Cisco CCNP - Quiz Questions with Answers

1.0 Architecture

1.

Which of the following uses Cisco's Locator/ID Separation Protocol (LISP)?

Control Plane

Data Plane

Policy Plane

Management Plane

Correct answer: Control Plane

Cisco's Software Defined (SD)-Access includes three different planes:

- ***Control Plane:** Uses the Locator/ID Separation Protocol (LISP), which uses a central Map Server (MS) to track remote destination data, enabling routers to only manage local routes and ask the MS for remote routes.*
- ***Data Plane:** Uses Virtual Extensible Local Area Network (VXLAN) to encapsulate traffic and perform tunneling while preserving the original Ethernet packet header. This enables the protocol to support overlays at Layers 2 and 3 and work on Internet Protocol (IP)-based networks that incorporate network segmentation and group-based policy.*
- ***Policy Plane:** Uses Cisco TrustSec Scalable Group Tags (SGTs) to encode information about groups, and these tags are used to apply corporate policies.*

Cisco SD-Access doesn't have a management plane.

2.

SSO is a protocol that allows synchronization between which of the following if a router contains multiple of them?

- **RP**
- RIB
- FIB
- FP
- PIR

Correct answer: RP

Stateful Switchover (SSO) enables a router that has multiple Router Processors (RPs) to synchronize the router configuration, Layer 2 protocol state information, and line card operation from the active to the standby RP.

An RP is responsible for the routing table/Routing Information Base (RIB), control plane, and, in centralized forwarding architectures, the adjacency table and Forwarding Information Base (FIB).

FP and PIR are fabricated terms.

3.

Which of the following are types of memory specific to Cisco devices? (Choose two.)

- **CAM**
- **TCAM**
- RAM
- TRAM

Content Addressable Memory (CAM) holds the Media Access Control (MAC) address table and uses specialized search techniques to enable addresses to be found faster than with Random Access Memory (RAM).

Ternary Content Addressable Memory (TCAM) is memory on a Cisco switch that allows multiple different fields to be used to evaluate a packet. It's used for Layer 2 and 3 searching and returns 0, 1, or X (don't care).

RAM is a common type of memory, and TRAM is a fabricated term.

4.

An organization is implementing a network that covers a single floor of one building. Which of the following layers of the three-tier network model are most likely to be consolidated into one? (Choose two.)

- **Distribution**
- **Core**
- Access
- Edge

The three-tier network design includes three layers:

- ***Access Layer:*** *Network edge where end-user devices (computers, IoT devices, mobile, etc.) are connected to the network.*
- ***Distribution Layer:*** *A distribution layer switch connects to access-layer switches from a building, floor, etc. This layer creates a boundary between the access and core layers, creating a boundary for the Spanning Tree Protocol (STP) and a summarization point for Internet Protocol (IP) routing information.*
- ***Core Layer:*** *Connects multiple distribution-layer switches together to support multi-site networks. This includes physically distributed sites or corporate networks with data centers, cloud deployments, etc.*

The organization will likely consolidate the distribution and core layers since there are not multiple sites to consolidate.

Edge is not a layer in the three-tier network model.

5.

Which of the following QoS mechanisms uses buffering in an attempt to avoid dropping excess traffic?

Shaping

Marking

Policing

Congestion avoidance

Correct answer: Shaping

Quality of Service (QoS) provides priority to certain types of network traffic, reducing latency, jitter, and packet loss for them. Common components of QoS include:

- ***Classification and Marking:*** *Dividing network traffic into classes based on its purpose and importance to the business. After traffic is classified, it is marked to allow QoS policies to be applied to it.*
- ***Policing:*** *Policing helps to enforce traffic rates by transmitting or remarking inbound or outbound traffic that complies with the rates and dropping or marking down traffic that exceeds it.*
- ***Shaping:*** *Shaping implements a buffer for egress traffic that holds traffic exceeding the traffic rate until the rate drops to the defined level. If traffic is below the desired traffic rate, then egress traffic isn't buffered.*
- ***Congestion Avoidance:*** *Congestion avoidance attempts to proactively prevent network congestion by strategically dropping packets.*

6.

A central Map Server (MS) is a critical component of which plane of Cisco SD-Access?

Control Plane

Data Plane

Policy Plane

Management Plane

Orchestration Plane

Correct answer: Control Plane

Cisco's Software-Defined Access (SD-Access) includes three different planes:

- ***Control Plane:** Uses the Locator/ID Separation Protocol (LISP), which uses a central Map Server (MS) to track remote destination data, enabling routers to only manage local routes and ask the MS for remote routes.*
- ***Data Plane:** Uses Virtual Extensible Local Area Network (VXLAN) to encapsulate traffic and perform tunneling while preserving the original Ethernet packet header. This enables the protocol to support overlays at Layers 2 and 3 and work on Internet Protocol (IP)-based networks that incorporate network segmentation and group-based policy.*
- ***Policy Plane:** Uses Cisco TrustSec Scalable Group Tags (SGTs) to encode information about groups, and these tags are used to apply corporate policies.*

Management and orchestration are not valid planes for SD-Access.

7.

Which of the following elements of QoS uses traffic classifications to re-prioritize or drop traffic exceeding defined rates?

Policing

Shaping

Marking

Congestion Management

Classification

Correct answer: Policing

Quality of Service (QoS) provides priority to certain types of network traffic, reducing latency, jitter, and packet loss for them. Common components of QoS include:

- ***Classification and Marking:** Dividing network traffic into classes based on its purpose and importance to the business. After traffic is classified, it is marked to allow QoS policies to be applied to it.*
- ***Policing:** Policing helps to enforce traffic rates by transmitting or remarking inbound or outbound traffic that complies with the rates and dropping or marking down traffic that exceeds it.*
- ***Shaping:** Shaping implements a buffer for egress traffic that holds traffic exceeding the traffic rate until the rate drops to the defined level. If traffic is below the desired traffic rate, then egress traffic isn't buffered.*
- ***Congestion Management:** Congestion management buffers excess traffic, and then removes packets from the queue via various algorithms.*
- ***Congestion Avoidance:** Congestion avoidance attempts to proactively prevent network congestion by proactively dropping packets.*

8.

Which of the following is a type of memory designed to store the MAC address table and enable rapid searches of it on a Cisco switch?

CAM

RIB

FIB

CEF

Correct answer: CAM

Content Addressable Memory (CAM) holds the Media Access Control (MAC) address table and uses specialized search techniques to enable addresses to be found faster than with Random Access Memory (RAM).

Cisco Express Forwarding (CEF) is a packet-switching protocol developed by Cisco and used by default on the majority of Cisco platforms.

The Routing Information Base (RIB) is Cisco's term for a routing table, which stores information on how to reach different devices or networks.

The Forwarding Information Base (FIB) (part of software CEF) stores the next-hop information for each network destination and is used to make Internet Protocol (IP) destination prefix-based decisions about how to route packets.

9.

Which of the following Cisco SD-WAN solutions work together to define and enforce policies on the SD-WAN network? (Choose two.)

vManage

vSmart

vBond

vAnalytics

Edge devices

Cisco's Software-Defined Wide Area Network (SD-WAN) solutions include several different components:

- *vManage is used for centralized visibility and management, including policy definitions.*
- *vSmart implements and enforces the policies created in vManage.*
- *vBond performs tasks such as authentication, Network Address Translation (NAT) detection, and load balancing.*
- *vAnalytics is an optional analytics service.*
- *SD-WAN edge devices make up the data plane and are responsible for forwarding traffic between locations via various media.*

10.

Which of the following deployment models are not an option with a Cisco lightweight AP? (Choose two.)

- **Autonomous**
- **Controllerless**
- Centralized
- Distributed
- Cloud-based

Autonomous deployments require standalone Access Points (APs), while most Cisco wireless APs are lightweight and require a Wireless Local Area Network (LAN) Controller (WLC) to operate. Some APs have an Embedded Wireless Controller (EWC) that allows them to be used in controllerless deployments.

Controller-based wireless deployments have a standalone WLC that can be deployed under the following models:

- ***Centralized:** The WLC is placed in a centralized location, such as the network core, enabling it to support many APs. WLC is likely near commonly-used resources (Internet, data center, etc.) and enables centralized policy enforcement.*
- ***Distributed:** Multiple WLCs are located alongside each switch in the access layer. This design makes sense for geographically distributed sites.*
- ***Cloud-Based:** Similar to the centralized model, except the WLC is located in a public or private cloud, rather than on-premises.*

11.

Cisco TrustSec Scalable Group Tags (SGTs) are used by which plane of Cisco's SD-Access?

Policy Plane

Control Plane

Data Plane

Management Plane

Orchestration Plane

Correct answer: Policy Plane

Cisco's Sofware-Defined Access (SD-Access) includes three different planes:

- ***Control Plane:** Uses the Locator/ID Separation Protocol (LISP), which uses a central Map Server (MS) to track remote destination data, enabling routers to only manage local routes and ask the MS for remote routes.*
- ***Data Plane:** Uses Virtual Extensible Local Area Network (VXLAN) to encapsulate traffic and perform tunneling while preserving the original Ethernet packet header. This enables the protocol to support overlays at Layers 2 and 3 and work on Internet Protocol (IP)-based networks that incorporate network segmentation and group-based policy.*
- ***Policy Plane:** Uses Cisco TrustSec Scalable Group Tags (SGTs) to encode information about groups, and these tags are used to apply corporate policies.*

Management and orchestration are not valid SD-Access planes.

12.

Which of the following BEST describes a protocol that offers load balancing across multiple different routers called Active Virtual Forwarders (AVFs)?

