PassQueen
a.k.a dumpsqueen.com

Implementing Cisco SD-WAN Solutions (300-415 ENSDWI)

Cisco 300-415

Version Demo

Total Demo Questions: 59

Total Premium Questions: 593

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Architecture	57
Topic 2, Controller Deployment	84
Topic 3, Router Deployment	132
Topic 4, Policies	132
Topic 5, Security and Quality of Service	92
Topic 6, Management and Operations	96
Total	593

QUESTION NO: 1

Which two route attributes are considered by OMP when selecting the best route to a destination prefix? (Choose two.)

A. administrative distance

B. OMP route preference

C. site ID

D. TLOC color

E. VPN name

ANSWER: A B

Explanation:

Administrative distance and **OMP route preference** are considered during OMP route selection. OMP first compares the administrative distance associated with candidate routes. A route with a lower administrative distance is preferred. When candidate routes have the same administrative distance, OMP preference is one of the subsequent attributes used to select the more desirable route.

The site ID identifies the site from which a route originated and is used for loop prevention, but it is not a best-path preference value. The TLOC color identifies the transport type, such as mpls or biz-internet, and can be used in policy, but it is not itself an OMP route-selection attribute. Cisco documents OMP route selection in the Cisco SD-WAN OMP overview.

QUESTION NO: 2

Refer to the exhibit.

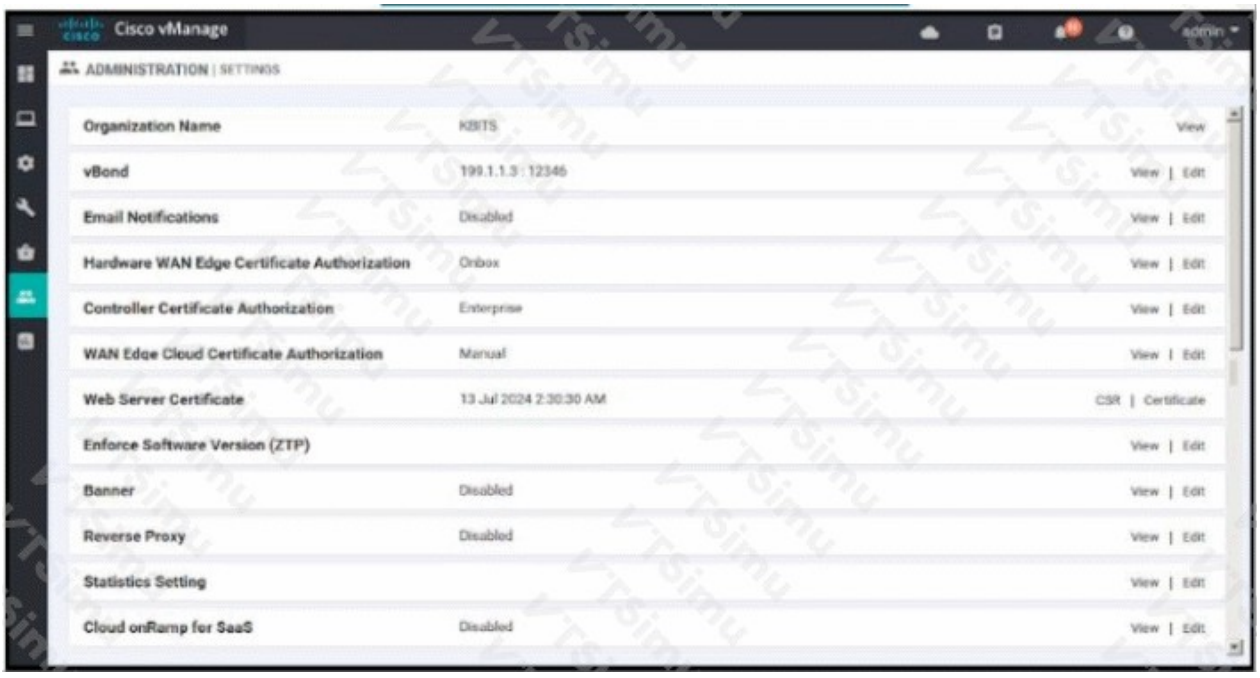

Which two configurations are needed to get the WAN Edges registered with the controllers when certificates are used? (Choose two)

A. Generate a CSR manually within vManage server

B. Generate a CSR manually on the WAN Edge

C. Request a certificate manually from the Enterprise CA server

D. Install the certificate received from the CA server manually on the WAN Edge

E. Install the certificate received from the CA server manually on the vManage

ANSWER: A E

Explanation:

Generate a CSR manually within vManage server and install the certificate received from the CA server manually on the vManage are required steps in an Enterprise CA certificate deployment. vManage must first create a certificate-signing request containing its identity and public key. That CSR is submitted to the organization's certificate authority, which validates and signs it. The resulting signed vManage certificate is then installed on vManage, binding the CA-validated identity to the vManage public key.

This establishes vManage as a trusted component of the SD-WAN control plane under the enterprise PKI. With the controller certificate installed and the enterprise trust chain configured, vManage can participate in authenticated control-plane communication and centrally manage the certificate-related onboarding process for WAN Edge devices. The signed certificate also lets peers validate that they are communicating with the legitimate vManage instance rather than an untrusted system. Cisco documents the Enterprise CA workflow as generating the vManage CSR and installing the CA-signed certificate on vManage before proceeding with the remaining SD-WAN certificate operations. See Cisco's SD-WAN Certificates Deployment Guide and the Cisco SD-WAN certificate configuration documentation.

QUESTION NO: 3

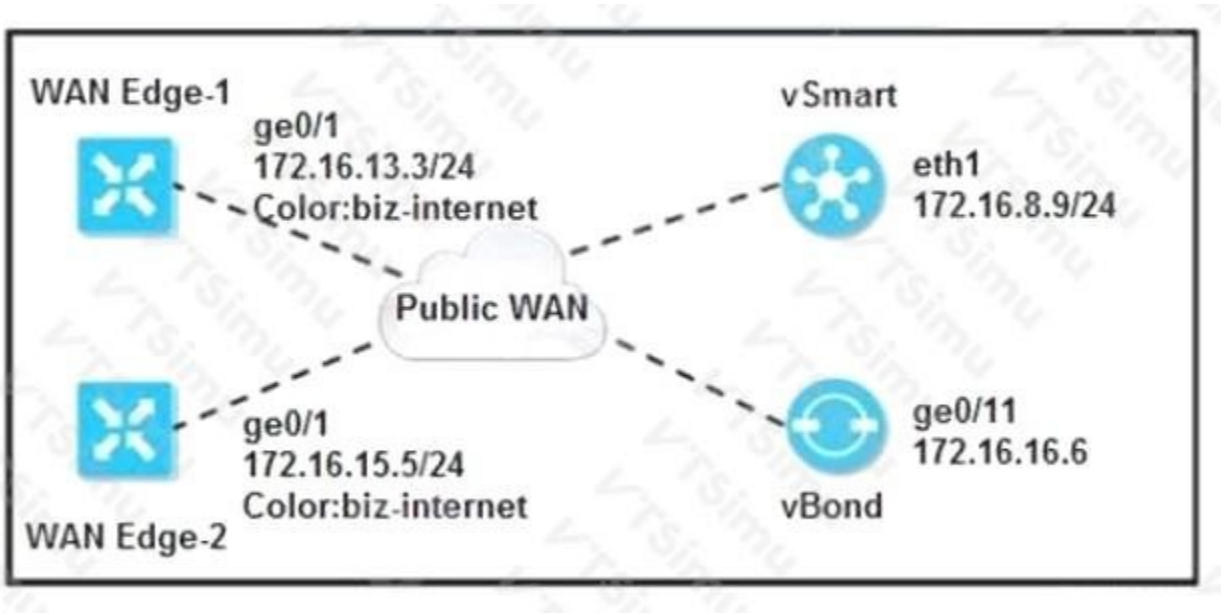

Refer to the exhibit. The tunnel interface configuration on both WAN Edge routers is:

```
vpn 0
 interface ge0/1
  tunnel-interface
   encapsulation ipsec
   color biz-internet
   allow-service dhcp
   allow-service dns
   allow-service icmp
   no allow-service sshd
   no allow-service ntp
   no allow-service stun
  !
  no shutdown
```

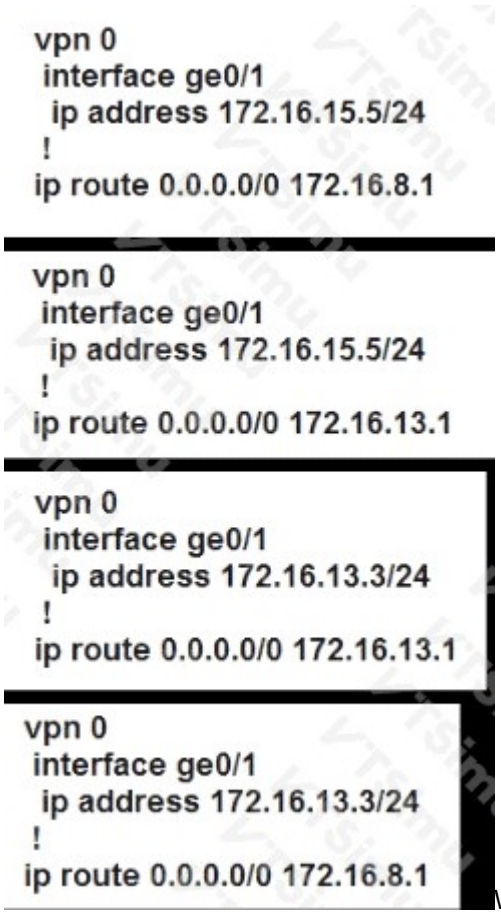

Which configuration for WAN Edge routers will connect to the Internet? A.

B. C.

D.

Answer: B

Explanation:

A. C.
D.

B. B.

ANSWER: B

Explanation:

The configuration identified as “B.” is correct. For a Cisco SD-WAN WAN Edge to establish Internet connectivity while participating in the SD-WAN overlay, its WAN-facing transport interface must have valid underlay IP reachability and be configured as a tunnel interface with the appropriate tunnel properties. The transport tunnel provides the connection between the WAN Edge and SD-WAN control-plane components, while the underlay-facing interface must be able to reach its configured default gateway or otherwise have a usable route toward the Internet. This enables the WAN Edge to resolve and contact controllers, form control connections, and subsequently establish IPsec data-plane tunnels to other WAN Edge devices. Cisco SD-WAN uses the transport tunnel configuration to associate the physical WAN circuit with a transport VPN and a TLOC; correct interface addressing and routing are therefore essential prerequisites for Internet-based transport operation. Cisco’s configuration guidance describes WAN tunnel interfaces as the mechanism used to enable SD-WAN transport connectivity and control-plane communication. See the Cisco tunnel-interface configuration guide and the Cisco SD-WAN systems and interfaces configuration guide.

QUESTION NO: 4

Refer to the exhibit.

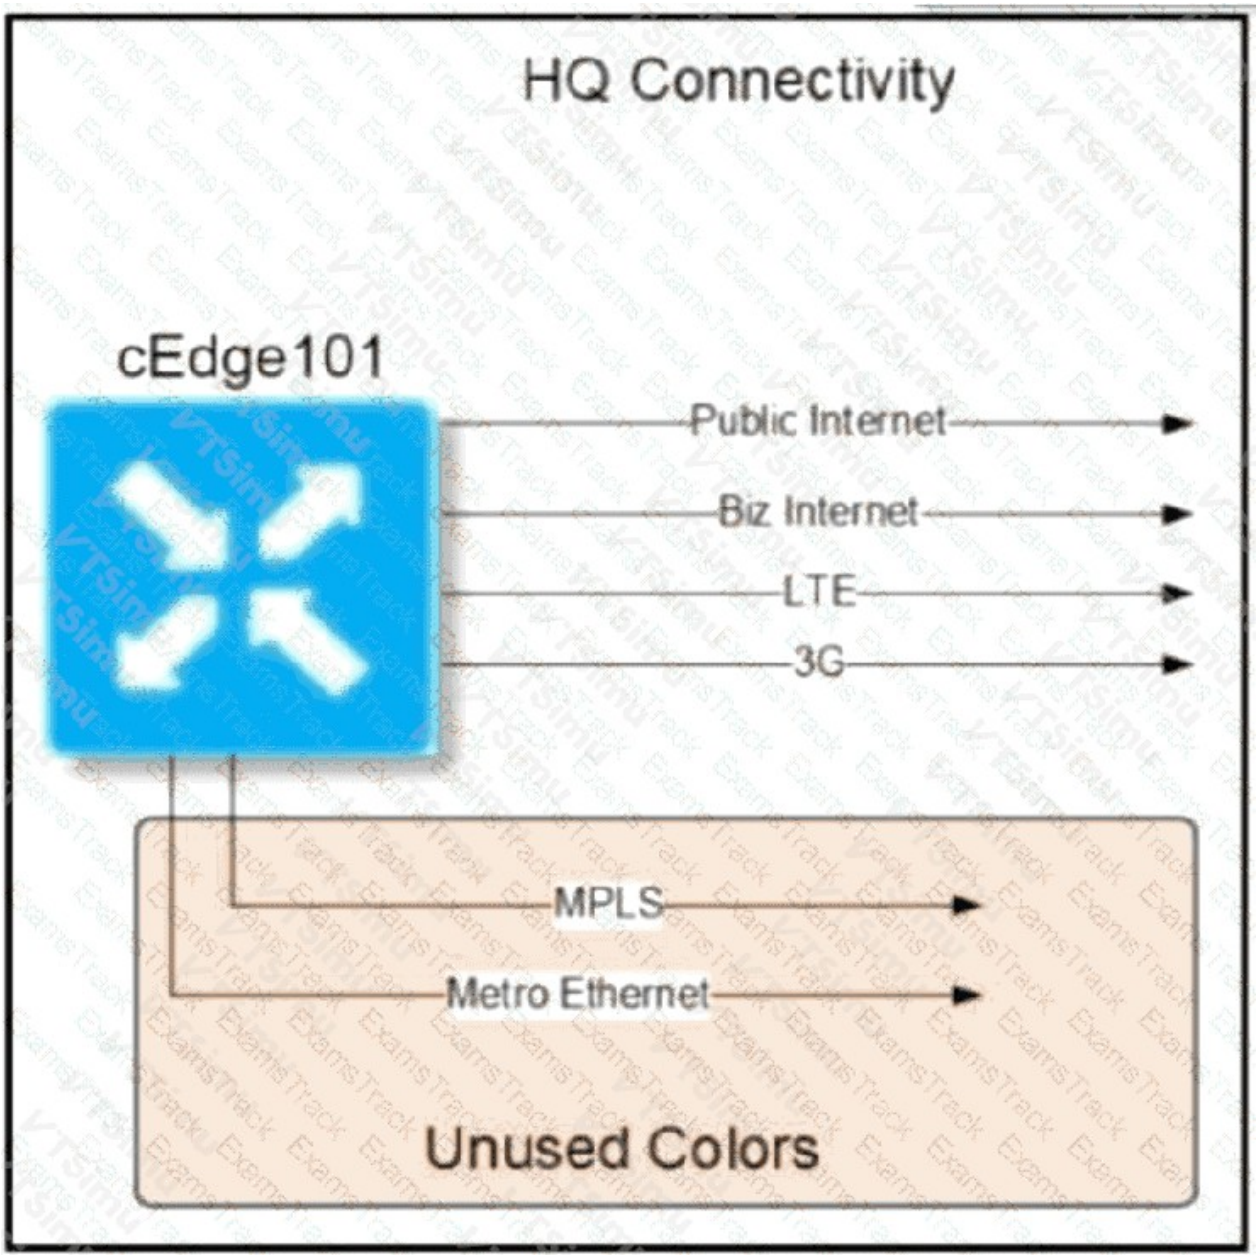

cEdge101 has six possible routes to connect to spokes. However, only four routes are currently in use. Which CLI configuration ensures that all six routes are used?

A. cEdge101# omp | send-path-limit 16

B. cEdge101# omp | send-backup-paths | send-path-limit 16

C. vsmart# omp | send-backup-paths

D. vsmart# omp | no shutdown | send-path-limit 16 | send-backup-paths | graceful-restart

ANSWER: D

Explanation:

vsmart# omp | no shutdown | send-path-limit 16 | send-backup-paths | graceful-restart is correct because the vSmart controller controls the OMP routes that it advertises to WAN edge devices. OMP advertises a maximum of four equal-cost paths for a route by default, so a higher `send-path-limit` is required when cEdge101 must receive and install all six usable paths. Configuring `send-path-limit 16` on vSmart raises the advertisement limit above the six required routes.

In addition, `send-backup-paths` instructs vSmart to advertise eligible backup OMP paths in addition to the preferred paths. This is needed when the available paths include routes that are not selected as the primary best paths but must still be delivered to the edge so they can participate in multipath forwarding. With these settings in the active OMP configuration, cEdge101 receives the complete set of six paths and can use them when the routes meet the normal equal-cost forwarding requirements. The included `no shutdown` keeps the OMP process enabled, while graceful restart preserves OMP routing continuity during a control-plane restart. Cisco documents OMP path advertisement behavior and the relevant configuration controls in its SD-WAN OMP documentation and Cisco SD-WAN Command Reference.

QUESTION NO: 5

What are the two advantages of configuration groups in a Cisco SD-WAN deployment? (Choose two.)

A. Individual devices are associated with a configuration group and a device template.

B. Individual devices are added to multiple groups.

C. Individual devices are grouped based on a shared configuration.

D. A subset of devices is identified with tags.

E. An individual device has multiple tag rules.

ANSWER: C D

Explanation:

Configuration groups provide a scalable way to deploy and maintain common Cisco SD-WAN Manager configuration across devices that have the same intended role or policy requirements. Therefore, **Individual devices are grouped based on a shared configuration.** A group can contain shared feature configuration and associated device-specific settings, allowing an administrator to define the common baseline once and consistently apply it to the devices assigned to that group. This reduces repeated configuration work and helps maintain standardization as the deployment grows.

Configuration groups also support tags to distinguish devices that need a variation within an otherwise common configuration. Thus, **A subset of devices is identified with tags.** Tags let the administrator target a subset of the group when defining configuration values or applying relevant settings, without having to create an entirely separate configuration group for every small difference. This combination of a common group baseline and targeted tag-based differentiation makes configuration management more efficient while retaining the flexibility needed for device or site-specific requirements. Cisco documents configuration groups and tag-based device targeting in its Cisco SD-WAN Configuration Groups documentation and the Cisco Catalyst SD-WAN API documentation.

QUESTION NO: 6

When a WAN Edge device joins the SD-WAN overlay, which Cisco SD-WAN components orchestrates the connection between the WAN Edge device and a vSmart controller?

A. vManage

B. vBond

C. OMP

D. APIC-EM

ANSWER: B

Explanation:

vBond is correct because it acts as the Cisco SD-WAN orchestration component during WAN Edge onboarding and control-plane establishment. After a WAN Edge device authenticates to the overlay, it contacts the vBond orchestrator. vBond validates the device identity and coordinates the discovery of the appropriate controllers, including vSmart controllers. It

supplies the WAN Edge with controller connection information and helps determine reachable transport addresses, supporting operation across NAT and firewall boundaries. This orchestration process enables the WAN Edge to establish its secure DTLS or TLS control connections to the vSmart controller, through which the device receives centralized routing and policy information. In a multitenant or geographically distributed deployment, vBond can also direct devices toward suitable controllers based on configured controller lists and reachability. Cisco describes vBond as the orchestration plane component responsible for bringing WAN Edge routers into the SD-WAN fabric and facilitating their controller connections. See Cisco's SD-WAN overview and its vBond configuration documentation.

QUESTION NO: 7

A network administrator needs a WAN Edge router to retain a local TLOC route when the associated tunnel interface is operational but BFD sessions to remote peers are down. Which tunnel-interface setting must be configured?

A. no-detect

B. restrict

C. shutdown

D. max-control-connections

ANSWER: A

Explanation:

no-detect is correct. By default, a WAN Edge router can withdraw its local TLOC route when BFD detects loss of reachability through the associated tunnel interface. Configuring `no-detect` prevents this BFD-based TLOC withdrawal behavior, allowing the local TLOC to remain advertised while the tunnel interface itself remains operational.

This setting is used only when the design requires the TLOC to remain available despite the absence of BFD sessions. A tunnel-interface shutdown disables the interface, and restrict controls whether remote devices can automatically establish tunnels to the TLOC. Cisco documents tunnel-interface and TLOC behavior in the Cisco SD-WAN TLOC documentation.

QUESTION NO: 8

Refer to the exhibit.

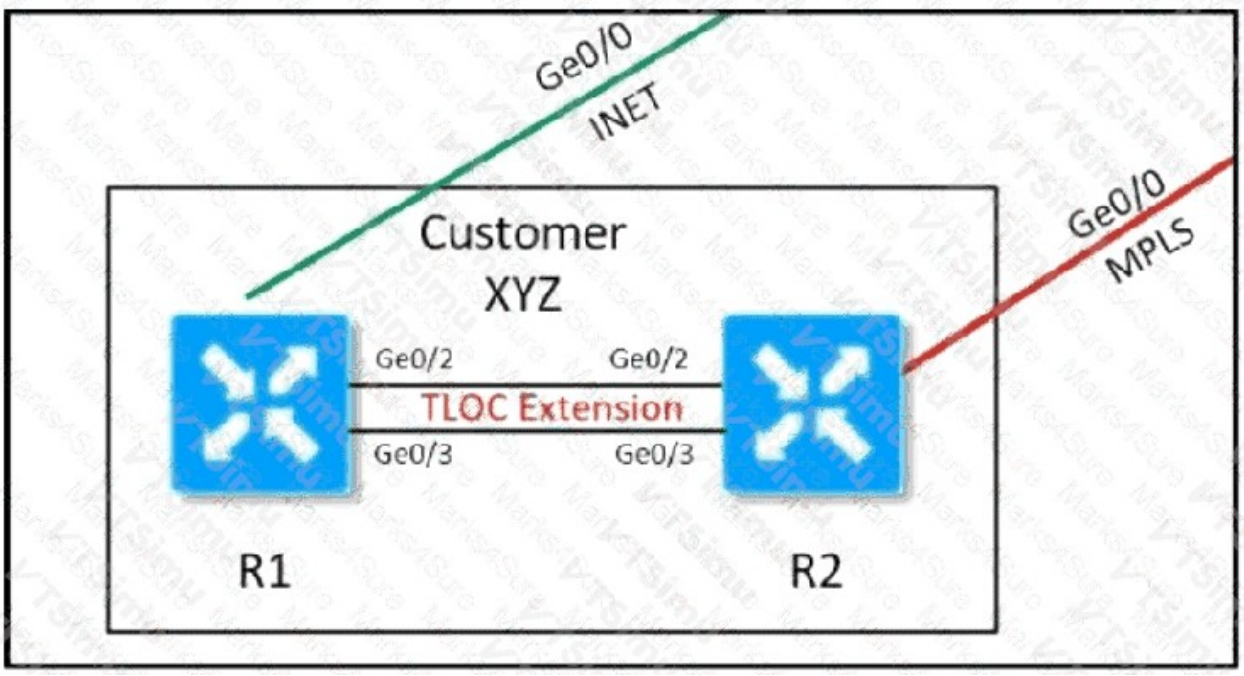

Customer XYZ cannot provison dual connectivity on both Its routers due to budget constratnts but wants to use tnth RI and R2 interface for users behind them for load toward the hub site Which configurauon achieves this objectives?

