

Question 1

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


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

Answer: B

Question 2

Refer to the exhibit.

An engineer must ensure that all traffic leaving AS 200 will choose Link 2 as the exit point. Assuming that all BGP neighbor relationships have been formed and that the attributes have not been changed on any of the routers, which configuration accomplish task?



- A. R4(config-router)#bgp default local-preference 200**
- B. R3(config-router)#neighbor 10.1.1.1 weight 200
- C. R3(config-router)#bgp default local-preference 200
- D. R4(config-router)#neighbor 10.2.2.2 weight 200

Answer: A

Explanation

Local preference is an indication to the AS about which path has preference to exit the AS in order to reach a certain network. A path with a higher local preference is preferred. The default value for local preference is 100.

Unlike the weight attribute, which is only relevant to the local router, local preference is an attribute that routers exchange in the same AS. The local preference is set with the "bgp default local-preference *value*" command.

In this case, both R3 & R4 have exit links but R4 has higher local-preference so R4 will be chosen as the preferred exit point from AS 200.

Question 3

Which protocol infers that a YANG data model is being used?

- A. SNMP
- B. REST
- C. **RESTCONF**
- D. NX-API

Answer: C

Explanation

YANG (Yet another Next Generation) is a data modeling language for the definition of data sent over network management protocols such as the NETCONF and RESTCONF.

Question 4

Which configuration restricts the amount of SSH that a router accepts 100 kbps?

```
A.
class-map match-all CoPP_SSH
match access-group name CoPP_SSH
!
Policy-map CoPP_SSH
class CoPP_SSH
police cir 100000
exceed-action drop
!
!
!
Interface GigabitEthernet0/1
ip address 209.165.200.225 255.255.255.0
ip access-group CoPP_SSH out
duplex auto
speed auto
```

```

media-type rj45
service-policy input CoPP_SSH
!
ip access-list extended CoPP_SSH
permit tcp any any eq 22
!

B.
class-map match-all CoPP_SSH
match access-group name CoPP_SSH
!
Policy-map CoPP_SSH
class CoPP_SSH
police cir CoPP_SSH
exceed-action drop
!
!
!
Interface GigabitEthernet0/1
ip address 209.165.200.225 255.255.255.0
ip access-group ... out
duplex auto
speed auto
media-type rj45
service-policy input CoPP_SSH
!
Ip access-list extended CoPP_SSH
deny tcp any any eq 22
!

```

```

C.
class-map match-all CoPP_SSH
match access-group name CoPP_SSH
!
Policy-map CoPP_SSH
class CoPP_SSH
police cir 100000
exceed-action drop
!
!
!
Control-plane
service-policy input CoPP_SSH
!
Ip access-list extended CoPP_SSH
deny tcp any any eq 22
!

```

```

D.
class-map match-all CoPP_SSH
match access-group name CoPP_SSH
!
Policy-map CoPP_SSH
class CoPP_SSH
police cir 100000
exceed-action drop
!
!
!
Control-plane transit
service-policy input CoPP_SSH
!
Ip access-list extended CoPP_SSH
permit tcp any any eq 22
!


```

Answer: C

Explanation

CoPP protects the route processor on network devices by treating route processor resources as a separate entity with its own ingress interface (and in some implementations, egress also). CoPP is used to police traffic that is destined to the route processor of the router such as:

- + routing protocols like OSPF, EIGRP, or BGP.
- + Gateway redundancy protocols like HSRP, VRRP, or GLBP.
- + Network management protocols like telnet, SSH, SNMP, or RADIUS.



Therefore we must apply the CoPP to deal with SSH because it is in the management plane. CoPP must be put under “control-plane” command.

Question 5


What NTP stratum level is a server that is connected directly to an authoritative time source?

- A. Stratum 0
- B. Stratum 1**
- C. Stratum 14
- D. Stratum 15

Answer: B

Explanation

The stratum levels define the distance from the reference clock. A reference clock is a stratum 0 device that is assumed to be accurate and has little or no delay associated with it. Stratum 0 servers cannot be used on the network but they are directly connected to computers which then operate as stratum-1 servers. A stratum 1 time server acts as a primary network time standard.



A stratum 2 server is connected to the stratum 1 server; then a stratum 3 server is connected to the stratum 2 server and so on. A stratum 2 server gets its time via NTP packet requests from a stratum 1 server. A stratum 3 server gets its time via NTP packet requests from a stratum-2 server... A stratum server may also peer with other stratum servers at the same level to provide more stable and robust time for all devices in the peer group (for example a stratum 2 server can peer with other stratum 2 servers).