GLBP

FHRP

VRRP

HSRP

Correct answer: GLBP

First-Hop Redundancy Protocols (FHRP) help ensure network resiliency by creating a Virtual Internet Protocol (VIP) gateway linked to multiple physical gateways. If a gateway goes down, then the device's traffic will be sent via another gateway. The three main FHRPs include:

- ***Hot Standby Router Protocol (HSRP):*** *A protocol developed by Cisco that creates a virtual IP and Media Access Control (MAC) address usually held by the active router. If the active router fails, a standby router takes over these addresses and acts as the gateway.*
- ***Virtual Router Redundancy Protocol (VRRP):*** *Industry standard protocol that operates similarly to HSRP but names the routers "master" and "backup". This protocol allows preemption by default and uses a particular MAC address structure for the VIP gateway.*
- ***Gateway Load Balancing Protocol (GLBP):*** *Offers both redundancy and load balancing. The network has up to four Active Virtual Forwarders (AVFs) responsible for forwarding traffic for their assigned hosts and a single Active Virtual Gateway (AVG) that responds to Address Resolution Protocol (ARP) requests with the virtual MAC of the assigned AVF. Failure of the AVG or an AVF causes another system to take over its role.*

13.

Which of the following are benefits of SD-WAN? (Choose two.)

Centralize device configuration and management

Provide seamless connectivity to public cloud

Optimize user experience for on-prem applications

Ensure optimal usage of a single transport medium

Software-Defined Wide Area Network (SD-WAN) is designed to offer centralized device configuration and management, and simplifies expansion to the public cloud.

It optimizes the user experience for cloud applications and the usage of multiple different transport media.

14.

Which of the following are planes defined within Cisco's SD-WAN solution? (Choose two.)

- **Control**
- **Data**
- **Management**
- Policy

Cisco’s Software-Defined Wide Area Network (SD-WAN) solution implements various planes, including the control and data planes. The control plane is implemented using vSmart, which is responsible for defining the network topology and advertising routes and data policies to SD-WAN edge devices. These SD-WAN edge devices make up the data plane and are responsible for forwarding traffic between locations via various media.

Cisco’s vManage Network Management System (NMS) and vBond orchestrator implement the management and orchestration planes respectively.

Policy isn't a plane in SD-WAN.

15.

Which of the following are planes defined in Cisco's SD-Access solution? (Choose three.)

- **Control**
- **Data**
- **Policy**
- Management
- Orchestration

SD-Access includes three different planes:

- ***Control Plane:*** *Uses the Locator/ID Separation Protocol (LISP), which uses a central Map Server (MS) to track remote destination data, enabling routers to only manage local routes and ask the MS for remote routes.*
- ***Data Plane:*** *Uses Virtual Extensible Local Area Network (VXLAN) to encapsulate traffic and perform tunneling while preserving the original Ethernet packet header. This enables the protocol to support overlays at Layers 2 and 3 and work on Internet Protocol (IP)-based networks that incorporate network segmentation and group-based policy.*
- ***Policy Plane:*** *Uses Cisco TrustSec Scalable Group Tags (SGTs) to encode information about groups, and these tags are used to apply corporate policies.*

Management and orchestration planes are part of Software-Defined Wide Area Network (SD-WAN), not SD-Access.

16.

Which of the following is a system that allows synchronization of router configuration and other data between redundant components on a router?

- **SSO**
- RP
- RIB
- FIB

Correct answer: SSO

Stateful Switchover (SSO) enables a router that has multiple Router Processors (RPs) to synchronize the router configuration, Layer 2 protocol state information, and line card operation from the active to the standby RP.

An RP is responsible for the routing table or Routing Information Base (RIB), control plane, and (in centralized forwarding architectures) the adjacency table and Forwarding Information Base (FIB).

17.

Which of the following Cisco wireless deployment models uses an EWC?

Controllerless

Centralized

Distributed

Cloud-Based

Autonomous

Correct answer: Controllerless

Cisco wireless Access Points (APs) are lightweight and require a Wireless Local Area Network (LAN) Controller (WLC) to operate, which it connects to via a pair of Control and Provisioning of Wireless Access Points (CAPWAP) tunnels.

Controllerless deployments integrate the WLC into an AP, which is called an Embedded Wireless Controller (EWC). These can be used in distributed deployments as well.

Controller-based wireless deployments have a standalone WLC that can be deployed under the following models:

- ***Centralized:** The WLC is placed in a centralized location, such as the network core, enabling it to support many APs. WLC is likely near commonly-used resources (Internet, data center, etc.) and enables centralized policy enforcement.*
- ***Distributed:** Multiple WLCs are located alongside each switch in the access layer. This design makes sense for geographically distributed sites.*
- ***Cloud-Based:** Similar to the centralized model, except the WLC is located in a public or private cloud, rather than on-premises.*

Autonomous deployments use standalone APs.

18.

Which of the following traffic policing algorithms states that all traffic exceeding the CIR is downgraded in priority?

- **Single-Rate, Two-Color**
- Single-Rate, Three-Color
- Two-Rate, Two-Color
- Two-Rate, Three Color

Correct answer: Single-Rate, Two-Color

Cisco supports three types of traffic policing algorithms:

- ***Single-Rate, Two-Color:*** *Traffic has a single Committed Information Rate (CIR), and traffic exceeding that rate (using up all tokens in the bucket) is downgraded in priority.*
- ***Single-Rate, Three-Color:*** *This is a two-bucket algorithm in which traffic for which there are no tokens available in the first bucket might use tokens in a second bucket designed to handle temporary bursts (which is filled using excess tokens from the first bucket). This traffic is usually marked down (but can be dropped), and any traffic beyond that is likely dropped (but can be marked down).*
- ***Two-Rate, Three-Color:*** *Introduces a Peak Information Rate (PIR) in addition to the CIR, which defines the rate at which tokens are added to the second bucket. Otherwise operates similarly to Single-Rate Three-Color.*

Two-rate, two-color is not a traffic policing algorithm.

19.

An organization is designing a wireless network for a company that has several different sites. Which of the following is the best deployment model for this approach if it's using lightweight APs (no EWC)?

Distributed

Centralized

Cloud-based

Controllerless

Autonomous

Correct answer: Distributed

Cisco wireless Access Points (APs) are lightweight and require a Wireless Local Area Network (LAN) Controller (WLC) to operate, which it connects to via a pair of Control And Provisioning of Wireless Access Points (CAPWAP) tunnels. Controller-based wireless deployments have a standalone WLC that can be deployed under the following models:

- ***Centralized:** The WLC is placed in a centralized location, such as the network core, enabling it to support many APs. WLC is likely near commonly-used resources (Internet, data center, etc.) and enables centralized policy enforcement.*
- ***Distributed:** Multiple WLCs are located alongside each switch in the access layer. This design makes sense for geographically-distributed sites.*
- ***Cloud-Based:** Similar to the centralized model, but WLC is located in a public or private cloud, rather than on-premises.*

Controllerless deployments integrate the WLC into an AP, which is called an Embedded Wireless Controller (EWC). These can be used in distributed deployments as well.

Autonomous deployments use standalone APs, not lightweight ones.

20.

Which of the following is an industry-standard alternative to a Cisco protocol that offers preemption by default and has a defined structure for VIP gateway MAC addresses?

VRRP

HSRP

FHRP

GLBP

NHRP

Correct answer: VRRP

First-Hop Redundancy Protocols (FHRP) help ensure network resiliency by creating a Virtual Internet Protocol (VIP) gateway linked to multiple physical gateways. If a gateway goes down, then the devices's traffic will be sent via another gateway. The three main FHRPs include:

- ***Hot Standby Router Protocol (HSRP):*** *Protocol developed by Cisco that creates a virtual IP and Media Access Control (MAC) address usually held by the active router. If the active router fails, a standby router takes over these addresses and acts as the gateway.*
- ***Virtual Router Redundancy Protocol (VRRP):*** *Industry standard protocol that operates similarly to HSRP but names the routers "master" and "backup". This protocol allows preemption by default and uses a particular MAC address structure for the VIP gateway.*
- ***Gateway Load Balancing Protocol (GLBP):*** *Offers both redundancy and load balancing. The network has up to four Active Virtual Forwarders (AVFs) responsible for forwarding traffic for their assigned hosts and a single Active Virtual Gateway (AVG) that responds to Address Resolution Protocol (ARP) requests with the virtual MAC of the assigned AVF. Failure of the AVG or an AVF causes another system to take over its role.*

NHRP is a fabricated term.

21.

Which of the following network models are most dependent on network virtualization technology? (Choose two.)

- **Fabric**
- **Cloud**
- Two-tier
- Three-tier

Fabric and cloud environments are implemented using virtual overlays on physical networks. This implements network infrastructure in software rather than in hardware.

Two-tier and three-tier networks are usually built with physical appliances.

22.

Which of the following methods of locating devices can't be used for devices outside of the organization's control but are connected to the network?

RFID tags

Interference detection

Measuring Wi-Fi probe requests

Measuring RSS

Signal triangulation

Correct answer: RFID Tags

Location services track device locations within a wireless network, which is useful for asset tracking. Some methods for implementing it include:

- ***RFID Tags:** Radio-Frequency Identifier (RFID) tags attached to devices can be used to map devices' locations as they attach to Wi-Fi or send out probe requests. If an organization doesn't control a device, it can't put an RFID tag on it.*
- ***Measuring Received Signal Strength:** If a client is in an open space, measuring Received Signal Strength (RSS) from different Access Points (APs) can help to triangulate its location. Cisco devices use a Radio Frequency (RF) calibration template that is generated using an RF scanner and accurately measures attenuation and signal propagation within a space, enabling it to be used in other places than an empty room.*
- ***Probe Requests:** Devices with Wi-Fi active will send out probe requests to APs on all supported wireless channels, enabling triangulation.*
- ***Interference Detection:** Devices that create signal interference, such as cordless phones and wireless cameras, can be located via spectrum analysis and identifying locations experiencing interference.*

Signal triangulation is part of measuring RSS or probe requests, making it a viable option.