A)

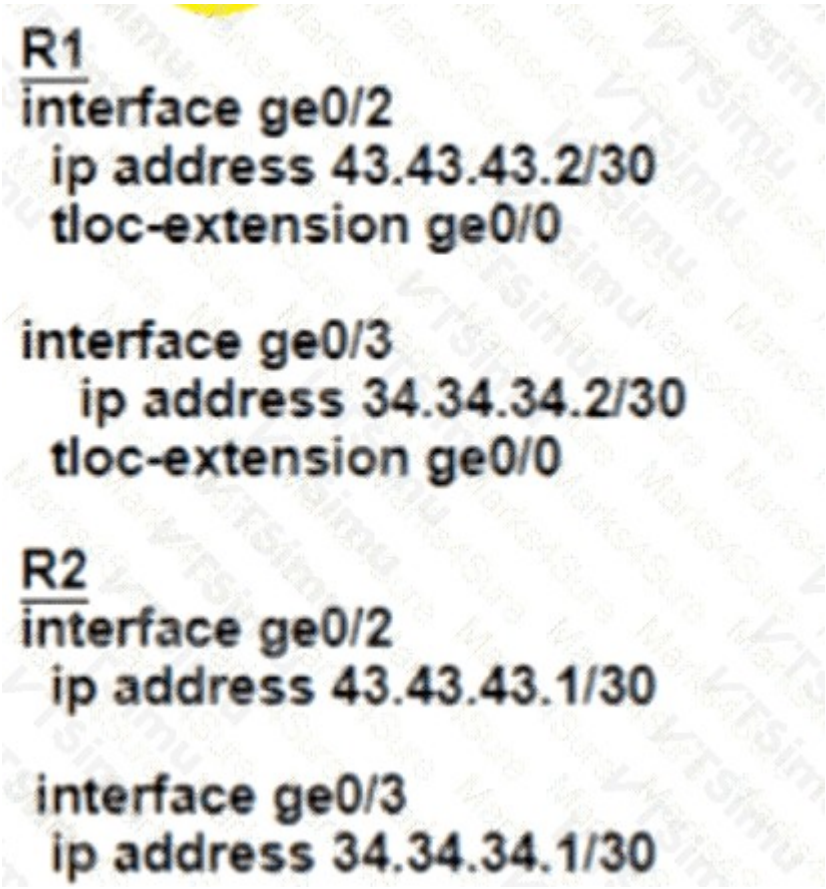

B)

```
R1
interface ge0/2
 ip address 43.43.43.2/30
 tloc-extension ge0/0

interface ge0/3
 ip address 34.34.34.1/30
 tunnel-interface
  color mpls

R2
interface ge0/2
 ip address 43.43.43.1/30
 tunnel-interface
    color public-internet

interface ge0/3
 ip address 34.34.34.2/30
 tloc-extension ge0/0
```

C)

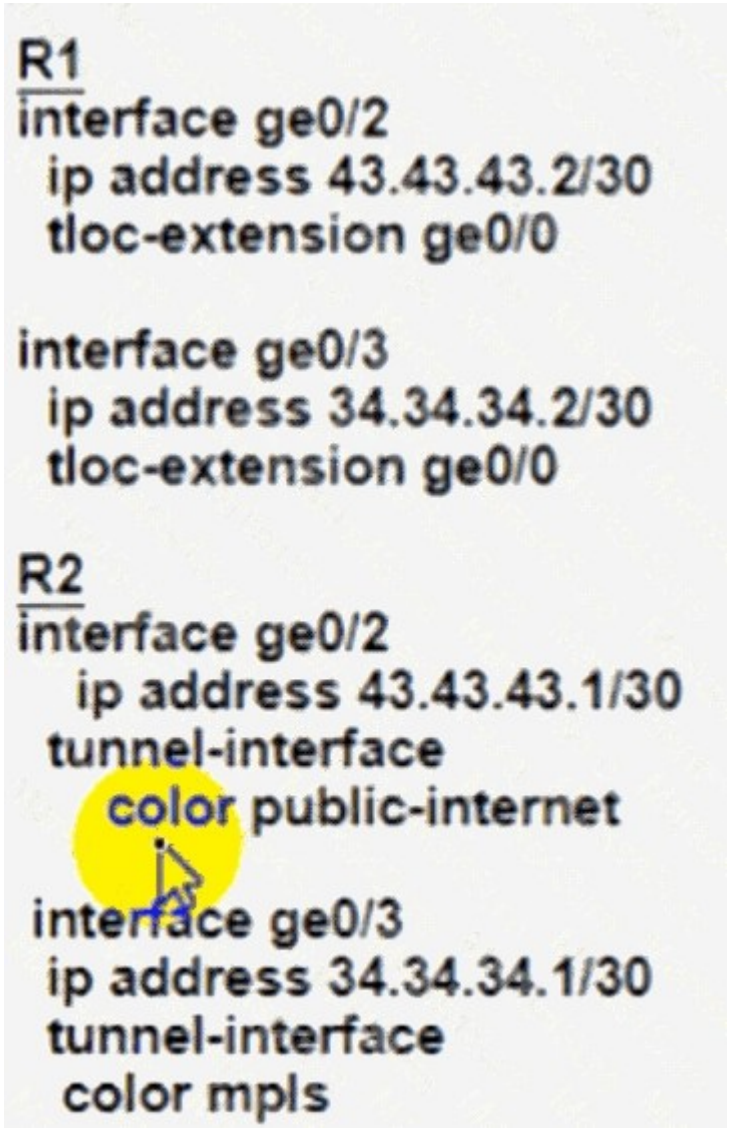
R1
interface ge0/2
ip address 43.43.43.2/30
tloc-extension ge0/0
interface ge0/3
ip address 34.34.34.2/30
tloc-extension ge0/0
R2
interface ge0/2
ip address 43.43.43.1/30
tunnel-interface
color public-internet
interface ge0/3
ip address 34.34.34.1/30
tunnel-interface
color mpls

D)

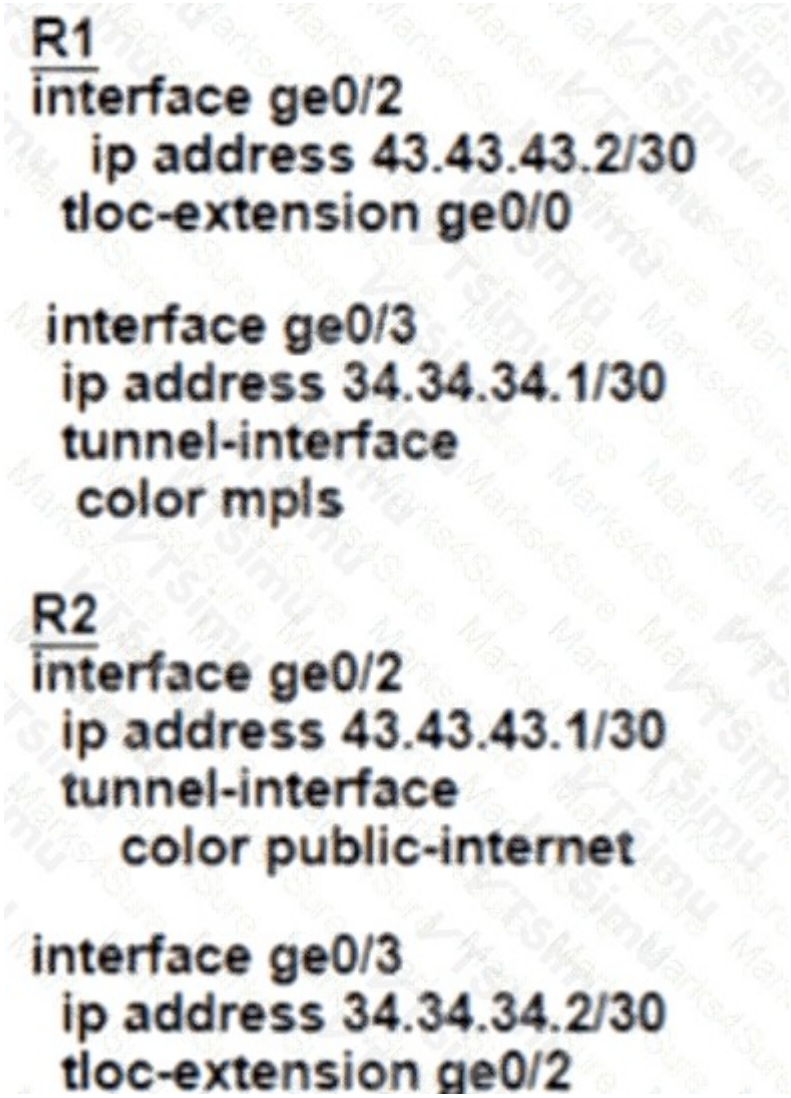

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: A

Explanation:

The selected configuration is correct because it implements a TLOC-extension design between the two SD-WAN edge routers. TLOC extension is intended for sites where each router has only one independent WAN transport, but the site still requires both routers to use both available transports. Each router retains its locally connected WAN circuit and advertises its own transport locator, while the inter-router extension connection lets it reach and use the peer router's WAN transport as an extended TLOC.

This design provides transport redundancy without requiring a second provider circuit on each router. It also allows SD-WAN policy and path selection to distribute application traffic toward the hub across the available WAN paths, including a locally connected transport and a transport reached through the peer. The extension interfaces must be placed in VPN 0 and be reachable directly between the routers so that the control plane and data plane can use the extended transport correctly. This arrangement is specifically useful when cost constraints prevent dual-homing each individual router, while the site still needs resilient, load-sharing connectivity.

Cisco's TLOC-extension configuration guidance is available in the Cisco TLOC Extension configuration document and the Cisco SD-WAN System Interfaces guide.

QUESTION NO: 9

Which actions must be taken to allow certain departments to require firewall protection when interacting with data center network without including other departments? (Choose two.)

A. Use classification policing and marking

B. Apply data policies at vEdge.

C. Deploy a service-chained firewall service per VPN

D. The regional hub advertises the availability of the firewall service

E. Advertise lo vSmart controllers

ANSWER: C D

Explanation:

Deploying a service-chained firewall service per VPN provides the required departmental separation. In Cisco SD-WAN, a VPN represents an isolated routing domain, so associating the firewall service with only the VPNs used by departments that require inspection ensures that their data-center-bound traffic is steered through the firewall. Other departmental VPNs do not need to be included in that service chain and can retain their independent forwarding behavior. This approach allows security policy, inspection, and routing treatment to be applied at the intended VPN scope rather than globally.

The regional hub must also advertise the availability of the firewall service so that the SD-WAN control plane can distribute the service information to the sites and VPNs that need it. Service chaining depends on control-plane advertisements to identify the service location and the transport path toward it. Once the service is advertised, relevant remote sites can resolve the service and steer matching traffic toward the regional hub firewall before it reaches the data center. Cisco documents service chaining as the mechanism for redirecting selected VPN traffic through network services while maintaining VPN segmentation. See the Cisco SD-WAN Security Configuration Guide and the Cisco SD-WAN solution documentation.

QUESTION NO: 10

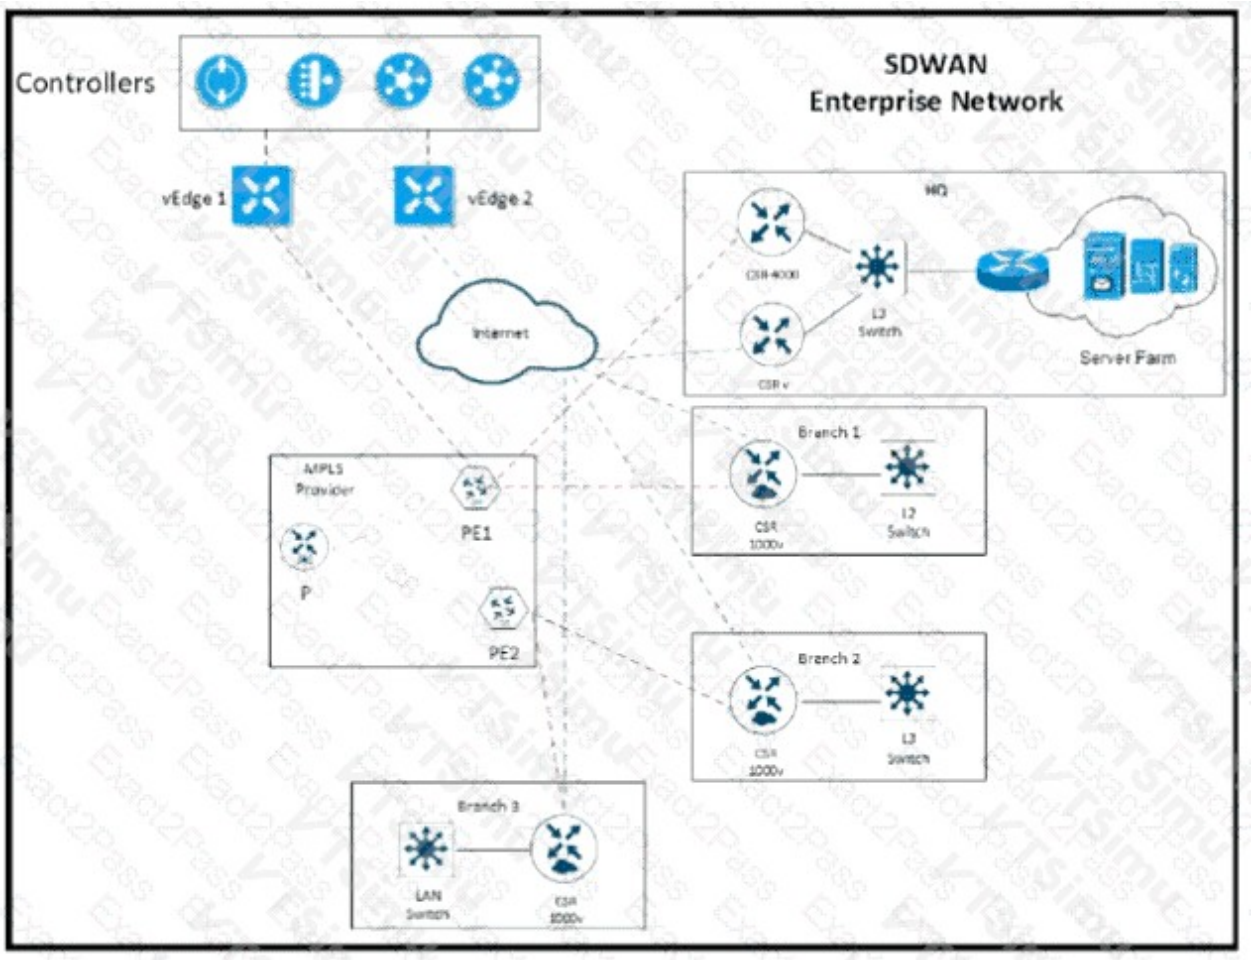

Refer to the exhibit The network team must configure EIGRP peering at HQ with devices in the service VPN connected to WAN Edge CSRv. CSRv is currently configured with

```
viptela-system:system
 device-model      vedge-cloud
 host-nam                  CSRv
 location          CA
 system-ip                 10.4.4.4
 domain-id                 1
 site-id           1
 organization-name ABC
 clock timezone US/Newyork
 vbond 10.10.0.1 port 12346

omp
 no shutdown
 graceful-restart
```

Which configuration on the WAN Edge meets the requirement?

Which configuration on the WAN Edge meets the requiremnet

A)

```
vpn 0
 router
  eigrp
   address-family ipv4-unicast
    maximum-path paths 4
    network 10.0.0.0/28
    redistribute omp

interface ge0/0
 description "*** Internet Circuit ***"
 ip address 209.165.200.229/30
 tunnel-interface
  encapsulation ipsec
  color public-internet
  allow-service all

vpn 512
interface eth0
 shutdown
```

B)

```
vpn 0
vpn 1
interface ge0/0
 description "*** Internet Circuit ***"
 ip address 209.165.200.229/30
 tunnel-interface
  encapsulation ipsec
  color public-internet
  allow-service all

 router
  eigrp
   address-family ipv4-unicast
    maximum-path paths 4
    network 10.0.0.0/28
    redistribute omp

interface ge0/1
 ip address 10.0.0.1/28

vpn 512
interface eth0
 shutdown
```

C)

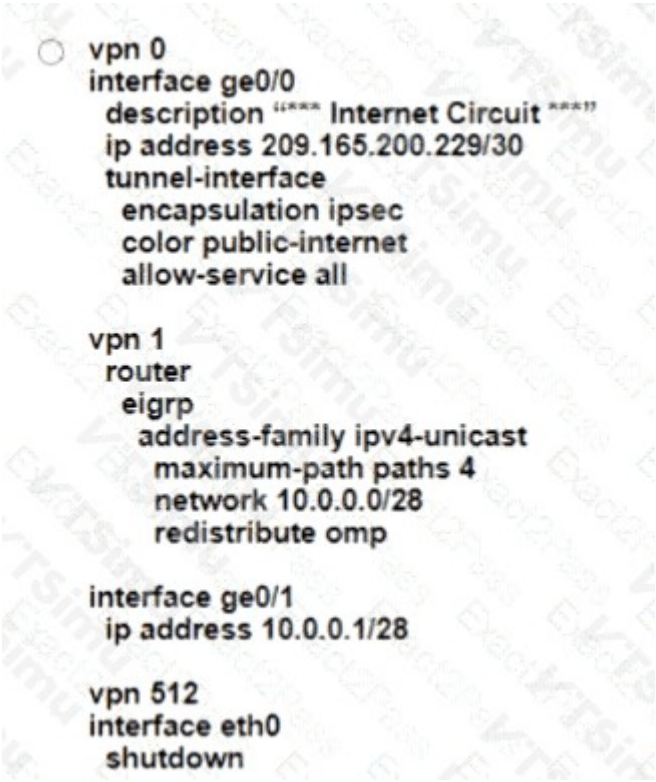

D)

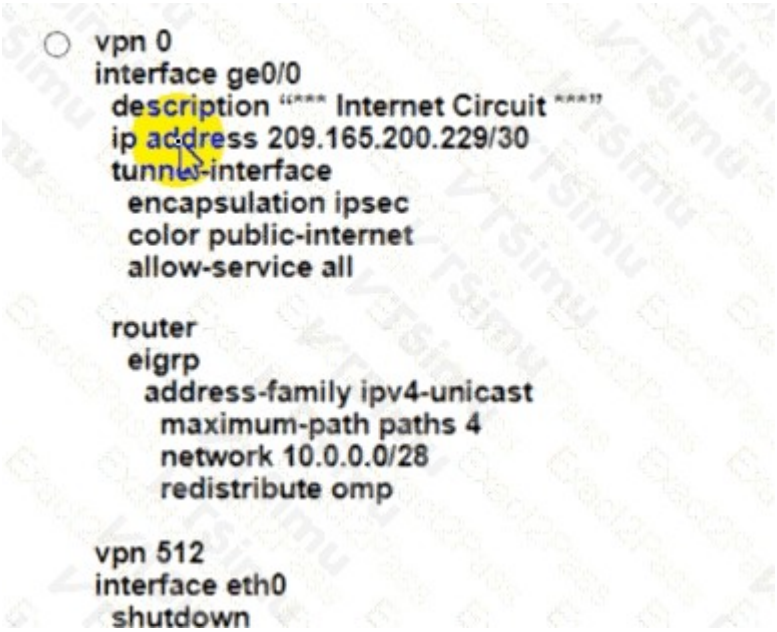

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: C

Explanation:

Option C is correct because EIGRP peering with a device connected to a Cisco SD-WAN WAN Edge must be configured within the relevant service VPN, rather than in VPN 0, which is reserved for transport connectivity. The EIGRP process must use the required autonomous-system value and enable the Layer 3 interface that faces the HQ EIGRP neighbor. This allows

the WAN Edge to exchange EIGRP routing information locally in the service VPN and makes those learned routes available to the SD-WAN routing infrastructure according to the configured route-policy and redistribution behavior.

Cisco SD-WAN supports EIGRP as a service-side routing protocol on supported WAN Edge platforms. The configuration is applied under the service VPN configuration hierarchy, with EIGRP parameters and interface participation defined for that VPN. This preserves the intended separation between the SD-WAN transport underlay and enterprise service-side routing. The configuration shown in Option C places the EIGRP elements in the appropriate service-VPN context and supplies the information needed for the CSRv WAN Edge to establish the required adjacency with the HQ-side devices. Cisco's configuration guidance for service-side routing protocols is available in the Cisco SD-WAN Routing Configuration Guide and the Cisco SD-WAN System and Interfaces Configuration Guide.

QUESTION NO: 11

Refer to the exhibit vManage and vBond have an issue establishing a connection to vSmart Which two actions does the administrator take to fix the issue? (Choose two)

A. Install the certificate received from the certificate server.

B. Manually resync vManage and vBond

C. Reconfigure the vSmart from CLI with the proper Hostname & System IP

D. Delete and re-add vSmart Click Generate and validate CSR

E. Request a certificate from the certificate server based on the CSR for the vSmart

ANSWER: A E

Explanation:

"Request a certificate from the certificate server based on the CSR for the vSmart" and "Install the certificate received from the certificate server." are the required steps when the vSmart controller does not yet have a valid signed controller certificate. The certificate-signing request contains the controller identity and public-key information that must be signed by the enterprise certificate authority or Cisco certificate server. Requesting the certificate using that CSR produces the signed certificate that establishes trust for the vSmart.

After the certificate authority returns the signed certificate, it must be installed through vManage for the appropriate vSmart. Installation binds the trusted certificate to the controller identity and allows the controller to authenticate securely as part of the SD-WAN control plane. Once that certificate workflow is complete, the vSmart can establish authenticated DTLS or TLS control connections with vBond and vManage, allowing it to participate in the SD-WAN overlay. Cisco documents controller certificate management and the CSR-to-signed-certificate installation workflow in its SD-WAN installation and configuration guides and Cisco SD-WAN developer documentation.

QUESTION NO: 12

How is the scalability of the Manage increased in Cisco SD-WAN Fabric?

A. Increase the bandwidth of the WAN link connected to the vManage

B. Increase licensing on the vManage

C. Deploy more than one vManage controllers on different physical server

D. Deploy multiple vManage controllers in a cluster

ANSWER: D

Explanation:

Deploy multiple vManage controllers in a cluster is correct. In Cisco SD-WAN, vManage is the centralized management plane component responsible for configuration, monitoring, software management, and operational visibility across the

fabric. A vManage cluster combines multiple vManage instances into one logical management system, allowing the cluster to distribute management workload and support a larger number of WAN Edge devices than a standalone instance. The cluster architecture also provides high availability: vManage instances synchronize relevant data, and the system can continue providing management services if an individual cluster member becomes unavailable. Cisco supports forming vManage clusters with multiple nodes, with the appropriate cluster configuration and shared storage/database design required for the deployed release and platform. This approach is specifically intended to increase both scale and resilience for larger SD-WAN deployments, rather than merely increasing the resources or network bandwidth of one server. For sizing limits, Cisco publishes validated scale guidance that varies by software version, hardware or virtual-machine profile, and enabled management features. See Cisco's vManage cluster documentation and the Cisco SD-WAN release notes and scalability information.

QUESTION NO: 13

What is the procedure to upgrade all Cisco SD-WAN devices to a recent version?

A. The upgrade is performed for a group of WAN Edge devices first to ensure data-plane availability when other controllers are being updated.

B. Upgrade and reboot are performed first on vManage, then on vBond, then on vSmart, and finally on the Cisco WAN Edge devices.

C. The upgrade is performed first on vManage, then on WAN Edge devices, then on vBond, and finally on vSmart. The reboot must start from WAN Edge devices.

D. Upgrade and reboot are performed first on vBond, then on vSmart, and finally on the Cisco WAN Edge devices.

ANSWER: B

Explanation:

"Upgrade and reboot are performed first on vManage, then on vBond, then on vSmart, and finally on the Cisco WAN Edge devices." is correct because Cisco SD-WAN software upgrades must preserve management-plane and control-plane compatibility throughout the process. vManage is upgraded first because it orchestrates the lifecycle workflow and must be able to manage the target software versions on the remaining fabric components. The vBond orchestrators are then upgraded, followed by the vSmart controllers, so that the control-plane elements are updated before the WAN Edge routers that depend on them for orchestration and route-policy distribution.

