Cisco Certified Network Professional Enterprise (CCNP)

Exam 350-401 Implementing and Operating Cisco Enterprise Network Core Technologies (ENCOR)

Demo Questions

CISCO

CISCO CERTIFIED CCNP

(350-401) Implementing Cisco Enterprise Network Core Technologies

注意：題目有 ★ 標記只代表近期出過的題目，★ 數量越多，代表過去多個月出的次數越多，但同學必須溫習“整份”題目才可提升考試合格機會。 (★=Jun, ★=Jul, ★=Aug, ★=Sep)

QUESTION 1 ★★

What is a characteristic of a traditional WAN?

A. low complexity and high overall solution scale

B. unified data plane and control plane

C. centralized reachability, security, and application policies

D. operates over DTLS and TLS authenticated and secured tunnels

Answer: B

- **Traditional WANs** are built on legacy architectures where the **control plane (routing decisions)** and the **data plane (traffic forwarding)** are tightly coupled and not separated.
- This makes them less flexible compared to modern SD-WAN solutions, which separate control and data planes for scalability and centralized policy enforcement.

Why the other options are incorrect

A. Low complexity and high scale → Traditional WANs are actually complex and difficult to scale.

C. Centralized reachability, security, and application policies → This is a feature of SD-WAN, not traditional WAN.

D. Operates over DTLS/TLS tunnels → Again, this is SD-WAN, which uses secure overlay tunnels.

QUESTION 2 ★★

Which network design principle should be followed to improve Layer 2 stability in an enterprise campus design?

A. Extend only required user VLANs throughout the aggregation layer.

B. Extend only required user VLANs throughout the access layer.

C. Extend all user VLANs throughout the core layer.

D. Extend all user VLANs throughout the access layer.

Answer: A

Why A is correct:

Limiting VLAN extension to only those needed at the aggregation layer reduces spanning-tree complexity, minimizes broadcast domains, and improves Layer 2 stability. This design aligns with hierarchical campus best practices.

Why others are incorrect:

B. Access layer only → The access layer is already the VLAN boundary; extending VLANs here doesn’t improve stability.

C. Core layer all VLANs → Extending all VLANs into the core increases STP complexity and instability.

D. Access layer all VLANs → Extending all VLANs across access switches enlarges broadcast domains and risks instability.

QUESTION 3

In VXLAN, what is the source of encapsulated and de-encapsulated traffic?

A. virtual tunnel endpoint

B. network virtual interface

C. router interface on the VXLAN underlay

D. loopback 0

Answer: A

- In VXLAN, the VTEP is responsible for encapsulating Layer 2 Ethernet frames into VXLAN packets and de-encapsulating them back to Layer 2 when they arrive.
- Each VTEP typically resides on a physical or virtual device and uses an IP interface (often a loopback) for tunnel endpoints, but the actual source of encapsulation/de-encapsulation is the VTEP itself.

Why the other options are incorrect

B. Network virtual interface → Provides logical connectivity but does not perform VXLAN encapsulation.

C. Router interface on the VXLAN underlay → Carries VXLAN traffic but does not encapsulate/de-encapsulate it.

D. Loopback 0 → Often used as the VTEP's IP address, but the loopback interface itself is not the encapsulation source.

QUESTION 4

Which feature of a virtual machine allows it to run applications from different operating systems on the same physical hardware?

A. hypervisor

B. virtual CPU

C. virtual disk

D. virtual-machine monitor

Answer: A

- The **hypervisor** (also called a virtual-machine monitor) is the key feature that enables virtualization.
- It abstracts the physical hardware and allows multiple operating systems to run concurrently on the same physical machine by creating and managing virtual machines.
- Each VM can run its own OS and applications independently, sharing the underlying hardware resources.

Why the other options are incorrect

B. Virtual CPU → Provides processing resources to a VM, but doesn't enable multiple OS environments.

C. Virtual disk → Provides storage for a VM, but not the mechanism for running different OSes.

D. Virtual-machine monitor → This is essentially another name for a hypervisor, but in Cisco/IT exam contexts, the precise answer is hypervisor.

QUESTION 5 ★★

Which component does Cisco Threat Defense use to measure bandwidth, application performance, and utilization?

A. TrustSec

B. NetFlow

C. Cisco Umbrella

D. Advanced Malware Protection for Endpoints

Answer: B

- NetFlow is a Cisco-developed protocol that collects IP traffic information and monitors network flow.
- It provides detailed visibility into bandwidth usage, application performance, and traffic patterns, which is essential for threat defense and network optimization.
- By analyzing NetFlow data, administrators can detect anomalies, identify bottlenecks, and improve overall security posture.

Why not the other options?

- TrustSec: Focuses on segmentation and policy enforcement using security group tags, not bandwidth measurement.
- Cisco Umbrella: A cloud-delivered security service that provides DNS-layer protection, blocking malicious domains before connections are made. It doesn't measure bandwidth or utilization.
- Advanced Malware Protection (AMP) for Endpoints: Provides malware detection and remediation on endpoints. It's about threat detection, not traffic performance monitoring.

QUESTION 6 ★★★

What is a benefit of Cisco TrustSec in a multilayered LAN network design?

