action to set the next hop: Policy maps with ePBR then take detailed actions based on the set statements configured. You can configure an ePBR policy map with the class map and set the next hop.

Apply the service policy on the interface: Finally, you apply the ePBR policy map to the interface. For example, you can apply the policy map to a GigabitEthernet interface. Reference :=

Implementing Enhanced Policy Based Routing - Cisco

Cisco Catalyst SD-WAN Policies Configuration Guide, Cisco IOS XE

How to configure PBR - Cisco Community

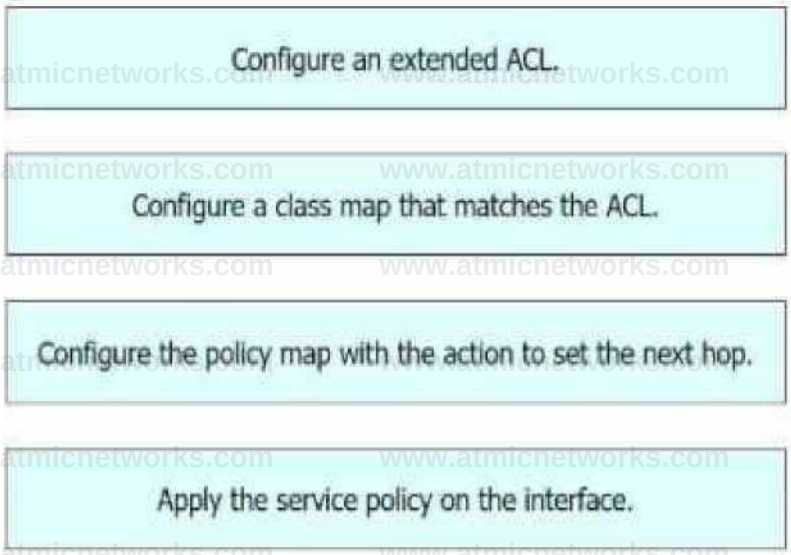

Question: 23

Which feature is unique to Cisco SD-WAN IPsec tunnels compared to native IPsec VPN tunnels?

A. real-time dynamic path selection
B. tunneling protocols
C. end-to-end encryption
D. authentication mechanisms

Answer: A

Explanation:

Cisco SD-WAN IPsec tunnels are different from native IPsec VPN tunnels in several ways. One of the unique features of Cisco SD-WAN IPsec tunnels is that they support real-time dynamic path selection, which means that they can automatically choose the best path for each application based on the network conditions and policies. This feature improves the performance, reliability, and efficiency of the network traffic. Native IPsec VPN tunnels, on the other hand, do not have this capability and rely on static routing or manual configuration to select the path for each tunnel. This can result in suboptimal performance, increased latency, and higher costs. Reference := Traditional IPsec Versus Cisco SD-WAN IPsec, SD-WAN vs IPsec VPN's - What's the difference?, SD-WAN vs. VPN: How Do They Compare?, Traditional IPSEC Versus SD-WAN IPSEC

Question: 24

Which approach does a centralized internet gateway use to provide connectivity to SaaS applications?

A. A cloud-based proxy server routes traffic from the on-premises infrastructure to the SaaS provider data center.
B. Internet traffic from the on-premises infrastructure is routed through a centralized gateway that provides access controls for SaaS applications.
C. VPN connections are used to provide secure access to SaaS applications from the on-premises infrastructure.
D. A dedicated, private connection is established between the on-premises infrastructure and the SaaS provider data center using colocation services.

Answer: B

Explanation:

A centralized internet gateway is a network design that routes all internet-bound traffic from the onpremises infrastructure through a single point of egress, typically located at the data center or a regional hub1. This approach allows the enterprise to apply consistent security policies and access controls for SaaS applications, as well as optimize the bandwidth utilization and performance of the WAN links2. A centralized internet gateway can use various technologies to provide connectivity to SaaS applications, such as proxy servers, firewalls, web filters, and WAN optimizers3. However, a cloud-based proxy server (option A) is not a part of the centralized internet gateway, but rather a separate service that can be used to route traffic from the on-premises infrastructure to the SaaS provider data center4. VPN connections (option C) and dedicated, private connections (option D) are also not related to the centralized internet gateway, but rather alternative ways of providing secure and reliable access to SaaS applications from the on-premises infrastructure5. Therefore, the correct answer is option B, which describes the basic function of a centralized internet gateway. Reference := 1: Designing and Implementing Cloud Connectivity (ENCC) v1.0, Module 1: Cloud Connectivity Overview, Lesson 1: Cloud Connectivity Concepts, Topic: Centralized Internet Gateway 2: Cloud OnRamp for SaaS, Cisco IOS XE Catalyst SD-WAN Release 17.3.1a and Later, Topic: Centralized Internet Gateway 3: Architect and optimize your internet traffic with Azure routing preference, Microsoft Azure Blog, Topic: Routing via the premium Microsoft global network 4: What is SaaS? Software as a Service, Microsoft Azure, Topic: How SaaS works 5: How an application gateway works, Microsoft Learn, Topic: Application gateway components