After the management and controller components are operating on the intended release, the Cisco WAN Edge devices can be upgraded in staged groups. Staging edge upgrades helps maintain forwarding availability for sites with redundant WAN Edge routers and allows validation of the new release before wider deployment. Before beginning, administrators should confirm the Cisco-supported upgrade paths and release interoperability, upload the required images to vManage, and take configuration and operational backups. Cisco's SD-WAN upgrade guidance documents this controller-first sequence and emphasizes careful planning for high availability and rollback readiness. See Cisco SD-WAN Software Upgrade Best Practices.

QUESTION NO: 14

Which two WAN Edge devices should be deployed in a cloud? (Choose two.)

A. vEdge 5000v

B. ASR 1000v

C. CSR 1000v

D. vEdge 100wm

E. vEdge cloud

ANSWER: C E

Explanation:

CSR 1000v and **vEdge cloud** are virtual WAN Edge platforms designed for deployment in cloud environments. CSR 1000v is Cisco's virtual-router form factor and can operate as a Cisco SD-WAN WAN Edge in public-cloud and virtualized infrastructures. It provides the WAN Edge functions needed to establish the secure SD-WAN overlay, connect cloud workloads to branch and data-center sites, and apply centralized SD-WAN policy.

vEdge cloud is the cloud-hosted virtual vEdge appliance. It is intended to run as a virtual machine in supported cloud or virtualized environments and provides the same fundamental SD-WAN overlay role: it authenticates with the Cisco SD-WAN controllers, forms secure control and data-plane tunnels, and provides connectivity between cloud networks and the rest of the SD-WAN fabric. These virtual platforms are therefore appropriate when WAN Edge connectivity must be placed close to applications and resources hosted in cloud infrastructure, rather than at a physical branch location.

Cisco's platform documentation identifies virtual-cloud WAN Edge form factors and their deployment roles. See the Cisco vEdge Cloud Router documentation and the Cisco CSR 1000v data sheet.

QUESTION NO: 15

Which combination of platforms are managed by vManage?

A. ISR4321, ASR1001, ENCS, ISRv

B. ISR4351, ASR1002HX, vEdge2000, vEdge Cloud

C. ISR4321, ASR1001, Nexus, ENCS

D. ISR435l, ASRl009, vEdge2000, CSR1000v

ANSWER: B

Explanation:

ISR4351, ASR1002HX, vEdge2000, vEdge Cloud is correct because Cisco vManage (now called Cisco Catalyst SD-WAN Manager in current documentation) centrally manages both Cisco IOS XE SD-WAN edge routers and legacy vEdge devices. The ISR4351 is an ISR 4000 Series platform that can operate as a Cisco IOS XE SD-WAN (cEdge) router, while the ASR1002-HX is an ASR 1000 Series platform supported as a high-performance IOS XE SD-WAN edge router. vEdge 2000 is a physical legacy vEdge platform, and vEdge Cloud is the virtualized legacy vEdge offering for cloud deployments. This mix demonstrates vManage's role as the centralized management plane across physical branch devices, higher-capacity aggregation or hub routers, and virtual cloud WAN edges. From vManage, an administrator can onboard devices, push templates and configuration, apply centralized policy, monitor operational status, and manage the SD-WAN fabric consistently across these supported platform families. Cisco's SD-WAN compatibility documentation identifies the supported IOS XE SD-WAN and vEdge hardware and virtual platforms, while Cisco's SD-WAN FAQ describes vManage as the centralized management component for the solution. See the Cisco SD-WAN Compatibility Matrix and the Cisco SD-WAN FAQ.

QUESTION NO: 16

What are the two requirements for plug-and-play provisioning on Cisco IOS XE SD-WAN devices? (Choose two.)

A. The gateway router for the WAN Edge device must be able to reach devicehelper.cisco.com.

B. The gateway router for the WAN Edge device must be able to reach public DNS servers.

C. The gateway router for the WAN Edge device must be able to reach ztp.viptela.com.

D. Devices at branch offices must be able to reach the Cisco SD-WAN vSmart controller at the headquarters site.

E. The WAN Edge device must have a valid certificate.

ANSWER: A E

Explanation:

For Cisco IOS XE SD-WAN Plug and Play provisioning, the WAN Edge must have outbound network connectivity through its default gateway to Cisco's Plug and Play service, `devicehelper.cisco.com`. During initial startup, the device obtains basic IP connectivity, uses the Cisco PnP service to identify its assigned controller and configuration, and then begins its automated onboarding workflow. Therefore, the gateway path used by the WAN Edge must permit the device's access to that Cisco-hosted service.

The WAN Edge device must also possess a valid device certificate. Cisco uses the device identity certificate to authenticate the device securely during onboarding and to establish trust before the device receives its SD-WAN configuration and joins the overlay. On supported IOS XE SD-WAN platforms, this identity is typically based on the manufacturer-installed SUDI certificate, unless the deployment uses an alternative supported certificate process. These two conditions provide both the network path to initiate PnP and the cryptographic identity required to authorize the device. Cisco's PnP overview is available at Cisco Plug and Play documentation, and Cisco's SD-WAN documentation is available at Cisco SD-WAN support documentation.

QUESTION NO: 17

Which two types of SGT propagation are supported by Cisco TrustSec? (Choose two.)

A. reconciliation

B. SXP

C. offline tagging

D. key chain

E. inline tagging

ANSWER: B E

Explanation:

Cisco TrustSec supports propagation of Security Group Tags through **SXP** and **inline tagging**. SXP, or Security Group Tag Exchange Protocol, provides IP-to-SGT binding exchange between TrustSec-capable devices. This allows the receiving device to learn the security-group classification associated with an endpoint IP address even when the network path, intermediate devices, or endpoint attachment does not preserve a tag in packet headers. SXP is therefore useful for extending group-based policy across parts of the network that cannot transport TrustSec tags natively.

Inline tagging carries the SGT directly with the traffic frame by using the TrustSec Security Group Tag field in the Cisco metadata header. Network devices that support TrustSec inline tagging can read the tag as traffic traverses the network and enforce Security Group Access Control List policy based on the source and destination security groups. Together, these mechanisms allow TrustSec policy to remain based on identity and group membership rather than solely on IP addresses or network location. Cisco documents both inline tagging and SXP as supported SGT propagation mechanisms in its TrustSec configuration guidance. See Cisco TrustSec SGT Propagation Configuration Guide and Cisco TrustSec overview.

QUESTION NO: 18

An engineer configures an application-aware routing policy for a group of sites The locations depend on public and private transports The policy does not work as expected when one of the transports does not perform properly This policy is configured:

which configuration completes the policy so that it works for all locations?

A)

B)

C)

D)

A. Option A

B. Option B

C. Option C

D. Option D

E. Configure the intended primary transport color as the preferred color and the alternate public or private transport color as the backup preferred color in the application-aware routing policy.

ANSWER: E

Explanation:

Configure the intended primary transport color as the preferred color and the alternate public or private transport color as the backup preferred color in the application-aware routing policy. Application-aware routing evaluates the configured SLA class continuously for matching application traffic. When the path associated with the preferred color no longer meets the SLA requirements, the policy can select a tunnel using the backup preferred color, provided that the alternate path meets the same SLA criteria. This gives sites using both public and private underlay transports a deterministic failover path instead of relying solely on the primary transport. The preferred and backup colors should correspond to the transport colors actually available at the participating sites, such as a private MPLS color and a public Internet color. Cisco SD-WAN uses these color preferences together with BFD-derived path measurements and the SLA thresholds to steer traffic dynamically while preserving the application policy intent. For Cisco's description of application-aware routing and SLA-based path selection, see Cisco SD-WAN Centralized Policy Configuration Guide and Cisco SD-WAN Tunnel Interface Configuration Guide.

QUESTION NO: 19

Refer to the exhibit.

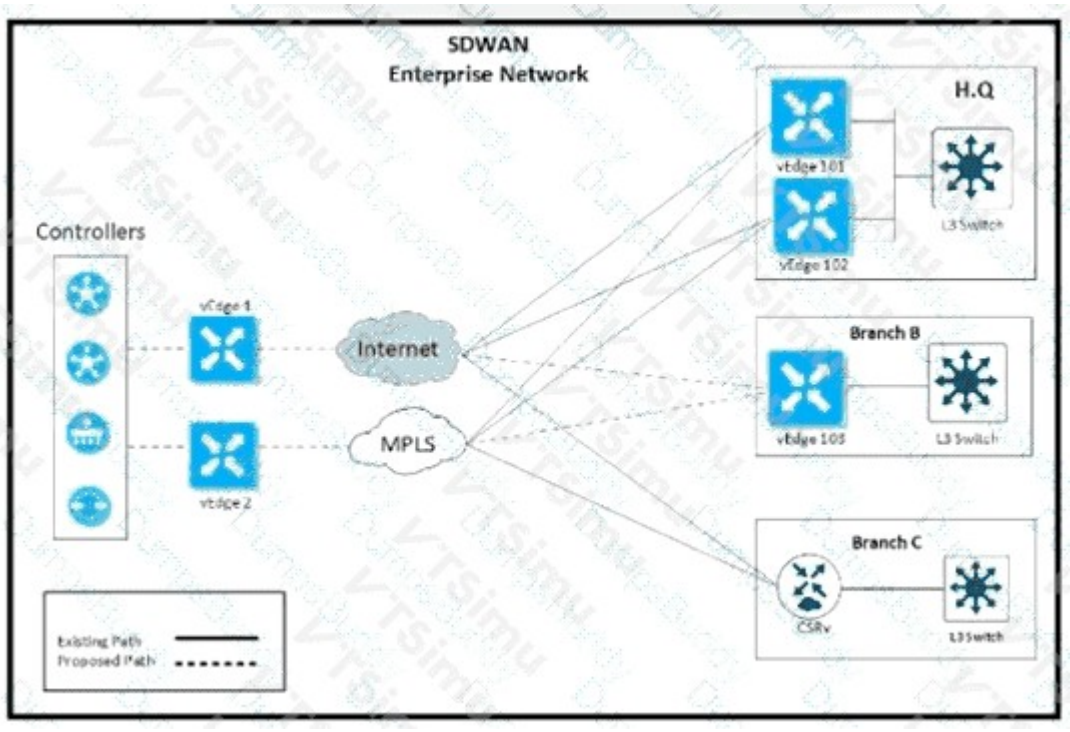

The network team must configure branch B WAN Edge device 103 to establish dynamic full-mesh IPsec tunnels between all colors with branches over MPLS and Internet circuits. The branch ts configured with:

```
viptela-system:system
 device-model        vedge-cloud
 host-nam                    vEdge-Cloud
 location            CA
 system-ip                   10.4.4.4
 domain-id                   1
 site-id             1
 organization-name ABC
 clock timezone US/Newyork
 vbond 10.10.0.1 port 12346
```

```
omp
 no shutdown
 graceful-restart

vpn 512
interface eth0
 ip address 10.0.0.1/24
 no shutdown
```

Which configuration meets the requirement?

A)

```
vpn 0
interface ge0/0
 description "*** Internet Circuit ***"
 ip address 209.165.200.225/30
 tunnel-interface
  encapsulation ipsec
  color public-internet
  allow-service all

interface ge0/1
 description "*** MPLS ***"
 ip address 10.10.0.101/30
 tunnel-interface
  encapsulation ipsec
  color mpls
  allow-service all
```

B)

```
vpn 0
interface ge0/0
 description "*** Internet Circuit ***"
 ip address 209.165.200.225/30
 tunnel-interface
  encapsulation ipsec
  color public-internet restrict
  allow-service all

interface ge0/1
 description "*** MPLS ***"
 ip address 10.10.0.101/30
 tunnel-interface
  encapsulation ipsec
  color mpls restrict
  allow-service all
```

C)

```
vpn 0
interface ge0/0
  description "*** Internet Circuit ***"
  ip address 209.165.200.225/30
  tunnel-interface
    encapsulation ipsec
    color public-internet restrict
    allow-service all

interface ge0/1
  description "** MPLS ***"
  ip address 10.10.0.101/30
  tunnel-interface
    encapsulation ipsec
    color mpls restrict
    allow-service all
```

D)

```
vpn 0
interface ge0/0
  description "*** Internet Circuit ***"
  ip address 209.165.200.225/30
  tunnel-interface
    encapsulation ipsec
    color public-internet
    allow-service all

vpn 1
interface ge0/1
  description "** MPLS ***"
  ip address 10.10.0.101/30
  tunnel-interface
    encapsulation ipsec
    color mpls
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: A

Explanation:

The selected configuration is correct because it enables the WAN Edge to build IPsec data-plane tunnels through each applicable transport tunnel interface while preserving the MPLS and Internet transport colors. In Cisco SD-WAN, a tunnel interface represents a TLOC. The color assigned to that tunnel interface identifies the transport domain advertised to vSmart controllers, and the controller uses the advertised TLOC information to distribute tunnel endpoint information to eligible remote WAN Edges. This allows WAN Edges to establish dynamic direct IPsec tunnels, creating the required full-mesh connectivity rather than forcing traffic through a hub.

The configuration also permits tunnel formation across the required transport colors. Therefore, branches can use MPLS-to-MPLS, Internet-to-Internet, and any policy-permitted cross-color paths as determined by centralized control policy and TLOC reachability. The tunnel interface must be enabled on the WAN-facing VPN 0 transport interfaces, because SD-WAN control and data-plane tunnel establishment occurs through VPN 0 transport locators. Cisco documents the relationship among tunnel interfaces, colors, TLOC advertisements, and IPsec tunnel creation in its tunnel-interface configuration guide and Cisco SD-WAN overview.

QUESTION NO: 20

An engineer must redirect traffic from branch users through a firewall hosted at a data center before it reaches a private application. Which two configurations are required to implement this service insertion design? (Choose two.)

A. Configure a service route for the firewall.

B. Configure a centralized data policy with a service action.

C. Configure a localized control policy on each branch WAN Edge router.

D. Configure a BFD template on the firewall interface.

E. Configure the firewall as a vBond Orchestrator.

ANSWER: A B

Explanation:

A **service route** is required so that the Cisco SD-WAN overlay advertises reachability to the firewall service. Service routes identify a network service that is connected to a WAN Edge router, allowing other WAN Edge routers to learn how to reach the service insertion point through OMP.

A **centralized data policy** is also required to classify the intended traffic and apply the service action. The policy can match traffic by criteria such as source prefix, destination prefix, application, protocol, or DSCP value, and then steer matching packets to the advertised service. Together, the service route provides service reachability and the centralized data policy redirects selected traffic through that service. Cisco documents service chaining and centralized data-policy processing in the Cisco SD-WAN Policies Configuration Guide.

QUESTION NO: 21

How many vCPUs and how much RAM are recommended to run the vSmart controller on the KVM server for 251 to 1000 devices in software version 20.4.x?

A. 4vCPUs. 16 GB

B. 4 vCPUs. 8 GB

C. 8vCPUs. 16 GB

D. 2vCPUs.4GB

ANSWER: B

Explanation:

4 vCPUs. 8 GB is the recommended KVM resource allocation for a Cisco vSmart controller supporting a deployment scale of 251 through 1000 devices in the Cisco SD-WAN 20.4.x release family. Controller virtual-machine sizing is based on the number of managed WAN edge devices because the vSmart controller must maintain control-plane sessions, distribute routing and policy information, process OMP updates, and support IPsec tunnel orchestration for the fabric. Allocating four virtual CPUs provides the processing capacity expected for this medium-scale device range, while 8 GB of RAM provides the memory required by the controller services and their control-plane state. These values should be treated as the baseline recommendation when creating the vSmart virtual machine on a KVM hypervisor; the host itself must also have sufficient available CPU, memory, storage, and network resources to avoid contention with other workloads. Cisco documents controller virtual-machine deployment and platform requirements in its SD-WAN virtual-machine installation documentation. See the Cisco SD-WAN Virtual Machine Installation Guide and the Cisco vSmart Controller support documentation.

QUESTION NO: 22

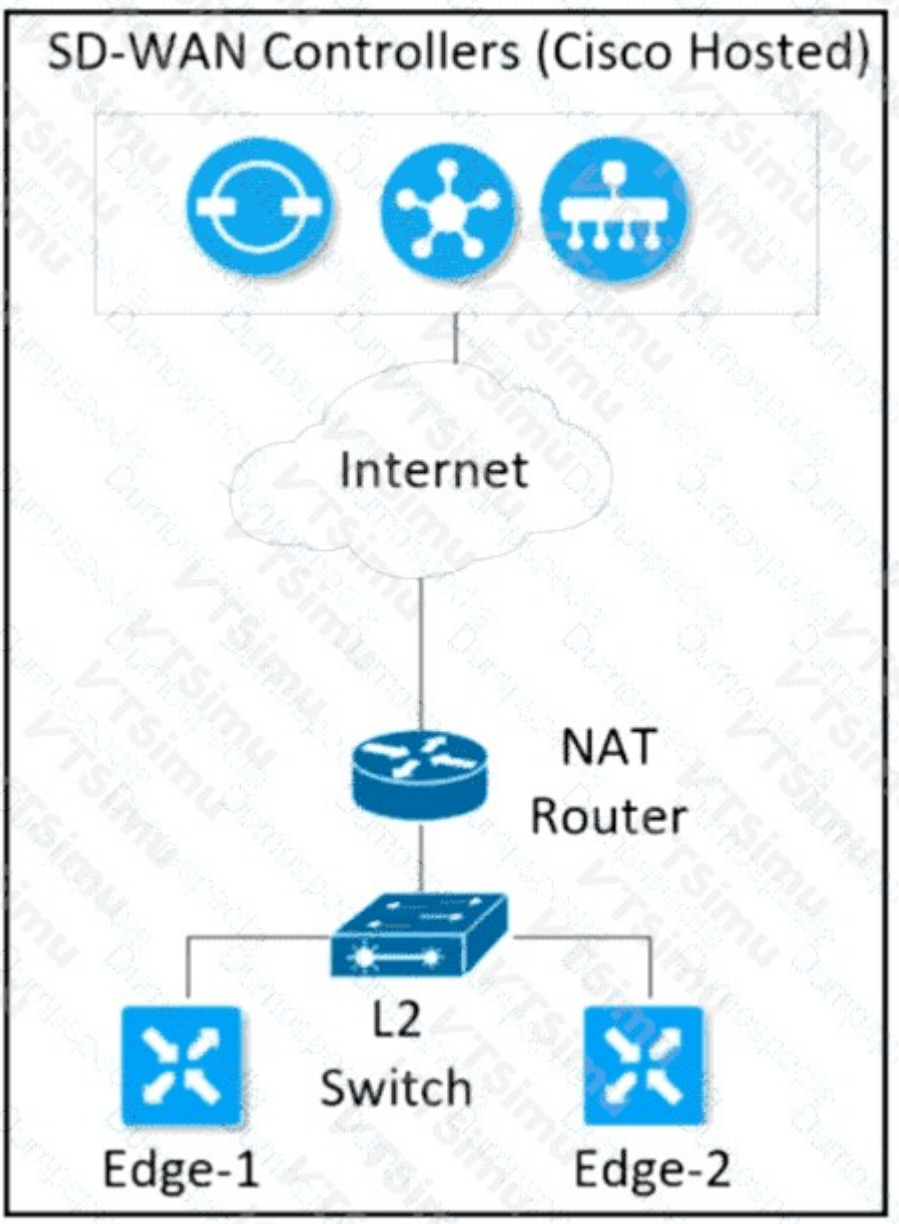

Refer to the exhibit Which configuration must the engineer use to form underlay connectivity for the Cisco SD-WAN network?

A)

```
R1
vpn 512
interface eth0
ip address 10.0.0.21/24
no shutdown
!
ip route 0.0.0.0/0 10.0.0.254

R2
vpn 512
interface eth0
ip address 10.0.0.2/24
no shutdown
!
ip route 0.0.0.0/0 10.0.0.254
```

B)

```
R1
vpn 10
interface ge0/2
ip address 10.10.10.9/29
no shutdown
!
ip route 0.0.0.0/0 10.10.10.11

R2
vpn 10
interface ge0/2
ip address 10.10.10.10/29
no shutdown
!
ip route 0.0.0.0/0 10.10.10.11
```

C)

```
R1
vpn 0
interface 10ge0/0
ip address 10.50.0.2/29
no shutdown
!
ip route 0.0.0.0/0 10.50.0.3

R2
vpn 0
interface 10ge0/0
ip address 10.50.0.1/29
no shutdown
!
ip route 0.0.0.0/0 10.50.0.3
```

D)

```
R1
vpn 0
interface ge0/0
ip address 10.50.0.2/30
tunnel-interface
!
ip route 0.0.0.0/0 10.50.0.4

R2
vpn 0
interface ge0/0
ip address 10.50.0.3/30
tunnel-interface
!
ip route 0.0.0.0/0 10.50.0.4
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: C

Explanation:

The configuration shown in Option C is the required underlay configuration because Cisco SD-WAN transport connectivity is established through a WAN-facing tunnel interface in VPN 0. VPN 0 is the transport VPN and carries the control-plane and data-plane connectivity between WAN Edge devices and the SD-WAN fabric. The transport interface must have reachable IP addressing and routing toward the WAN, and it must be configured as a tunnel interface with the appropriate transport color. The color identifies the transport type and allows Cisco SD-WAN to classify and build tunnels across that WAN circuit. The tunnel-interface configuration also enables the secure IPsec-based SD-WAN overlay tunnels used by the fabric. Once the WAN Edge has IP reachability to the controllers through this transport interface, it can establish its control connections and exchange routing and policy information. This separation of VPN 0 transport connectivity from service-side VPNs is a core Cisco SD-WAN design principle: service VPNs provide user and application reachability, while VPN 0 supplies the underlay path used to reach the SD-WAN control plane and remote transport endpoints. Cisco documents the transport-interface and VPN 0 requirements in its SD-WAN interface configuration guide and summarizes the transport and overlay architecture in the Cisco SD-WAN Design Guide.

QUESTION NO: 23

```
PEER    PEER    PEER            SITE       DOMAIN
PEER            PRIVATE PEER               PUBLIC                                  LOCAL   REMOTE      REPEAT
INSTANCE TYPE       PROTOCOL SYSTEM IP     ID          ID      PRIVATE IP          PORT    PUBLIC IP   PORT REMOTE
COLOR    STATE              ERROR          ERROR       COUNT DOWNTIME
-------------------------------------------------------------------------------------------------------------------
0        vbond      dtls     0.0.0.0        0           0       192.168.0.231       12346   192.168.0.231 12346 default
tear_down    CRTREJSER NOERR  9          2019-06-01T19:06:32+0200
```

Refer to the exhibit. An engineer is troubleshooting tear down of control connections even though a valid CertificateSerialNumber is entered. Which two actions resolve the issue? (Choose two.)

A. Enter a valid product ID (model) on the PNP portal

B. Match the serial number file between the controllers

C. Remove the duplicate IP in the network

D. Restore network reachability for the controller

E. Enter a valid serial number on the controllers for a given device

ANSWER: A C

Explanation:

Enter a valid product ID (model) on the PNP portal and remove the duplicate IP in the network. Cisco SD-WAN device authorization uses the device identity information supplied through Plug and Play, including the chassis/serial identity and the product ID (PID). A certificate serial number by itself is therefore not sufficient for successful authorization if the associated device record does not contain a valid PID matching the onboard device. Correcting the PnP record lets the controllers validate the device identity and admit its control-plane connections.

The network also must provide a unique reachable transport endpoint for each SD-WAN device. A duplicate IP address can cause return traffic, DTLS/TLS control sessions, and tunnel endpoint identification to be directed ambiguously. That ambiguity can cause an established control connection to be torn down even after certificate validation succeeds. Removing the duplicate address restores a consistent path and unique device reachability, allowing the vBond-orchestrated control connection to remain established. Cisco's control-connection troubleshooting guidance includes checking certificate/authorization information as well as transport reachability and addressing conditions: Cisco: Troubleshoot SD-WAN Control Connections. Cisco also documents PnP onboarding and device identity requirements at Cisco SD-WAN Plug and Play Configuration Guide.