NTP uses the concept of a stratum to describe how many NTP hops away a machine is from an authoritative time source. A stratum 1 time server typically has an authoritative time source (such as a radio or atomic clock, or a Global Positioning System (GPS) time source) directly attached, a stratum 2 time server receives its time via NTP from a stratum 1 time server, and so on.

Reference: <https://www.cisco.com/c/en/us/td/docs/routers/asr920/configuration/guide/b-sm/16-6-1/b-sm-xe-16-6-1-asr920/b-sm-time-calendar-set.html>

Question 6


How does QoS traffic shaping alleviate network congestion?

- A. It drops packets when traffic exceeds a certain bitrate.
- B. It buffers and queue packets above the committed rate.**
- C. It fragments large packets and queues them for delivery.
- D. It drops packets randomly from lower priority queues.

Answer: B

Explanation

Traffic shaping retains excess packets in a queue and then schedules the excess for later transmission over increments of time. The result of traffic shaping is a smoothed packet output rate.



Question 7

An engineer is describing QoS to a client. Which two facts apply to traffic policing? (Choose two)

- A. Policing adapts to network congestion by queuing excess traffic
- B. Policing should be performed as close to the destination as possible
- C. Policing drops traffic that exceeds the defined rate**
- D. Policing typically delays the traffic, rather than drops it
- E. Policing should be performed as close to the source as possible**

Answer: C E

Explanation

Traffic policing propagates bursts. When the traffic rate reaches the configured maximum rate (or committed information rate), excess traffic is dropped (or remarked). The result is an output rate that appears as a saw-tooth with crests and troughs.

Unlike traffic shaping, traffic policing does not cause delay.

Classification (which includes traffic policing, traffic shaping and queuing techniques) should take place at the network edge. It is recommended that classification occur as close to the source of the traffic as possible.

Also according to this [Cisco link](#), "policing traffic as close to the source as possible".

Question 8

What mechanism does PIM use to forward multicast traffic?

- A. PIM sparse mode uses a pull model to deliver multicast traffic
- B. PIM dense mode uses a pull model to deliver multicast traffic
- C. PIM sparse mode uses receivers to register with the RP
- D. PIM sparse mode uses a flood and prune model to deliver multicast traffic

Answer: A

Explanation

PIM dense mode (PIM-DM) uses a push model to flood multicast traffic to every corner of the network. This push model is a brute-force method of delivering data to the receivers. This method would be efficient in certain deployments in which there are active receivers on every subnet in the network. PIM-DM initially floods multicast traffic throughout the network. Routers that have no downstream neighbors prune the unwanted traffic. This process repeats every 3 minutes.

PIM Sparse Mode (PIM-SM) uses a pull model to deliver multicast traffic. Only network segments with active receivers that have explicitly requested the data receive the traffic. PIM-SM distributes information about active sources by forwarding data packets on the shared tree. Because PIM-SM uses shared trees (at least initially), it requires the use of an RP. The RP must be administratively configured in the network.

Answer C seems to be correct but it is not, PIM spare mode uses sources (not receivers) to register with the RP. Sources register with the RP, and then data is forwarded down the shared tree to the receivers.

Reference: Selecting MPLS VPN Services Book, page 193

Question 9

Which two namespaces does the LISP network architecture and protocol use?
(Choose two)

- A. TLOC
- B. RLOC
- C. DNS
- D. VTEP
- E. EID

Answer: B E

Explanation

Locator ID Separation Protocol (LISP) is a network architecture and protocol that implements the use of two namespaces instead of a single IP address:

- + Endpoint identifiers (EIDs)—assigned to end hosts.
- + Routing locators (RLOCs)—assigned to devices (primarily routers) that make up the global routing system.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_lisp/configuration/xr-3s/irl-xr-3s-book/irl-overview.html

Question 10

Which First Hop Redundancy Protocol should be used to meet a design requirements for more efficient default bandwidth usage across multiple devices?

- A. GLBP
- B. LCAP
- C. HSRP
- D. VRRP


Answer: A

Explanation

The main disadvantage of HSRP and VRRP is that only one gateway is elected to be the active gateway and used to forward traffic whilst the rest are unused until the active one fails. Gateway Load Balancing Protocol (GLBP) is a Cisco proprietary protocol and performs the similar function to HSRP and VRRP but it supports load balancing among members in a GLBP group.

Question 11

Refer to the exhibit.