Question: 25

```
crypto keyring keyring-vpn-000001
pre-shared-key address 192.10.10.10 key secretkey01
!
interface Tunnel1
 ip address 20.20.20.21 255.255.255.252
 tunnel destination 192.10.10.10
!
crypto ikev2 keyring AWS_Keyring
 peer AWS_Peer
 [                                ]
   pre-shared-key local awssecretkey01
   pre-shared-key remote awssecretkey02
!
```

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 onpremises 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. Reference:
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

Question: 26

DRAG DROP
An engineer must configure a CLI add-on feature template in Cisco vManage for enhanced policybased routing (ePBR) for IPv4. These configurations were deleted:

- licensing config enable false
- licensing config privacy hostname true
- licensing config privacy version false
- licensing config utility utility-enable true

Drag and drop the steps from the left onto the order on the right to complete the configuration.

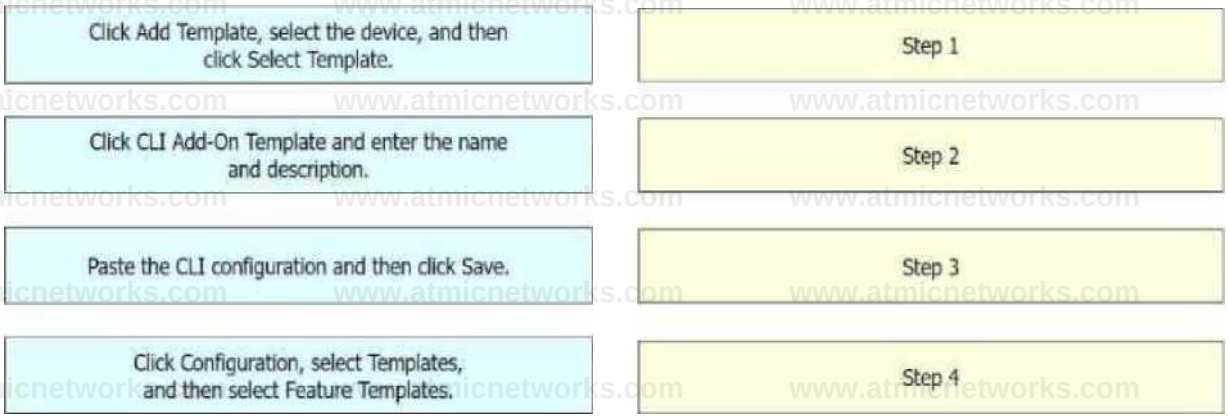

Answer:

Explanation:

Step 1 = Click Configuration, select Templates, and then select Feature Templates. Step 2 = Click Add Template, select the device, and then click Select Template. Step 3 = Click CLI Add-On Template and enter the name and description. Step 4 = Paste the CLI configuration and then click Save.

The process of configuring a CLI add-on feature template in Cisco vManage for enhanced policybased routing (ePBR) for IPv4 involves several steps1234.
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 Select Template: In this step, you add a new template for the device1.
Click CLI Add-On Template and enter the name and description: After setting up the template, you select the CLI Add-On Template option, and then enter the name and description for the template1. Paste the CLI configuration and then click Save: Finally, you paste the CLI configuration into the template and save the changes1.
Reference :=
CLI Add-On Feature Templates - Cisco
Cisco Catalyst SD-WAN Systems and Interfaces Configuration Guide, Cisco IOS XE Catalyst SD-WAN Release 17.x - CLI Add-On Feature Templates
Cisco SD-WAN vSmart CLI Template - NetworkLessons.com
CLI Templates for Cisco XE SD-WAN Routers

Question: 27

DRAG DROP

An engineer must configure an AppGoE service node for WAN optimization for applications that are hosted in the cloud using Cisco vManage for C8000V or C8500L-8S4X devices. Drag and drop the steps from the left onto the order on the right to complete the configuration.