QUESTION NO: 24

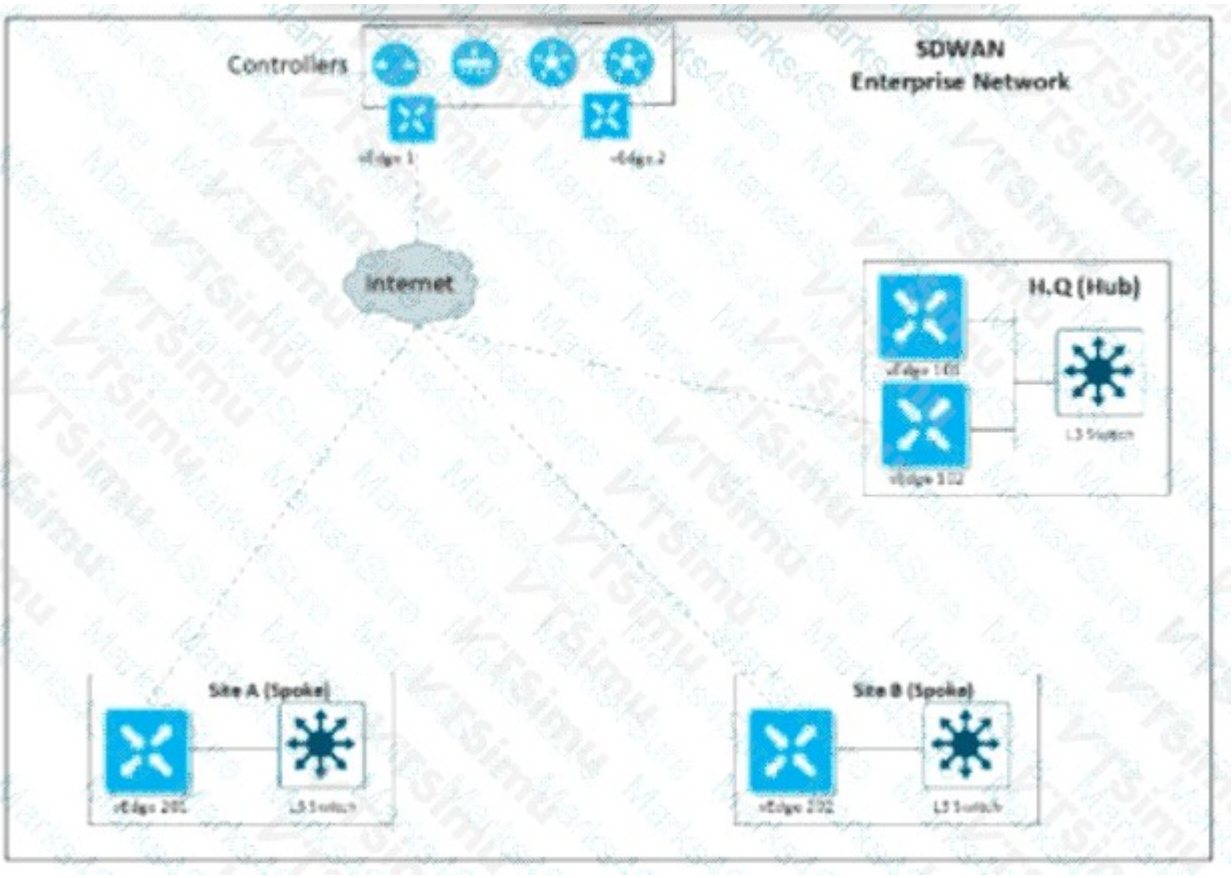

Refer to the exhibit An engineer must configure a QoS policy between me hub and site A (spoke) over a standard internet circuit where traffic shaping is adjusted automatically based on evaiiabk» bandwidth Which configuration meets the requirement?

Hub

```
sdwan
 interface GigabitEthernet1
 qos-adaptive
  period 90
  downstream 8000
  downstream range 6000 10000
  upstream 8000
  upstream range 4000 16000
  exit
 tunnel-interface
  encapsulation ipsec weight 1
  color public-internet
  no last-resort-circuit
  vmanage-connection-preference 5
  allow-service all
```

Site A (Spoke)

```
sdwan
 interface GigabitEthernet1
 qos-adaptive
  period 40
  downstream 10000
  downstream range 8000 12000
  upstream 9000
  upstream range 5000 16000
  exit
 tunnel-interface
  encapsulation ipsec weight 1
  color mpls
  no last-resort-circuit
  vmanage-connection-preference 5
  allow-service all
```

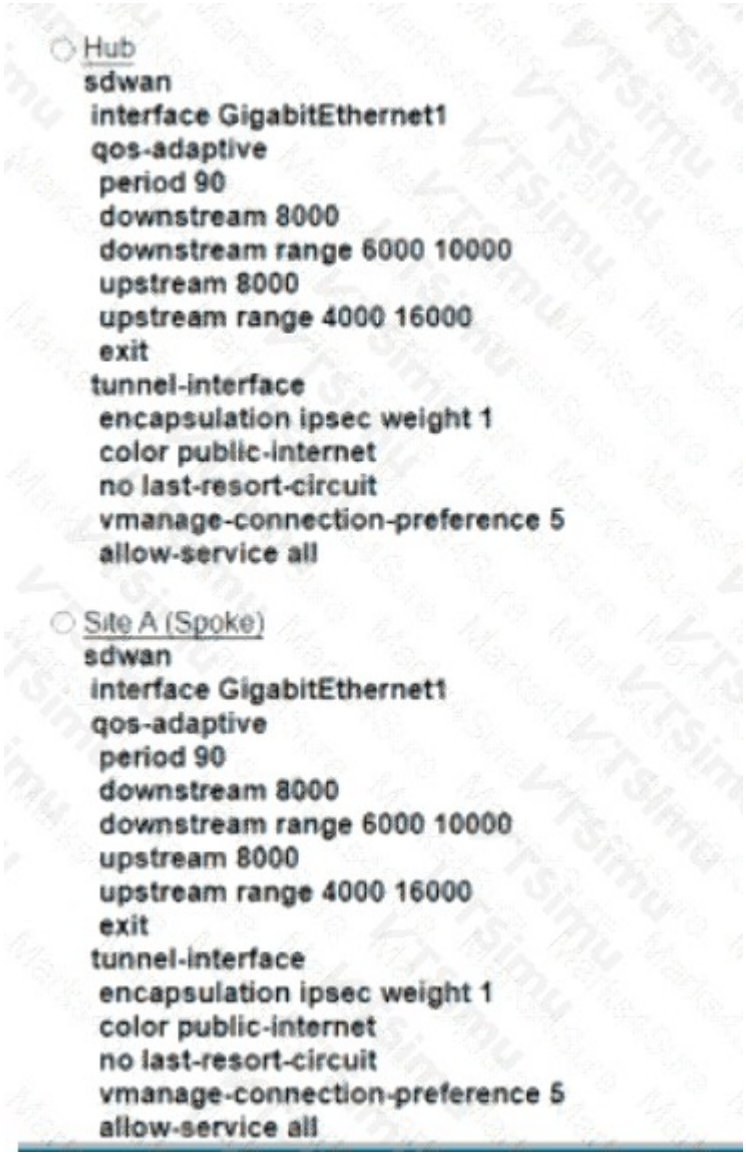

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: D

Explanation:

The configuration shown in Option D is correct because it applies adaptive QoS shaping for the Internet transport path. Adaptive QoS is designed for circuits whose usable bandwidth can vary, which is common on public Internet links. Rather than relying on one static shaper value, the router uses the available-bandwidth information learned for the tunnel and dynamically adjusts the shaping rate. This ensures that QoS queues and bandwidth allocations operate against the bandwidth currently available on the WAN path, helping avoid persistent congestion and excessive packet loss when capacity decreases.

For a hub-to-spoke SD-WAN path, the adaptive shaping behavior must be associated with the relevant tunnel/interface and QoS policy so that egress traffic is shaped according to the path's changing capacity. The QoS classes can then schedule important traffic within that dynamically calculated rate. Cisco documents adaptive QoS as a mechanism that adapts to changing transport bandwidth and integrates with SD-WAN tunnel performance and bandwidth information. See Cisco's SD-WAN QoS Configuration Guide and the Cisco SD-WAN Interfaces Configuration Guide.

QUESTION NO: 25

Which two advanced security features are available on the Cisco SD-WAN WAN Edge (vEdge) device? (Choose two.)

A. URL filtering

B. snort intrusion prevention system

C. Cisco Umbrella DNS Security

D. Cisco AMP and AMP Threat Grid

E. Enterprise Firewall

ANSWER: B E

Explanation:

snort intrusion prevention system and **Enterprise Firewall** are the advanced security capabilities implemented on Cisco SD-WAN vEdge WAN Edge devices. The enterprise firewall provides stateful, zone-based traffic inspection and policy enforcement at the WAN edge. It can apply rules using Layer 3/Layer 4 information and application-aware criteria, allowing security policy to be enforced locally as traffic enters, leaves, or traverses a site.

The Snort intrusion prevention system adds deep-packet inspection and signature-based detection and prevention of known threats. Cisco SD-WAN integrates Snort IPS with the WAN Edge security policy framework, so inspected traffic can be blocked according to configured IPS signatures and security policies. Together, these functions provide both firewall control and threat inspection directly on the vEdge appliance, reducing the need to steer all site traffic to a separate security device for these protections. Cisco's SD-WAN security documentation describes firewall and intrusion-prevention services as WAN Edge security capabilities; see Cisco SD-WAN Zone-Based Firewall Configuration Guide and Cisco SD-WAN Intrusion Prevention System Configuration Guide.

QUESTION NO: 26

Which two vRoute attributes should be matched or set in vSmart policies and modified by data policies? (Choose two.)

A. site ID

B. preference

C. VPN

D. TLOC

E. origin

ANSWER: C D

Explanation:

VPN and **TLOC** are the relevant attributes for a vSmart centralized data policy. A data policy can match traffic within a particular VPN, allowing the policy to apply forwarding treatment only to packets associated with the intended VPN service. VPN matching is fundamental because Cisco SD-WAN maintains separate routing and forwarding domains for VPNs, including service VPNs and the transport VPN.

A data policy can also set a TLOC as the preferred forwarding destination for matching traffic. A TLOC identifies a reachable WAN transport endpoint through its system IP address, transport color, and encapsulation type. Setting the TLOC directs matching traffic toward a specified transport path or tunnel endpoint, enabling traffic engineering decisions such as steering selected applications or prefixes over a desired WAN circuit. Together, matching the VPN and setting the TLOC lets vSmart apply targeted, VPN-aware forwarding behavior while retaining centralized policy control. Cisco documents VPN matching and the `set tloc` action as supported components of centralized data policies. See the Cisco centralized data-policy documentation and the Cisco SD-WAN configuration documentation.

QUESTION NO: 27

What are the two functions of vSmart? (Choose two)

A. It orchestrates connectivity between WAN Edge routers using policies to create network topology

B. It ensures that valid WAN Edge routers can build the control pane connectivity

C. It uses TLOCs to uniquely identify the circuit interface to control plane and data plane information

D. It validates that the WAN Edge trying to join the overlay is authorized to join.

E. It builds control plane connections with WAN Edge routers using TLS or DTLS

ANSWER: A E

Explanation:

The vSmart controller is Cisco SD-WAN's centralized control-plane component. It uses Overlay Management Protocol information received from WAN Edge routers to calculate and distribute routing and transport-location information, and it applies centralized control policies. This allows it to orchestrate the intended overlay topology and control which routes, TLOCs, and service paths are advertised among WAN Edge routers. Consequently, using policies to create the network topology is a core vSmart responsibility.

vSmart also establishes secure control-plane sessions directly with WAN Edge routers. These controller connections use TLS or DTLS, depending on the device software and configured transport protocol. Across those secure sessions, vSmart exchanges OMP updates and policy information that enable WAN Edge routers to learn reachable prefixes and available transport tunnels. This centralized policy and route distribution is what lets the SD-WAN fabric operate consistently across sites. Cisco describes vSmart as the controller that manages the control plane and enforces centralized policies in its Cisco SD-WAN Overview. Additional details on controller and WAN Edge control connections are available in the Cisco SD-WAN Controllers documentation.

QUESTION NO: 28

An engineer configured a data policy called ROME-POLICY. Which configuration allows traffic flow from the Rome internal network toward other sites?

A. apply-policy site-list Rome data-policy ROME-POLICY from-tunnel

B. apply-policy site-list Rome data-policy ROME-POLICY from-service

C. site-list Rome control-policy ROME-POLICY in

D. site-list Rome control-policy ROME-POLICY out

ANSWER: B

Explanation:

`apply-policy site-list Rome data-policy ROME-POLICY from-service` is correct because traffic originating on the Rome internal network enters the WAN edge through a service-side interface. The `from-service` direction attaches the centralized data policy to traffic received from local service networks before that traffic is forwarded into the SD-WAN fabric. This is the appropriate attachment direction when the policy must evaluate, permit, steer, or otherwise act on packets travelling from a branch LAN toward remote sites across overlay tunnels.

In Cisco SD-WAN, centralized data policies are applied to a site list and use an attachment direction that identifies the packet entry point relative to the WAN edge. Service-side traffic is traffic received from locally connected networks, whereas tunnel-side traffic is received from the SD-WAN overlay. Therefore, selecting `from-service` makes ROME-POLICY effective for Rome-originated internal traffic as it starts its path toward other sites. Cisco documents centralized data-policy attachment and direction behavior in its Configuring Data Policy guide and the Cisco IOS XE SD-WAN configuration documentation.

QUESTION NO: 29

Which two criteria ate supported to filter traffic on a Cisco Umbrella Cloud-delivered firewall? (Choose two)

A. tunnels

B. site ID

C. URL

D. geolocation

E. protocol

ANSWER: A E

Explanation:

Cisco Umbrella's cloud-delivered firewall policy rules can use tunnel identity as a traffic-selection criterion. A tunnel represents a protected network location or site that forwards its traffic to Umbrella, so selecting tunnels lets an administrator scope a firewall policy to traffic received through specific configured tunnels. This allows different security controls to be applied to separate branch offices, networks, or traffic paths without affecting all Umbrella-protected users and locations.

Protocol is also a supported firewall-rule criterion. Firewall rules are designed to control layer 3 and layer 4 traffic, and protocol matching enables policies to target traffic such as TCP, UDP, ICMP, or other supported IP protocols. Protocol conditions can be combined with source and destination parameters, ports where applicable, and action settings to create precise controls over permitted or blocked network communications. Cisco documents tunnel selection under policy identities and protocol as a firewall rule condition in its firewall-policy configuration guidance. See Cisco Umbrella: Add a Firewall Policy and Cisco Umbrella Firewall Policy documentation.

QUESTION NO: 30

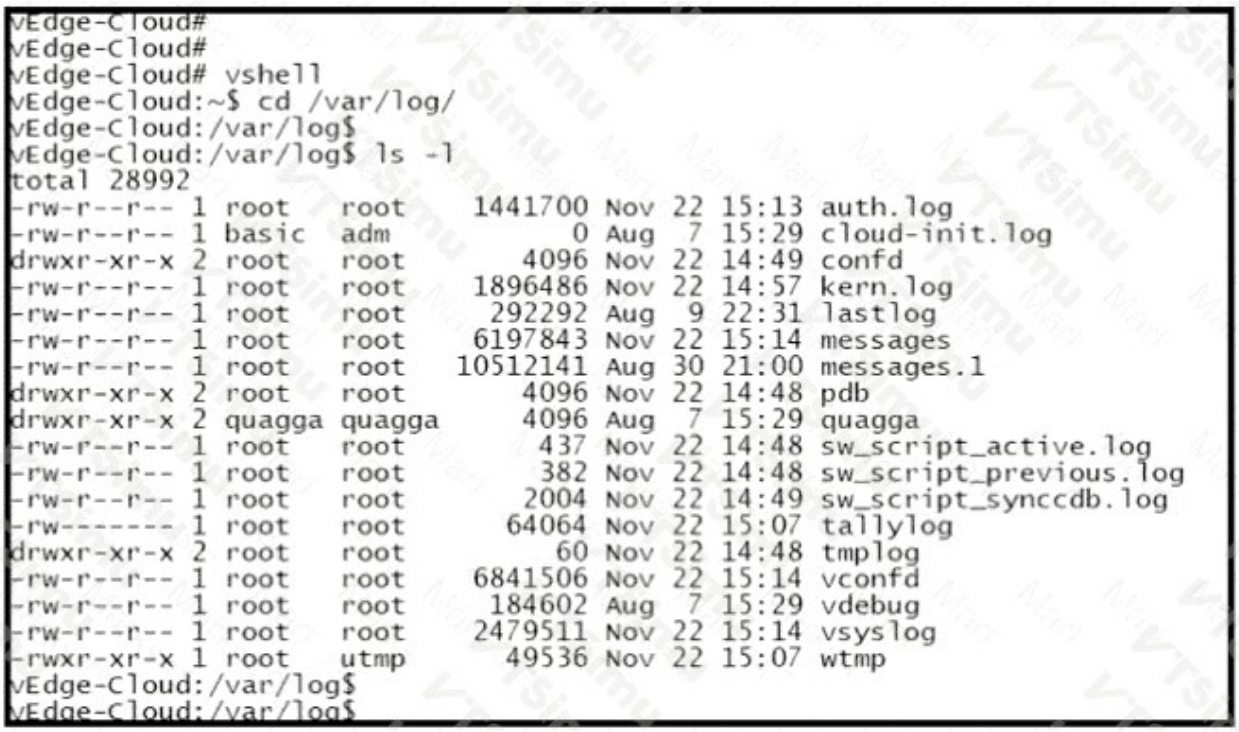

Refer to the exhibit. Which configuration stops Netconf CLI logging on WAN Edge devices during migration?

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

The configuration shown in **Option B** is correct because it uses the negated NETCONF CLI logging command, which disables the logging of CLI commands processed through the NETCONF interface. During a WAN Edge migration, this prevents NETCONF-originated CLI activity from being recorded while retaining the NETCONF service needed for management and migration operations. This is important because stopping the logging function is different from disabling NETCONF itself: WAN Edge devices can continue to accept required management transactions without generating the associated CLI log entries. Cisco's SD-WAN command reference documents the NETCONF CLI logging control and its negated form for disabling this logging behavior. See the Cisco SD-WAN Command Reference and Cisco's SD-WAN installation and configuration guides for WAN Edge configuration context.

QUESTION NO: 31

Which two algorithms authenticate a user when configuring SNMPv3 monitoring on a WAN Edge router? (Choose two.)

A. AES-256

B. MD5

C. SHA-2

D. AES-128

E. SHA-1

ANSWER: B E

Explanation:

SNMPv3 uses the User-based Security Model (USM) to provide user authentication, optional privacy encryption, and message-integrity protection. On Cisco SD-WAN WAN Edge routers, the supported SNMPv3 user authentication protocols for this configuration are MD5 and SHA-1. When an SNMPv3 user is created, one of these algorithms is selected to derive and validate the authentication digest associated with that user's configured authentication password. The receiving device uses the same configured authentication method to verify that an SNMP message originated from the expected SNMPv3 user and that it was not altered in transit. This authentication function is distinct from SNMPv3 privacy, which encrypts SNMP payload data when privacy is enabled. Cisco's SD-WAN SNMP configuration documentation identifies `md5` and `sha` as the

authentication choices for an SNMPv3 user; in this context, `sha` denotes SHA-1. See the Cisco SD-WAN SNMP Configuration Guide and Cisco's SNMPv3 configuration documentation.

QUESTION NO: 32

Which two prerequisites must be met before the Cloud onRamp for IaaS is initiated on vManage to expand to the AWS cloud? (Choose two.)

A. Attach an OSPF feature template to the AWS cloud Edge router template.

B. Attach the "AmazonCreateVPC" and "AmazonProvisionEC2" permission policy to the IAM account.

C. Subscribe to the SD-WAN Edge router AMI in the AWS account.

D. Attach a device template to the cloud WAN Edge router to be deployed in the AWS.

E. Preprovision the transit VPC in the AWS region.

ANSWER: C D

Explanation:

Before initiating AWS Cloud onRamp for IaaS from vManage, the AWS account must be subscribed to the Cisco SD-WAN Edge router Amazon Machine Image (AMI) in AWS Marketplace. This subscription grants the account access to the licensed cloud-router image that vManage uses when it orchestrates the deployment of WAN Edge instances in AWS. The AMI is therefore a prerequisite to automated instance creation.

A device template must also be attached to the cloud WAN Edge router that will be deployed in AWS. The template supplies the router's intended Cisco SD-WAN configuration, including its system and transport settings and any required feature templates. Cloud onRamp uses this template-driven configuration model so that a newly provisioned cloud router can be onboarded into the SD-WAN overlay and operate consistently with the rest of the fabric. The Cloud onRamp workflow automates AWS infrastructure and instance deployment, but it depends on an available Marketplace image and a prepared WAN Edge device template before deployment begins. Cisco's SD-WAN configuration documentation is available from Cisco SD-WAN installation and configuration guides, and Cisco publishes the cloud router image through AWS Marketplace.

QUESTION NO: 33

Which command disables the logging of syslog messages to the local disk?

A. no system logging disk local

B. system logging server remote

C. no system logging disk enable

D. system logging disk disable

ANSWER: C

Explanation:

The command `no system logging disk enable` disables local disk logging for syslog messages in Cisco SD-WAN system configuration. Cisco SD-WAN uses the `system logging disk` configuration hierarchy to control whether syslog output is retained on the device's local persistent storage. The `enable` leaf explicitly activates this capability; applying `no` to that leaf removes the enabled setting and turns off disk logging.

This setting is useful when local storage consumption must be minimized, when log retention is handled by an external syslog collector, or when operational policy requires logs to remain off the device. Disabling disk logging does not inherently prevent the device from generating log events or from using separately configured logging destinations. Those behaviors are

controlled through their own logging configuration settings. The command should be entered in the appropriate Cisco SD-WAN configuration context and committed through the platform's normal configuration workflow.

Cisco documents disk logging as part of the system logging configuration and identifies the disk `enable` setting as the control for writing syslog messages locally. See Cisco's System Logging documentation and the Cisco SD-WAN CLI reference.

QUESTION NO: 34

How many concurrent sessions does a vManage REST API have before it invalidates the least recently used session if the maximum concurrent session number is reached?

A. 150

B. 200

C. 250

D. 300

ANSWER: C

Explanation:

250 is correct. Cisco vManage maintains REST API login sessions after a client authenticates and receives a session cookie/token for API requests. The platform limits concurrent REST API sessions to 250. This limit prevents unbounded session consumption while still allowing automation platforms, administrators, and integrations to make parallel API requests. When a new login request would exceed the supported concurrent-session maximum, vManage invalidates the least recently used existing session. The client associated with that invalidated session must authenticate again before it can continue making authenticated API calls. For operational automation, applications should reuse a valid authenticated session where possible, explicitly log out when work is complete, and be prepared to reauthenticate if a session is no longer accepted. These practices reduce unnecessary session churn and avoid disruption when several scripts or services access the vManage REST API simultaneously. Cisco documents vManage REST API authentication and session handling in its Cisco SD-WAN API overview. Additional Cisco SD-WAN API documentation is available through the Cisco DevNet SD-WAN documentation portal.

QUESTION NO: 35

Which API call retrieves a list of all devices in the network?

A. https://vmanage_IP_address/dataservice/system/device/{{model}}

B. http://vmanage_IP_address/dataservice/system/device/{{model}}

C. http://vmanage_IP_address/api-call/system/device/{{model}}

D. https://vmanage_IP_address/api-call/system/device/{{model}}

E. https://vmanage_IP_address/dataservice/system/device

ANSWER: E

Explanation:

`https://vmanage_IP_address/dataservice/system/device` is the Cisco Catalyst SD-WAN vManage REST API endpoint that retrieves the network-wide device list. The request is sent to the vManage data-service API namespace, identified by `/dataservice`, and it uses HTTPS, which is the supported protocol for authenticated vManage API access. A successful GET request returns device inventory information for the devices known to vManage, such as their identifiers, device type, reachability, software version, site information, and operational status. This is the appropriate starting endpoint

when an automation workflow needs to discover every device before selecting a particular device or applying additional filtering.