A. Policy can be applied on a hop-by-hop basis

B. There is no requirement to run IEEE 802.1X when TrustSec is enabled on a switch port

C. Application flows between hosts on the LAN to remote destinations can be encrypted

D. Policy or ACLs are not required

Answer: A

- Cisco TrustSec provides software-defined segmentation in a multilayered LAN design. Instead of relying on complex IP-based ACLs, TrustSec uses Security Group Tags (SGTs) to classify traffic.
- These tags travel with the packets, allowing switches, routers, and firewalls to enforce consistent policies at every hop in the network.
- This hop-by-hop enforcement is the key benefit: it simplifies policy management, improves scalability, and strengthens security across the LAN.

Why the other options are incorrect

B. No requirement to run IEEE 802.1X → TrustSec often integrates with 802.1X for identity-based access; it doesn't eliminate the need.

C. Application flows encrypted → That's handled by MACsec or IPsec, not TrustSec itself.

D. Policy or ACLs not required → TrustSec reduces ACL complexity but doesn't eliminate the need for policy enforcement.

QUESTION 7

Which two methods are used to assign security group tags to the user in a Cisco TrustSec architecture? (Choose two)

A. DHCP

B. web authentication

C. policy routing

D. modular QoS

E. IEEE 802.1x

Answer: B, E

Web authentication (B)

- Used when 802.1X is not available (e.g., devices without supplicants).
- The user is redirected to a web portal for authentication, and once validated, an SGT is assigned.

IEEE 802.1X (E)

- Provides identity-based access control at the switch port.
- When a user authenticates, Cisco ISE assigns a Security Group Tag (SGT), which is then carried through the network for segmentation and policy enforcement.

Why the other options are incorrect

A. DHCP: While DHCP is used for IP address assignment, it is not a standard method for assigning SGTs in the TrustSec framework.

C. policy routing: Policy-based routing (PBR) is used to make forwarding decisions based on specific criteria (like source IP), not for initial user-to-SGT classification.

D. modular QoS: Modular QoS CLI (MQC) is used for traffic marking and queuing to manage bandwidth and priority, not for identity-based security tagging.

QUESTION 8

An engineer must protect their company against ransomware attacks. Which solution allows the engineer to block the execution stage and prevent file encryption?

A. Use Cisco Firepower and block traffic to TOR networks.

B. Use Cisco Firepower with Intrusion Policy and snort rules blocking SMB exploitation.

C. Use Cisco AMP deployment with the Exploit Prevention engine enabled.

D. Use Cisco AMP deployment with the Malicious Activity Protection engine enabled.

Answer: D

- The MAP engine in Cisco AMP (Cisco Secure Endpoint) is specifically designed to stop ransomware at the execution stage.
- It continuously monitors endpoint processes for suspicious behaviors such as:
  - Rapid encryption of multiple files
  - Mass renaming or deletion of files
  - Other ransomware-like activity
- When detected, MAP immediately terminates the malicious process, preventing file encryption from spreading across the system.

Why the other options are less effective

A. Exploit Prevention engine → Focuses on blocking memory-based exploits (buffer overflows, code injection). It prevents initial compromise but does not directly stop ransomware encryption.

B. Firepower blocking TOR networks → Helps reduce command-and-control traffic but cannot stop ransomware already executing locally.

C. Firepower Intrusion Policy with SMB rules → Prevents lateral movement (e.g., WannaCry spreading via SMB) but does not block encryption once a host is infected.

QUESTION 9

```
R1
ntp peer 192.168.100.1
ntp source interface GigabitEthernet0/0
interface GigabitEthernet 0/1
ip address 192.168.1.1 255.255.255.248
ntp broadcast
```

Refer to the exhibit. What is the result of the configuration?

A. It configures NTP in client mode.

B. It configures NTP in symmetric active mode.

C. This router will work as an authoritative NTP server.

D. This router is allowed to send and receive NTP broadcast packets.

Answer: B

- The command **`ntp peer 192.168.100.1`** puts the router into symmetric active mode with the specified peer.
  - In this mode, both devices can act as **server or client**, synchronizing time with each other.
  - This is different from client mode (**`ntp server`**) or authoritative server mode (**`ntp master`**).
- The additional command **`ntp broadcast`** enables the router to send NTP broadcast packets on the interface, but by itself that does not mean the router is automatically receiving them. To receive, you'd need **`ntp broadcast client`**.

Why the other options are incorrect

A. Client mode → Requires **`ntp server <IP>`** configuration, not ntp peer.

B. Authoritative NTP server → Requires **`ntp master`**, which is not present.

D. Send/receive broadcast packets → The config only enables sending broadcasts. Receiving requires **`ntp broadcast client.`**

QUESTION 10

```
hostname CPE
!
ip dhcp excluded-address 192.168.10.0 192.168.10.10
!
ip dhcp pool LAN
 network 192.168.10.0 255.255.255.0
 default-router 192.168.10.1
 dns-server 192.0.2.254
!
interface Loopback0
 ip address 192.168.255.1 255.255.255.255
!
interface GigabitEthernet0/1
 description => LAN <=
 ip address 192.168.10.1 255.255.255.0
!
--
CPE# debug ip dhcp server packet
DHCP server packet debugging is on.
CPE#
*Sep 11 11:00:12.820: DHCPD: inconsistent relay information.
*Sep 11 11:00:12.820: DHCPD: relay information option exists, but giaddr is zero.
CPE#
```

Refer to the exhibit. The CPE router acts as a DHCP server for the locally attached LAN. After DHCP snooping is enabled on the switch where the DHCP clients are connected, clients are unable to obtain their configuration from the DHCP server. What is the cause of this issue?