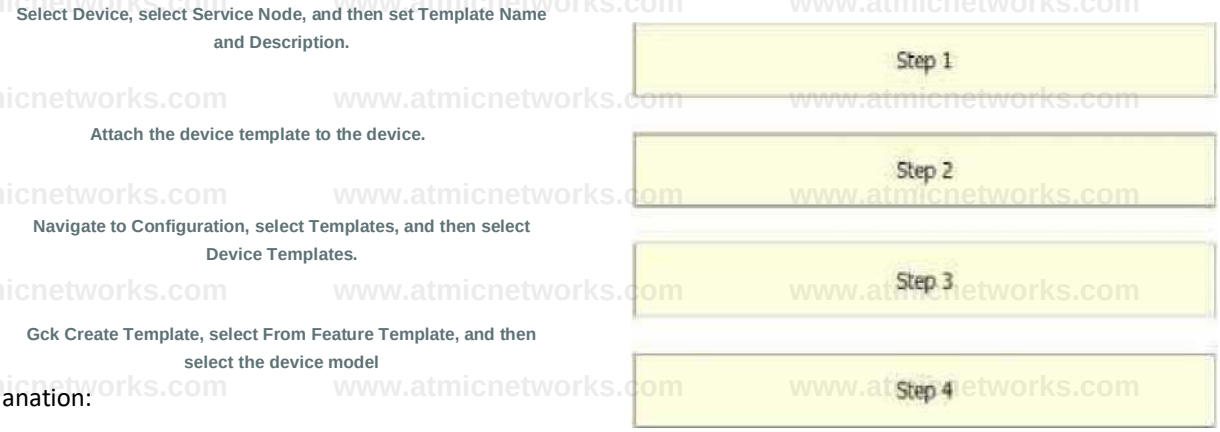

Explanation:

Answer:

Step 1 = Navigate to Configuration, select Templates, and then select Device Templates. Step 2 = Click Create Template, select From Feature Template, and then select the device model. Step 3 = Select Device, select Service Node, and then set Template Name and Description. Step 4 = Attach the device template to the device.

The process of configuring an AppGoE service node for WAN optimization for applications that are hosted in the cloud using Cisco vManage for C8000V or C8500L-8S4X devices involves several steps12.

Navigate to Configuration, select Templates, and then select Device Templates: This is the first step where you navigate to the Templates section in the Configuration menu of Cisco vManage1.

Click Create Template, select From Feature Template, and then select the device model: In this step, you create a new template for the device model from the feature template1.

Select Device, select Service Node, and then set Template Name and Description: After setting up the template, you select the device and the service node, and then set the template name and description1.

Attach the device template to the device: Finally, you attach the created device template to the device1.

Reference :=

AppQoE - Step-by-Step Configuration - Cisco Community

Cisco Catalyst SD-WAN AppQoE Configuration Guide, Cisco IOS XE Catalyst SD-WAN Release 17.x

Question: 28

hostname R2

interface GigabitEthemet0/0 ip address 10.30.10.1 255.255.255.0 duplex auto speed auto

!

interface GigabitEthernetO/l ip address 10.20.10.1 255.255.255.0 duplex auto speed auto i

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

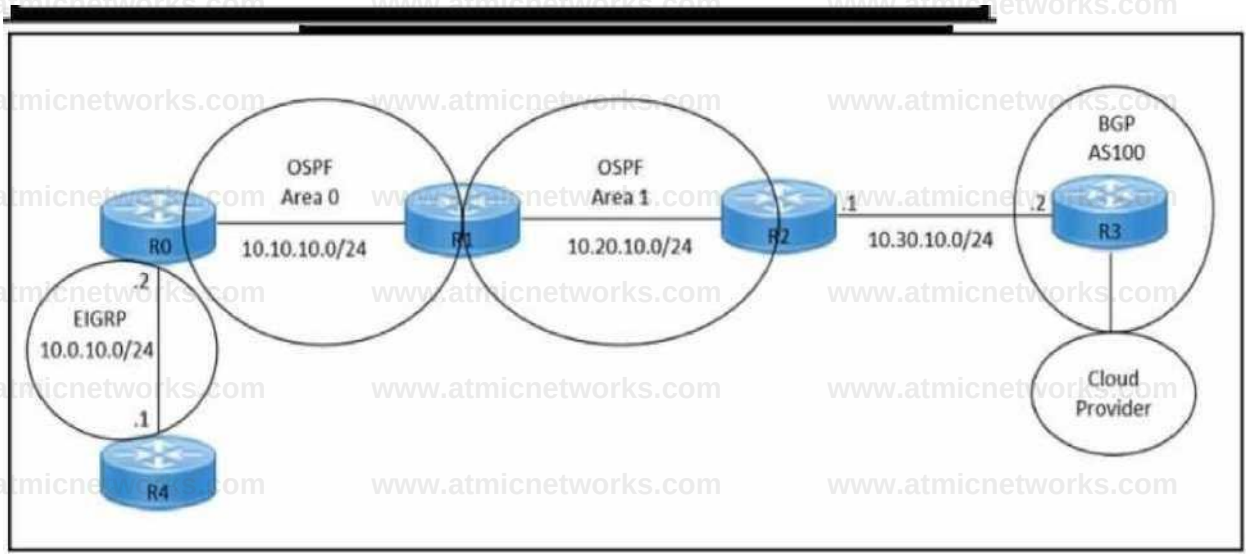

Refer to the exhibits. An engineer must redistribute only the 10.0.10.0/24 network into BGP to connect an on-premises network to a public cloud provider. These routes are currently redistributed:

•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. Reference :=

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

Question: 29

```
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
```

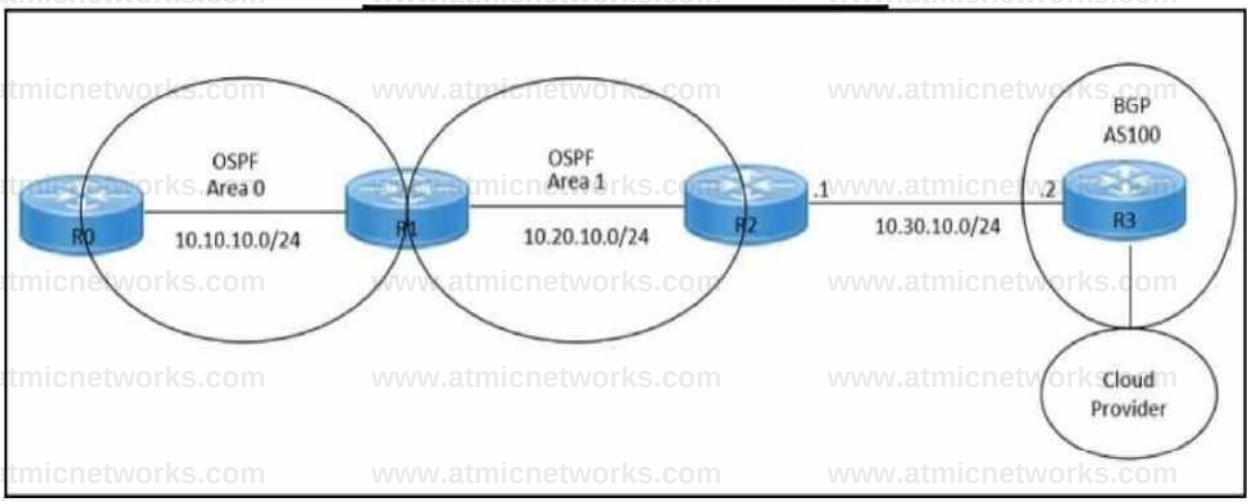

Refer to the exhibits. An engineer must redistribute OSPF internal routes into BGP to connect an onpremises network to a cloud provider. Which two commands should the engineer run on router R2? (Choose two.)

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

Answer: A D

Explanation:

To redistribute OSPF internal routes into BGP for connecting an on-premises network to a cloud provider, the engineer should run the commands “router bgp 100” and “redistribute ospf 1” on router R2. The command “router

bgp 100" is used to create a BGP routing process with AS number 100. The command "redistribute ospf 1" is used to redistribute OSPF routes from process ID 1 into BGP. Reference: = I need to access the specific content of Designing and Implementing Cloud Connectivity (ENCC) v1.0 from Cisco's official resources to provide exact references. However, I don't have direct access to external databases or resources, including the Cisco ENCC course materials. I recommend referring to the ENCC course materials for the most accurate and detailed information. Please note that this answer is based on general networking principles and may not reflect the specific content of the ENCC course. Always refer to the official course materials for the most accurate information.