The absence of a model-specific path component is important because the requested result is the complete device inventory rather than devices limited to one platform model. Cisco's API documentation identifies the Device List operation under the System APIs and documents the REST resource used to obtain this inventory. Authentication must be established first, typically by obtaining a vManage session or token as required by the deployment, and the request must be issued by a user with suitable permissions. See Cisco's Get Device List API documentation and the Cisco Catalyst SD-WAN API documentation.

QUESTION NO: 36

A Cisco SD-WAN customer has a requirement to calculate the SHA value for files as they pass through the device to see the returned disposition and determine if the file is good, unknown or malicious. The customer also wants to perform real-time traffic analysis and generate alerts when threats are detected Which two Cisco SD-WAN solutions meet the requirements? (Choose two.)

A. Cisco Trust Anchor Module

B. Cisco Threat Grid

C. Cisco Snort IPS

D. Cisco AMP

E. Cisco Secure Endpoint

ANSWER: C D

Explanation:

Cisco AMP and **Cisco Snort IPS** meet these requirements when deployed through the Cisco SD-WAN security/Unified Threat Defense capabilities. Cisco AMP for Networks performs file inspection and generates a SHA-256 value for an observed file. It uses Cisco's cloud-based file reputation intelligence to return a disposition, such as clean (good), unknown, or malicious. This gives the branch device an actionable verdict for files traversing the WAN edge and supports malware protection based on known file reputation and retrospective security intelligence.

Cisco Snort IPS provides the real-time network-traffic inspection component. The Snort intrusion prevention engine evaluates traffic against configured intrusion rules and threat signatures. It can detect suspicious or malicious activity as flows pass through the device and generate alerts and logs for identified threats. Together, AMP supplies the file hash and file-disposition workflow, while Snort IPS supplies continuous packet and protocol analysis for network-based threat detection. Cisco documents AMP and IPS as security services available on supported IOS XE SD-WAN edge platforms. See Cisco's Unified Threat Defense configuration guide and the Cisco Advanced Malware Protection overview.

QUESTION NO: 37

Refer to exhibit. An engineer is troubleshooting tear of control connection even though a valid CertificateSerialNumber is entered. Which two actions resolve Issue? (Choose two)

A. Restore network reachability on the controller.

B. Enter a valid serial cumber on the controller for a given device

C. Enter a valid product ID (mode) on the PNP portal.

D. Match the serial number file between the controller

E. Remove the duplicate IP in the network

ANSWER: B C

Explanation:

A Cisco SD-WAN device must be authorized with the serial number that is contained in, and presented by, its device certificate during control-plane authentication. Therefore, entering a valid serial number on the controller for the specific device ensures that the controller's authorized-device inventory has the matching certificate identity. This allows the controller to validate the device when it attempts to establish its secure control connections.

Entering a valid product ID (model) on the PNP portal is also required because Cisco Plug and Play associates the device's serial number with its hardware product identifier. The product ID must correspond to the actual device platform so that the correct certificate and device identity information are made available for onboarding and validation. With both the device serial number authorized in the controller inventory and the correct product ID associated in the PnP portal, the certificate identity can be validated consistently and control connections can be established. Cisco documents device authorization and certificate validation in its Cisco SD-WAN certificate configuration guide and PnP device onboarding in the Cisco Plug and Play documentation.

QUESTION NO: 38

Refer to the exhibit.

The engineer must assign community tags to 3 of its 74 critical server networks as soon as that are advertised to BGP peers. These server networks must not be advertised outside AS. Which configuration fulfill this requirement?

A)

```
policy
 route-policy eBGP_Community_Policy
   sequence 1
    match
     address Community_Prefix
    action accept
     set
      community 999:65000 local-as
  default-action reject
lists
 prefix-list Community_Prefix
  ip-prefix 20.20.20.0/24
  ip-prefix 21.21.21.0/24
  ip-prefix 22.22.22.0/24
```

B)

C)

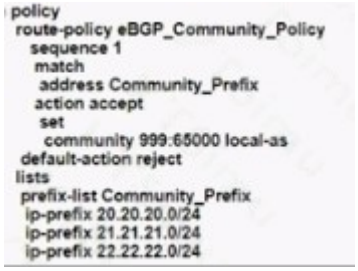

D)

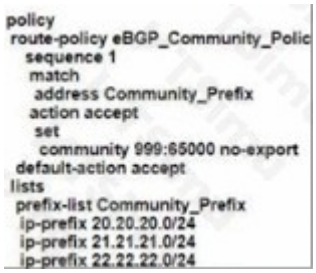

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: D

Explanation:

The correct configuration identifies the three critical server prefixes with a prefix list, matches that prefix list in a route map, and applies the well-known `no-export` BGP community while the routes are being advertised to BGP peers. This provides targeted policy control: only the selected server networks receive the community, rather than all BGP routes. The route map must be applied in the outbound direction to the relevant BGP neighbor so that the community is attached at advertisement time. In addition, the BGP process must be configured to send communities to the peer; otherwise, the community attribute is not propagated with the update.

The `no-export` community is a well-known, transitive BGP community defined as `0xFFFFFF01`. A BGP speaker that receives a route carrying this community can advertise it to peers within its own autonomous system, including confederation handling as defined by the standard, but must not advertise it to external BGP peers. Therefore, assigning `no-export` to the designated server prefixes satisfies the requirement to keep those networks from being advertised outside the AS while preserving their internal reachability. Cisco documents BGP community policy configuration at Cisco IOS XE BGP Community Configuration Guide; the community's standard behavior is defined in RFC 1997.

QUESTION NO: 39

What are the two advantages of deploying cloud-based Cisco SD-WAN controllers? (Choose two.)

A. centralized control and data plane

B. infrastructure as a service

C. management of SLA

D. centralized raid storage of data

E. distributed authentication policies

ANSWER: A C

Explanation:

Cloud-based Cisco SD-WAN controllers provide a centralized operational model for controlling the WAN data plane. Cisco SD-WAN separates the control and data planes: controllers establish and distribute policy, routing, segmentation, and tunnel information, while WAN Edge devices forward application traffic according to that centrally delivered intent. Hosting the controller components in the cloud retains this centralized control capability while avoiding the need to operate the controller infrastructure in a customer data center.

Management of SLA is also a key advantage because Cisco SD-WAN continuously measures path characteristics such as loss, latency, and jitter. vManage provides centralized visibility into those measurements, and SD-WAN policy can use application-aware routing to select paths that satisfy the defined SLA requirements. This enables operations teams to monitor WAN health centrally and apply consistent performance policy across sites. Cisco documents the controller-based architecture and centralized management capabilities in its SD-WAN architecture overview and documents SLA-based path selection in its application-aware routing guidance. See Cisco: What Is SD-WAN? and Cisco IOS XE SD-WAN Application-Aware Routing.

QUESTION NO: 40

Which two image formats are supported for controller codes? (Choose two.)

A. .nxos

B. .qcow2

C. .iso

D. .ova

E. .bin

ANSWER: B D

Explanation:

Controller code images for Cisco SD-WAN virtual controller deployment support the **.qcow2** and **.ova** formats. A QCOW2 image is the virtual-disk format used by KVM/QEMU-based hypervisors, making it suitable when a controller is deployed in KVM-oriented virtual environments. An OVA image is an Open Virtual Appliance package, commonly used to deploy virtual appliances through VMware vSphere and other platforms that import OVF/OVA packages. These formats provide the virtual-machine image packaging required for Cisco SD-WAN controller nodes, such as Cisco vManage, Cisco vSmart Controller, and Cisco vBond Orchestrator, in their respective supported virtualized deployments. The image selected must match the target hypervisor and the Cisco SD-WAN release and platform requirements. Cisco's SD-WAN documentation catalog provides the installation and configuration guidance for controller deployment, while Cisco's SD-WAN product documentation describes the controller components and supported deployment models. See the Cisco SD-WAN installation and configuration guides and the Cisco SD-WAN solution documentation.

QUESTION NO: 41

Which pathway under Monitor > Network > Select Device is used to verify service insertion configuration?

A. System Status

B. ACL Logs

C. Real Time

D. Events

ANSWER: C

Explanation:

Real Time is the correct pathway because service insertion is verified through current operational information collected from the selected SD-WAN device. In Cisco vManage, the Real Time area under a device's monitoring view provides live, device-level operational status and counters for services and forwarding features. This enables an administrator to confirm that the service-insertion feature is active and to inspect the status of the relevant service chain or service node from the perspective of that device.

Service insertion verification is an operational task rather than merely a review of historical notifications or general device inventory. The real-time monitoring workflow is intended to expose the current state of configured SD-WAN functions, helping validate that traffic is being directed through the intended service path and that the device has the required operational information available. Cisco's SD-WAN documentation describes monitoring device operational data through vManage and provides the platform documentation needed to navigate these device-monitoring capabilities. See the Cisco SD-WAN documentation search for service insertion monitoring and the Cisco SD-WAN developer documentation.

QUESTION NO: 42

Which website allows access to visualize the geography screen from vManager using the internet?

A. *.openstreetmap.org

B. *.fullstreetmaps.org

C. *.callstreelmaps.org

D. *.globaistreetmaps.org

ANSWER: A

Explanation:

`*.openstreetmap.org` is the required internet domain for displaying the geographic map view in Cisco vManage. The vManage Geography screen uses OpenStreetMap map data and tiles to render the visual topology and geographic locations of SD-WAN devices. Therefore, the vManage server must be able to reach the OpenStreetMap service through any intervening firewall, proxy, DNS policy, or internet-access control. The wildcard notation permits access to OpenStreetMap subdomains that may be used to retrieve map tiles and related mapping resources. This connectivity is relevant to the GUI map display only; it enables the background geographic content on which vManage presents device and site information. Cisco SD-WAN deployment planning should account for required external services where geographic visualization is expected, while security teams should allow only the required HTTPS/DNS access according to organizational policy. Cisco's SD-WAN installation and configuration documentation is available through the Cisco SD-WAN configuration guide collection. OpenStreetMap publishes information about the tile infrastructure and its access expectations in its tile usage policy.

QUESTION NO: 43

Which component of the Cisco SD-WAN architecture oversees the control plane of overlay network to establish, adjust, and maintain the connections between the WAN Edge devices that form the Cisco SD-WAN fabric?

A. APIC-EM

B. vSmart

C. vManage

D. vBond

ANSWER: B

Explanation:

vSmart is correct because it is the Cisco SD-WAN controller responsible for the overlay control plane. WAN Edge devices establish secure control-plane connections to vSmart controllers, which distribute routing, topology, policy, and service information through the Overlay Management Protocol (OMP). Using this centralized control-plane information, vSmart determines which routes and transport tunnel information are advertised to WAN Edge devices and applies centralized control policies that influence reachability and path selection across the SD-WAN fabric. This allows the overlay to adapt as WAN Edge devices, links, routes, and policy requirements change, while the WAN Edge devices establish the corresponding IPsec data-plane tunnels for traffic forwarding. Cisco identifies the vSmart controller as the component that manages the control plane and enables the scalable, policy-driven connectivity model used by Cisco SD-WAN. For additional architectural detail, see the Cisco SD-WAN Control Plane Overview and Cisco's SD-WAN architecture overview.

QUESTION NO: 44

Which two requirements must be met for DNS inspection when integrating with cisco umbrella? (Choose two)

A. Upload the WAN Edge serial allow list to the Umbrella portal.

B. Attach security policy to the device template.

C. Configure the Umbrella token on the vManage

D. Create and attach a System feature template with the Umbrella registration credentials.

E. Register and configure the vManage public IP and serial number in the Umbrella portal.

ANSWER: B C

Explanation:

Configuring the Umbrella token on vManage is required because the token securely associates the SD-WAN management system with the organization's Cisco Umbrella tenant. vManage uses this Umbrella API token to obtain and manage the registration information that allows WAN Edge devices to use Umbrella DNS-layer security. The token is configured globally in vManage before the Umbrella DNS-security settings can be deployed to managed devices.

Attaching a security policy to the device template is also required because DNS inspection and Umbrella DNS security are policy-driven services in Cisco SD-WAN. The policy contains the DNS-security configuration that directs eligible DNS requests through the Umbrella integration and applies the intended inspection behavior. Associating that policy with the applicable device template ensures that the configuration is delivered to the selected WAN Edge routers when the template is attached and configured. This connects the global Umbrella tenant integration with actual traffic enforcement on the edge devices. Cisco documents the vManage-based Umbrella integration and its deployment workflow in the Cisco SD-WAN Security Configuration Guide and describes DNS-layer policy functionality in the Cisco Umbrella DNS-Layer Security documentation.

QUESTION NO: 45

Which two virtualized environments are available for a company to install the controllers using the on-premises model? (Choose two)

A. VMware vSphere ESXi

B. VMware Workstation

C. kernel-based virtual machine

D. OpenStack

E. Microsoft Hyper-V

ANSWER: A D

Explanation:

VMware vSphere ESXi and OpenStack are the supported virtualized environments identified for deploying Cisco SD-WAN controller components in the on-premises deployment model. In this model, the control-plane and management-plane functions are delivered as virtual machines, allowing an organization to host the controller infrastructure in its own data center rather than consuming a Cisco-hosted service. VMware vSphere ESXi provides the conventional enterprise hypervisor platform on which the controller virtual-machine images can be deployed and operated. OpenStack provides a private-cloud virtualization environment in which the corresponding controller images can be instantiated and managed through OpenStack services. Selecting either of these environments supports an on-premises architecture while retaining the core Cisco SD-WAN controller roles required for orchestration, management, and control-plane operation. The platform and version requirements must still be verified against Cisco's applicable compatibility documentation before deployment, because supported releases, image formats, resource sizing, and lifecycle guidance can vary by Cisco SD-WAN release. Cisco's deployment documentation and compatibility information are available through Cisco SD-WAN installation and configuration guides and the Cisco SD-WAN compatibility matrix.

QUESTION NO: 46

Which two mechanisms are used by vManage to ensure that the certificate serial number of the WAN Edge router that is needed to authenticate is listed in the WAN Edge Authorized Señal Number Hst' (Choose two)

A. Synchronize to the PnP

B. Manually upload it to vManage

C. The devices register to vManage directly as the devices come online

D. The vManage is shipped with the list

E. Synchronize to the Smart Account

ANSWER: B E

Explanation:

Manually uploading the WAN Edge serial-number file to vManage and synchronizing with the Smart Account are supported mechanisms for populating the authorized serial-number list used during WAN Edge onboarding. The authorized serial-number list establishes that a device is entitled to join the Cisco SD-WAN overlay before vManage completes device authorization and configuration. In a manual workflow, an administrator obtains the WAN Edge authorization/serial-number information and uploads the file into vManage. This approach is useful when an organization needs direct control of the inventory data imported into the controller. In the Smart Account workflow, vManage synchronizes with Cisco Smart Software Manager through the Cisco Smart Account, allowing the controller to retrieve the eligible WAN Edge inventory associated with that account. This reduces manual data handling and helps ensure that serial numbers for purchased and claimed devices are available to the SD-WAN environment. After the serial number is present in the authorized list, the WAN Edge can be authenticated and authorized as part of the secure bring-up process. Cisco documents WAN Edge authorization and serial-number management in its certificate authentication guidance and its Cisco Smart Software Manager documentation.

QUESTION NO: 47 - (DRAG DROP)

DRAG DROP

Drag and drop the attributes from the left that make each transport location unique onto the right. Not all options are used.

Select and Place:

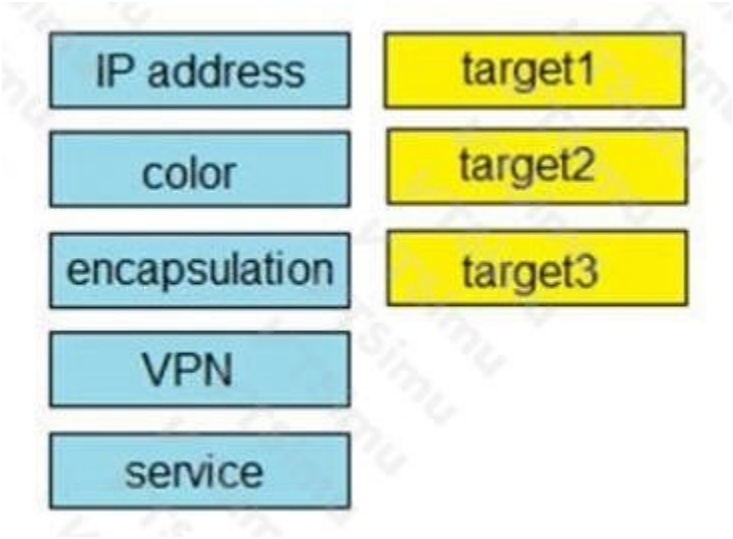

ANSWER:

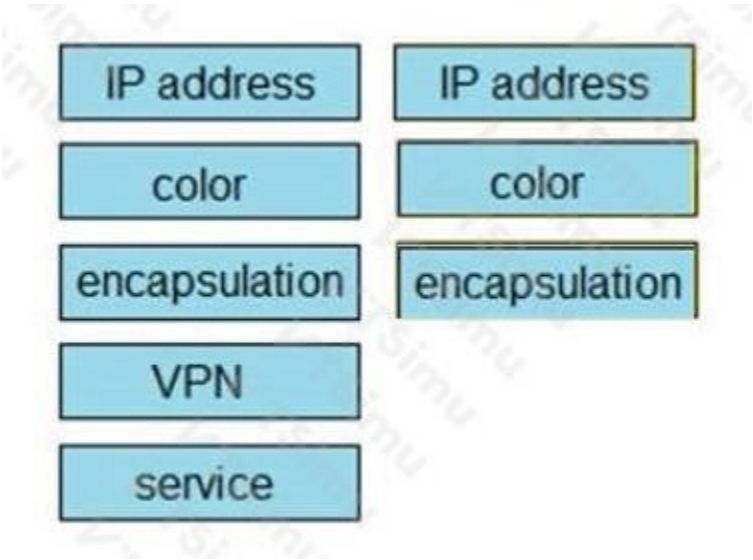

Explanation:

A Cisco SD-WAN transport location is represented by a Transport Locator, commonly called a TLOC. The unique identity of that TLOC is built from three attributes: the device's IP address, the transport color, and the encapsulation type. The IP address identifies the WAN edge device that advertises and owns the transport locator. Transport color identifies the logical transport category and its intended role, such as a public internet, MPLS, LTE, or private WAN transport. Encapsulation identifies the tunnel protocol used for the transport connection, such as IPsec or GRE. Together, these three values form the TLOC tuple used by the SD-WAN control plane to distinguish reachable transport endpoints and to construct WAN tunnels and data-plane paths.

The placed mapping is therefore correct: IP address, color, and encapsulation are the attributes that make each transport location unique. This identity is significant when OMP advertises transport-related routes and when WAN Edge devices select valid tunnel endpoints across the SD-WAN fabric. Cisco documentation describes a TLOC as the combination of system IP, color, and encapsulation, and explains its role in identifying a WAN Edge transport endpoint. See Cisco's SD-WAN TLOC documentation and the Cisco SD-WAN Design Guide for additional implementation context.

QUESTION NO: 48 - (DRAG DROP)

DRAG DROP

Drag and drop the actions from the left into the correct sequence on the right to create a data policy to direct traffic to the Internet exit.

Select and Place:

ANSWER:

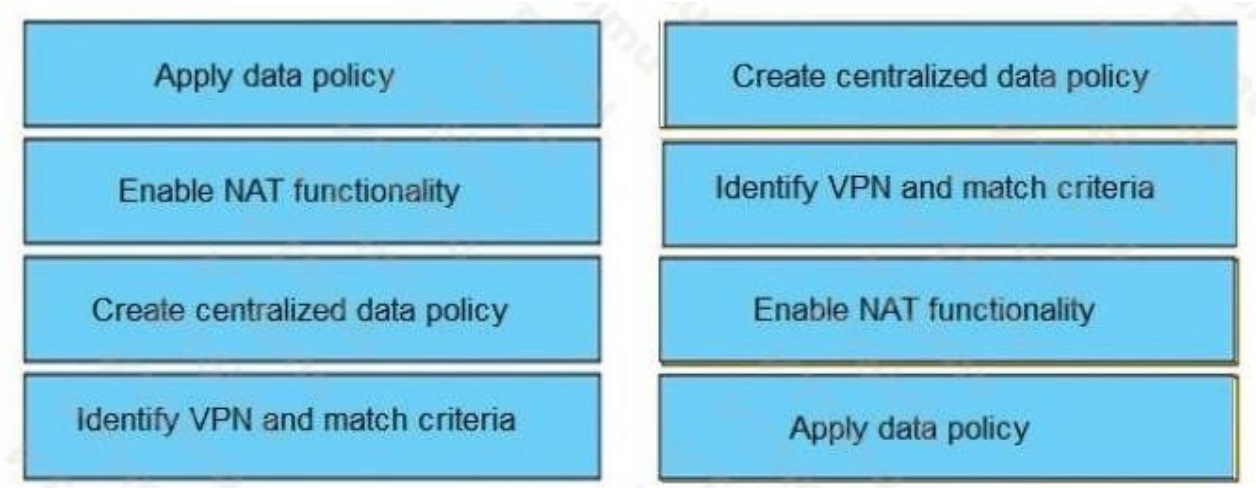

Explanation:

Creating the centralized data policy first establishes the policy container in Cisco SD-WAN Manager where traffic-handling logic is defined. Within that policy, identifying the VPN and match criteria determines exactly which traffic is eligible for Internet-exit handling. The VPN identifies the service-side traffic context, while match conditions can narrow the policy to the desired destinations, protocols, source prefixes, or other traffic attributes. This ensures that the Internet-exit behavior is applied only to the intended flows.

After the relevant traffic is identified, NAT functionality is enabled for that matched traffic. NAT is essential for an Internet exit because private service-VPN addresses generally require translation before packets are sent toward the public Internet. The centralized data-policy NAT action provides the required translation behavior so traffic can leave through the selected Internet-facing path while maintaining the intended SD-WAN policy control.

Applying the data policy is the final step because it activates the completed policy on the appropriate WAN Edge devices, sites, or VPN scope. Until the policy is applied, the configured match and NAT logic remains only a definition and does not affect forwarding behavior. This sequence follows the centralized-policy configuration model described in Cisco's vEdge as a NAT Device documentation and aligns with the normal workflow of defining policy, building its traffic sequence, specifying NAT treatment, and deploying the policy.

QUESTION NO: 49

Which multicast component is irrelevant when defining a multicast replicator outside the local network without any multicast sources or receivers?

A. PIM interfaces

B. TLOC

C. overlay BFD

D. OMP

ANSWER: A

Explanation:

PIM interfaces is correct because Protocol Independent Multicast is needed on LAN-facing interfaces that connect to multicast sources, multicast receivers, or multicast-capable downstream routers. PIM supplies the LAN multicast-control signaling used to learn receiver interest and build multicast forwarding state toward locally attached endpoints. A WAN Edge selected solely as a Cisco SD-WAN multicast replicator can be located away from the local LANs where multicast traffic originates or is consumed. In that dedicated placement, it receives multicast traffic over the SD-WAN overlay and replicates it toward remote multicast sites; it has no directly attached multicast source or receiver for which LAN PIM adjacency or PIM join processing is required. Therefore, no PIM interface is relevant to defining that off-LAN replicator role. Cisco SD-WAN multicast uses a replicator to copy multicast packets across overlay tunnels, while PIM remains a LAN-side function where multicast hosts and routers are attached. This separation permits a WAN Edge with suitable overlay reachability and capacity to act as the replication point without participating in local PIM on behalf of nonexistent local multicast endpoints. See Cisco's Cisco SD-WAN Multicast Configuration Guide and the multicast configuration documentation.