A network engineer is configuring OSPF between router R1 and router R2. The engineer must ensure that a DR/BDR election does not occur on the Gigabit Ethernet interfaces in area 0. Which configuration set accomplishes this goal?

- A.
R1 (config-if) #interface Gi0/0
R1 (config-if) #ip ospf network point-to-point
R2 (config-if) #interface Gi0/0
R2 (config-if) #ip ospf network point-to-point
- B.
R1 (config-if) #interface Gi0/0
R1 (config-if) #ip ospf network broadcast
R2 (config-if) #interface Gi0/0
R2 (config-if) #ip ospf network broadcast
- C.
R1 (config-if) #interface Gi0/0
R1 (config-if) #ip ospf database-filter all out
R2 (config-if) #interface Gi0/0
R2 (config-if) #ip ospf database-filter all out
- D.
R1 (config-if) #interface Gi0/0
R1 (config-if) #ip ospf priority 1
R2 (config-if) #interface Gi0/0
R2 (config-if) #ip ospf priority 1

Answer: A

Explanation

Broadcast and Non-Broadcast networks elect DR/BDR while Point-to-point/multipoint do not elect DR/BDR. Therefore we have to set the two Gi0/0 interfaces to point-to-point or point-to-multipoint network to ensure that a DR/BDR election does not occur.

Question 12

What are two reasons why broadcast radiation is caused in the virtual machine environment? (Choose two)

- A. vSwitch must interrupt the server CPU to process the broadcast packet
- B. The Layer 2 domain can be large in virtual machine environments**
- C. Virtual machines communicate primarily through broadcast mode**
- D. Communication between vSwitch and network switch is broadcast based
- E. Communication between vSwitch and network switch is multicast based

Answer: B C

Explanation

Broadcast radiation is the accumulation of broadcast and multicast traffic on a computer network. Extreme amounts of broadcast traffic constitute a broadcast storm.

The amount of broadcast traffic you should see within a broadcast domain is directly proportional to the size of the broadcast domain. Therefore if the layer 2 domain in virtual machine environment is too large, broadcast radiation may occur -> VLANs should be used to reduce broadcast radiation.


Also if virtual machines communicate via broadcast too much, broadcast radiation may occur.

Another reason for broadcast radiation is using a trunk (to extend VLANs) from the network switch to the physical server.

Note about the structure of virtualization in a hypervisor:

Hypervisors provide virtual switch (vSwitch) that Virtual Machines (VMs) use to communicate with other VMs on the same host. The vSwitch may also be connected to the host's physical NIC to allow VMs to get layer 2 access to the outside world.

Each VM is provided with a virtual NIC (vNIC) that is connected to the virtual switch. Multiple vNICs can connect to a single vSwitch, allowing VMs on a physical host to communicate with one another at layer 2 without having to go out to a physical switch.



Although vSwitch does not run Spanning-tree protocol but vSwitch implements other loop prevention mechanisms. For example, a frame that enters from one VMNIC is not going to go out of the physical host from a different VMNIC card.

Question 13


A company plans to implement intent-based networking in its campus infrastructure. Which design facilitates a migrate from a traditional campus design to a programmer fabric designer?

- A. Layer 2 access
- B. three-tier
- C. two-tier**
- D. routed access

Answer: C

Explanation

Intent-based Networking (IBN) transforms a hardware-centric, manual network into a controller-led network that captures business intent and translates it into policies that can be automated and applied consistently across the network. The goal is for the network to continuously monitor and adjust network performance to help assure desired business outcomes. IBN builds on software-defined networking (SDN). SDN usually uses spine-leaf architecture, which is typically deployed as two layers: spines (such as an aggregation layer), and leaves (such as an access layer).



Question 14

When a wireless client roams between two different wireless controllers, a network connectivity outage is experienced for a period of time. Which configuration issue would cause this problem?

- A. Not all of the controllers in the mobility group are using the same mobility group name
- B. Not all of the controllers within the mobility group are using the same virtual interface IP address**
- C. All of the controllers within the mobility group are using the same virtual interface IP address
- D. All of the controllers in the mobility group are using the same mobility group name

Answer: B

Explanation


A prerequisite for configuring Mobility Groups is "All controllers must be configured with the same virtual interface IP address". If all the controllers within a mobility group are not using the same virtual interface, inter-controller roaming may appear to work, but the handoff does not complete, and the client loses connectivity for a period of time. -> Answer B is correct.