Question: 30

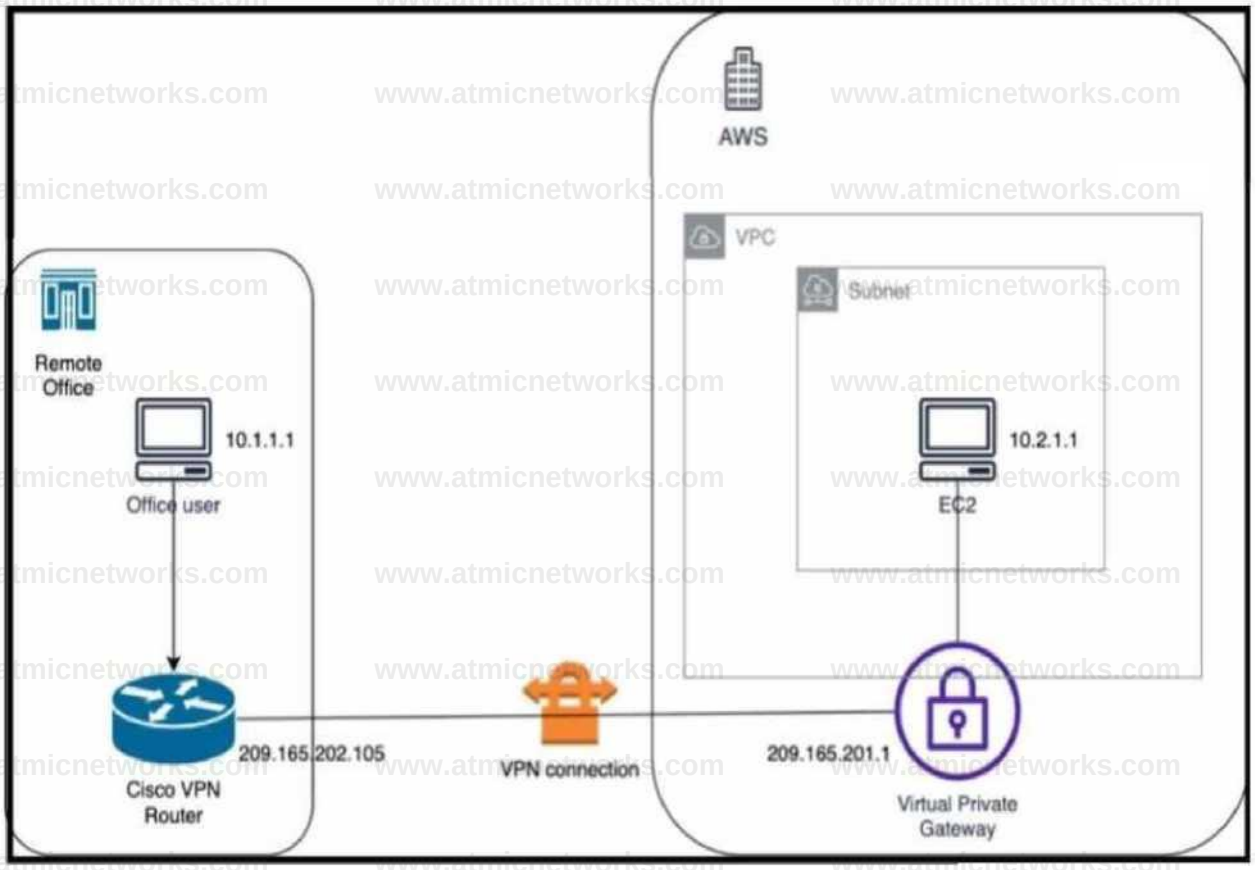

Refer to the exhibit. An engineer successfully brings up the site-to-site VPN tunnel between the remote office and the AWS virtual private gateway, and the site-to-site routing works correctly. However, the end-to-end ping between the office user PC and the AWS EC2 instance is not working. Which two actions diagnose the loss of connectivity? (Choose two.)

A. Check the network security group rules on the host VNET.
B. Check the security group rules for the host VPC.
C. Check the IPsec SA counters.
D. On the Cisco VPN router, configure the IPsec SA to allow ping packets.
E. On the AWS private virtual gateway, configure the IPsec SA to allow ping packets.

Answer: B, C

Explanation:

The end-to-end ping between the office user PC and the AWS EC2 instance is not working because either the security group rules for the host VPC are blocking the ICMP traffic or the IPsec SA counters are showing errors or drops. To diagnose the loss of connectivity, the engineer should check both the security group rules and the IPsec SA counters. The network security group rules on the host VNET are not relevant because they apply to Azure, not AWS. The IPsec SA configuration on the Cisco VPN router and the AWS private virtual gateway are not likely to be the cause of the problem because the site-to-site VPN tunnel is already up and the site-to-site routing works correctly. Reference := Designing and Implementing Cloud Connectivity (ENCC, Track 1 of 5), Module 3: Configuring IPsec VPN from Cisco IOS XE to AWS, Lesson 3: Verify IPsec VPN Connectivity
Security for VPNs with IPsec Configuration Guide, Cisco IOS XE, Chapter: IPsec VPN Overview, Section: IPsec Security Association
AWS Documentation, User Guide for AWS VPN, Section: Security Groups for Your VPC

Question: 31

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. Reference :=

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

Question: 32

A company with multiple branch offices wants a suitable connectivity model to meet these network architecture requirements:

- high availability
- quality of service (QoS)
- multihoming

- specific routing needs

Which connectivity model meets these requirements?