QUESTION NO: 50

Which two sets of identifiers does OMP carry when it advertises TLOC routes between WAN Edge routers? (Choose two.)

A. TLOC public and private address, carrier, and preference

B. source and destination IP address, MAC, and site ID

C. system IP address, link color, and encapsulation

D. VPN ID, local site network, and BGP next-hop IP address

E. TLOC public and private address, tunnel ID, and performance

ANSWER: A C

Explanation:

OMP uses Transport Location (TLOC) routes to distribute WAN transport-tunnel reachability between WAN Edge routers. The canonical TLOC identity is the combination of the WAN Edge system IP address, the tunnel link color, and the encapsulation type. This tuple lets remote WAN Edge routers identify a specific transport endpoint and determine how an IPsec data-plane tunnel can be formed to it. Therefore, *system IP address, link color, and encapsulation* is a required set of TLOC-identifying information.

An advertised TLOC route also carries transport-address information for the endpoint, including its private and public addresses as applicable, together with TLOC attributes used for path selection and transport differentiation. The carrier attribute identifies the transport provider associated with a tunnel, while preference influences the relative desirability of that TLOC when multiple viable transport paths exist. Thus, *TLOC public and private address, carrier, and preference* correctly identifies information OMP advertises with TLOC reachability. Cisco's OMP documentation describes TLOC routes as advertisements of transport locations and their associated attributes, while Cisco's SD-WAN routing documentation describes the system-IP/color/encapsulation TLOC tuple. See Cisco OMP Configuration Guide and Cisco TLOC documentation.

QUESTION NO: 51

What are two benefits of installing Cisco SD-WAN controllers on cloud-hosted services? (Choose two.)

A. utilizes well-known cloud services such as Azure. AWS. and GCP

B. accelerates Cisco SD-WAN deployment

C. allows integration of the WAN Edge devices In the cloud

D. installs the controllers in two cloud regions in a primary and backup setup

E. automatically Implements zone-based firewalling on the controllers

ANSWER: B D

Explanation:

accelerates Cisco SD-WAN deployment is correct because cloud hosting eliminates much of the time and operational effort associated with acquiring hardware, preparing a data center, installing controller software, and building the underlying controller infrastructure. Cloud-hosted controller instances can be provisioned rapidly, allowing the organization to bring up vManage, vSmart, and vBond services and begin onboarding WAN Edge devices sooner. This is especially valuable for organizations that need to deploy SD-WAN quickly or do not want to operate controller infrastructure in their own facilities.

installs the controllers in two cloud regions in a primary and backup setup is also correct because a cloud deployment supports geographic redundancy and disaster recovery. Placing controller services across separate cloud regions reduces the impact of a regional service disruption and helps preserve management and control-plane availability. Cisco SD-WAN controller design emphasizes high availability and redundancy for critical controller roles; cloud regions provide suitable, separated locations for implementing that design. The precise controller sizing and redundancy model must still follow the organization's scale and availability requirements. Cisco describes controller roles and deployment considerations in its SD-WAN controller documentation and its Cisco SD-WAN overview.

QUESTION NO: 52 - (SIMULATION)

Configure individual VRFs for each customer according to the topology to achieve these goals :

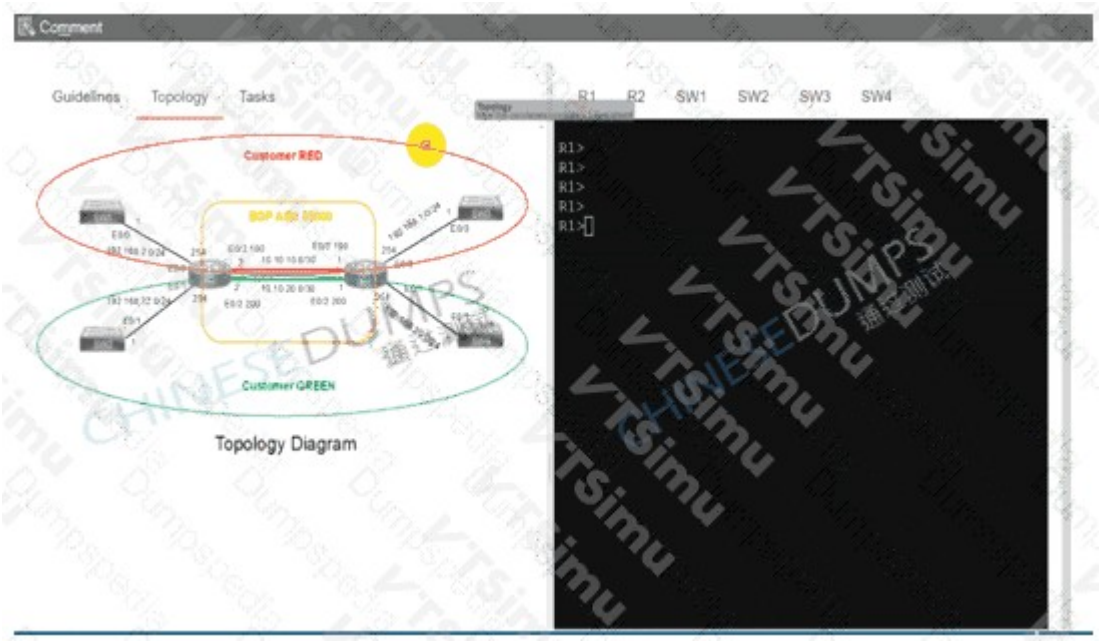


Guidelines
Topology
Tasks
Configure individual VRFs for each customer according to the topology to achieve these goals.
1. VRF "cu-red" has interfaces on routers R1 and R2. Both routers are preconfigured with IP addressing, VRFs, and BGP. Do not use the BGP network statement for advertisement.
2. VRF "cu-green" has interfaces on routers R1 and R2.
3. BGP on router R1 populates VRF routes between router R1 and R2.
4. BGP on router R2 populates VRF routes between router R1 and R2.
5. LAN to LAN is reachable between SW1 and SW3 for VRF "cu-red" and between SW2 and SW4 for VRF "cu-green". All switches are preconfigured.
R1
R2
SW1
SW2
SW3
SW4
CHINESEDUMPS
通过测试

R1


R1
R2
SW1
SW2
SW3
SW4
R1>
R1>
R1>
R1>
R1>en
R1#sh run
Building configuration...
Current configuration : 1353 bytes
!
version 15.8
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname R1
!
boot-start-marker
boot-end-marker
!
!
!
no aaa new-model
!
!
CHINESEDUMPS
通过测试

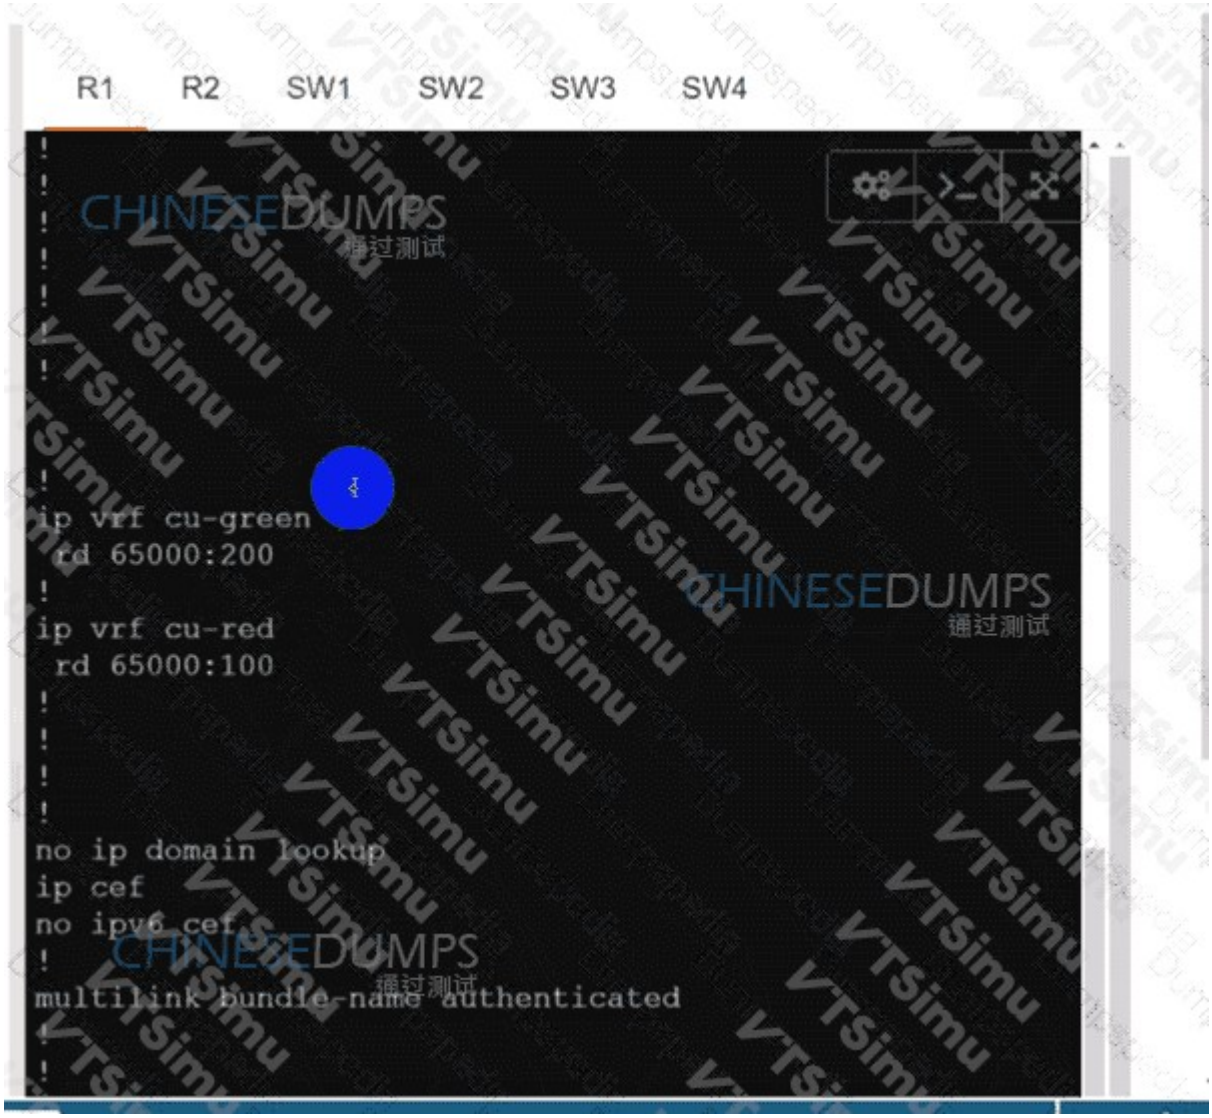
R1
R2
SW1
SW2
SW3
SW4
ip vrf cu-green
rd 65000:200
ip vrf cu-red
rd 65000:100
no ip domain lookup
ip cef
no ipv6 cef
multilink bundle-name authenticated

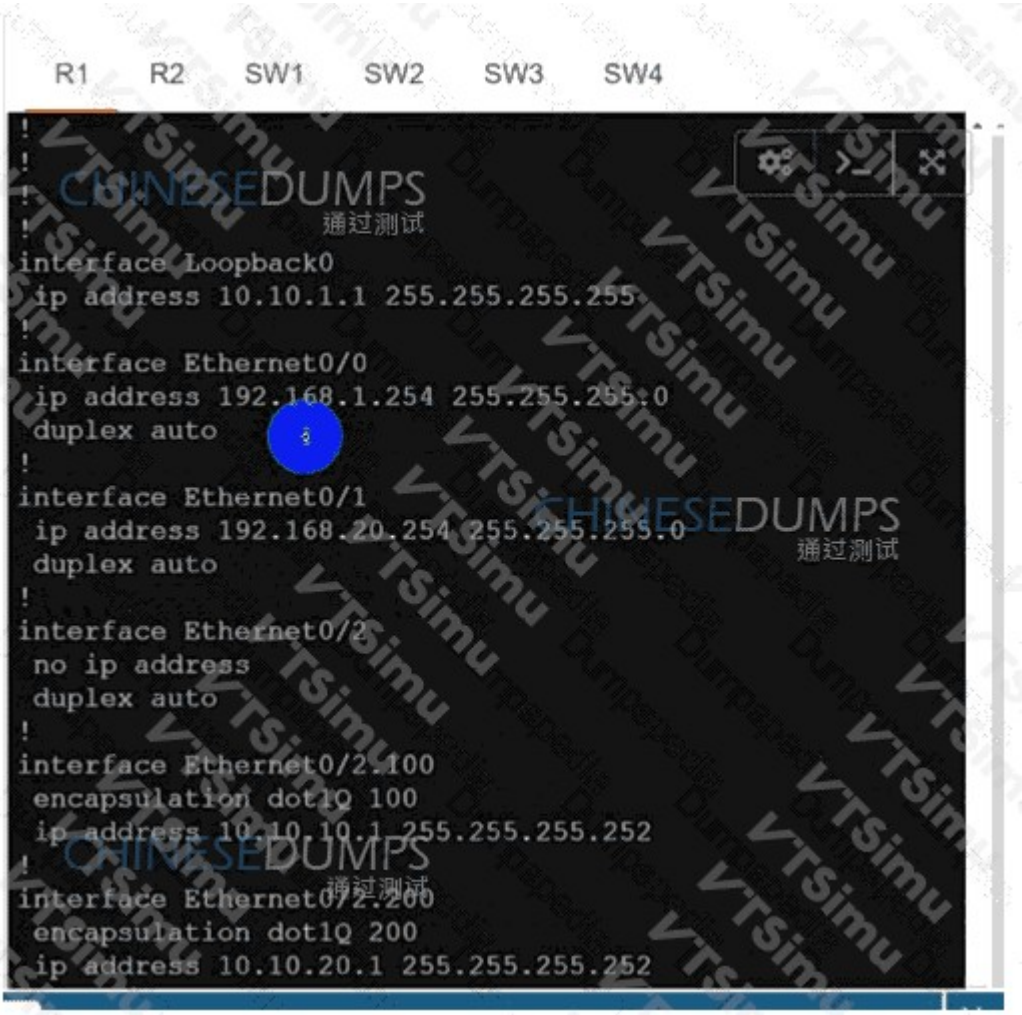
R1
R2
SW1
SW2
SW3
SW4
!
!
!
!
interface Loopback0
ip address 10.10.1.1 255.255.255.255
!
interface Ethernet0/0
ip address 192.168.1.254 255.255.255.0
duplex auto
!
interface Ethernet0/1
ip address 192.168.20.254 255.255.255.0
duplex auto
!
interface Ethernet0/2
no ip address
duplex auto
!
interface Ethernet0/2.100
encapsulation dot1Q 100
ip address 10.10.10.1 255.255.255.252
!
interface Ethernet0/2.200
encapsulation dot1Q 200
ip address 10.10.20.1 255.255.255.252

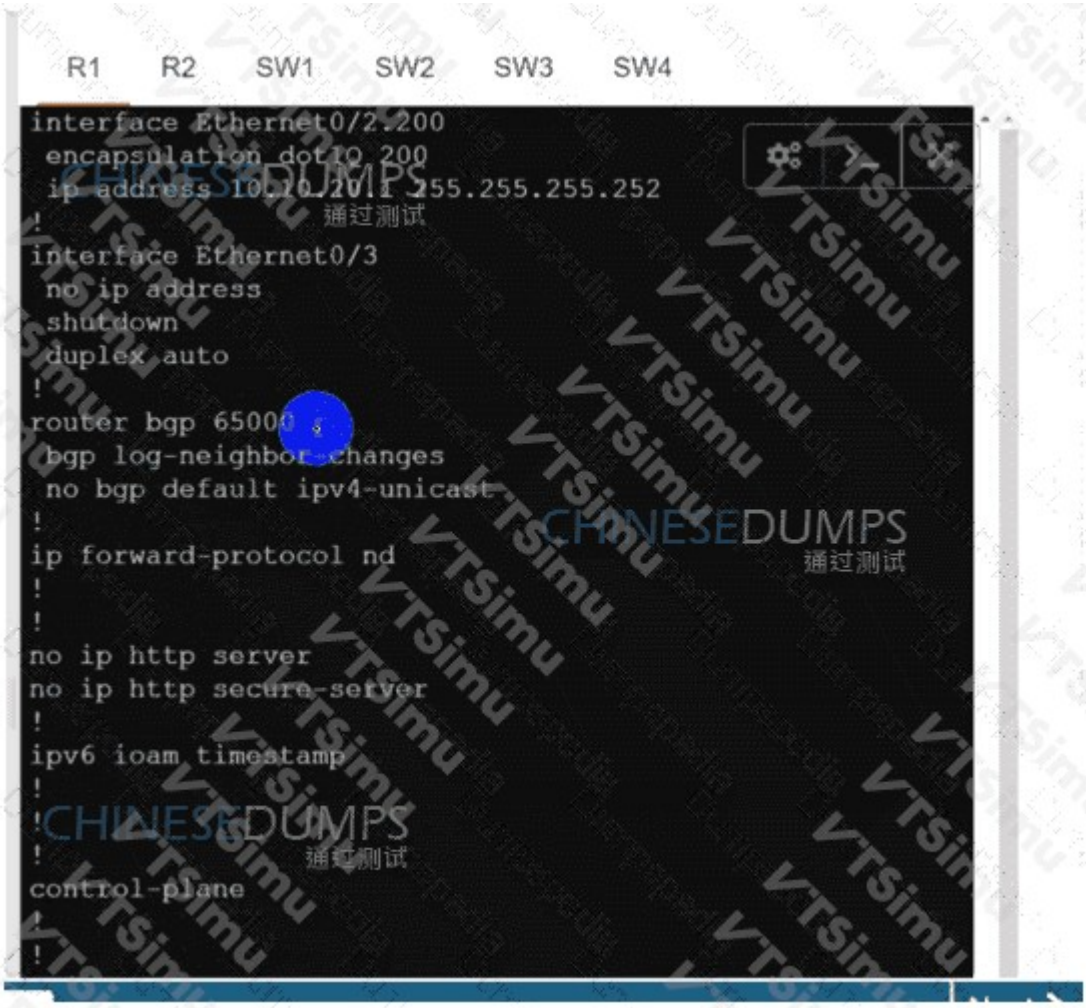
R1
R2
SW1
SW2
SW3
SW4
interface Ethernet0/2.200
encapsulation dot1Q 200
ip address 10.10.20.1 255.255.255.252
!
interface Ethernet0/3
no ip address
shutdown
duplex auto
!
router bgp 65000
bgp log-neighbor-changes
no bgp default ipv4-unicast
!
ip forward-protocol nd
!
!
no ip http server
no ip http secure-server
!
ipv6 ioam timestamp
!
!
control-plane
!
!

R2


R1
R2
SW1
SW2
SW3
SW4
R2>en
R2#Show run
Building configuration...
Current configuration : 1353 bytes
!
version 15.8
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname R2
!
boot-start-marker
boot-end-marker
!
!
!
no aaa new-model
!
!
!
clock timezone PST -8 0
mmi polling-interval 60
no mmi auto-configure
CHINESEDUMPS
通过测试


R1
R2
SW1
SW2
SW3
SW4
!
!
!
!
!
!
!
ip vrf cu-green
rd 65000:200
!
ip vrf cu-red
rd 65000:100
!
!
!
!
no ip domain lookup
ip cef
no ipv6 cef
!
multilink bundle-name authenticated
!
!


R1
R2
SW1
SW2
SW3
SW4
!
!
!
!
interface Loopback0
ip address 10.10.2.2 255.255.255.255
!
interface Ethernet0/0
ip address 192.168.2.254 255.255.255.0
duplex auto
!
interface Ethernet0/1
ip address 192.168.22.254 255.255.255.0
duplex auto
!
interface Ethernet0/2
no ip address
duplex auto
!
interface Ethernet0/2.100
encapsulation dot1Q 100
ip address 10.10.10.2 255.255.255.252
!
interface Ethernet0/2.200
encapsulation dot1Q 200
ip address 10.10.20.2 255.255.255.252

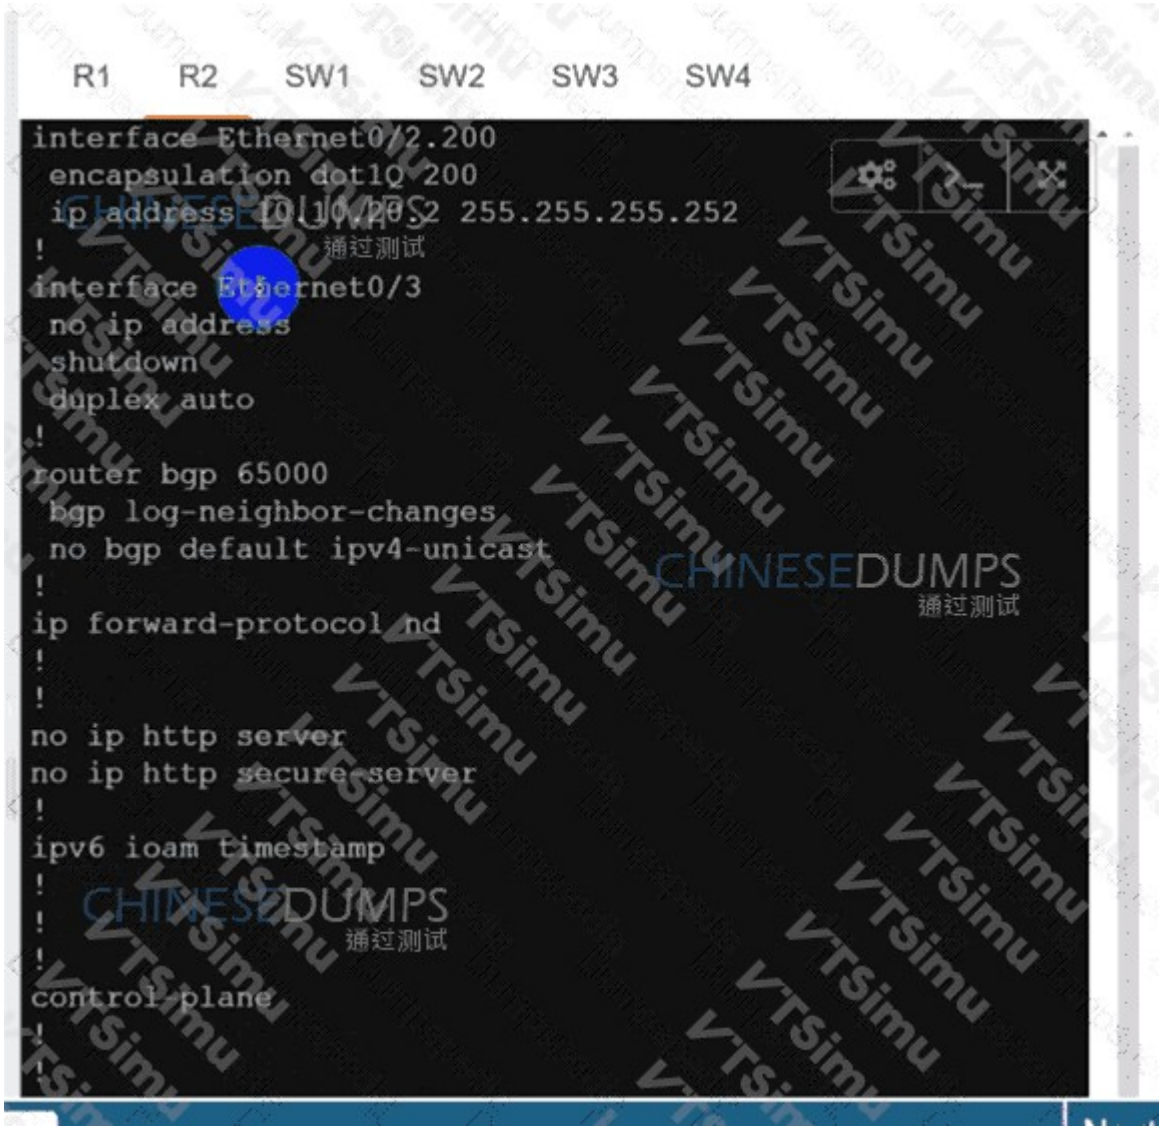
R1
R2
SW1
SW2
SW3
SW4
interface Ethernet0/2.200
encapsulation dot1Q 200
ip address 10.10.20.2 255.255.255.252
!
interface Ethernet0/3
no ip address
shutdown
duplex auto
!
router bgp 65000
bgp log-neighbor-changes
no bgp default ipv4-unicast
!
ip forward-protocol nd
!
!
no ip http server
no ip http secure-server
!
ipv6 ioam timestamp
!
!
!
control-plane
!
!
CHINESEDUMPS
通过测试

SW1


R1
R2
SW1
SW2
SW3
SW4
SW1>en
SW1#sh run
Building configuration...
Current configuration : 942 bytes
!
! Last configuration change at 04:43:09 PST Sat May 7 20
22
!
version 15.2
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
!
hostname SW1
!
boot-start-marker
boot-end-marker
!
!
!
no aaa new-model
clock timezone PST -8 0
!