Reference: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/mobility_groups.html

Answer A is not correct because when the client moves to a different mobility group (with different mobility group name), that client would be connected (provided that the new connected controller had information about this client in its mobility list already) or drop (if the new connected controller has not had information about this client in its mobility list). For more information please read the note below.

Note:

A mobility group is a set of controllers, identified by the same mobility group name that defines the realm of seamless roaming for wireless clients. By creating a mobility group, you can enable multiple controllers in a network to dynamically share information and forward data traffic when inter-controller or inter-subnet roaming occurs. Controllers in the same mobility group can share the context and state of client devices as well as their list of access points so that they do not consider each other's access points as rogue devices.



Let's take an example:

The controllers in the ABC mobility group share access point and client information with each other. The controllers in the ABC mobility group do not share the access point or client information with the XYZ controllers, which are in a different mobility group. Therefore if a client from ABC mobility group moves to XYZ mobility group, and the new connected controller does not have information about this client in its mobility list, that client will be dropped.

Note: Clients may roam between access points in different mobility groups if the controllers are included in each other's mobility lists.

Question 15

Which algorithms are used to secure REST API from brute attacks and minimize the impact?

- A. SHA-512 and SHA-384
- B. MD5 algorithm-128 and SHA-384
- C. SHA-1, SHA-256, and SHA-512
- D. PBKDF2, BCrypt, and SCrypt**

Answer: D

Explanation

One of the best practices to secure REST APIs is using password hash. Passwords must always be hashed to protect the system (or minimize the damage) even if it is compromised in some hacking attempts. There are many such hashing algorithms which can prove really effective for password security e.g. PBKDF2, bcrypt and scrypt algorithms.

Other ways to secure REST APIs are: Always use HTTPS, Never expose information on URLs (Usernames, passwords, session tokens, and API keys should not appear in the URL), Adding Timestamp in Request, Using OAuth, Input Parameter Validation.

Reference: <https://restfulapi.net/security-essentials/>

We should not use MD5 or any SHA (SHA-1, SHA-256, SHA-512...) algorithm to hash password as they are not totally secure.

Note: A brute-force attack is an attempt to discover a password by systematically trying every possible combination of letters, numbers, and symbols until you discover the one correct combination that works.

Question 16

What is the role of the RP in PIM sparse mode?

- A. The RP responds to the PIM join messages with the source of requested multicast group**
- B. The RP maintains default aging timeouts for all multicast streams requested by the receivers
- C. The RP acts as a control-plane node and does not receive or forward multicast packets
- D. The RP is the multicast that is the root of the PIM-SM shared multicast distribution tree

Answer: A

Question 17

A network administrator is preparing a Python script to configure a Cisco IOS XE-based device on the network. The administrator is worried that colleagues will make changes to the device while the script is running. Which operation of the client manager in prevent colleague making changes to the device while the script is running?

- A. m.lock (config='running')
- B. m.lock (target='running')**
- C. m.freeze (target='running')
- D. m.freeze(config='running')

Answer: B

Explanation

The example below shows the usage of lock command:

```
def demo(host, user, names):
```

```
With manager. Connect(host=host, port=22, username=user) as m:
    With m.locked(target='running'):
        for n in names:
            m.edit_config (target='running', config=template % n)
```

The command “m.locked (target='running')” causes a lock to be acquired on the running datastore.

Question 18

What are two device roles in Cisco SD-Access fabric? (Choose two)


- A. core switch
- B. vBond controller
- C. edge node
- D. access switch
- E. border node

Answer: C E

Explanation

There are five basic device roles in the fabric overlay:

- + **Control plane node:** This node contains the settings, protocols, and mapping tables to provide the endpoint-to-location (EID-to-RLOC) mapping system for the fabric overlay.
- + **Fabric border node:** This fabric device (for example, core layer device) connects external Layer 3 networks to the SDA fabric.
- + **Fabric edge node:** This fabric device (for example, access or distribution layer device) connects wired endpoints to the SDA fabric.
- + **Fabric WLAN controller (WLC):** This fabric device connects APs and wireless endpoints to the SDA fabric.
- + **Intermediate nodes:** These are intermediate routers or extended switches that do not provide any sort of SD-Access fabric role other than underlay services.



Question 19

Drag and drop the LISP components from the left onto the function they perform on the right. Not all options are used.