A. hub-and-spoke topology using MPLS with static routing and dedicated bandwidth for QoS
B. star topology with internet-based VPN connections and BGP for routing
C. hybrid topology that combines MPLS and SD-WAN
D. fully meshed topology with SD-WAN technology using dynamic routing and prioritized traffic for QoS

Answer: D

Explanation:

A fully meshed topology with SD-WAN technology using dynamic routing and prioritized traffic for QoS meets the network architecture requirements of the company. A fully meshed topology provides high availability by eliminating single points of failure and allowing multiple paths between branch offices. SD-WAN technology enables multihoming by supporting multiple transport options, such as MPLS, internet, LTE, etc. SD-WAN also provides QoS by applying policies to prioritize traffic based on application, user, or network conditions. Dynamic routing allows the SD-WAN solution to adapt to changing network conditions and optimize the path selection for each traffic type. A fully meshed topology with SD-WAN technology can also support specific routing needs, such as segment routing, policy-based routing, or application-aware routing. Reference:
Designing and Implementing Cloud Connectivity (ENCC) v1.0
[Cisco SD-WAN Design Guide]
[Cisco SD-WAN Configuration Guide]

Question: 33

A company has multiple branch offices across different geographic locations and a centralized data center. The company plans to migrate Its critical business applications to the public cloud infrastructure that is hosted in Microsoft Azure. The company requires high availability, redundancy, and low latency for its business applications. Which connectivity model meets these requirements?

A. ExpressRoute with private peering using SDCI
B. hybrid connectivity with SD-WAN
C. AWS Direct Connect with dedicated connections
D. site-to-site VPN with Azure VPN gateway

Answer: A

Explanation:

The connectivity model that meets the requirements of high availability, redundancy, and low latency for the company's business applications is ExpressRoute with private peering using SDCI.
ExpressRoute is a service that provides a dedicated, private, and high-bandwidth connection between the customer's on-premises network and Microsoft Azure cloud network1.
Private peering is a type of ExpressRoute circuit that allows the customer to access Azure services that are hosted in a virtual network, such as virtual machines, storage, and databases2.
SDCI (Secure Data Center Interconnect) is a Cisco solution that enables secure and scalable connectivity between multiple data centers and cloud providers, using technologies such as MPLS, IPsec, and SD-WAN3.
By using ExpressRoute with private peering and SDCI, the company can achieve the following benefits:
High availability: ExpressRoute circuits are redundant and resilient, and can be configured with multiple service

providers and locations for failover and load balancing1. SDCI also provides high availability by using dynamic routing protocols and encryption mechanisms to ensure optimal and secure path selection3.
Redundancy: ExpressRoute circuits can be paired together to form a redundant connection between the customer's network and Azure4. SDCI also supports redundancy by allowing multiple connections between data centers and cloud providers, using different transport technologies and service levels3.
Low latency: ExpressRoute circuits offer lower latency than public internet connections, as they bypass the congestion and variability of the internet1. SDCI also reduces latency by using MPLS and SD-WAN to optimize the performance and quality of service for the traffic between data centers and cloud providers3.

Reference:

What is Azure ExpressRoute?
Azure ExpressRoute peering
Cisco Secure Data Center Interconnect
ExpressRoute circuit and routing domain

Question: 34

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 or internal 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. Reference := FedRAMP Authorization Boundary

Guidance document1

Question: 35

DRAG DROP
Refer to the exhibit.