A. The DHCP server drops DHCP packets carrying Option 82 and an empty relay agent IP address.
B. The excluded DHCP range contains the subnet address of the entire LAN network.
C. The configuration of Gi0/1 is missing the **`ip helper-address 192.168.255.1`** command.
D. The IP address of the DHCP server is in the excluded DHCP range.

Answer: A

- When DHCP snooping is enabled on a switch, the switch inserts Option 82 (relay agent information) into DHCP requests.
- In this case, the debug output shows:
  `DHCPD: inconsistent relay information.`
  `DHCPD: relay information option exists, but giaddr is zero.`
  This means the DHCP server received packets with Option 82 present but with a **zero giaddr (gateway IP address)**, which is invalid.
- As a result, the DHCP server **drops the packets**, preventing clients from obtaining their IP configuration.

Why the other options are incorrect:

- B. The excluded range is 192.168.10.0 – 192.168.10.10. This only excludes the first 11 addresses, not the entire subnet. The rest of the pool is still valid.
- C. An ip helper-address is required when the DHCP server is remote. Here, the DHCP server is local on the CPE router, so no helper address is needed.
- D. The DHCP server address is 192.168.10.1. The excluded range is 192.168.10.0 – 192.168.10.10, which does include .1. However, excluding the server's own IP is normal practice to prevent it from being assigned to clients. This does not cause the issue.

QUESTION 11 ★★

Why would an engineer use policy-based routing?

A. to use multicast for video traffic transmission

B. to route video traffic using the shortest path to the destination

C. to route video traffic through a dedicated link

D. to perform traffic shaping to avoid video traffic delay

Answer: C

- Policy-Based Routing (PBR) allows engineers to override the normal routing table decisions and forward traffic based on policies (like source, destination, protocol, or application type).
- In this case, PBR can be used to force video traffic through a dedicated link, ensuring predictable performance and separation from other traffic.

Why the other options are incorrect

- Use multicast for video traffic transmission → Multicast is a separate technology for efficient one-to-many delivery, not related to PBR.
- Route video traffic using the shortest path to the destination → That's the job of standard routing protocols (OSPF, EIGRP, etc.), not PBR.
- Perform traffic shaping to avoid video traffic delay → Traffic shaping is a QoS function, not PBR.

QUESTION 12 ★★★

Which feature allows HSRP to failover from the active route processor to the standby route processor without loss of data or path change?

A. HSRP tracking

B. preemption

C. stateful switchover

D. IP SLA tracking

Answer: C

- **Stateful Switchover (SSO)** allows HSRP to failover from the active route processor to the standby route processor **without losing data or changing paths.**
- It synchronizes the state information between processors so that the standby can immediately take over with minimal disruption.
- This ensures continuous packet forwarding and session preservation during a failover event.

Why the other options are incorrect

A. HSRP tracking → Monitors interface or object states to influence HSRP priority, but does not guarantee seamless failover.

B. Preemption → Allows a higher-priority router to take over as active, but can cause brief disruption.

D. IP SLA tracking → Monitors reachability and performance, used with HSRP for decision-making, but not for seamless failover.

QUESTION 13

```
R1#show running-config interface fastEthernet 0/0
interface FastEthernet0/0
 ip address 192.168.1.251 255.255.255.0
 duplex auto
 speed auto
 vrrp 1 ip 192.168.1.254
 vrrp 1 timers advertise 3
 vrrp 1 priority 120
end

R2#show running-config interface fastEthernet 0/0
interface FastEthernet0/0
 ip address 192.168.1.252 255.255.255.0
 duplex auto
 speed auto
 vrrp 1 ip 192.168.1.254
 vrrp 1 priority 110
end

R3#show running-config interface fastEthernet 0/0
interface FastEthernet0/0
 ip address 192.168.1.253 255.255.255.0
 duplex auto
 speed auto
 vrrp 1 ip 192.168.1.254
 vrrp 1 priority 100
end
```

Refer to the exhibit. An engineer must configure the backup routers to achieve these results:

- R1 is the primary, and R2 and R3 are backup routers.
- R1 has a primary advertisement interval of 3 seconds.

Which command must be added to the backup routers to complete the configuration?

A. **`vrrp 1 timers learn`**

B. **`vrrp 1 timers advertise 3`**

C. **`vrrp 1 preempt`**

D. **`vrrp 1 preempt delay minimum 3`**

Answer: A

- R1 is configured with **`vrrp 1 timers advertise 3`**, meaning it sends VRRP advertisements every 3 seconds.
- Backup routers (R2 and R3) must **learn this interval from the primary** to stay synchronized.
- The command **`vrrp 1 timers learn`** ensures that backup routers adopt the advertisement interval set by the primary router.