23.

Which of the following is a superset of the others (meaning it is a general term where the others are specific protocols)?

FHRP

VRRP

HSRP

GLBP

NHRP

Correct answer: FHRP

First-Hop Redundancy Protocols (FHRP) help ensure network resiliency by creating a Virtual Internet Protocol (VIP) gateway linked to multiple physical gateways. If a gateway goes down, then the devices's traffic will be sent via another gateway. The three main FHRPs include:

- ***Hot Standby Router Protocol (HSRP):*** *Protocol developed by Cisco that creates a virtual IP and Media Access Control (MAC) address usually held by the active router. If the active router fails, a standby router takes over these addresses and acts as the gateway.*
- ***Virtual Router Redundancy Protocol (VRRP):*** *Industry standard protocol that operates similarly to HSRP but names the routers "master" and "backup". This protocol allows preemption by default and uses a particular MAC address structure for the VIP gateway.*
- ***Gateway Load Balancing Protocol (GLBP):*** *Offers both redundancy and load balancing. The network has up to four Active Virtual Forwarders (AVFs) responsible for forwarding traffic for their assigned hosts and a single Active Virtual Gateway (AVG) that responds to Address Resolution Protocol (ARP) requests with the virtual MAC of the assigned AVF. Failure of the AVG or an AVF causes another system to take over its role.*

NHRP is a fabricated term.

24.

Which of the following are potential search results for TCAM? (Choose three.)

- **0**
- **1**
- **X**
- N/A

Ternary Content Addressable Memory (TCAM) is memory on a Cisco switch that allows multiple different fields to be used to evaluate a packet. It's used for Layer 2 and 3 searching and returns 0, 1, or X (don't care).

25.

Which of the following involves buffering excess traffic which can help prevent it from being dropped? (Choose two.)

- **Shaping**
- **Congestion Mitigation**
- Congestion Avoidance
- Policing
- Marking

Quality of Service (QoS) provides priority to certain types of network traffic, reducing latency, jitter, and packet loss for them. Common components of QoS include:

- ***Classification and Marking:*** *Dividing network traffic into classes based on its purpose and importance to the business. After traffic is classified, it is marked to allow QoS policies to be applied to it.*
- ***Policing:*** *Policing helps to enforce traffic rates by transmitting or remarking inbound or outbound traffic that complies with the rates and dropping or marking down traffic that exceeds it.*
- ***Shaping:*** *Shaping implements a buffer for egress traffic that holds traffic exceeding the traffic rate until the rate drops to the defined level. If traffic is below the desired traffic rate, then egress traffic isn't buffered.*
- ***Congestion Management:*** *Congestion management buffers excess traffic, and then removes packets from the queue via various algorithms.*
- ***Congestion Avoidance:*** *Congestion avoidance attempts to proactively prevent network congestion by selectively dropping packets.*

26.

Where on a Cisco switch can you find the MAC addresses of each directly connected next-hop IP address?

AIB

RIB

FIB

CAM

TCAM

Correct answer: AIB

An adjacency table or the Adjacency Information Base (AIB) contains the next-hop Media Access Control (MAC) addresses for each directly connected next-hop Internet Protocol (IP) address as well as the MAC address of the egress interface. It's populated by the Address Resolution Protocol (ARP) table or similar sources.

The Forwarding Information Base (FIB), part of the software Cisco Express Forwarding (CEF), stores the next-hop information for each network destination and is used to make IP destination prefix-based decisions about how to route packets.

Content Addressable Memory (CAM) holds the MAC address table and uses specialized search techniques to enable addresses to be found faster than with Random Access Memory (RAM).

The Routing Information Base (RIB) is Cisco's term for a routing table, which stores information on how to reach different devices or networks.

Ternary Content Addressable Memory (TCAM) is memory on a Cisco switch that allows multiple different fields to be used to evaluate a packet. It's used for Layer 2 and 3 searching and returns 0, 1, or X (don't care).

27.

Which of the following algorithms has tokens in a bucket that are allocated for handling bursty traffic? (Choose two.)

Two-Rate, Three Color

Single-Rate, Three Color

Single-Rate, Two Color

Two-Rate, Two Color

Cisco supports three types of traffic policing algorithms:

- ***Single-Rate, Two-Color:*** *Traffic has a single Committed Information Rate (CIR), and traffic exceeding that rate (using up all tokens in the bucket) is downgraded in priority.*
- ***Single-Rate, Three Color:*** *This is a two-bucket algorithm in which traffic for which there are no tokens available in the first bucket might use tokens in a second bucket designed to handle temporary bursts (which is filled using excess tokens from the first bucket). This traffic is usually marked down (but can be dropped), and any traffic beyond that is likely dropped (but can be marked down).*
- ***Two-Rate, Three-Color:*** *Introduces a Peak Information Rate (PIR) in addition to the CIR, which defines the rate at which tokens are added to the second bucket. Otherwise operates similarly to Single-Rate, Three-Color.*

Two-rate, two-color is not a traffic policing algorithm.

28.

In the three-tier network model, which of the following layers is MOST likely to contain a wireless AP?

Access

Distribution

Core

Edge

Correct answer: Access

The three-tier network design includes three layers:

- ***Access Layer:*** *Network edge where end-user devices (computers, Internet of Things (IoT) devices, mobile, etc.) are connected to the network. This is the layer where wireless Access Points (APs) would be deployed.*
- ***Distribution Layer:*** *A distribution layer switch connects to access-layer switches from a building, floor, etc. This layer creates a boundary between the access and core layers, creating a boundary for the Spanning Tree Protocol (STP) and a summarization point for Internet Protocol (IP) routing information.*
- ***Core Layer:*** *Connects multiple distribution-layer switches together to support multi-site networks. This includes physically distributed sites or corporate networks with data centers, cloud deployments, etc.*

Edge is not a layer of the three-tier network model.

29.

Which of the following is a Cisco-specific congestion management protocol?

CQ

WRR

FIFO

CBWFQ

Correct answer: CQ

Congestion management buffers excess traffic and then removes packets from the queue via various algorithms. Some legacy examples include:

- ***First-In First-Out (FIFO):** The system operates a single queue, and the oldest packet in the queue is removed first.*
- ***Weighted Round Robin (WRR):** Adds support for prioritization to round robin where each queue is assigned bandwidth based on its weight.*
- ***Custom Queueing (CQ):** Cisco WRR protocol with 16 queues and FIFO ordering within a queue.*
- ***Class-Based Weighted Fair Queueing (CBWFQ):** Modern algorithm where up to 256 queues are created (one per traffic class) that are serviced based on the class's assigned bandwidth. Uses traffic descriptors to classify traffic, and classified traffic is assigned bandwidth, weight, maximum packet limit, and queue limit. Packets exceeding the queue limit are dropped.*

30.

Which of the following is not one of the Management Layer Components of Cisco SD-Access?

Administration

Design

Policy

Provision

Correct answer: Administration

Administration is not a part of Cisco SD-Access at any layer.

The Management Layer consists of the Cisco DNA Center GUI, Base Automation, Design, Policy, Provision, and Assurance.

31.

Which of the following parts of a Cisco switch contains information used to make decisions on where to send packets based on the destination IP's prefix?

FIB

AIB

RIB

TCAM

CAM

Correct answer: FIB

The Forwarding Information Base (FIB) stores the next-hop information for each network destination and is used to make Internet Protocol (IP) destination prefix-based decisions about how to route packets.

The Routing Information Base (RIB) is Cisco's term for a routing table, which stores information on how to reach different devices or networks.

Ternary Content Addressable Memory (TCAM) is memory on a Cisco switch that allows multiple different fields to be used to evaluate a packet. It's used for Layer 2 and 3 searching and returns 0, 1, or X (don't care).

An adjacency table or the Adjacency Information Base (AIB), which is part of software Cisco Express Forwarding (CEF), contains the next-hop Media Access Control (MAC) addresses for each directly connected next-hop IP address as well as the MAC address of the egress interface. It's populated by the Address Resolution Protocol (ARP) table or similar sources.

Content Addressable Memory (CAM) holds the MAC address table and uses specialized search techniques to enable addresses to be found faster than with Random Access Memory (RAM).

32.

Which of the following is true of the client density for a particular wireless AP?

It increases as its antenna's service area increases

It increases as the number of APs increases

If it increases, then performance increases

It increases as client activity increases

Correct answer: It increases as its antenna's service area increases

Client density is the number of devices connected to an Access Point (AP). Organizations can limit client density by deploying more APs and choosing antennas with a smaller coverage area.

With more clients and more active clients, performance degrades.

33.

A campus fabric solution must be managed via which of the following to be considered SD-Access? (Choose one.)

Cisco DNA Center

APIs

CLI

NETCONF

A campus fabric overlay solution must be managed by Cisco DNA Center to be considered Software Defined (SD)-Access.

Solutions managed by Command-Line Interface (CLI) or Application Programming Interfaces (APIs) via Network Configuration (NETCONF) protocol are considered a campus fabric solution but not "SD-Access."

34.

In which of the following might you expect to find MAC addresses stored? (Choose two).

- **AIB**
- **CAM**
- CEF
- RIB
- FIB

An adjacency table or the Adjacency Information Base (AIB) contains the next-hop Media Access Control (MAC) addresses for each directly connected next-hop Internet Protocol (IP) address as well as the MAC address of the egress interface. It's populated by the Address Resolution Protocol (ARP) or similar sources.

Content Addressable Memory (CAM) holds the MAC address table and uses specialized search techniques to enable addresses to be found faster than with Random Access Memory (RAM).

Cisco Express Forwarding (CEF) is a packet-switching protocol developed by Cisco and used by default on the majority of Cisco platforms.