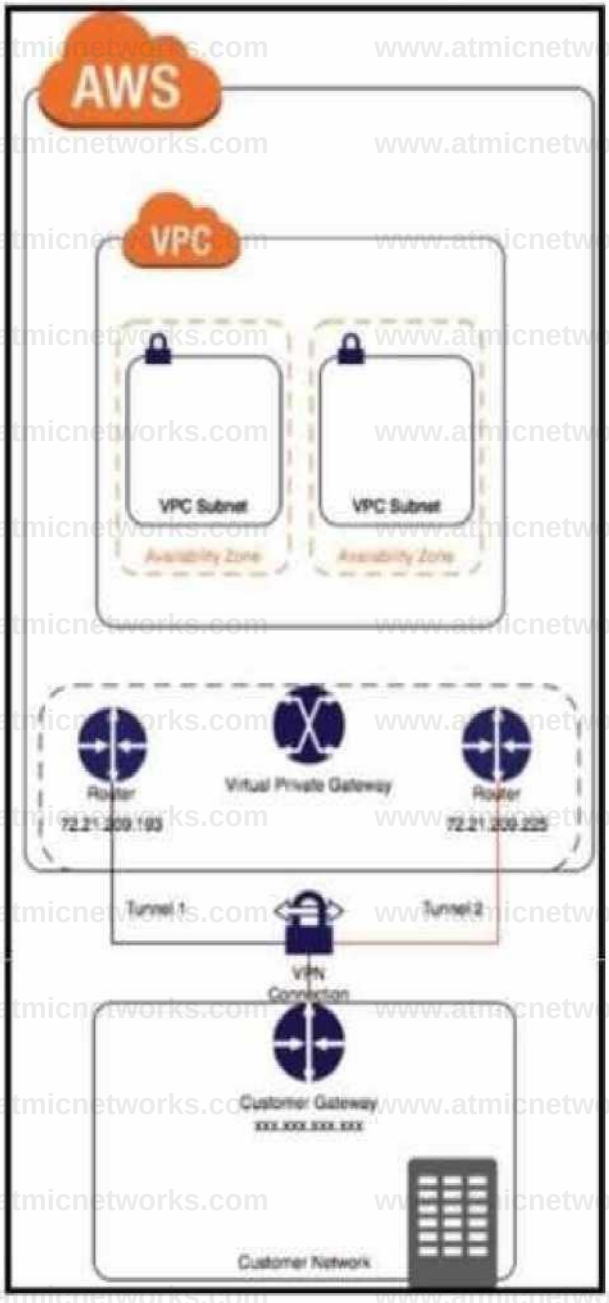

Drag and drop the steps from the left onto the order on the right to configure a site-to-site VPN connection between an on-premises Cisco IOS XE router and Amazon Web Services (AWS).

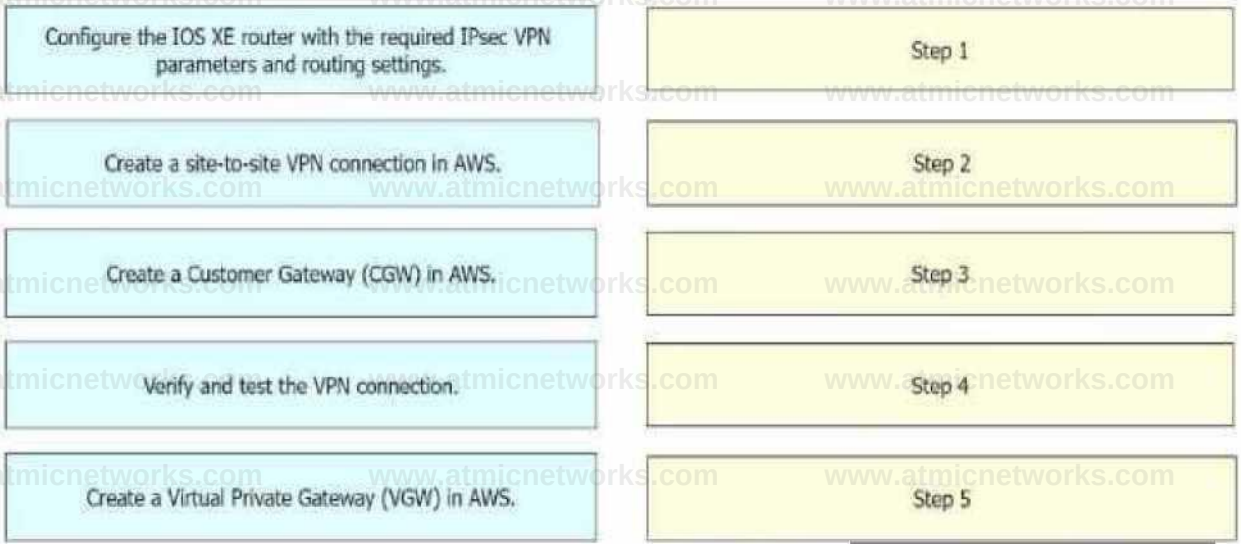

Answer:

Explanation:

Step 1 = Create a Customer Gateway (CGW) in AWS. Step 2 = Create a Virtual Private Gateway (VGW) in AWS. Step 3 = Create a site-to-site VPN connection in AWS. Step 4 = Configure the IOS XE router with the required IPsec VPN parameters and routing settings. Step 5 = Verify and test the VPN connection.

The process of configuring a site-to-site VPN connection between an on-premises Cisco IOS XE router and Amazon Web Services (AWS) involves several steps12.

Create a Customer Gateway (CGW) in AWS: This is the first step where you define the public IP address of your on-premises Cisco IOS XE router in AWS1.

Create a Virtual Private Gateway (VGW) in AWS: This involves creating a VGW and attaching it to the VPC in AWS1.

Create a site-to-site VPN connection in AWS: After setting up the CGW and VGW, you then create a site-to-site VPN connection in AWS. This involves specifying the CGW, VGW, and the static IP prefixes for your on-premises network1.

Configure the IOS XE router with the required IPsec VPN parameters and routing settings: After the AWS side is set up, you configure the on-premises Cisco IOS XE router with the required IPsec VPN parameters and routing settings2.

Verify and test the VPN connection: Finally, you verify and test the VPN connection to ensure that it is working correctly12.

Reference :=

Configure IOS-XE Site-to-Site VPN Connection to Amazon Web Services - Cisco Community

SD-WAN Configuration Example: Site-to-site (LAN to LAN) IPSec between vEdge and Cisco IOS - Cisco Community

Question: 36

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.

Answer:

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. Reference :=

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

Question: 37

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 IKEv1 to IKEv2 in Cisco vManage. Drag and drop the steps from the left onto the order on the right to complete the

configuration.