LISP map resolver	accepts LISP encapsulated map requests
LISP proxy ETR	learns of EID prefix mapping entries from an ETR
LISP route reflector	receives traffic from LISP sites and sends it to non-LISP sites
LISP ITR	receives packets from site-facing interfaces
LISP map server	

Answer:

- + **accepts LISP encapsulated map requests:** **LISP map resolver**
- + **learns of EID prefix mapping entries from an ETR:** **LISP map server**
- + **receives traffic from LISP sites and sends it to non-LISP sites:** **LISP proxy ETR**
- + **receives packets from site-facing interfaces:** **LISP ITR**

Explanation

ITR is the function that maps the destination EID to a destination RLOC and then encapsulates the original packet with an additional header that has the source IP address of the ITR RLOC and the destination IP address of the RLOC of an Egress Tunnel Router (ETR). After the encapsulation, the original packet become a LISP packet.

ETR is the function that receives LISP encapsulated packets, decapsulates them and forwards to its local EIDs. This function also requires EID-to-RLOC mappings so we need to point out an "map-server" IP address and the key (password) for authentication.

A LISP proxy ETR (PETR) implements ETR functions on behalf of non-LISP sites. A PETR is typically used when a LISP site needs to send traffic to non-LISP sites but the LISP site is connected through a service provider that does not accept no routable EIDs as packet sources. PETRs act just like ETRs but for EIDs that send traffic to destinations at non-LISP sites.

Map Server (MS) processes the registration of authentication keys and EID-to-RLOC mappings. ETRs sends periodic Map-Register messages to all its configured Map Servers.

Map Resolver (MR): a LISP component which accepts LISP Encapsulated Map Requests, typically from an ITR, quickly determines whether or not the destination IP address is part of the EID namespace

Question 20

Drag and Drop the descriptions from the left onto the routing protocol they describe on the right.

summaries can be created anywhere in the IGP topology	OSPF
uses areas to segment a network	
DUAL algorithm	
summarizes can be created in specific parts of the IGP topology	EIGRP

Answer:

OSPF:

- + uses areas to segment a network
- + summarizes can be created in specific parts of the IGP topology

EIGRP:

- + summaries can be created anywhere in the IGP topology
- + DUAL algorithm

Explanation

Unlike OSPF where we can summarize only on ABR or ASBR, in EIGRP we can summarize anywhere. Manual summarization can be applied anywhere in EIGRP domain, on every router, on every interface via the `ip summary-address eigrp as-number address mask [administrative-distance]` command (for example: `ip summary-address eigrp 1 192.168.16.0 255.255.248.0`). Summary route will exist in routing table as long as at least one more specific route will exist. If the last specific route will disappear, summary route also will fade out. The metric used by EIGRP manual summary route is the minimum metric of the specific routes.

Question 21

Which component handles the orchestration plane of the Cisco SD-WAN?

- A. vBond
- B. vSmart
- C. vManage
- D. vEdge

Answer: A

Explanation

+ Orchestration plane (vBond) assists in securely onboarding the SD-WAN WAN Edge routers into the SD-WAN overlay. The vBond controller, or orchestrator, authenticates and authorizes the SD-WAN components onto the network. The vBond orchestrator takes an added responsibility to distribute the list of vSmart and vManage controller information to the WAN Edge routers. vBond is the only device in SD-WAN that requires a public IP address as it is the first point of contact and authentication for all SD-WAN components to join the SD-WAN fabric. All other components need to know the vBond IP or DNS information.

Question 22

Which two entities are Type 1 hypervisors? (Choose two)

- A. Oracle VM Virtual Box
- B. Microsoft Hyper-V
- C. VMware server
- D. VMware ESX
- E. Microsoft Virtual PC

Answer: B D

Explanation

A bare-metal hypervisor (Type 1) is a layer of software we install directly on top of a physical server and its underlying hardware. There is no software or any operating system in between, hence the name bare-metal hypervisor. A Type 1 hypervisor is proven in providing excellent performance and stability since it does not run inside Windows or any other operating system. These are the most common type 1 hypervisors:

- + VMware vSphere with ESX/ESXi
- + KVM (Kernel-Based Virtual Machine)
- + Microsoft Hyper-V
- + Oracle VM
- + Citrix Hypervisor (formerly known as Xen Server)

Question 23

Which access point mode allows a supported AP to function like a WLAN client would, associating and identifying client connectivity issues?