The Forwarding Information Base (FIB) (part of software CEF) stores the next-hop information for each network destination and is used to make IP destination prefix-based decisions about how to route packets.

The Routing Information Base (RIB) is Cisco's term for a routing table, which stores information on how to reach different devices or networks.

35.

Which of the following Cisco solutions uses the Overlay Management Protocol (OMP) to learn and disseminate routing information?

vSmart

vBond

vAnalytics

vManage

Correct answer: vSmart

Cisco's Software-Defined Wide Area Network (SD-WAN) solution implements four planes:

- ***Control:** The control plane is implemented using vSmart, which is responsible for defining the network topology and uses the Overlay Management Protocol (OMP) to advertise routes to SD-WAN edge devices.*
- ***Management:** Cisco's vManage Network Management System (NMS) implements the management plane, providing single-pane-of-glass visibility and control.*
- ***Orchestration:** vBond implements the orchestration plane and performs tasks such as authentication, Network Address Translation (NAT) detection, and load balancing.*
- ***Data:** These SD-WAN edge devices make up the data plane and are responsible for forwarding traffic between locations via various media.*

vAnalytics is an optional analytics service.

36.

Which of the following components of Cisco's SD-WAN solution performs authentication, NAT detection, and load balancing?

vBond

vManage

vSmart

vAnalytics

Edge devices

Correct answer: vBond

Cisco’s Software-Defined Wide Area Network (SD-WAN) solution implements four planes:

- ***Orchestration:*** *vBond implements the orchestration plane and performs tasks such as authentication, Network Address Translation (NAT) detection, and load balancing.*
- ***Control:*** *The control plane is implemented using vSmart, which is responsible for defining the network topology and advertising routes and data policies to SD-WAN edge devices.*
- ***Data:*** *These SD-WAN edge devices make up the data plane and are responsible for forwarding traffic between locations via various media.*
- ***Management:*** *Cisco’s vManage Network Management System (NMS) implements the management plane, providing single-pane-of-glass visibility and control.*

vAnalytics is an optional analytics service.

37.

Which of the following device roles in Cisco's SD-Access connect wired or wireless endpoints to the SD-Access fabric? (Choose two.)

- **Fabric edge node**
- **Fabric WLC**
- Control plane node
- Intermediate node
- Fabric border node

Fabric edge nodes and fabric Wireless LAN Controllers (WLCs) connect wired and wireless endpoints to the Software Defined (SD)-Access fabric, respectively.

The control plane node contains settings, protocols, and mapping tables for the fabric overlay.

Fabric border nodes connect the SD-Access fabric to external Layer 3 networks.

Intermediate nodes have no SD-Access function beyond underlay services.

38.

Network segmentation is easiest to implement in which of the following network designs?

Fabric

Three-tier

Two-tier

Hierarchical

Data center

Correct answer: Fabric

Fabric networks create a virtual overlay on top of physical network architecture. This makes it easier to implement network segmentation or redesign the network architecture without a physical redesign.

The three-tier network design includes three layers:

- ***Access Layer:** Network edge where end-user devices (computers, IoT devices, mobile, etc.) are connected to the network.*
- ***Distribution Layer:** A distribution layer switch connects to access-layer switches from a building, floor, etc. This layer creates a boundary between the access and core layers, creating a boundary for the Spanning Tree Protocol (STP) and a summarization point for Internet Protocol (IP) routing information.*
- ***Core Layer:** Connects multiple distribution-layer switches together to support multi-site networks. This includes physically distributed sites or corporate networks with data centers, cloud deployments, etc.*

A two-tier network combines the distribution and core layers into a single layer.

Data center and hierarchical are not network designs.

39.

At which layer of SD-Access is traffic encapsulated and tunneled over the network?

Data Plane

Control Plane

Policy Plane

Orchestration Plane

Correct answer: Data Plane

Software Defined (SD)-Access includes three different planes:

- ***Control Plane:*** *Uses the Locator/ID Separation Protocol (LISP), which uses a central Map Server (MS) to track remote destination data, enabling routers to only manage local routes and ask the MS for remote routes.*
- ***Data Plane:*** *Uses Virtual Extensible Local Area Network (VXLAN) to encapsulate traffic and perform tunneling while preserving the original Ethernet packet header. This enables the protocol to support overlays at Layers 2 and 3 and work on Internet Protocol (IP)-based networks that incorporate network segmentation and group-based policy.*
- ***Policy Plane:*** *Uses Cisco TrustSec Scalable Group Tags (SGTs) to encode information about groups, and these tags are used to apply corporate policies.*

SD-Access doesn't have an orchestration plane.

40.

Virtual Extensible Local Area Network (VXLAN) implements which plane of Cisco SD-Access?

- **Data plane**
- Control plane
- Management plane
- Policy plane
- Network plane

Correct answer: Data plane

Virtual Extensible Local Area Network (VXLAN) implements the data plane in Cisco Software Defined (SD)-Access. The other two planes are the control and policy planes.

SD-Access doesn't include a management or network plane.

41.

Which of the following are layers in the three-tier network model? (Choose three.)

- **Access**
- **Distribution**
- **Core**
- Edge
- Management

The three-tier network design includes three layers:

- ***Access Layer:*** *Network edge where end-user devices (computers, Internet of Things (IoT) devices, mobile, etc.) are connected to the network.*
- ***Distribution Layer:*** *A distribution layer switch connects to access-layer switches from a building, floor, etc. This layer creates a boundary between the access and core layers, creating a boundary for the Spanning Tree Protocol (STP) and a summarization point for Internet Protocol (IP) routing information.*
- ***Core Layer:*** *Connects multiple distribution-layer switches together to support multi-site networks. This includes physically distributed sites or corporate networks with data centers, cloud deployments, etc.*

Edge and management are not part of the three-tier model.

42.

Which of the following traffic policing algorithms provide the option to downgrade or drop excess traffic? (Choose two.)

- **Single-Rate, Three Color**
- **Two-Rate, Three Color**
- Single-Rate, Two Color
- Two-Rate, Two Color
- Single-Rate, Single Color

Cisco supports three types of traffic policing algorithms:

- ***Single-Rate, Two-Color:*** *Traffic has a single Committed Information Rate (CIR), and traffic exceeding that rate (using up all tokens in the bucket) is downgraded in priority.*
- ***Single-Rate, Three Color:*** *This is a two-bucket algorithm in which traffic for which there are no tokens available in the first bucket might use tokens in a second bucket designed to handle temporary bursts (which is filled using excess tokens from the first bucket). This traffic is usually marked down (but can be dropped), and any traffic beyond that is likely dropped (but can be marked down).*
- ***Two-Rate, Three-Color:*** *Introduces a Peak Information Rate (PIR) in addition to the CIR, which defines the rate at which tokens are added to the second bucket. Otherwise operates similarly to Single-Rate Three-Color.*

Two-rate, two color and single-rate, single color are not traffic policing algorithms.

43.

Which of the following traffic policing algorithms has both a CIR and PIR?

- **Two-Rate, Three-Color**
- Single-Rate, Two-Color
- Single-Rate, Three-Color
- Two-Rate, Two-Color
- Two-Rate, Single-Color

Correct answer: Two-Rate, Three-Color

Cisco supports three types of traffic policing algorithms:

- ***Single-Rate, Two-Color:*** *Traffic has a single Committed Information Rate (CIR), and traffic exceeding that rate (using up all tokens in the bucket) is downgraded in priority.*
- ***Single-Rate, Three-Color:*** *This is a two-bucket algorithm in which traffic for which there are no tokens available in the first bucket might use tokens in a second bucket designed to handle temporary bursts (which is filled using excess tokens from the first bucket). This traffic is usually marked down (but can be dropped), and any traffic beyond that is likely dropped (but can be marked down).*
- ***Two-Rate, Three-Color:*** *Introduces a Peak Information Rate (PIR) in addition to the CIR, which defines the rate at which tokens are added to the second bucket. Otherwise operates similarly to Single-Rate, Three-Color.*

Two-Rate, Two-Color and Two-Rate, Single-Color aren't traffic policing algorithms.

44.

Which of the following levels of QoS policy for WLANs is intended for voice traffic?

Platinum

Gold

Silver

Bronze

Correct answer: Platinum

Cisco offers four different Quality of Service (QoS) policies for Wireless Local Area Networks (WLANs), including:

- ***Platinum:** Platinum QoS is often intended for voice traffic and has an 802.1p tag of 5 and a DSCP value of 46 (EF).*
- ***Gold:** Gold QoS is for video traffic with an 802.1p tag of 4 and a DSCP value of 34 (AF41).*
- ***Silver:** Silver is the default and implements “best effort” handling with an 802.1p tag and DSCP value of 0.*
- ***Bronze:** Bronze implements background handling with an 802.1p tag of 1 and a DSCP value of 10 (AF11).*

45.

Where on a Cisco switch would you look for the MAC address of the egress interface?

AIB

RIB

FIB

CEF

Correct answer: AIB

An adjacency table or the Adjacency Information Base (AIB), which is part of the software Cisco Express Forwarding (CEF), contains the next-hop Media Access Control (MAC) addresses for each directly-connected next-hop Internet Protocol (IP) address as well as the MAC address of the egress interface. It's populated by the Address Resolution Protocol (ARP) table or similar sources.

The Forwarding Information Base (FIB) (part of software CEF) stores the next-hop information for each network destination and is used to make IP destination prefix-based decisions about how to route packets.

The Routing Information Base (RIB) is Cisco's term for a routing table, which stores information on how to reach different devices or networks.

CEF is a packet-switching protocol developed by Cisco and used by default on the majority of Cisco platforms.

46.

The Network Control Platform (NCP) and Network Data Platform (NDP) are part of which layer of Cisco SD-Access?

- **Controller**
- Network
- Physical
- Management

Correct answer: Controller

The Network Control Platform (NCP) and Network Data Platform (NDP) are part of the Cisco DNA Center and make up the bulk of the Controller Layer of SD-Access.