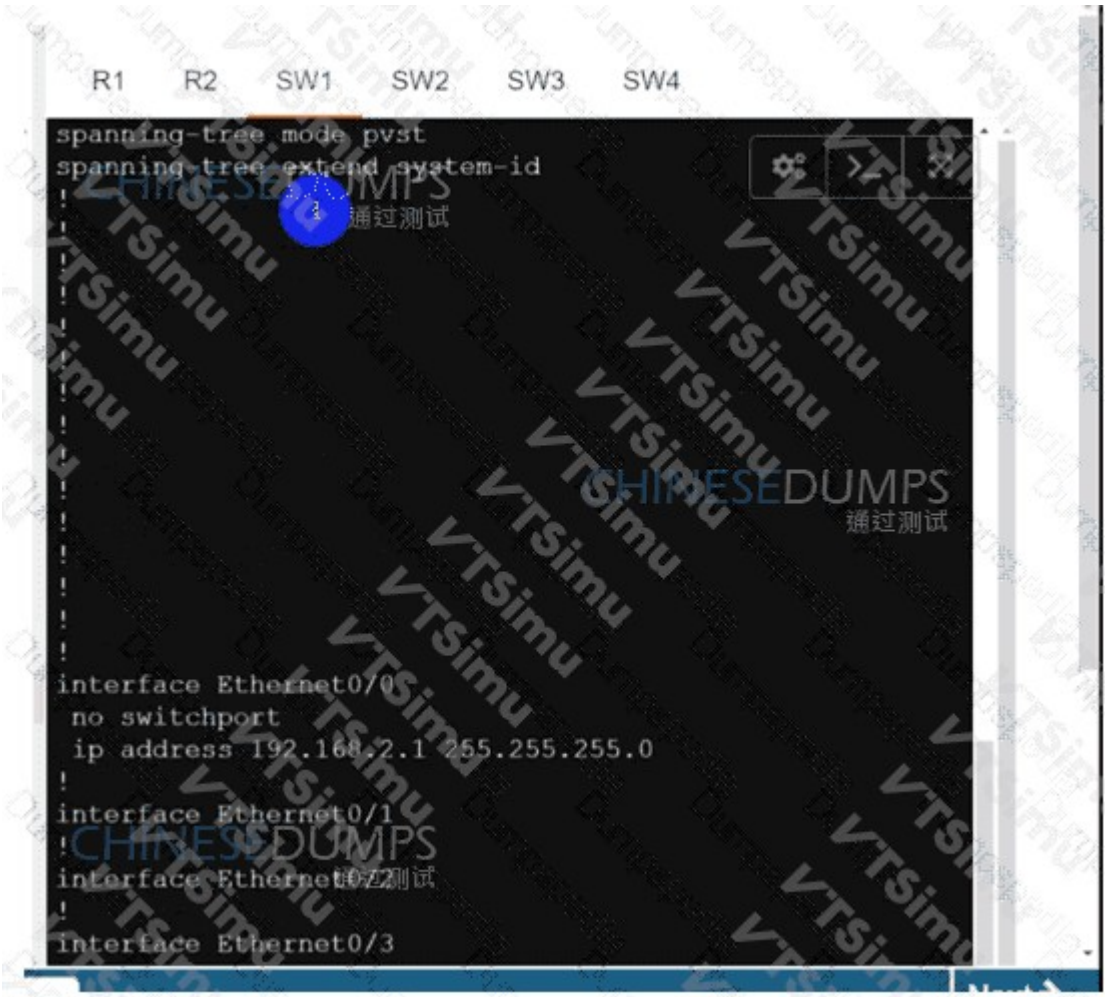
R1
R2
SW1
SW2
SW3
SW4
spanning-tree mode pvst
spanning-tree extend system-id
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Ethernet0/0
 no switchport
 ip address 192.168.2.1 255.255.255.0
!
interface Ethernet0/1
!
interface Ethernet0/2
!
interface Ethernet0/3

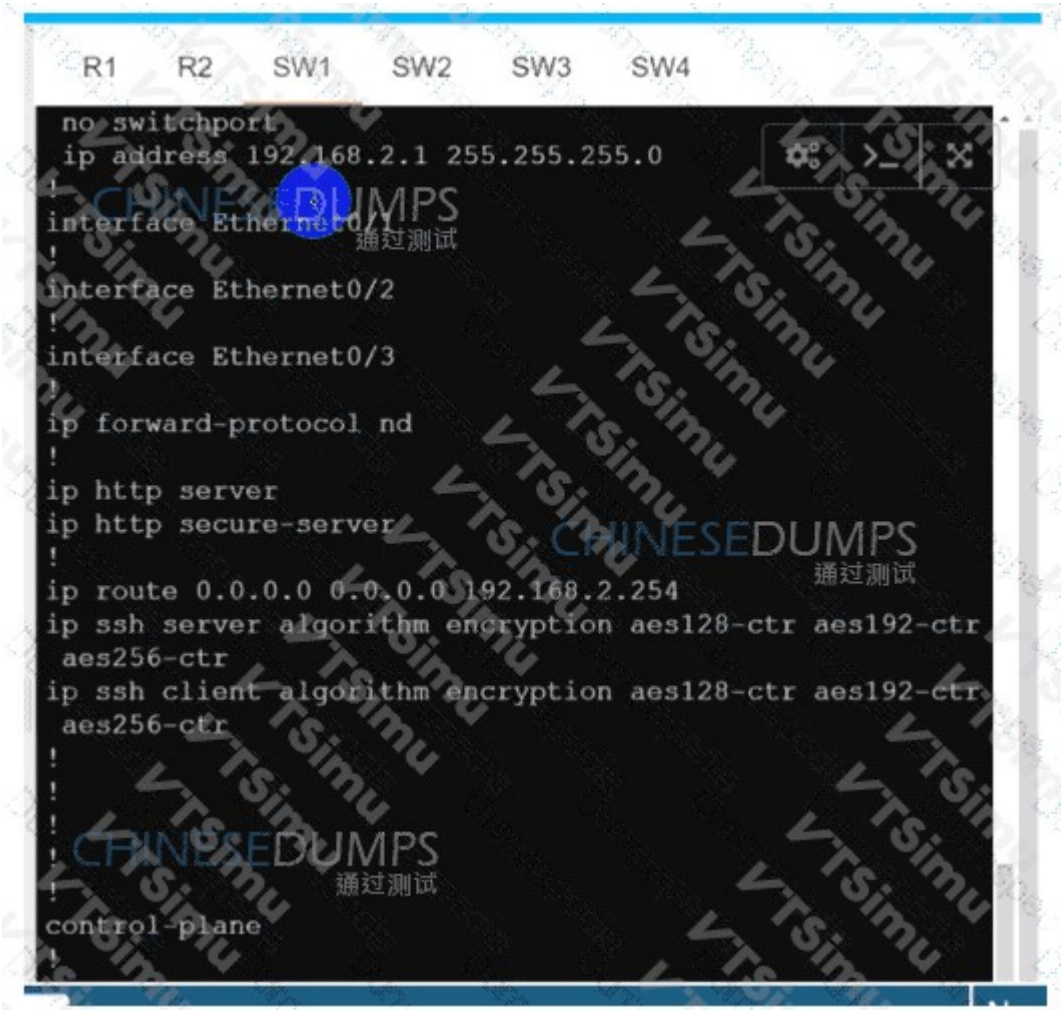
R1
R2
SW1
SW2
SW3
SW4
no switchport
ip address 192.168.2.1 255.255.255.0
!
interface Ethernet0/1
!
interface Ethernet0/2
!
interface Ethernet0/3
!
ip forward-protocol nd
!
ip http server
ip http secure-server
!
ip route 0.0.0.0 0.0.0.0 192.168.2.254
ip ssh server algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
ip ssh client algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
!
!
!
!
!
control-plane
!

SW2

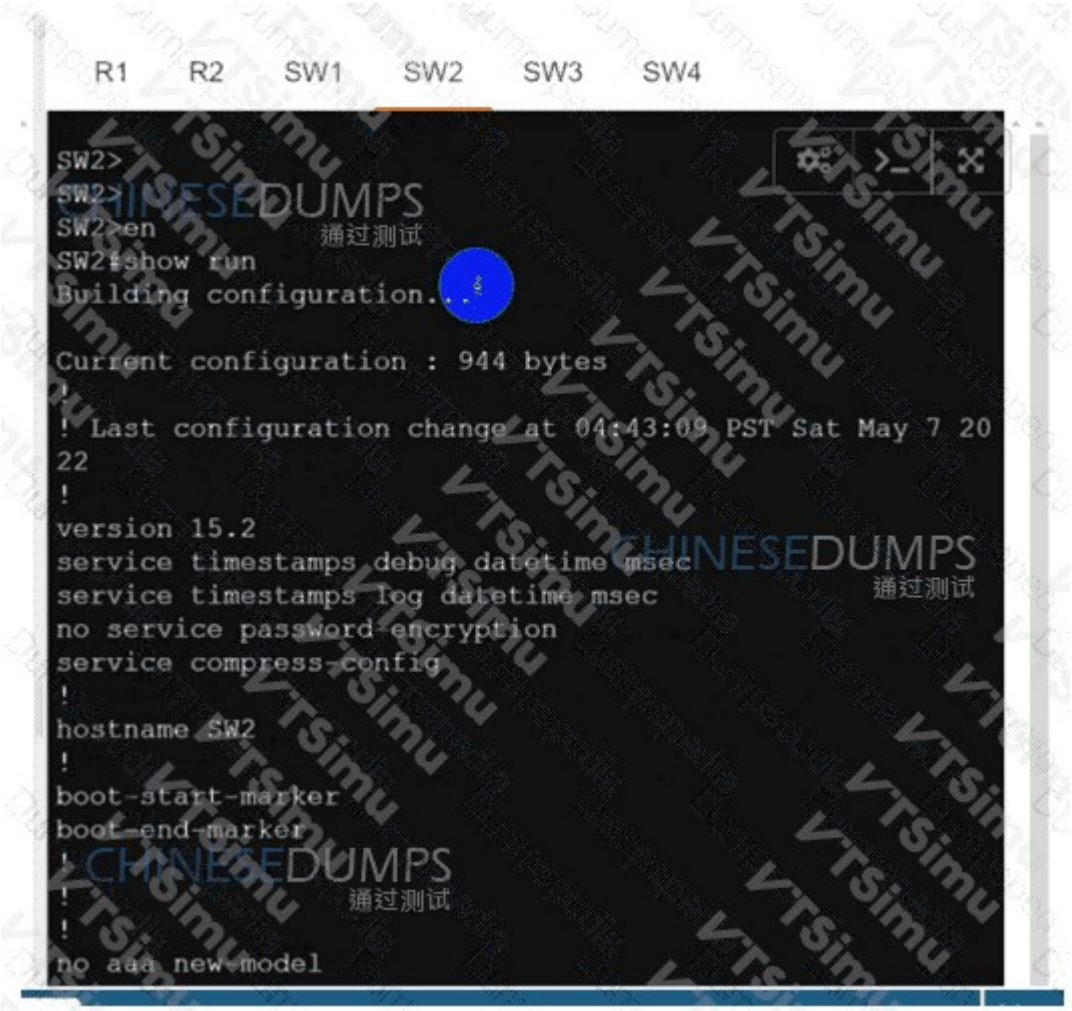
R1
R2
SW1
SW2
SW3
SW4
SW2>
SW2>
SW2>en
SW2#show run
Building configuration...
Current configuration : 944 bytes
!
! Last configuration change at 04:43:09 PST Sat May 7 20
22
!
version 15.2
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
!
hostname SW2
!
boot-start-marker
boot-end-marker
!
!
!
no aaa new-model


R1
R2
SW1
SW2
SW3
SW4
spanning-tree mode pvst
spanning-tree extend system-id
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Ethernet0/0
!
interface Ethernet0/1
no switchport
ip address 192.168.22.1 255.255.255.0
!
interface Ethernet0/2
!
interface Ethernet0/3

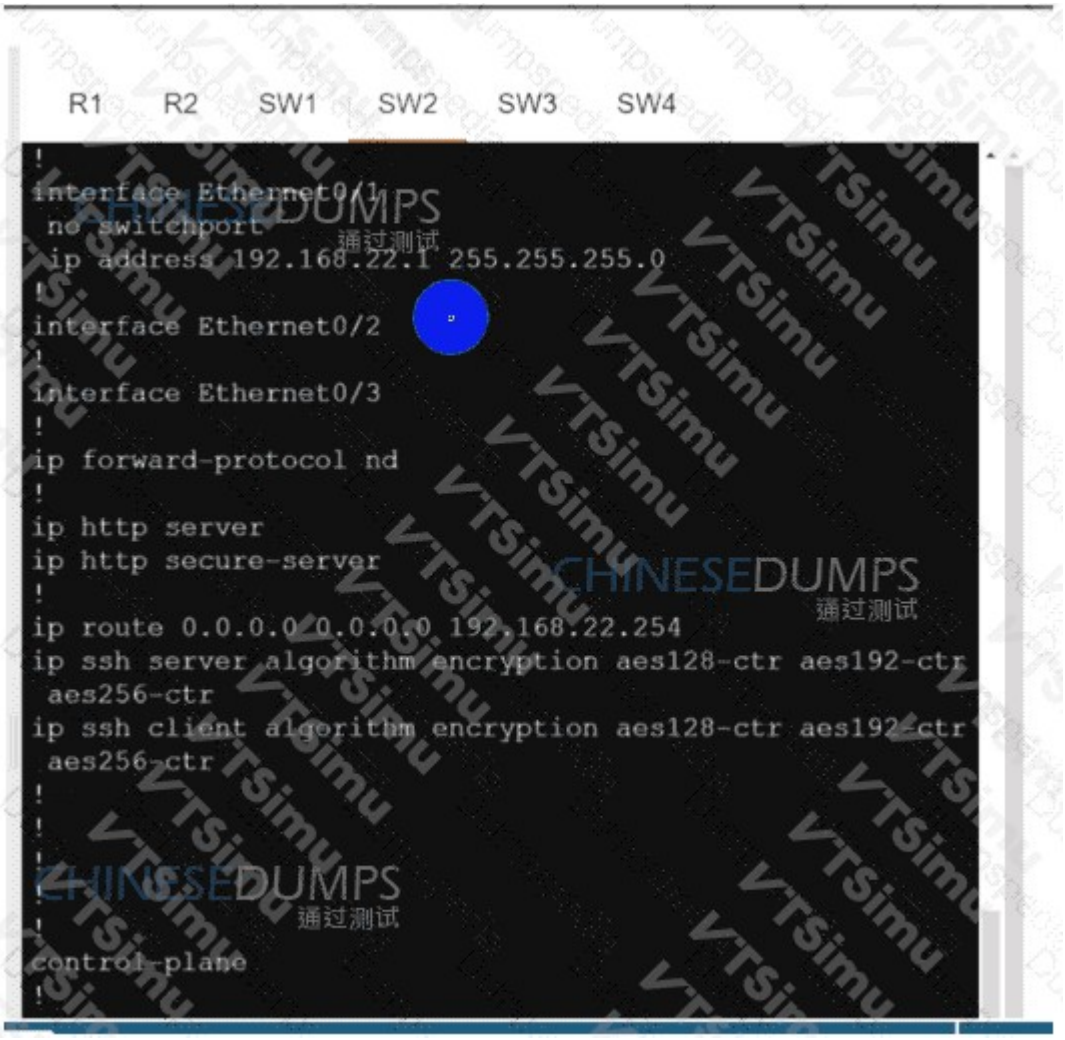
R1
R2
SW1
SW2
SW3
SW4
!
interface Ethernet0/1
no switchport
ip address 192.168.22.1 255.255.255.0
!
interface Ethernet0/2
!
interface Ethernet0/3
!
ip forward-protocol nd
!
ip http server
ip http secure-server
!
ip route 0.0.0.0 0.0.0.0 192.168.22.254
ip ssh server algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
ip ssh client algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
!
!
!
!
control-plane
!

SW3

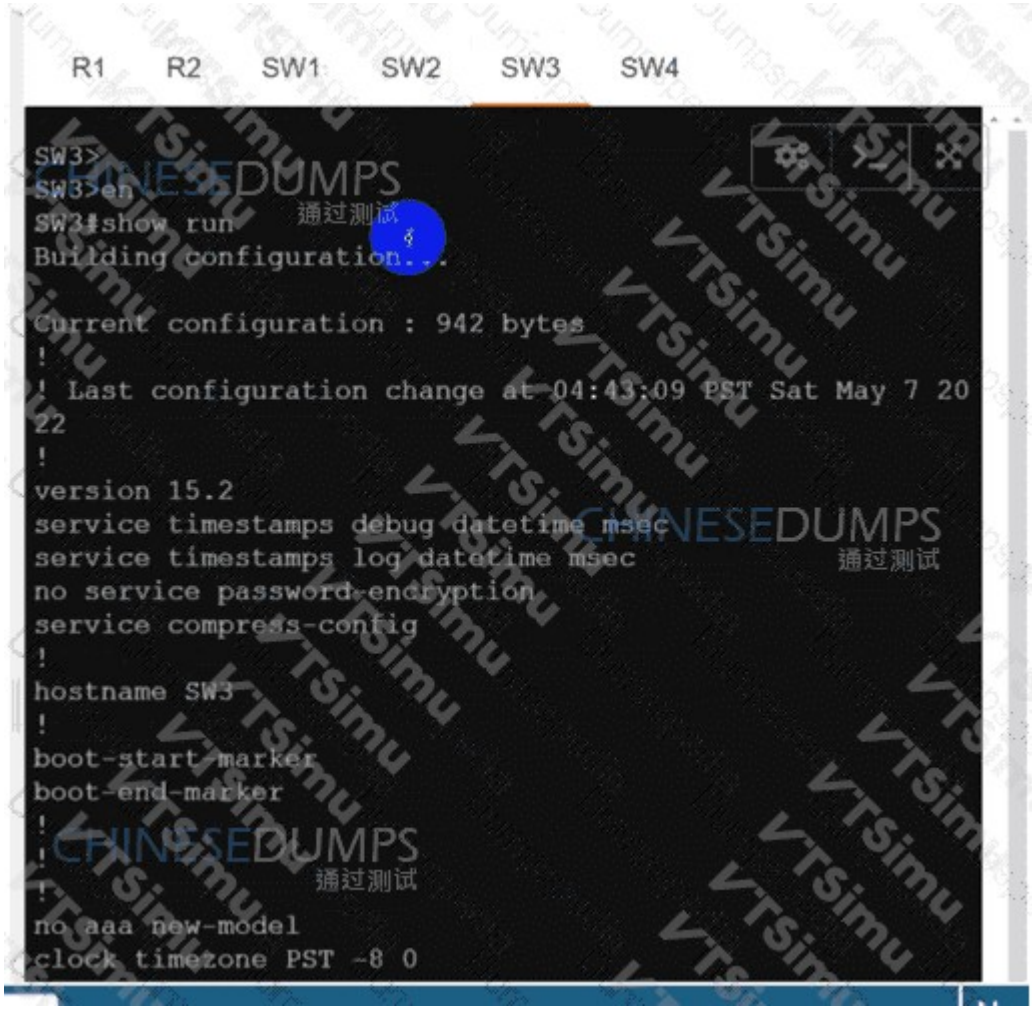
R1
R2
SW1
SW2
SW3
SW4
SW3>
SW3>en
SW3#show run
Building configuration...
Current configuration : 942 bytes
!
! Last configuration change at 04:43:09 PST Sat May 7 20
22
!
version 15.2
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
!
hostname SW3
!
boot-start-marker
boot-end-marker
!
!
!
no aaa new-model
clock timezone PST -8 0

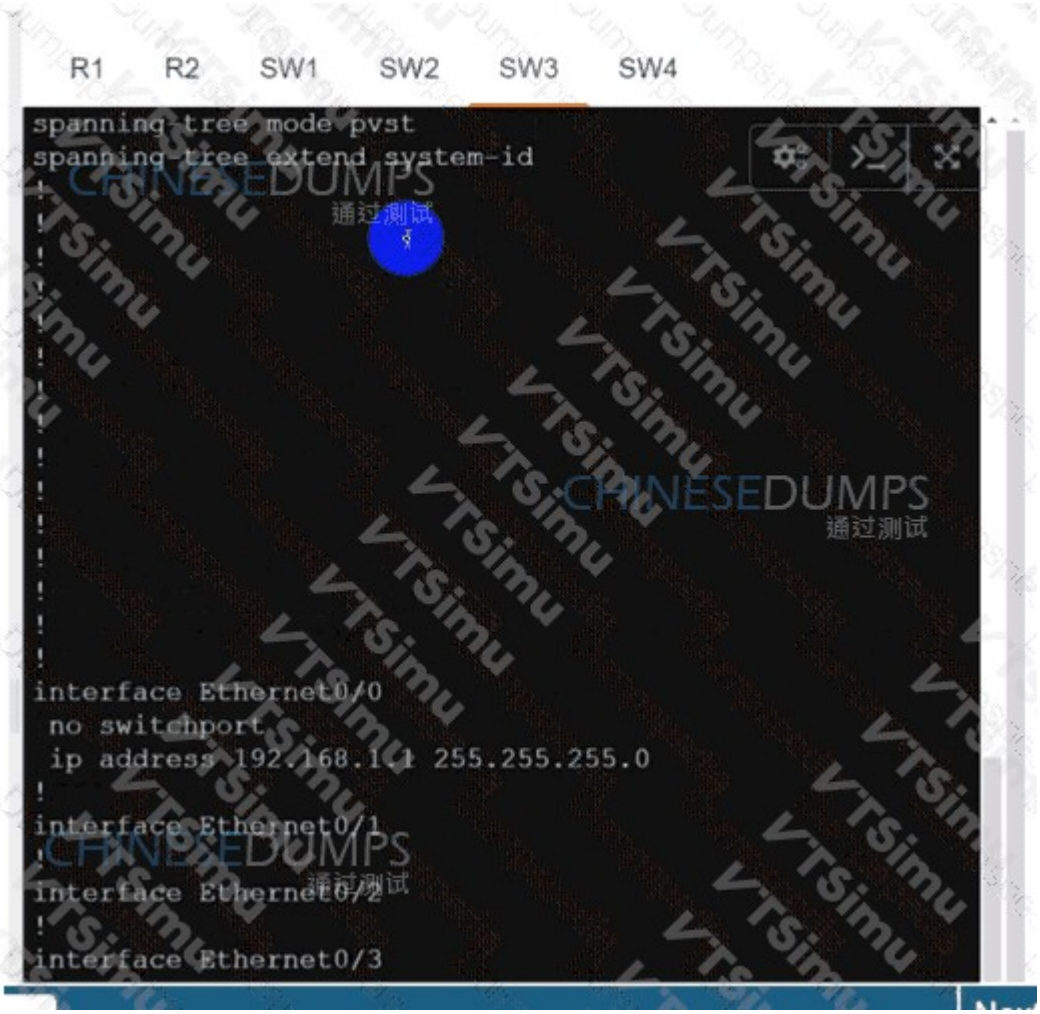
R1
R2
SW1
SW2
SW3
SW4
spanning-tree mode pvst
spanning-tree extend system-id
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Ethernet0/0
 no switchport
 ip address 192.168.1.1 255.255.255.0
!
interface Ethernet0/1
!
interface Ethernet0/2
!
interface Ethernet0/3
CHINESEDUMPS
通过测试