- A. client mode
- B. SE-connect mode
- C. sensor mode
- D. sniffer mode**

Answer: D


Explanation

An lightweight AP (LAP) operates in one of six different modes:

- + Local mode (default mode): measures noise floor and interference, and scans for intrusion detection (IDS) events every 180 seconds on unused channels
- + Flex Connect, formerly known as Hybrid Remote Edge AP (H-REAP), mode: allows data traffic to be switched locally and not go back to the controller. The Flex Connect AP can perform standalone client authentication and switch VLAN traffic locally even when it's disconnected to the WLC (Local Switched). Flex Connect AP can also tunnel (via CAPWAP) both user wireless data and control traffic to a centralized WLC (Central Switched).
- + Monitor mode: does not handle data traffic between clients and the infrastructure. It acts like a sensor for location-based services (LBS), rogue AP detection, and IDS
- + Rogue detector mode: monitor for rogue APs. It does not handle data at all.
- + Sniffer mode: run as a sniffer and captures and forwards all the packets on a particular channel to a remote machine where you can use protocol analysis tool (Wireshark, Airopoek, etc) to review the packets and diagnose issues. Strictly used for troubleshooting purposes.
- + Bridge mode: bridge together the WLAN and the wired infrastructure together.

Question 24

Refer to the exhibit.



An engineer must deny Telnet traffic from the loopback interface of router R3 to the loopback interface of router R2 during the weekend hours. All other traffic between the loopback interfaces of routers R3 and R2 must be allowed at all times. Which command accomplish this task?

A.

```
R3(config)#time-range WEEKEND
```

```
R3(config-time-range)#periodic Saturday Sunday 00:00 to 23:59
```

```
R3(config)#access-list 150 deny tcp host 10.3.3.3 host 10.2.2.2 eq 23 time-range WEEKEND
```

```
R3(config)#access-list 150 permit ip any any time-range WEEKEND
```

```
R3(config)#interface Gi0/1
```

```
R3(config-if)#ip access-group 150 out
```

B.

```
R1(config)#time-range WEEKEND
```

```
R1(config-time-range)#periodic Friday Sunday 00:00 to 00:00
```

```
R1(config)#access-list 150 deny tcp host 10.3.3.3 host 10.2.2.2 eq 23 time-range WEEKEND
```

```
R1(config)#access-list 150 permit ip any any
```

```
R1(config)#interface Gi0/1
```

```
R1(config-if)#ip access-group 150 in
```

C.

```
R1(config)#time-range WEEKEND
```

```
R1(config-time-range)#periodic weekend 00:00 to 23:59
```

R1(config)#access-list 150 deny tcp host 10.3.3.3 host 10.2.2.2 eq 23 time-range WEEKEND

R1(config)#access-list 150 permit ip any any

R1(config)#interface Gi0/1

R1(config-if)#ip access-group 150 in

D.

R3(config)#time-range WEEKEND

R3(config-time-range)#periodic weekend 00:00 to 23:59

R3(config)#access-list 150 deny tcp host 10.3.3.3 host 10.2.2.2 eq 23 time-range WEEKEND

R3(config)#access-list 150 permit ip any any time-range WEEKEND

R3(config)#interface Gi0/1

R3(config-if)#ip access-group 150 out

Answer: C

Explanation

We cannot filter traffic that is originated from the local router (R3 in this case) so we can only configure the ACL on R1 or R2. "Weekend hours" means from Saturday morning through Sunday night so we have to configure: "periodic weekend 00:00 to 23:59".

Note: The time is specified in 24-hour time (hh:mm), where the hours range from 0 to 23 and the minutes range from 0 to 59.

Question 25

Which tool is used in Cisco DNA Center to build generic configurations that are able to be applied on device with similar network settings?

A. Command Runner

B. Template Editor

C. Application Policies

D. Authentication Template

Answer: B

Explanation

Cisco DNA Center provides an interactive editor called Template Editor to author CLI templates. Template Editor is a centralized CLI management tool to help design a set of device configurations that you need to build devices in a branch. When you have a site, office, or branch that uses a similar set of devices and configurations, you can use Template Editor to build generic configurations and apply the configurations to one or more devices in the branch.

Reference: https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-3/user_guide/b_cisco_dna_center_ug_1_3/b_cisco_dna_center_ug_1_3_chapter_0111.html

Question 26

A client device roams between access points located on different floors in an atrium. The access points joined to the same controller and configuration in local mode. The access points are in different IP addresses, but the client VLAN in the group same. What type of roam occurs?

- A. inter-controller
- B. inter-subnet**
- C. intra-VLAN
- D. intra-controller

Answer: B

Explanation

Mobility, or roaming, is a wireless LAN client's ability to maintain its association seamlessly from one access point to another securely and with as little latency as possible. Three popular types of client roaming are:

Intra-Controller Roaming: Each controller supports same-controller client roaming across access points managed by the same controller. This roaming is transparent to the client as the session is sustained, and the client continues using the same DHCP-assigned or client-assigned IP address.