Cisco Identity Services Engine (ISE) is the other Controller Layer component.

Network, Physical, and Management are the other three SD-Access Layers.

47.

An organization is attempting to manage client density to ensure good network performance. Which of the following is a good choice? (Choose two.)

- **Increasing the number of access points**
- **Choosing antennas with smaller coverage areas**
- Choosing antennas with larger coverage areas
- Decreasing the number of access points

Client density is the number of devices connected to an Access Point (AP). With more clients and more active clients, performance degrades. Organizations can limit client density by deploying more APs and choosing antennas with a smaller coverage area. This combination ensures that adequate APs are available and that clients are unlikely to all connect to the same AP.

48.

An organization is implementing a two-tier network design because all of its systems are located on a single floor of a building. Which of the following layers of the hierarchical model is it most likely to consolidate? (Choose two).

- **Distribution**
- **Core**
- Access
- Edge
- Hub

The three-tier network design includes three layers:

- *Access Layer: Network edge where end-user devices (computers, Internet of Things (IoT) devices, mobile, etc.) are connected to the network.*
- *Distribution Layer: A distribution layer switch connects to access-layer switches from a building, floor, etc. This layer creates a boundary between the access and core layers, creating a boundary for the Spanning Tree Protocol (STP) and a summarization point for Internet Protocol (IP) routing information.*
- *Core Layer: Connects multiple distribution-layer switches together to support multi-site networks. This includes physically distributed sites or corporate networks with data centers, cloud deployments, etc.*

A two-tier network combines the distribution and core layers into a single layer.

Edge and hub are not layers in the three-tier model.

49.

In a Cisco router, where is routing information stored? (Choose two.)

- **RIB**
- **FIB**
- AIB
- TCAM
- CAM

The Routing Information Base (RIB) is Cisco's term for a routing table, which stores information on how to reach different devices or networks.

The Forwarding Information Base (FIB) stores the next-hop information for each network destination and is used to make Internet Protocol (IP) destination prefix-based decisions about how to route packets.

An adjacency table or the Adjacency Information Base (AIB), which is part of software CEF, contains the next-hop Media Access Control (MAC) addresses for each directly connected next-hop IP address as well as the MAC address of the egress interface. It's populated by the Address Resolution Protocol (ARP) or similar sources.

Content Addressable Memory (CAM) holds the MAC address table and uses specialized search techniques to enable addresses to be found faster than with Random Access Memory (RAM).

Ternary Content Addressable Memory (TCAM) is memory on a Cisco switch that allows multiple different fields to be used to evaluate a packet. It's used for Layer 2 or 3 searching and returns 0, 1, or X (don't care).

50.

Cisco's DNA Center is part of which of the following layers of the Cisco SD-Access architecture? (Choose three.)

- **Controller Layer**
- **Management Layer**
- **Physical Layer**
- Network Layer

Cisco DNA Center is part of the Controller and Physical layers of the Cisco SD-Access architecture, and its Graphical User Interface (GUI) is part of the Management layer.

Cisco DNA Center doesn't appear in the Network layer.

51.

Which of the following are open-source protocols that describe how a network can have multiple, redundant gateways? (Choose two.)

VRRP

GLBP

HSRP

FHRP

RGBP

First-Hop Redundancy Protocols (FHRP) help ensure network resiliency by creating a Virtual Internet Protocol (VIP) gateway linked to multiple physical gateways. If a gateway goes down, then the device's traffic will be sent via another gateway. The three main FHRPs include:

- ***Hot Standby Router Protocol (HSRP)**: Protocol developed by Cisco that creates a virtual IP and Media Access Control (MAC) address usually held by the active router. If the active router fails, a standby router takes over these addresses and acts as the gateway.*
- ***Virtual Router Redundancy Protocol (VRRP)**: Industry standard protocol that operates similarly to HSRP but names the routers "master" and "backup". This protocol allows preemption by default and uses a particular MAC address structure for the VIP gateway.*
- ***Gateway Load Balancing Protocol (GLBP)**: Offers both redundancy and load balancing. The network has up to four Active Virtual Forwarders (AVFs) responsible for forwarding traffic for their assigned hosts and a single Active Virtual Gateway (AVG) that responds to Address Resolution Protocol (ARP) requests with the virtual MAC of the assigned AVF. Failure of the AVG or an AVF causes another system to take over its role.*

RGBP is a fabricated term.

52.

Which of the following Cisco tools manages the network topology and advertises routes?

vSmart

vManage

Edge devices

vOrchestrator

vSphere

Correct answer: vSmart

Cisco's Software-Defined Wide Area Network (SD-WAN) solution implements various planes, including the control and data planes. The control plane is implemented using vSmart, which is responsible for defining the network topology, advertising routes, and data policies for SD-WAN edge devices.

The SD-WAN edge devices make up the data plane and are responsible for forwarding traffic between locations via various media. Cisco's vManage Network Management System (NMS) and vBond orchestrator implement the management and orchestration planes respectively.

vSphere is a VMware hypervisor.

53.

Which of the following elements of Cisco's SD-WAN solution offers centralized visibility and configuration control for the entire solution?

vManage

vControl

vBond

vAnalytics

Correct answer: vManage

Cisco’s Software-Defined Wide Area Network (SD-WAN) solution implements four planes:

- ***Management:** Cisco’s vManage Network Management System (NMS) implements the management plane, providing single-pane-of-glass visibility and control.*
- ***Orchestration:** vBond implements the orchestration plane and performs tasks such as authentication, Network Address Translation (NAT) detection, and load balancing.*
- ***Control:** The control plane is implemented using vSmart, which is responsible for defining the network topology, advertising routes, and data policies for SD-WAN edge devices.*
- ***Data:** These SD-WAN edge devices make up the data plane and are responsible for forwarding traffic between locations via various media.*

vAnalytics is an optional analytics service.

54.

Which of the following network models is MOST dependent on virtualization technology?

- **Fabric**
- Hierarchical
- Two-tier
- Three-tier

Correct answer: Fabric

Fabric networks create a virtual overlay on top of physical network architecture. This makes it easier to implement network segmentation or redesign the network architecture without a physical redesign.

The other network designs can be implemented at least partially using physical components.

55.

Which of the following is part of the Controller Layer of the Cisco SD-Access Architecture but not part of the DNA Center?

Identity Services Engine (ISE)

Network Control Platform (NCP)

Network Data Platform (NDP)

SD-Access Overlay Network

Base Automation

Correct answer: Identity Services Engine (ISE)

The Controller layer of the Cisco Software Defined (SD)-Access architecture consists of the Identity Services Engine (ISE), Network Control Platform (NCP), and Network Data Platform (NDP). The NCP and NDP are part of Cisco DNA Center but ISE is not.

The SD-Access Overlay Network is part of the Network layer, and Base Automation is part of the Management layer.

56.

Which of the following wireless deployment models may require more WLCs than the others? (Choose two.)

Distributed

Controllerless

Centralized

Cloud-based

Most Cisco wireless Access Points (APs) are lightweight and require a Wireless Local Area Network (LAN) Controller (WLC) to operate. Controller-based wireless deployments have a standalone WLC that can be deployed under the following models:

- ***Centralized**: The WLC is placed in a centralized location, such as the network core, enabling it to support many APs. WLC is likely near commonly-used resources (Internet, data center, etc.) and enables centralized policy enforcement.*
- ***Distributed**: Multiple WLCs are located alongside each switch in the access layer. This design makes sense for geographically distributed sites.*
- ***Cloud-Based**: Similar to the centralized model, except the WLC is located in a public or private cloud, rather than on-premises.*

Controllerless deployments integrate the WLC into an AP, which is called an Embedded Wireless Controller (EWC). These can be used in distributed deployments as well.

Distributed and controllerless distributed deployments may require many WLCs or EWCs since each AP may have its own.

57.

Which of the following is a Cisco algorithm that takes priority into account when deciding which packets to drop for congestion avoidance?

WRED

RED

Tail Drop

WRR

LLQ

Correct answer: WRED

Congestion avoidance attempts to proactively prevent network congestion. Some common methods include:

- ***Tail Drop:** Drops all types of traffic when queue output buffers are full. Can cause Transport Control Protocol (TCP) global synchronization where TCP streams slow down and speed up in sync, oscillating between underuse and congestion.*
- ***Random Early Detection (RED):** RED randomly drops packets when buffers are full, avoiding congestion and TCP global synchronization.*
- ***Weighted RED (WRED):** Cisco RED implementation uses weights to affect random dropping. Weights are based on Internet Protocol (IP) Precedence (IPP) or Differentiated Services Code Point (DSCP).*

Weighted Round Robin (WRR) and Low-Latency Queueing (LLQ) are congestion management algorithms.

58.

Which of the following contains next-hop information for each destination on the network?

FIB

RIB

AIB

TCAM

CAM

Correct answer: FIB

The Forwarding Information Base (FIB) stores the next-hop information for each network destination and is used to make Internet Protocol (IP) destination prefix-based decisions about how to route packets.

Content Addressable Memory (CAM) holds the Media Access Control (MAC) address table and uses specialized search techniques to enable addresses to be found faster than with Random Access Memory (RAM).

The Routing Information Base (RIB) is Cisco's term for a routing table, which stores information on how to reach different devices or networks.

Ternary Content Addressable Memory (TCAM) is memory on a Cisco switch that allows multiple different fields to be used to evaluate a packet. It's used for Layer 2 and 3 searching and returns 0, 1, or X (don't care).

An adjacency table or the Adjacency Information Base (AIB) contains the next-hop MAC addresses for each directly connected next-hop IP address as well as the MAC address of the egress interface. It's populated by the Address Resolution Protocol (ARP) table or similar sources.

59.