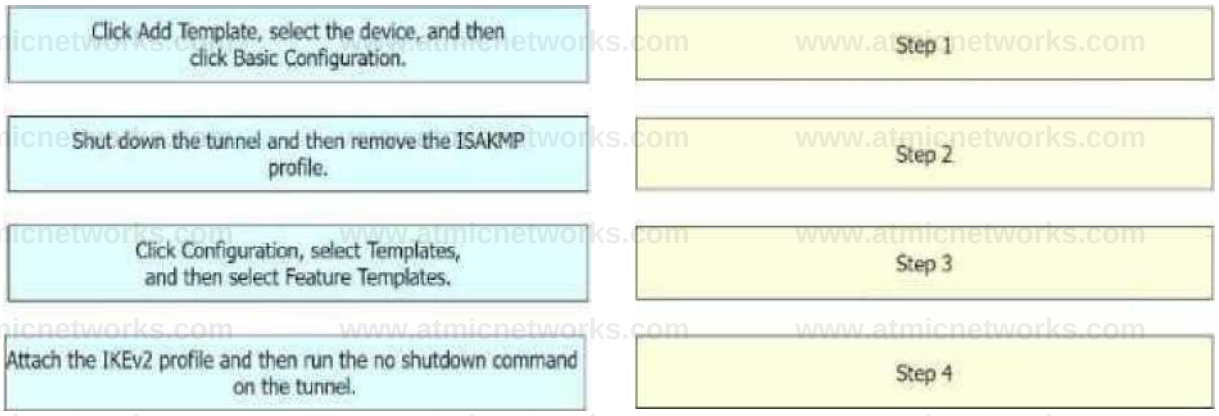

Answer:

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.

Reference :=

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

Question: 38

DRAG DROP

An engineer must use Cisco vManage to configure an application-aware routing policy Drag and drop the steps from the left onto the order on the right to complete the configuration.

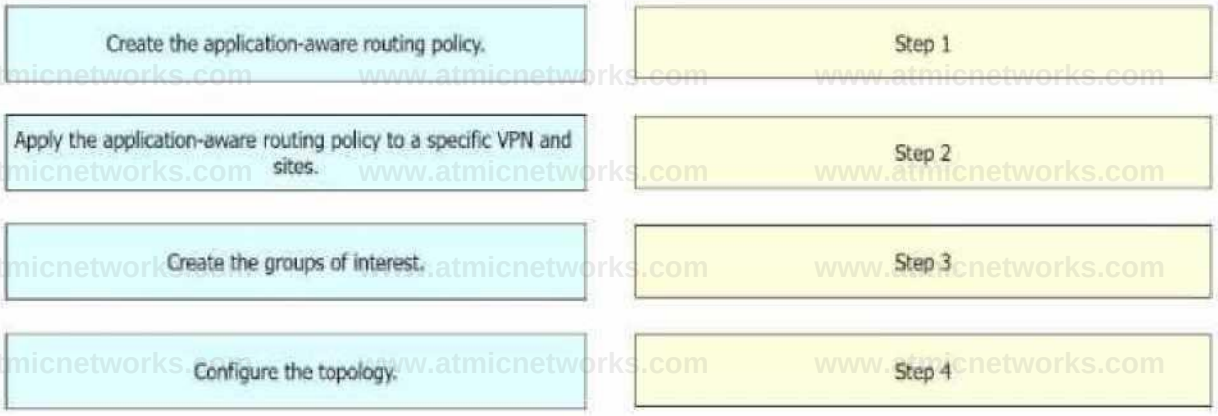

Answer:

Explanation:

Step 1 = Create the groups of interest. Step 2 = Configure the topology. Step 3 = Create the application-aware routing policy. Step 4 = Apply the application-aware routing policy to a specific VPN and sites.

The process of configuring an application-aware routing policy in Cisco vManage involves several steps12.
Create the groups of interest: This is the first step where you define the applications or groups that the policy will affect1.
Configure the topology: This involves setting up the network topology that the policy will operate within1.
Create the application-aware routing policy: After setting up the groups and topology, you then create the application-aware routing policy. This policy tracks network and path characteristics of the data plane tunnels between Cisco SD-WAN devices and uses the collected information to compute optimal paths for data traffic31.
Apply the application-aware routing policy to a specific VPN and sites: Finally, the created policy is applied to a specific VPN and sites. This allows the policy to affect the desired network traffic1. Reference :=
Designing and Implementing Cloud Connectivity (ENCC) v1.0
Learning Plan: Designing and Implementing Cloud Connectivity v1.0 (ENCC 300-440)
Information About Application-Aware Routing - Cisco
Configuring Application-Aware Routing (AAR) Policies | NetworkAcademy.io
Policies Configuration Guide, Cisco IOS XE SD-WAN Releases 16.11, 16.12