Inter-Controller Roaming: Multiple-controller deployments support client roaming across access points managed by controllers in the same mobility group and on the same subnet. This roaming is also transparent to the client because the session is sustained and a tunnel between controllers allows the client to continue using the same DHCP- or client-assigned IP address as long as the session remains active.

Inter-Subnet Roaming: Multiple-controller deployments support client roaming across access points managed by controllers in the same mobility group on different subnets. This roaming is transparent to the client because the session is sustained and a tunnel between the controllers allows the client to continue using the same DHCP-assigned or client-assigned IP address as long as the session remains active.

Reference: https://www.cisco.com/c/en/us/td/docs/wireless/controller/7-4/configuration/guides/consolidated/b_cg74_CONSOLIDATED/b_cg74_CONSOLIDATED_chapter_01100.html

In three types of client roaming above, only with Inter-Subnet Roaming the controllers are in different subnets.

Question 27

What does the LAP send when multiple WLCs respond to the CISCO_CAPWAP-CONTROLLER.localdomain hostname during the CAPWAP discovery and join process?

- A. broadcast discover request
- B. join request to all the WLCs
- C. unicast discovery request to each WLC
- D. Unicast discovery request to the first WLC that resolves the domain name**

Answer: D

Question 28

Refer to the exhibit.

```
vlan 222
  remote-span
!
vlan 223
  remote-span
!
monitor session 1 source interface FastEthernet0/1 tx
monitor session 1 source interface FastEthernet0/2 rx
monitor session 1 source interface port-channel 5
monitor session 1 destination remote vlan 222
!
```

What is the result when a technician adds the monitor session 1 destination remote vlan 233 command?

- A. The RSPAN VLAN is replaced by VLAN 223**
- B. RSPAN traffic is sent to VLANs 222 and 223
- C. An error is flagged for configuring two destinations
- D. RSPAN traffic is split between VLANs 222 and 223

Answer: A

Question 29


In an SD-Access solution what is the role of a fabric edge node?

- A. to connect external Layer 3- network to the SD-Access fabric
- B. to connect wired endpoint to the SD-Access fabric**
- C. to advertise fabric IP address space to external network
- D. to connect the fusion router to the SD-Access fabric

Answer: B

Explanation

+ Fabric edge node: This fabric device (for example, access or distribution layer device) connects wired endpoints to the SDA fabric.



Question 30

Refer to the exhibit.

```
access-list 1 permit 172.16.1.0 0.0.0.255
ip nat inside source list 1 interface gigabitethernet0/0 overload
```

The inside and outside interfaces in the NAT configuration of this device have been correctly identified. What is the effect of this configuration?

- A. dynamic NAT
- B. static NAT
- C. PAT
- D. NAT64

Answer: C

Explanation

The command "ip nat inside source list 1 interface gigabitethernet0/0 overload" translates all source addresses that pass access list 1, which means 172.16.1.0/24 subnet, into an address assigned to gigabitethernet0/0 interface. Overload keyword allows to map multiple IP addresses to a single registered IP address (many-to-one) by using different ports so it is called Port Address Translation (PAT).

Question 31

Which component of the Cisco Cyber Threat Defense solution provides user and flow context analysis?

- A. Cisco Firepower and FireSIGHT
- B. Cisco Stealth watch system**
- C. Advanced Malware Protection
- D. Cisco Web Security Appliance

Answer: B

Explanation

The goal of the Cyber Threat Defense solution is to introduce a design and architecture that can help facilitate the discovery, containment, and remediation of threats once they have penetrated into the network interior.

Cisco Cyber Threat Defense version 2.0 makes use of several solutions to accomplish its objectives:

- * NetFlow and the Lancope StealthWatch System
 - Broad visibility
 - User and flow context analysis
 - Network behavior and anomaly detection
 - Incident response and network forensics
- * Cisco FirePOWER and FireSIGHT
 - Real-time threat management
 - Deeper contextual visibility for threats bypassing the perimeters
 - URL control
- * Advanced Malware Protection (AMP)
 - Endpoint control with AMP for Endpoints
 - Malware control with AMP for networks and content
- * Content Security Appliances and Services
 - Cisco Web Security Appliance (WSA) and Cloud Web Security (CWS)
 - Dynamic threat control for web traffic
 - Outbound URL analysis and data transfer controls
 - Detection of suspicious web activity