Which of the following are Cisco-proprietary protocols for deploying redundant network gateways? (Choose two.)

HSRP

GLBP

FHRP

VRRP

First-Hop Redundancy Protocols (FHRP) help ensure network resiliency by creating a Virtual Internet Protocol (VIP) gateway linked to multiple physical gateways. If a gateway goes down, then the device's traffic will be sent via another gateway. The three main FHRPs include:

- ***Hot Standby Router Protocol (HSRP):*** *Protocol developed by Cisco that creates a virtual IP and Media Access Control (MAC) address usually held by the active router. If the active router fails, a standby router takes over these addresses and acts as the gateway.*
- ***Gateway Load Balancing Protocol (GLBP):*** *Cisco-proprietary protocol that offers both redundancy and load balancing. The network has up to four Active Virtual Forwarders (AVFs) responsible for forwarding traffic for their assigned hosts and a single Active Virtual Gateway (AVG) that responds to Address Resolution Protocol (ARP) requests with the virtual MAC of the assigned AVF. Failure of the AVG or an AVF causes another system to take over its role.*
- ***Virtual Router Redundancy Protocol (VRRP):*** *Industry standard protocol that operates similarly to HSRP but names the routers "master" and "backup". This protocol allows preemption by default and uses a particular MAC address structure for the VIP gateway.*

60.

The edge devices deployed within Cisco's SD-WAN network comprise which of the following?

- **Data Plane**
- Management Plane
- Control Plane
- Orchestration Plane
- Policy Plane

Correct answer: Data Plane

Cisco’s Software-Defined Wide Area Network (SD-WAN) solution implements various planes, including the control and data planes. The control plane is implemented using vSmart, which is responsible for defining the network topology and advertising routes and data policies to SD-WAN edge devices. These SD-WAN edge devices make up the data plane and are responsible for forwarding traffic between locations via various media.

Cisco’s vManage Network Management System (NMS) and vBond orchestrator implement the management and orchestration planes respectively.

The policy plane is part of Cisco's SD-Access.

61.

Which of the following algorithms for congestion management handles both real-time and non-real-time traffic?

LLQ

CBFWQ

WFQ

WRR

CQ

Correct answer: LLQ

Congestion management buffers excess traffic, and then removes packets from the queue via various algorithms. Some examples include:

- ***Weighted Round Robin (WRR):** Legacy algorithm that adds support for prioritization to round robin where each queue is assigned bandwidth based on its weight.*
- ***Custom Queueing (CQ):** Legacy Cisco WRR protocol with 16 queues and First In First Out (FIFO) ordering within a queue.*
- ***Weighted Fair Queueing (WFQ):** Legacy algorithms where interface bandwidth is divided among network flows so each flow gets bandwidth based on its weight.*
- ***Class-Based Weighted Fair Queueing (CBWFQ):** Up to 256 queues are created (one per traffic class) that are serviced based on the class's assigned bandwidth. Uses traffic descriptors to classify traffic, and classified traffic is assigned bandwidth, weight, maximum packet limit, and queue limit. Packets exceeding the queue limit are dropped.*
- ***Low-Latency Queueing (LLQ):** Combines Priority Queueing (PQ) with CBWFQ, creating a strict-priority queue that is serviced before all CBWFQ queues. This queue includes real-time traffic, such as voice traffic. This algorithm handles both real-time and non-real-time traffic.*

62.

Which of the following are part of the role of vBond in Cisco SD-WAN? (Choose three.)

Authentication

NAT Detection

Load Balancing

Route Advertisement

Policy Management

In Cisco Software-Defined Wide Area Network (SD-WAN), vBond is the orchestration service that performs authentication, Network Address Translation (NAT) detection, and load balancing.

vManage is responsible for policy management, and vSmart handles route advertisements.

63.

An organization wants to identify the locations of devices in a space by measuring RSS. However, the area contains walls, furniture, and other sources of attenuation.

Which of the following could help address this?

Creating an RF calibration template

Spacing APs equally through the space

Using three APs for triangulation

Scanning RFID frequencies

Correct answer: Creating an RF calibration template

If a client is in an open space, measuring Received Signal Strength (RSS) from different Access Points (APs) can help to triangulate its location. Cisco devices use a Radio Frequency (RF) calibration template that is generated using an RF scanner and accurately measures attenuation and signal propagation within a space, enabling it to be used in other places than an empty room.

Equally spacing APs and using three APs for triangulation won't work since the various attenuation sources will affect the perceived distance of a device from an AP.

RF Identification (RFID) scanning only works if devices have attached RFID tags.

64.

Which of the following is implemented by Cisco's vSmart?

- **Control plane**
- Management plane
- Orchestration plane
- Data plane

Correct answer: Control plane

Cisco’s Software-Defined Wide Area Network (SD-WAN) solution implements various planes, including the control and data planes. The control plane is implemented using vSmart, which is responsible for defining the network topology and advertising routes and data policies to SD-WAN edge devices.

These SD-WAN edge devices make up the data plane and are responsible for forwarding traffic between locations via various media.

Cisco’s vManage Network Management System (NMS) and vBond orchestrator implement the management and orchestration planes respectively.

65.

A department store wants to track consumers' movements through their store, but many devices aren't connected to the store's Wi-Fi. Which of the following is an option for accomplishing this?

Monitoring probe requests

Measuring RSS

Tracking RFID tags

Measuring interference

Correct answer: Monitoring probe requests

Location services track device locations within a wireless network, which is useful for asset tracking. Some methods for implementing it include:

- ***Measuring Received Signal Strength (RSS):** If a client is in an open space, measuring RSS from different Access Points (APs) can help to triangulate its location. Cisco devices use a Radio Frequency (RF) calibration template that is generated using an RF scanner and accurately measures attenuation and signal propagation within a space, enabling it to be used in other places than an empty room.*
- ***Probe Requests:** Devices with Wi-Fi active will send out Probe Requests to APs on all supported wireless channels, enabling triangulation.*
- ***RFID Tags:** Radio Frequency Identification (RFID) tags attached to devices can be used to map devices' locations as they attach to Wi-Fi or send out Probe Requests.*
- ***Interference Detection:** Devices that create signal interference, such as cordless phones and wireless cameras, can be located via spectrum analysis and identifying locations experiencing interference.*

66.

An organization's users primarily connect to Software as a Service (SaaS) apps. Which of the following wireless deployment models is the best fit?

Cloud-based

Centralized

Distributed

Controllerless

Correct answer: Cloud-based

Cisco wireless Access Points (APs) are lightweight and require a Wireless Local Area Network (LAN) Controller (WLC) to operate, which it connects to via a pair of Control and Provisioning of Wireless Access Points (CAPWAP) tunnels. Controller-based wireless deployments have a standalone WLC that can be deployed under the following models:

- ***Centralized:** The WLC is placed in a centralized location, such as the network core, enabling it to support many APs. WLC is likely near commonly-used resources (Internet, data center, etc.) and enables centralized policy enforcement.*
- ***Distributed:** Multiple WLCs are located alongside each switch in the access layer. This design makes sense for geographically distributed sites.*
- ***Cloud-Based:** Similar to the centralized model, but the WLC is located in a public or private cloud, rather than on-premises.*

Controllerless deployments integrate the WLC into an AP, which is called an Embedded Wireless Controller (EWC). These can be used in distributed deployments as well.

67.

Which of the following wireless network designs locates WLCs near resources that users are likely to access? (Choose two.)

Centralized

Cloud-Based

Distributed

Controllerless

Cisco wireless Access Points (APs) are lightweight and require a Wireless Local Area Network (LAN) Controller (WLC) to operate, which it connects to via a pair of Control And Provisioning of Wireless Access Points (CAPWAP) tunnels. Controller-based wireless deployments have a standalone WLC that can be deployed under the following models:

- ***Centralized:** The WLC is placed in a centralized location, such as the network core, enabling it to support many APs. WLC is likely near commonly-used resources (Internet, data center, etc.) and enables centralized policy enforcement.*
- ***Cloud-Based:** Similar to the centralized model, except the WLC is located in a public or private cloud, rather than on-premises. This means that WLCs are close to cloud-based resources that employees are likely to access.*
- ***Distributed:** Multiple WLCs are located alongside each switch in the access layer. This design makes sense for geographically distributed sites.*

Controllerless deployments integrate the WLC into an AP, which is called an Embedded Wireless Controller (EWC). These can be used in distributed deployments as well.

68.

Which of the following components of Cisco's SD-WAN solution is optional?

vAnalytics

vManage

vSmart

vBond

Edge devices

Correct answer: vAnalytics

Cisco's Software-Defined Wide Area Network (SD-WAN) has an optional vAnalytics solution.

The four main components include vManage, vSmart, vBond, and the edge devices that make up the network.

69.

An organization wants to have visibility into the location of its IT assets that can't connect to the corporate Wi-Fi. Which of the following could identify these assets? (Choose two.)

- **RFID tags**
- **Interference detection**
- Measuring RSS
- Monitoring probe requests
- Using traceroute

Radio Frequency Identification (RFID) tags attached to devices can be used to map devices' locations as the tags attach to Wi-Fi or send out Probe Requests. Interference detection can be used to identify devices that use other wireless protocols and interfere with Wi-Fi signals.

Measuring Received Signal Strength (RSS) and monitoring probe requests assume that Wi-Fi is enabled and active on the device, which is not always true.

Traceroute is a tool to identify the path that packets take over the network, not device location.

70.

Which of the following are modern congestion management algorithms? (Choose two.)