Why the other options are incorrect

B. vrrp 1 timers advertise 3 → Backups don't advertise unless they become primary, so this is unnecessary.

C. vrrp 1 preempt → Enables preemption but doesn't address advertisement interval synchronization.

D. vrrp 1 preempt delay minimum 3 → Controls preemption delay, not advertisement interval learning.

QUESTION 14 ★★★

What is the data policy in a Cisco Catalyst SD-WAN deployment?

A. list of ordered statements that define node configurations and authentication used within the SD-WAN overlay

B. set of statements that defines how data is forwarded based on IP packet information and specific VPNs

C. detailed database mapping several kinds of addresses with their corresponding location

D. group of services tested to guarantee devices and links liveliness within the SD-WAN overlay

Answer: B

- In Cisco Catalyst SD-WAN, policies are divided into **control policies** (affecting routing information) and **data policies** (affecting traffic forwarding).

  Data Policy

  - A **set of ordered statements** that match traffic flows based on IP header fields (source/destination IP, ports, DSCP, protocol, VPN ID).
  - Defines **forwarding actions** such as permit, deny, or redirect.
  - Can be applied **centrally** (via vSmart controllers) or **locally** (on WAN edge routers).
  - Used to enforce segmentation, traffic filtering, and application-aware routing.

Why the other options are incorrect

A. This describes **control policies**, not data policies.

C. This refers to **mapping systems** (like LISP in SD-Access), not SD-WAN data policies.

D. This describes **Bidirectional Forwarding Detection (BFD)** or liveliness checks, not data policies.

QUESTION 15 ★★

Which plane transports traffic between sites across the SD-WAN fabric?

A. IPsec

B. Management

C. Control

D. Data

Answer: D

- In Cisco SD-WAN, the **data plane** is responsible for transporting actual user traffic between sites across the SD-WAN fabric.
- This traffic is typically secured using **IPsec tunnels,** but the plane itself is the **data plane.**

Why the other options are incorrect

- **IPsec:** A security protocol used to encrypt traffic, but it is not itself a "plane." It operates within the data plane to secure traffic.
- **Management Plane:** Provides centralized configuration, monitoring, and management via vManage.
- **Control Plane:** Handles routing, policy, and control information exchange between vSmart, vBond, and edge devices.

QUESTION 16

Which protocol is implemented to establish secure control plane adjacencies between Cisco Catalyst SD-WAN nodes?

A. TLS
B. IPsec
C. IKE
D. ESP

Answer: A

- In Cisco SD-WAN, the **control plane** (used for exchanging routing, policy, and management information between vSmart, vBond, and vEdge/Catalyst Edge devices) is secured using TLS.
- TLS provides **mutual authentication, encryption, and integrity protection** for control plane adjacencies.
- The data plane, on the other hand, is typically secured using **IPsec** tunnels for actual user traffic.

Why the other options are incorrect

B. IPsec: Used to secure the data plane traffic, not the control plane adjacencies.

C. IKE (Internet Key Exchange): A protocol used to establish IPsec security associations, but not directly used for control plane adjacencies in Cisco SD-WAN.

D. ESP (Encapsulating Security Payload): A component of IPsec that provides encryption and authentication for data traffic, again not for control plane adjacencies.

QUESTION 17

When a wired client connects to an edge switch in a Cisco SD-Access fabric, which component decides whether the client has access to the network?

A. control-plane node
B. Identity Services Engine
C. edge node
D. RADIUS server

Answer: B

- In Cisco SD-Access, **ISE** provides centralized policy and identity management.
- When a wired client connects to an **edge node (switch)**, the edge node forwards authentication requests to ISE.
- ISE decides whether the client is allowed access, based on policies (such as 802.1X, MAB, or profiling).

Why the other options are incorrect

A. control-plane node → handles endpoint ID-to-location mapping, not access decisions.

C. edge node → enforces the decision but does not make it.

D. RADIUS server → is the protocol mechanism, but in SD-Access, ISE is the actual policy decision point.

QUESTION 8

Drag and drop the characteristics from the left onto the routing protocols they describe on the right.

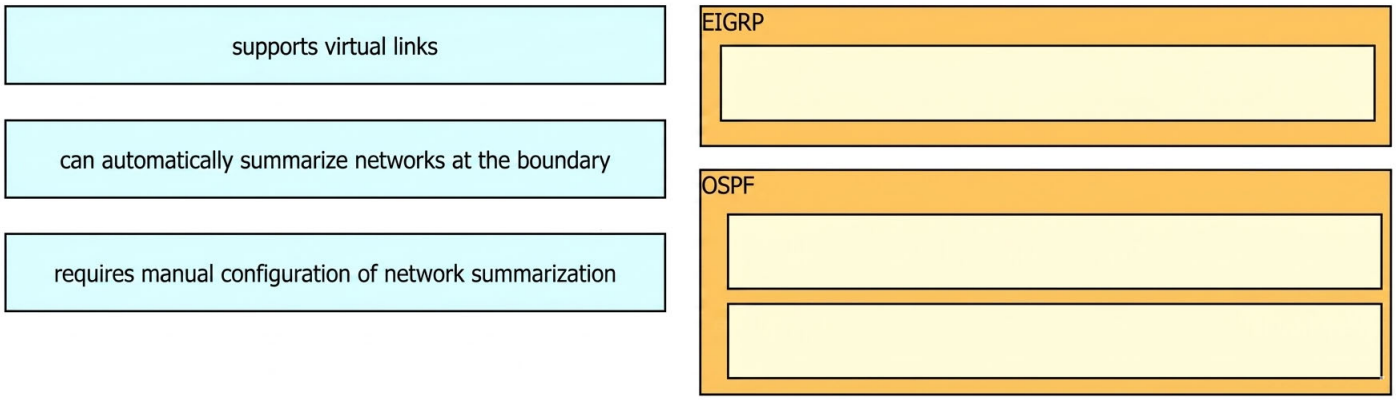

Answer:

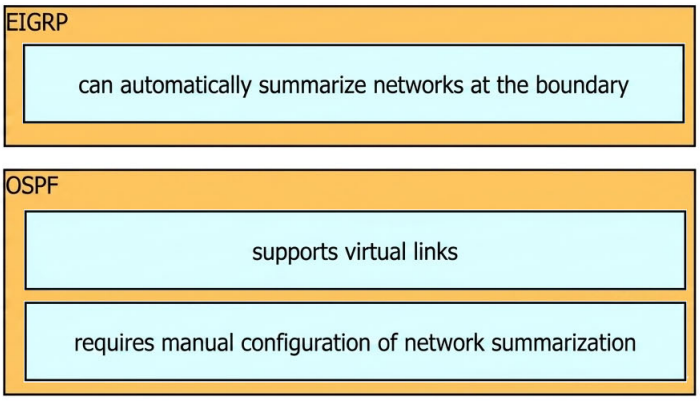

- **EIGRP** is a hybrid/advanced distance-vector protocol. By default (in legacy Cisco IOS versions), it **automatically summarizes routes** at major classful network boundaries (Class A, B, or C). If a router has subnets like 10.1.1.0/24 and 10.1.2.0/24, it advertises a single 10.0.0.0/8 route to its neighbors across a network boundary.

- **OSPF** demands a strict two-tier hierarchical design where all non-backbone areas (Area 1, Area 2, etc.) must physically connect to the backbone area (Area 0). If an area cannot physically connect to Area 0 due to network design limits or mergers, a logical tunnel called a **Virtual Link** is built through a transit area to restore back-end connectivity to Area 0. EIGRP operates on a flat system and has no concept of areas or virtual links.

- **OSPF** routers within the same area must maintain identical Link-State Databases (LSDB). Because of this strict consistency, **OSPF cannot automatically summarize routes**. If you want to condense OSPF routes to reduce routing table sizes, a network engineer must manually configure summarization explicitly at specific boundary routers:
  - ABRs (Area Border Routers) using the **`area range`** command.
  - ASBRs (Autonomous System Boundary Routers) using the **`summary-address`** command.

QUESTION 1

Drag and drop the characteristics from the left onto the routing protocols on the right.

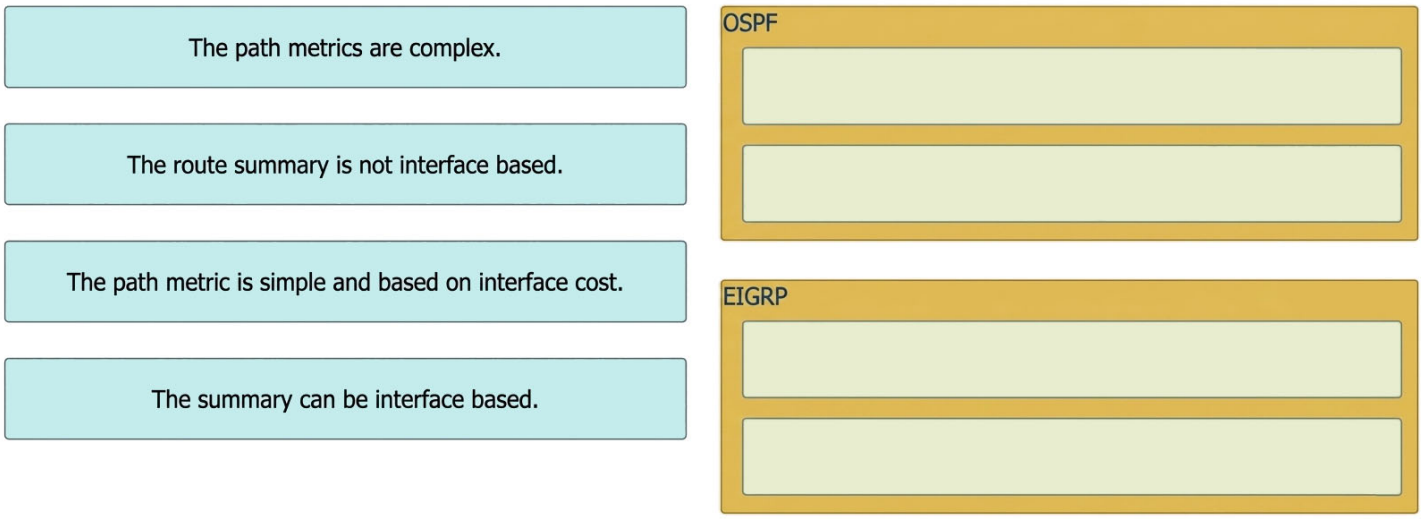

Answer:

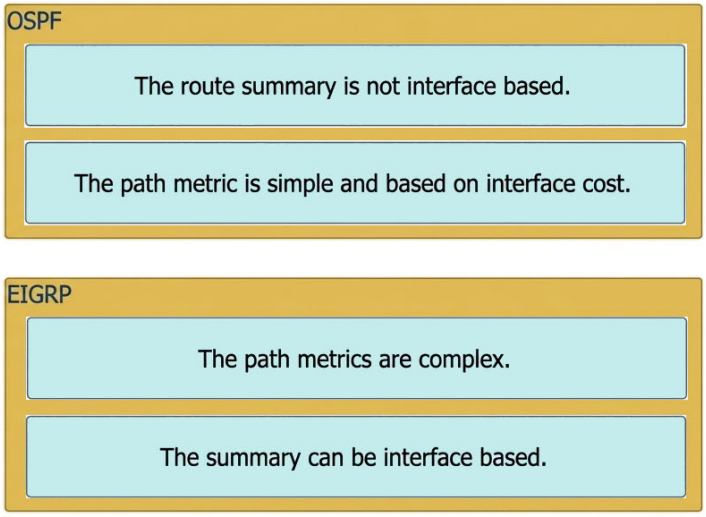

- **OSPF** uses a straightforward metric: cost, derived from interface bandwidth. Summarization is done at the area boundary, not per interface.
- **EIGRP** uses a composite metric (bandwidth, delay, reliability, load, MTU), making it more complex. It also supports manual summarization at the interface level, giving engineers flexibility.

QUESTION 2

Drag and drop the tools from the left onto the agent types on the right.

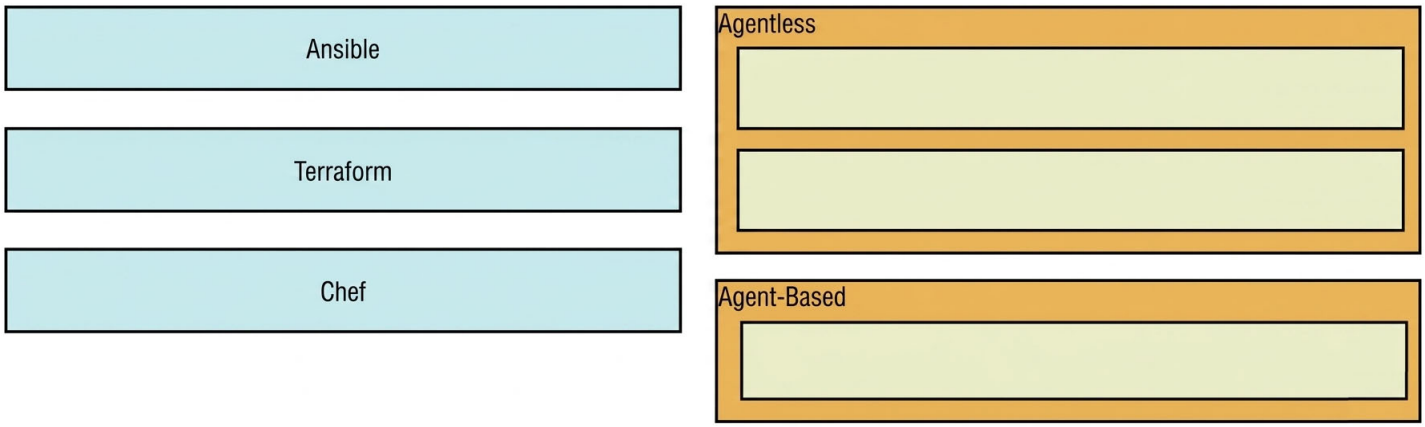

Answer:

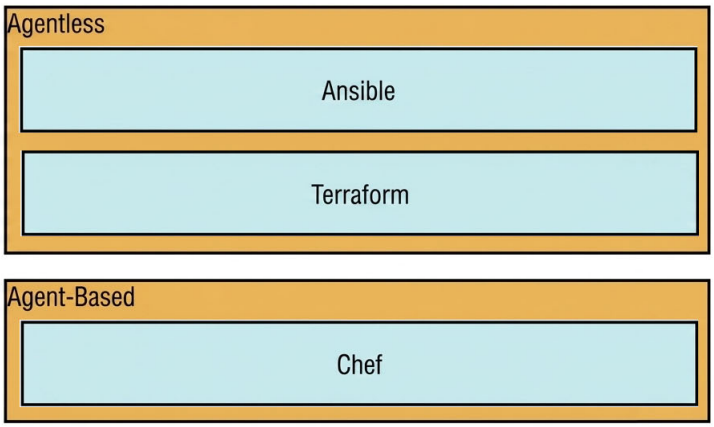

Classification of Tools

Agentless

- Ansible
- Terraform

These tools do not require agents installed on managed nodes. They typically use SSH or APIs to communicate directly with devices or cloud platforms.

Agent-Based

- Chef

Chef requires an agent (Chef client) installed on each managed node to pull configurations from the Chef server and apply them locally.