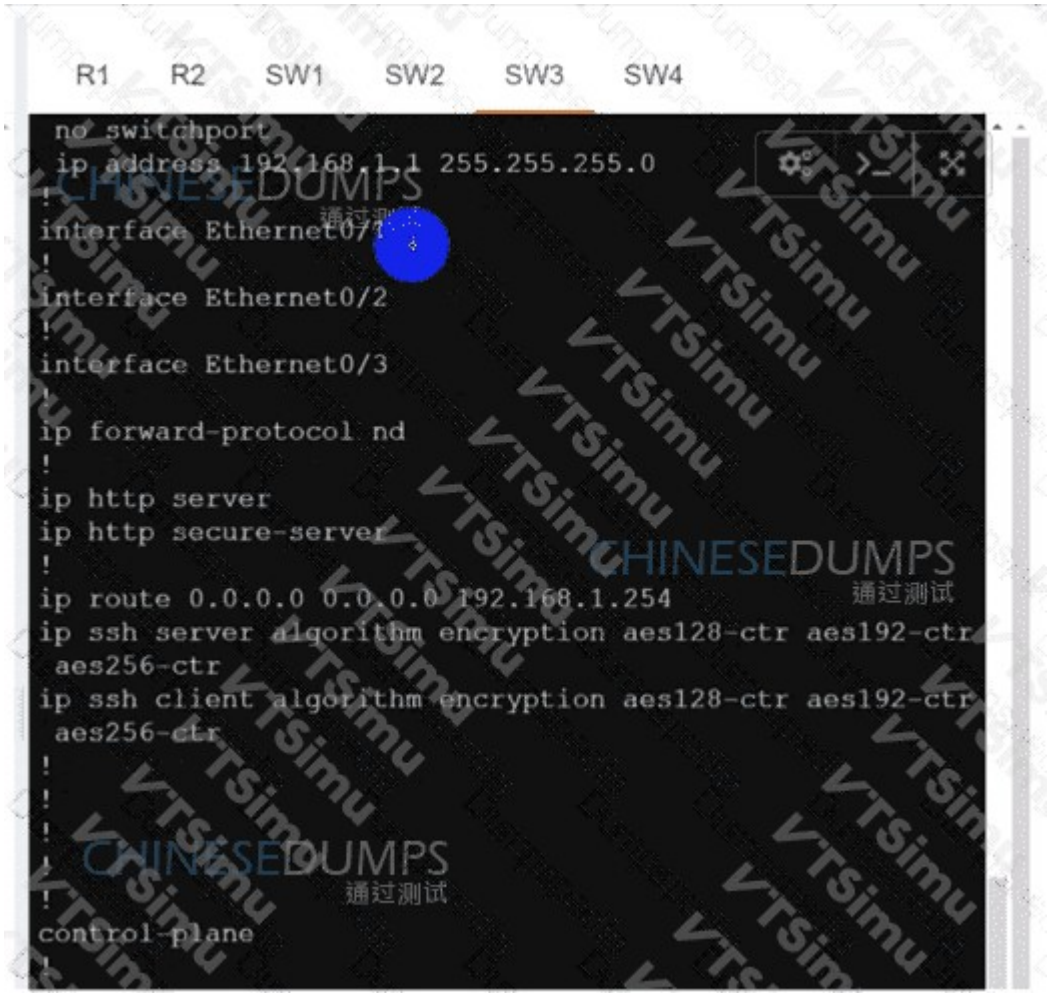
R1
R2
SW1
SW2
SW3
SW4
no switchport
ip address 192.168.1.1 255.255.255.0
!
interface Ethernet0/1
!
interface Ethernet0/2
!
interface Ethernet0/3
!
ip forward-protocol nd
!
ip http server
ip http secure-server
!
ip route 0.0.0.0 0.0.0.0 192.168.1.254
ip ssh server algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
ip ssh client algorithm encryption aes128-ctr aes192-ctr
aes256-ctr
!
!
!
!
!
control-plane
!


R1
R2
SW1
SW2
SW3
SW4
SW4>en
SW4#show run
Building configuration...
Current configuration : 944 bytes
!
! Last configuration change at 04:43:09 PST Sat May 7 20
22
!
version 15.2
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
!
hostname SW4
!
boot-start-marker
boot-end-marker
!
!
!
no aaa new-model
clock timezone PST -8 0
!


R1
R2
SW1
SW2
SW3
SW4
spanning-tree mode pvst
spanning-tree extend system-id
interface Ethernet0/0
interface Ethernet0/1
no switchport
ip address 192.168.20.1 255.255.255.0
interface Ethernet0/2
interface Ethernet0/3
CHINESEDUMPS
通过测试

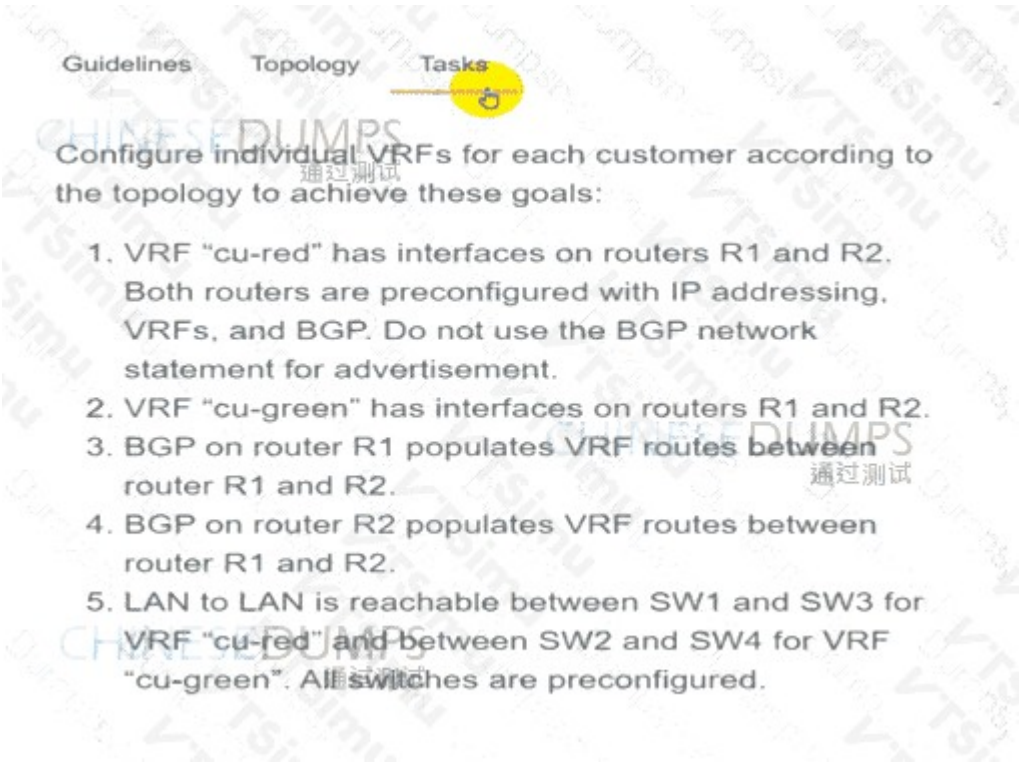

ANSWER: See the explanation for the answer

Explanation:

Configure the LAN-facing interfaces into their respective VRFs, add matching import/export route targets, enable MPLS on the two R1–R2 transport subinterfaces, and advertise the VRF-connected LAN routes through MP-BGP VPNv4. Do **not** use a BGP `network` statement.

R1

```
conf t
ip vrf cu-red
 route-target both 65000:100
!
ip vrf cu-green
 route-target both 65000:200
!
interface Ethernet0/0
 ip vrf forwarding cu-red
 ip address 192.168.1.254 255.255.255.0
 no shutdown
!
interface Ethernet0/1
 ip vrf forwarding cu-green
 ip address 192.168.20.254 255.255.255.0
 no shutdown
!
interface Ethernet0/2.100
 mpls ip
!
interface Ethernet0/2.200
 mpls ip
!
router bgp 65000
 neighbor 10.10.10.2 remote-as 65000
 neighbor 10.10.20.2 remote-as 65000
 address-family vpnv4
  neighbor 10.10.10.2 activate
  neighbor 10.10.10.2 send-community extended
  neighbor 10.10.20.2 activate
  neighbor 10.10.20.2 send-community extended
 exit-address-family
 address-family ipv4 vrf cu-red
  redistribute connected
 exit-address-family
 address-family ipv4 vrf cu-green
  redistribute connected
 exit-address-family
end
```

R2

```
conf t
ip vrf cu-red
 route-target both 65000:100
!
ip vrf cu-green
 route-target both 65000:200
!
interface Ethernet0/0
 ip vrf forwarding cu-red
 ip address 192.168.2.254 255.255.255.0
 no shutdown
!
interface Ethernet0/1
 ip vrf forwarding cu-green
 ip address 192.168.22.254 255.255.255.0
 no shutdown
!
interface Ethernet0/2.100
```

```
 mpls ip
!
interface Ethernet0/2.200
 mpls ip
!
router bgp 65000
 neighbor 10.10.10.1 remote-as 65000
 neighbor 10.10.20.1 remote-as 65000
 address-family vpnv4
  neighbor 10.10.10.1 activate
  neighbor 10.10.10.1 send-community extended
  neighbor 10.10.20.1 activate
  neighbor 10.10.20.1 send-community extended
 exit-address-family
 address-family ipv4 vrf cu-red
  redistribute connected
 exit-address-family
 address-family ipv4 vrf cu-green
  redistribute connected
 exit-address-family
end
```

The matching route targets keep the customers isolated:

`cu-red`: RT `65000:100`, connecting SW1 (`192.168.1.0/24`) to SW3 (`192.168.2.0/24`).
`cu-green`: RT `65000:200`, connecting SW2 (`192.168.22.0/24`) to SW4 (`192.168.20.0/24`).

Verify with `show ip bgp vpnv4 all summary`, `show ip route vrf cu-red`, and `show ip route vrf cu-green`. SW1 must reach SW3 but not the green subnets; SW2 must reach SW4 but not the red subnets.

QUESTION NO: 53

Refer to the exhibit. An enterprise decides to use the Cisco SD-WAN Cloud onRamp for SaaS feature and utilize H.Q site Biz iNET to reach SaaS Cloud for branch C. currently reaching SaaS Cloud directly. Which role must be assigned to devices at both sites in vManage Cloud Express for this solution to work?

A. H.Q to be added as Gateway and Branch as DIA.

B. Branch to be added as Client Sites and H.Q as DIA.

C. Branch to be added as DIA and H.Q as Client Site.

D. H.Q to be added as Gateway and Branch as Client Site.

ANSWER: D

Explanation:

H.Q to be added as Gateway and Branch as Client Site. is correct because Cloud OnRamp for SaaS Cloud Express uses gateway sites to provide the selected alternate SaaS egress path for other locations. In this scenario, the enterprise intends to use the HQ business-Internet connection for Branch C's SaaS access. Therefore, HQ is the site that supplies the preferred egress and must be configured as a gateway site. Branch C is the consuming location whose SaaS traffic can be steered across the SD-WAN fabric to that selected HQ gateway, so it must be configured as a client site.

Cloud Express continuously evaluates SaaS-path performance and can direct a client site's SaaS flows toward a suitable gateway based on the configured policy and measured application-path conditions. The client/gateway relationship is what enables Branch C to use HQ's business Internet circuit rather than relying solely on its own local Internet breakout. This role assignment also lets SD-WAN Manager distribute the required routing and traffic-steering intent to the participating WAN edge devices. Cisco documents Cloud OnRamp for SaaS deployment and Cloud Express configuration concepts in the Cisco SD-WAN Cloud OnRamp for SaaS Configuration Guide and its Cloud OnRamp overview.

QUESTION NO: 54

What is the role of the Session Traversal Utilities for NAT server provided by the vBond orchestrator?

A. It facilitates SD-WAN routers and controllers to discover their own mapped or translated IP addresses and port numbers

B. It prevents SD-WAN Edge routers from forming sessions with public transports among different service providers

C. It facilitates SD-WAN Edge routers to stay behind a NAT-enabled firewall while the transport addresses of the SD-WAN controller are unNAT-ed

D. It allows WAN Edge routers to form sessions among MPLS TLOCs using only public IP addresses

ANSWER: A

Explanation:

It facilitates SD-WAN routers and controllers to discover their own mapped or translated IP addresses and port numbers. In Cisco SD-WAN, the vBond orchestrator provides a STUN service as part of the onboarding and orchestration process. A WAN Edge device or controller that is located behind NAT sends a request to vBond, and vBond identifies the source public address and UDP port as seen from the Internet-facing side of the NAT device. This information is the device's NAT mapping, sometimes called its public or translated transport address.

Knowing that mapped address and port is essential for vBond to orchestrate secure control-plane connectivity between Cisco SD-WAN components. It enables vBond to provide peers with usable transport information and supports NAT traversal when establishing DTLS or TLS control connections. The mechanism is not limited to WAN Edge routers: controllers can also use the STUN service to learn their externally visible transport mapping when NAT is present. Cisco documents that vBond performs NAT discovery using STUN and communicates the discovered public IP address and port information during orchestration. See Cisco's NAT traversal configuration documentation and the vBond orchestrator documentation.

QUESTION NO: 55

Which VPN connects the transport-side WAN Edge interface to the underlay/WAN network?

A. VPN 1

B. VPN 0

C. VPN 512

D. VPN 511

ANSWER: B

Explanation:

VPN 0 is correct because it is the Cisco SD-WAN transport VPN. WAN Edge interfaces that face the WAN transport, such as Internet, MPLS, or other service-provider underlay networks, are configured in VPN 0. This VPN provides the underlay IP connectivity required for control-plane and data-plane tunnel establishment between WAN Edge devices. In particular, the transport interface receives reachability to controller addresses and remote tunnel endpoints through the WAN network, allowing the device to form secure IPsec tunnels that make up the SD-WAN overlay.

Cisco SD-WAN separates transport connectivity from service-side routing by using VPN segmentation. VPN 0 is reserved for the WAN-facing transport domain, while service VPNs carry user, branch, and application traffic. Placing the transport-side interface in VPN 0 ensures that the device can communicate across the provider underlay and establish its overlay connections independently of the customer service VPNs. Cisco also reserves specific VPN identifiers for internal purposes, but the fundamental transport-facing role belongs to VPN 0. See Cisco's transport interface documentation and the Cisco SD-WAN VPN overview.

QUESTION NO: 56

Which feature template configures OMP?

A.

Section	Parameter	Type	Variable/value
Basic configuration	Number of paths advertised per prefix	Global	16
Advertise	Radio	Global	Off
	Static	Global	Off

B.

Section	Parameter	Type	Variable/value
Server	Hostname/IP Address	Global	10.4.48.13
	VPN ID	Global	1
	Source interface	Global	loopback0

C.

Section	Parameter	Type	Variable/value
Authentication	Authentication Order	Drop-down	local
Local	User/admin/ Password	Device Specific	user_admin_passwd

D.

Section	Parameter	Type	Variable/value
Basic configuration	VPN	Global	512
	Name	Global	Management VPN
IPv4 Route	Prefix	Global	0.0.0.0/0
	Gateway	Radio button	Next Hop
	Next Hop	Device Specific	vpn512_mgt_next_hop_ip_addr

A. Option A

B. Option B

C. Option C

ANSWER: A

Explanation:

Option A is correct because it represents the Cisco SD-WAN feature template intended for Overlay Management Protocol configuration. OMP is the SD-WAN control-plane protocol that distributes routing, TLOC, service, and policy information between WAN Edge devices through the Cisco SD-WAN controllers. Its parameters are configured in the OMP feature-template context and can include advertisement behavior, graceful restart, hold time, and route-policy controls that influence the routes advertised into or accepted from the SD-WAN overlay.

In Cisco vManage configuration workflows, a device template combines multiple feature templates for the relevant device functions. The OMP-specific feature template is therefore the appropriate location for settings that control OMP route exchange, rather than a template used for system, interface, VPN, or transport configuration. This separation enables common OMP policy and control-plane behavior to be consistently attached to multiple devices through reusable templates. Cisco documents OMP as the protocol responsible for exchanging overlay routing information and describes its configuration

under the OMP configuration mode. See the Cisco SD-WAN OMP configuration guide and the Cisco OMP configuration documentation.

QUESTION NO: 57

Which two REST API functions are performed for Cisco devices in an overlay network? (Choose two)

A. distributing a Snort image among devices

B. attaching a device configuration template

C. managing connections for smart licensing

D. monitoring device certificates

E. querying a device and aggregating statistics

ANSWER: B E

Explanation:

Cisco SD-WAN Manager (formerly vManage) REST APIs support both configuration automation and operational monitoring of WAN-edge devices in the SD-WAN overlay. **attaching a device configuration template** is a configuration workflow exposed through the API: automation can select target devices, associate the required device template and feature variables, and initiate the template-attachment action. This enables consistent, repeatable provisioning without requiring an administrator to perform every operation in the graphical interface.

querying a device and aggregating statistics is an operational monitoring workflow. SD-WAN Manager REST endpoints return device inventory and status information as well as telemetry such as interface, tunnel, application, and system statistics. API consumers can query this data for selected devices and use the returned records to build reports, dashboards, alerts, or capacity-analysis processes. These two functions represent the core configuration and monitoring uses of the SD-WAN REST API for overlay devices. Cisco documents the available SD-WAN Manager API resources and workflows in the Cisco SD-WAN REST API documentation and describes template-based device provisioning in the Cisco SD-WAN Manager Configuration Guide.

QUESTION NO: 58

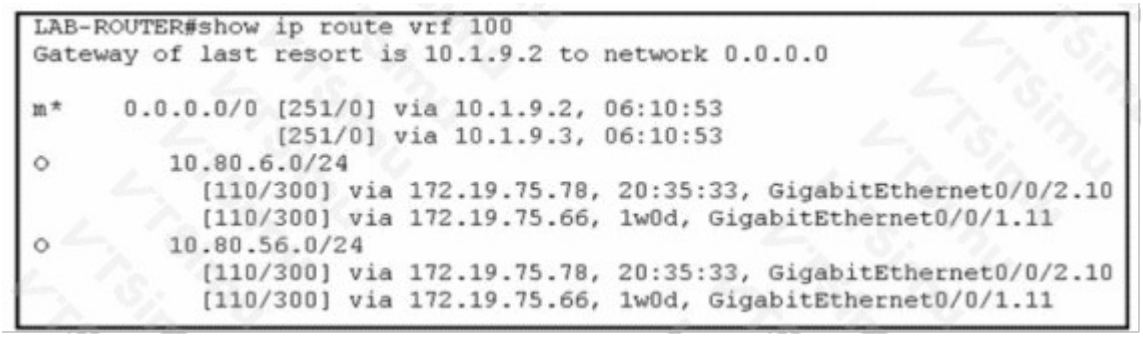

Refer to the exhibit. The network administrator has configured a centralized topology policy that results in the displayed routing table at a branch office. Which two configurations are verified by the output? [Choose two.)

A. The routing table is for the transport VPN.

B. The default route is learned via OMP.

C. This routing table is from a cEdge router.

D. The default route is configured locally.

E. The configured policy is adding a route tag of 300 to learned routes.

ANSWER: B C

Explanation:

The output verifies that **The default route is learned via OMP.** Cisco SD-WAN distributes reachable prefixes, including a default prefix when advertised, through the Overlay Management Protocol. In an IOS XE SD-WAN routing-table display, the route's protocol indicator identifies OMP as the source of the installed prefix. Thus, the displayed default route is an overlay-learned route rather than merely an unspecified gateway-of-last-resort entry. OMP routes are installed in the appropriate service VPN routing table after they are received from the SD-WAN control plane and accepted by policy.

The output also verifies that **This routing table is from a cEdge router.** The routing-table format and IOS XE route-code conventions shown in the exhibit are characteristic of a Cisco IOS XE SD-WAN device, commonly called a cEdge. cEdge routers use IOS XE routing-table presentation while participating in the Cisco SD-WAN overlay and learning routes through OMP. This combination of IOS XE-style operational output and an OMP-derived route identifies the platform type and route source together. Cisco documents OMP as the SD-WAN control-plane protocol used to advertise and learn service routes, and documents IOS XE SD-WAN routing behavior for cEdge devices. See Cisco OMP Configuration Guide and Cisco IOS XE SD-WAN Routing Configuration Guide.

QUESTION NO: 59

Refer to exhibit.

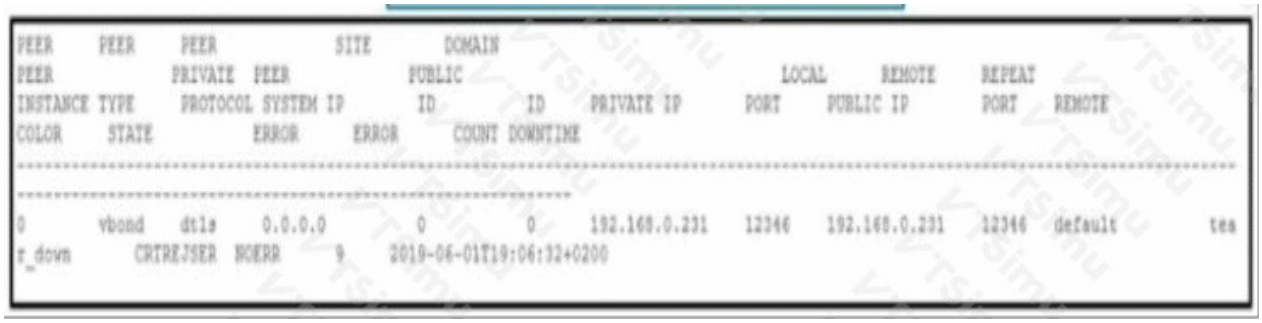

An engineer is troubleshooting tear down of control connections even though a valid Certificate Serial Number is entered Which two actions resolve the Issue? (Choose two)

A. Enter a valid serial number on the controllers for a given device

B. Remove the duplicate IP in the network.

C. Enter a valid product ID (model) on the PNP portal

D. Match the serial number file between the controllers

E. Restore network reachability for the controller

ANSWER: C D

Explanation:

Enter a valid product ID (model) on the PNP portal and **Match the serial number file between the controllers** resolve certificate authorization inconsistencies that can cause Cisco SD-WAN control connections to be established and then torn down. A WAN Edge device identity is represented by more than the chassis serial number: the Plug and Play record must associate the device with the correct product identifier and serial-number information. This lets the Cisco SD-WAN controller validate the device certificate and its hardware identity correctly during control-plane authentication.

The controller cluster must also use consistent authorized-device serial information. In particular, the controller responsible for orchestration and the management controller need matching serial-number authorization data so that the device receives a consistent validation result throughout the control-connection process. When the product identifier is valid in Plug and Play and the serial-number file is synchronized between controllers, certificate-based authentication can complete consistently and the control connection remains operational. Cisco's control-connection troubleshooting guidance discusses checking certificate and serial-file information: Cisco SD-WAN Routers Troubleshoot Control Connections. Cisco also documents WAN Edge authorization and certificate procedures in the Cisco SD-WAN Certificate Configuration Guide.