CBWFQ

LLQ

WRR

WFQ

Congestion management buffers excess traffic and then removes packets from the queue via various algorithms. Modern queuing algorithms include:

- ***Class-Based Weighted Fair Queueing (CBWFQ):*** *Up to 256 queues are created (one per traffic class) that are serviced based on the class's assigned bandwidth. Uses traffic descriptors to classify traffic, and classified traffic is assigned bandwidth, weight, maximum packet limit, and queue limit. Packets exceeding the queue limit are dropped.*
- ***Low-Latency Queueing (LLQ):*** *Combines Priority Queueing (PQ) with CBWFQ, creating a strict-priority queue that is serviced before all CBWFQ queues. This queue includes real-time traffic, such as voice traffic.*

Some legacy examples include:

- ***Weighted Round Robin (WRR):*** *Adds support for prioritization to round robin where each queue is assigned bandwidth based on its weight.*
- ***Weighted Fair Queueing (WFQ):*** *Interface bandwidth is divided among network flows so each flow gets bandwidth based on its weight.*

71.

Which of the following wireless deployment models are designed to accommodate Cisco's lightweight APs? (Choose three.)

- **Centralized**
- **Distributed**
- **Cloud-based**
- Controllerless
- Autonomous

Many Cisco wireless Access Points (APs) are lightweight and require a Wireless Local Area Network (LAN) Controller (WLC) to operate. Controller-based wireless deployments have a standalone WLC that can be deployed under the following models.

- ***Centralized**: The WLC is placed in a centralized location, such as the network core, enabling it to support many APs. WLC is likely near commonly-used resources (Internet, data center, etc.) and enables centralized policy enforcement.*
- ***Distributed**: Multiple WLCs are located alongside each switch in the access layer. This design makes sense for geographically distributed sites.*
- ***Cloud-Based**: Similar to a centralized model, but the WLC is located in a public or private cloud, rather than on-premises.*

Controllerless deployments integrate the WLC into an AP, which is called an Embedded Wireless Controller (EWC). These can be used in distributed deployments as well.

Autonomous deployments require standalone APs.

72.

Where on a Cisco switch is stored information about how to reach different networks and devices?

- **RIB**
- TCAM
- AIB
- FIB

Correct answer: RIB

The Routing Information Base (RIB) is Cisco's term for a routing table, which stores information on how to reach different devices or networks.

Ternary Content Addressable Memory (TCAM) is memory on a Cisco switch that allows multiple different fields to be used to evaluate a packet. It's used for Layer 2 and 3 searching and returns 0, 1, or X (don't care).

The Forwarding Information Base (FIB) stores the next-hop information for each network destination and is used to make Internet Protocol (IP) destination prefix-based decisions about how to route packets.

An adjacency table or the Adjacency Information Base (AIB), which is part of software CEF, contains the next-hop MAC addresses for each directly-connected next-hop IP address as well as the Media Access Control (MAC) address of the egress interface. It's populated by the Address Resolution Protocol (ARP) table or similar sources.

73.

Which of the following is responsible for forwarding traffic between locations in Cisco SD-WAN?

Edge devices

vSmart

vManage

vBond

Correct answer: Edge devices

Cisco's Software-Defined Wide Area Network (SD-WAN) solution implements various planes, including the control and data planes. The control plane is implemented using vSmart, which is responsible for defining the network topology and advertising routes and data policies to SD-WAN edge devices. These SD-WAN edge devices make up the data plane and are responsible for forwarding traffic between locations via various media.

Cisco's vManage Network Management System (NMS) and vBond orchestrator implement the management and orchestration planes respectively.

74.

Cisco supports two types of SD-WAN. Cisco SD-WAN was created by a company acquired by Cisco, and another was developed by another company. What are these two companies? (Choose two.)

Viptela

Meraki

Aryaka

Huawei

Cisco Software-Defined Wide Area Network (SD-WAN) was based on Viptela's solution. Cisco also supports Meraki SD-WAN.

Aryaka and Huawei are unrelated SD-WAN providers.

75.

Which of the following are the responsibilities of Cisco's vSmart tool? (Choose three.)

Define network topology

Advertise routes

Advertise data policies

Forward traffic between locations

Manage SD-WAN solution configurations

Cisco’s Software-Defined Wide Area Network (SD-WAN) solution implements various planes, including the control and data planes. The control plane is implemented using vSmart, which is responsible for defining the network topology and advertising routes and data policies to SD-WAN edge devices.

These SD-WAN edge devices make up the data plane and are responsible for forwarding traffic between locations via various media.

Cisco’s vManage Network Management System (NMS) manages SD-WAN solution configurations.

76.

Which of the following is the term used to describe the number of clients attached to a wireless AP?

- **Client density**
- User density
- Client capacity
- User capacity
- Device density

Correct answer: Client density

Client density is the number of devices connected to an Access Point (AP). With more clients and more active clients, performance degrades.

The other answers are fabricated terms.

2.0 Virtualization

77.

Which of the following VXLAN control plane options are preferred for data centers and private clouds? (Choose two.)

VXLAN with MP-BGP EVPN control plane

VXLAN with Multicast underlay

VXLAN with static unicast VXLAN tunnels

VXLAN with LISP control plane

Cisco devices support four Virtual eXtensible Local Area Network (VXLAN) control planes:

- *VXLAN with Multicast underlay*
- *VXLAN with static unicast VXLAN tunnels*
- *VXLAN with Multiprotocol Border Gateway Protocol Ethernet Virtual Private Network (MP-BGP EVPN) control plane*
- *VXLAN with Locator/ID Separation Protocol (LISP) control plane*

SD-Access is VXLAN with LISP control plane (preferred for campus environments). Multicast and MP-BGP EVPN are preferred for data center and private cloud environments.

78.

In Cisco's Locator/ID Separation Protocol (LISP), an xTR can fulfill the role of which of the following? (Choose two.)

- **ETR**
- **ITR**
- PETR
- PITR
- MR

In the Locator/ID Separation Protocol (LISP), a Tunnel Router (xTR) acts as an Ingress and Egress Tunnel Router (ITR/ETR).

A Proxy xTR (PxTR) can act as a Proxy ITR (PITR) or Proxy ETR (PETR).

Neither can act as a Map Resolver (MR).

79.

After typing "crypto isakmp policy priority" which of the following are valid commands? (Choose three.)

- **encryption**
- **authentication**
- **group**
- signature
- exchange

The "crypto isakmp policy priority" command begins the creation of an Internet Security Association and Key Management Protocol (ISAKMP) policy. After this command, the encryption, hash, authentication, and group commands can be used to choose algorithms and settings for this.

Signature and exchange are not valid commands in this context.

80.

Which of the following hypervisors needs to run on top of a client OS? (Choose two.)

- **VirtualBox**
- **VMware Fusion**
- VMware vSphere
- Citrix Hypervisor
- Red Hat KVM

A hypervisor is software that allows multiple Virtual Machines (VMs) to run on the same hardware. There are two types of hypervisors:

- ***Type 1:** A Type I, bare-metal, or native hypervisor runs directly on the device hardware with no Operating System (OS). Examples of Type 1 hypervisors include VMware vSphere, Citrix Hypervisor, and Red Hat Kernel-based Virtual Machine (KVM).*
- ***Type 2:** Type 2 hypervisors are software that runs within a host operating system. VirtualBox and VMware Fusion are examples of Type 2 hypervisors, which are typically used by client machines.*

81.

Which of the following exists for containers but not VMs?

- **Runtime**
- OS
- Applications
- Binaries
- Libraries

Correct answer: Runtime

Containers have a separate runtime and their own applications, binaries, and libraries.

Virtual Machines (VMs) have their own applications, binaries, libraries, and Operating System (OS). They don't have a runtime.

82.

Which of the following commands are legitimate commands used to set up or test a GRE tunnel? (Choose two.)

- **traceroute**
- **tunnel source**
- gre verify
- tunnel create

The traceroute command can be used to verify that traffic is flowing over a Generic Routing Encapsulation (GRE) tunnel. To do so, specify the destination and source addresses (e.g. traceroute 10.0.2.1 source 10.1.2.1) and check that the Internet Protocol (IP) address of the tunnel is included in the list of hops.

The tunnel source command is used to define the local end of the tunnel.

The commands tunnel create and gre verify are fabricated.

83.

Which of the following are true? (Choose two.)

Containers and VMs both have their own binaries and libraries.

Containers and VMs both use their host's NICs.

Containers and VMs both have their own OS.

Containers are a type of VM.

Containers and Virtual Machines (VMs) both have their own binaries and libraries, and both have network access via their host's Network Interface Cards (NICs).

However, containers are not a type of VM and are dependent on the host's kernel and OS.

84.

Which of the following commands correctly specifies the remote end of a GRE tunnel?

tunnel destination ip-address

tunnel source ip-address

gre destination ip-address

tunnel destination interface-id

gre source ip-address

Correct answer: tunnel destination ip-address

tunnel destination ip-address is the correct command to specify the remote end of a Generic Routing Encapsulation (GRE) tunnel. The source is the local end, and remote ends can't be specified using an interface-id.

The other answers are fabricated commands.

85.

Which of the following does a VM have that a container doesn't? (Choose two.)

- **Guest OS**
- **Hypervisor**
- Runtime
- Own Applications
- Own binaries and libraries

A Virtual Machine (VM) has a guest Operating System (OS) and a hypervisor, while containers don't.

Both have their own applications, binaries, and libraries.

Containers have a runtime.

86.

Which of the following is the correct command to create a GRE tunnel?

interface tunnel tunnel-number

tunnel create tunnel-number

tunnel gre tunnel-number

tunnel mode gre {ip | ipv6}

Correct answer: interface tunnel tunnel-number

Generic Routing Encapsulation (GRE) allows the creation of tunnels over Internet Protocol (IP) networks, which is useful for creating Virtual Private Networks (VPNs). GRE encapsulates an existing packet within a GRE packet, which includes an IP header that points to a remote endpoint where the packet will be de-encapsulated and forwarded to its destination.