QUESTION 3

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

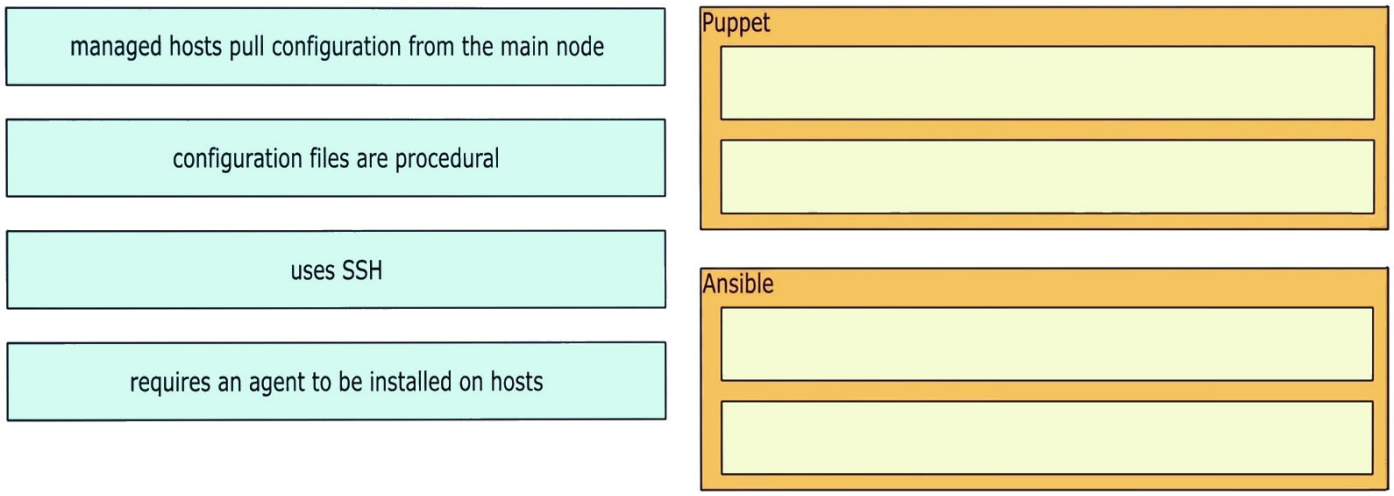

Answer:

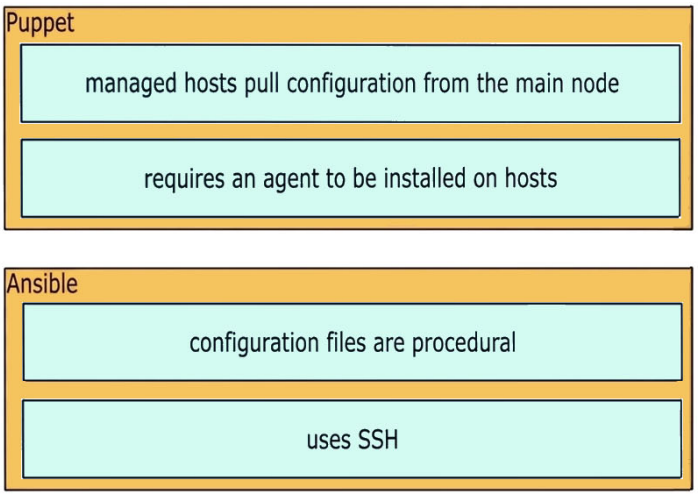

Simulation 1

Guidelines | Topology | Tasks

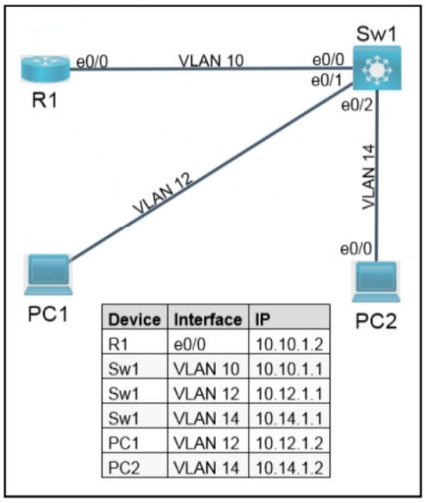

Device	Interface	IP
R1	e0/0	10.10.1.2
Sw1	VLAN 10	10.10.1.1
Sw1	VLAN 12	10.12.1.1
Sw1	VLAN 14	10.14.1.1
PC1	VLAN 12	10.12.1.2
PC2	VLAN 14	10.14.1.2

R1 | Sw1 | PC1 | PC2

```
R1#
```

Guidelines | Topology | Tasks

The Operations team started configuring several monitoring activities. Complete the configurations for the tasks below.

1. Enable Flexible NetFlow on R1 E0/0 in both directions using the pre-configured flow monitor.
2. Configure the switch port analyzer on Sw1 and mirror all VLAN 12 traffic to interface E1/3 using session number 12.
3. Configure a basic IP SLA ICMP echo operation on R1 to ping PC1 every 300 seconds.