GRE tunnels are defined using the following commands:

- ***interface tunnel tunnel-number:** Creates a new GRE tunnel.*
- ***tunnel source {ip-address|interface-id}:** Identifies the local end of the GRE tunnel.*
- ***tunnel destination ip-address:** Defines the remote end of the GRE tunnel.*
- ***ip address ip-address subnet-mask:** Assigns an IP address to the tunnel.*

The other commands are fabricated.

87.

Which of the following is not one of the key components of Cisco’s Locator/ID Separation Protocol (LISP)?

Management Plane

Control Plane

Routing Architecture

Data Plane

Correct answer: Management Plane

Cisco’s Locator/ID Separation Protocol (LISP) doesn't have a management plane. It has three main components:

- ***LISP Routing Architecture:*** *Traditionally, an IP address identifies a particular device at a specific location. LISP breaks it into separate Endpoint Identifiers (EIDs) and Routing Locators (RLOCs). If a device changes location, its EID stays the same while its RLOC changes.*
- ***LISP Control Plane:*** *LISP converts EIDs to RLOCs by making a map request to the Map Resolver (MR). This pull-based model is more efficient than Border Gateway Protocol (BGP) or Open Shortest Path First (OSPF), which pushes route data, even for unused routes.*
- ***LISP Data Plane:*** *An IP packet received from an EID is encapsulated in an Internet Protocol (IP)/User Datagram Protocol (UDP) packet using an IP address in the RLOC IP space.*

88.

Which of the following IPSec protocols offers data confidentiality?

ESP

IKE

ISAKMP

AH

Correct answer: ESP

Internet Protocol Security (IPSec) is a set of protocols used to create Virtual Private Networks (VPNs). Some key elements include:

- ***Authentication Header (AH):** The AH protects against replay attacks and ensures data integrity and peer authentication using digital signatures. It doesn't provide data confidentiality via encryption. It has a protocol number of 51.*
- ***Encapsulating Security Payload (ESP):** ESP provides data integrity, confidentiality, peer authentication, and replay protection. It has a protocol number of 50 and can carry packets either in tunnel mode (old packet's headers are encrypted and new IPSec headers are added) or transport mode (only the original packet payload is encrypted).*
- ***Internet Key Exchange (IKE):** IKE performs mutual authentication and allows the creation of Security Associations (SAs), which are tunnels that carry data and control plane traffic for IPsec. IKEv2 is the modern version of the protocol, which is more efficient, supports asymmetric authentication, and offers improved protocols and cryptographic algorithms.*

IKE is a particular implementation of the Internet Security Association and Key Management Protocol (ISAKMP).

89.

Which of the following is the correct command to assign an IP address on an interface for a VRF?

ip address ip-address subnet-mask

vrf address ip-address subnet mask

vrf address ipv4 ip-address subnet-mask

address ipv4 ip-address subnet-mask

address ipv6 ip-address subnet-mask

Correct answer: ip address ip-address subnet-mask

The command "ip address ip-address subnet-mask" assigns an Internet Protocol version 4 (IPv4) address to an interface for Virtual Routing and Forwarding (VRF).

Alternatively, IPv6 addresses can be assigned with "ipv6 address ipv6-address/prefix-length."

The other commands are fabricated.

90.

AES is the preferred algorithm option for which of the following IPsec protocols?

ESP

IKE

ISAKMP

AH

IKEv2

Correct answer: ESP

Internet Protocol Security (IPSec) is a set of protocols used for creating Virtual Private Networks (VPNs). Some key elements include:

- ***Authentication Header (AH):** The AH protects against replay attacks and ensures data integrity and peer authentication using digital signatures. It doesn't provide data confidentiality via encryption. It has a protocol number of 51.*
- ***Encapsulating Security Payload (ESP):** ESP provides data integrity, confidentiality, peer authentication, and replay protection. It has a protocol number of 50 and can carry packets either in tunnel mode (old packet's headers are encrypted and new IPsec headers are added) or transport mode (only the original packet payload is encrypted). The Advanced Encryption Standard (AES) is the preferred algorithm for data encryption.*
- ***Internet Key Exchange (IKE):** IKE performs mutual authentication and allows the creation of Security Associations (SAs), which are tunnels that carry data and control plane traffic for IPsec. Internet Key Exchange version two (IKEv2) is the modern version of the protocol, which is more efficient, supports asymmetric authentication, and offers improved protocols and cryptographic algorithms.*

IKE is a specific implementation of the Internet Security Association and Key Management Protocol (ISAKMP).

91.

Which of the following commands is used to associate a virtual router with a particular interface?

vrf forwarding

vrf definition

address-family

ip or ipv6

Correct answer: vrf forwarding

Virtual Routing and Forwarding (VRF) is:

- *Initialized with the **vrf definition vrf-name** command.*
- *Set to Internet Protocol version four (IPv4) or IP version six (IPv6) with the **address-family {ipv4\|ipv6}** command.*
- *Associated to a particular router interface with **vrf forwarding vrf-name** in that interface's configuration submodule (entered using the interface **interface-id** command).*
- *Assigned an IP address using the **ip address ip-address subnet-mask** and/or **ipv6 address ipv6-address/prefix-length** commands in the interface's configuration submodule.*

92.

Which of the following is true of VRFs on a router?

A router will have more routing tables than VRFs

A router will have an equal number of VRFs and routing tables

A router will always have a single routing table

All VRFs share a single routing table

A VRF may use multiple routing tables

Correct answer: A router will have more routing tables than VRFs

A router will have more routing tables than Virtual Router Functions (VRFs). This is because each VRF will have its own routing table, and a global routing table will exist for all traffic not flowing over a VRF.

93.

Different virtual routers created via VRF will have their own copies of all of the following, except:

Global routing table

Routing table

Router interface

Forwarding table

Correct answer: Global routing table

Virtual Route Forwarding (VRF) splits a single physical router into multiple virtual routers. Each virtual router has its own router interface, routing table, and forwarding table, keeping them isolated from one another. The router will also have a global routing table, which is the routing table for all traffic not assigned to a particular VRF.

94.

Which of the following types of hypervisors offers the greatest efficiency?

Type 1

Type 2

Type 3

Type 0

Correct answer: Type 1

A hypervisor is software that allows multiple Virtual Machines (VMs) to run on the same hardware. There are two types of hypervisors:

- ***Type 1:** A Type I, bare-metal, or native hypervisor runs directly on the device hardware with no operating system, making it the more efficient type of hypervisor. Examples of Type 1 hypervisors include VMware vSphere, Citrix Hypervisor, and Red Hat Kernel-based Virtual Machine (KVM).*
- ***Type 2:** Type 2 hypervisors are software that runs within a host operating system. VirtualBox and VMware Fusion are examples of Type 2 hypervisors, which are typically used by client machines.*

There are no Type 0 or Type 3 hypervisors.

95.

Which of the following is true of virtual switching?

vSwitches are used by both VMs and containers

Traffic can flow directly between vSwitches without going over the physical network

Multiple vSwitches can share a pNIC

A device can only operate a single vSwitch

Correct answer: vSwitches are used by both VMs and containers

Virtual switching is when software emulates a physical, layer-2 switch, just like a Virtual Machine (VM) emulates a physical computer.

A virtual switch (vSwitch) allows communication between VMs and the outside world via the computer's physical Network Interface Cards (pNICs). A computer can host multiple vSwitches, but a pNIC can only host one vSwitch, and all traffic between vSwitches needs to travel over the physical network (out one pNIC and in another).

96.

Which of the following will send a map registration packet in Cisco's Locator/ID Separation Protocol (LISP)?

ETR

ITR

PITR

PETR

MS

Correct answer: ETR

Cisco's Locator/ID Separation Protocol (LISP) Egress Tunnel Routers (ETRs) send a map register message to the Map Server (MS) to register Endpoint Identifiers (EIDs).

Proxy ETRs (PETRs) don't register EIDs, and Ingress Tunnel Routers (ITRs) and Proxy ITRs (PITRs) request EIDs, not register them.

The MS is the recipient of the map registration packet, not the sender.

97.

Which of the following types of hypervisors has its own OS?

Type 1

Type 2

Type 0

Type A

Type B

Correct answer: Type 1

A hypervisor is software that allows multiple Virtual Machines (VMs) to run on the same hardware. There are two types of hypervisors:

- ***Type 1:** A Type I, bare-metal, or native hypervisor runs directly on the device hardware with no Operating System (OS). Examples of Type 1 hypervisors include VMware vSphere, Citrix Hypervisor, and Red Hat Kernel-based Virtual Machine (KVM).*
- ***Type 2:** Type 2 hypervisors are software that runs within a host operating system. VirtualBox and VMware Fusion are examples of Type 2 hypervisors, which are typically used by client machines.*

Type 0, A, and B hypervisors do not exist.

98.

Which of the following are the benefits of VMs? (Choose three.)

- **Ability to host multiple VMs on a single system**
- **Ability to easily migrate VMs between systems**
- **Support for high-availability environments**
- Lower resource requirements than containers

Virtual Machines (VMs) allow multiple VMs to be hosted on a single system and to be migrated between systems. These help them to ensure the high availability of supported services.

However, VMs generally have higher resource requirements than containers.

99.

Which of the following are examples of a Type 2 hypervisor? (Choose two.)

VirtualBox

VMware Fusion

VMware vSphere

Red Hat KVM

A hypervisor is software that allows multiple Virtual Machines (VMs) to run on the same hardware. There are two types of hypervisors:

- ***Type 1:** A Type I, bare-metal, or native hypervisor runs directly on the device hardware with no operating system. Examples of Type 1 hypervisors include VMware vSphere, Citrix Hypervisor, and Red Hat Kernel-based Virtual Machine (KVM).*
- ***Type 2:** Type 2 hypervisors are software that runs within a host operating system. VirtualBox and VMware Fusion are examples of Type 2 hypervisors, which are typically used by client machines.*