R1 | Sw1 | PC1 | PC2

```
R1#
```

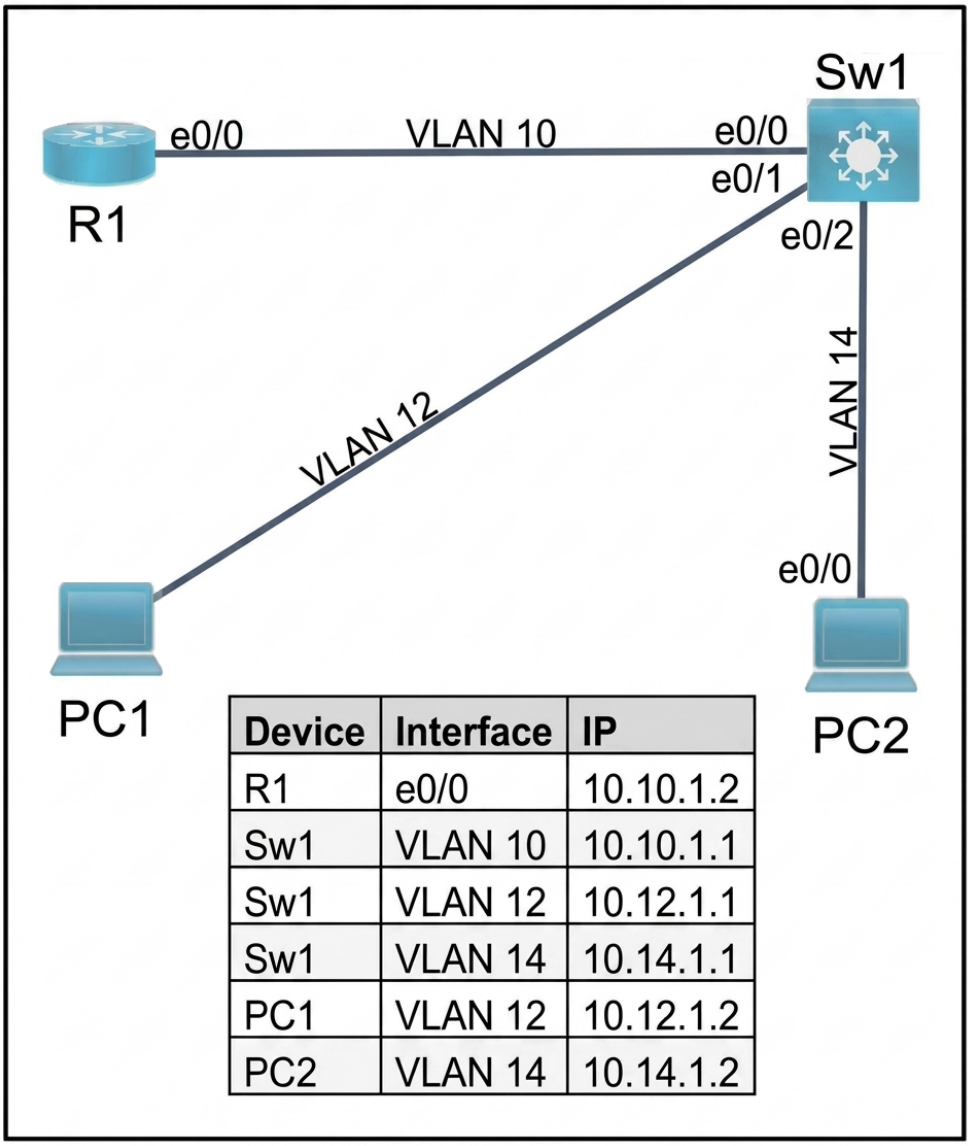

Device	Interface	IP
R1	e0/0	10.10.1.2
Sw1	VLAN 10	10.10.1.1
Sw1	VLAN 12	10.12.1.1
Sw1	VLAN 14	10.14.1.1
PC1	VLAN 12	10.12.1.2
PC2	VLAN 14	10.14.1.2

Task:

The Operations team started configuring several monitoring activities. Complete the configurations for the tasks below.

1. Enable Flexible NetFlow on R1 E0/0 in both directions using the pre-configured flow monitor.
2. Configure the switch port analyzer on Sw1 and mirror all VLAN 12 traffic to interface E1/3 using session number 12.
3. Configure a basic IP SLA ICMP echo operation on R1 to ping PC1 every 300 seconds.

Task Solutions:

Task 1：在 R1 E0/0 啟用雙向 Flexible NetFlow

```
R1# configure terminal
R1(config)# interface Ethernet0/0
R1(config-if)# ip flow monitor Monitor-R1Flow input
R1(config-if)# ip flow monitor Monitor-R1Flow output
R1(config-if)# end
```

Task 2：在 Sw1 上配置 SPAN 監控

將 VLAN 12 的流量鏡像到介面 E1/3，會話編號為 12：

```
Sw1# config t
Sw1(config)# monitor session 12 source vlan 12
Sw1(config)# monitor session 12 destination interface Ethernet1/3
Sw1(config)# end
```

Task 3：在 R1 上配置 IP SLA ICMP 回應操作

配置 R1 每 300 秒向 PC1（10.12.1.2）發送一次 Ping 偵測：

```
R1# config t
R1(config)# ip sla 1
R1(config-ip-sla)# icmp-echo 10.12.1.2
R1(config-ip-sla-echo)# frequency 300
R1(config-ip-sla-echo)# exit
R1(config)# ip sla schedule 1 life forever start-time now
R1(config)# end
```