- Cisco Email Security Appliance (ESA)
- Dynamic threat control for email traffic
- Detection of suspicious email activity
- * Cisco Identity Services Engine (ISE)
 - User and device identity integration with Lancope StealthWatch
 - Remediation policy actions using pxGrid

Reference: https://www.cisco.com/c/dam/en/us/td/docs/security/network_security/ctd/ctd2-0/design_guides/ctd_2-0_cvd_guide_jul15.pdf

Question 32

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

- A. Use Cisco AMP deployment with the Malicious Activity Protection engine enabled**
- B. Use Cisco AMP deployment with the Exploit Prevention engine enabled
- C. Use Cisco Firepower and block traffic to TOR networks
- D. Use Cisco Firepower with Intrusion Policy and snort rules blocking SMB exploitation

Answer: A

Explanation

Ransomware are malicious software that locks up critical resources of the users. Ransomware uses well-established public/private key cryptography which leaves the only way of recovering the files being the payment of the ransom, or restoring files from backups.

Cisco Advanced Malware Protection (AMP) for Endpoints Malicious Activity Protection (MAP) engine defends your endpoints by monitoring the system and identifying processes that exhibit malicious activities when they execute and stops them from running. Because the MAP engine detects threats by observing the behavior of the process at run time, it can generically determine if a system is under attack by a new variant of ransomware or malware that may have eluded other security products and detection technology, such as legacy signature-based malware detection. The first release of the MAP engine targets identification, blocking, and quarantine of ransomware attacks on the endpoint.

Reference: <https://www.cisco.com/c/dam/en/us/products/collateral/security/amp-for-endpoints/white-paper-c11-740980.pdf>

Question 33

Refer to the exhibit.

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 AMP deployment with the Malicious Activity Protection engine enabled**
- B. Use Cisco AMP deployment with the Exploit Prevention engine enabled
- C. Use Cisco Firepower and block traffic to TOR networks
- D. Use Cisco Firepower with Intrusion Policy and snort rules blocking SMB exploitation

Answer: A

Explanation

Ransomware are malicious software that locks up critical resources of the users. Ransomware uses well-established public/private key cryptography which leaves the only way of recovering the files being the payment of the ransom, or restoring files from backups.

Cisco Advanced Malware Protection (AMP) for Endpoints Malicious Activity Protection (MAP) engine defends your endpoints by monitoring the system and identifying processes that exhibit malicious activities when they execute and stops them from running. Because the MAP engine detects threats by observing the behavior of the process at run time, it can generically determine if a system is under attack by a new variant of ransomware or malware that may have eluded other security products and detection technology, such as legacy signature-based malware detection. The first release of the MAP engine targets identification, blocking, and quarantine of ransomware attacks on the endpoint.

Reference: <https://www.cisco.com/c/dam/en/us/products/collateral/security/amp-for-endpoints/white-paper-c11-740980.pdf>

Question 33

Refer to the exhibit.

WLANs > Edit 'LiveDemo'

General Security QoS Policy-Mapping Advanced

Layer 2 Layer 3 AAA Servers

Select AAA servers below to override use of default servers on this WLAN

Radius Servers

Radius Server Overwrite interface ☒ Enabled

Interface Priority WLAN

	Authentication Servers	Accounting Servers
	☒ Enabled	☒ Enabled
Server 1	None	None
Server 2	None	None
Server 3	None	None
Server 4	None	None
Server 5	None	None
Server 6	None	None

Assuming the WLC's interfaces are not in the same subnet as the RADIUS server, which interface would the WLC use as the source for all RADIUS-related traffic?

- A. the interface specified on the WLAN configuration
- B. any interface configured on the WLC
- C. the controller management interface
- D. the controller virtual interface

Answer: A

Question 34

Which benefit is offered by a cloud infrastructure deployment but is lacking in an on-premises deployment?

- A. efficient scalability
- B. virtualization
- C. storage capacity
- D. supported systems

Answer: A

Question 35

Wireless users report frequent disconnections from the wireless network. While troubleshooting a network engineer finds that after the user a disconnect, the connection reestablishes automatically without any input required. The engineer also notices these message logs.

```
AP 'AP2' is down Reason: Radio channel set. 6:54:04 PM
AP 'AP4' is down Reason: Radio channel set. 6:44:49 PM
AP 'AP7' is down Reason: Radio channel set. 6:34:32 PM
```

Which action reduces the user impact?

- A. increase the dynamic channel assignment interval
- B. increase BandSelect
- C. increase the AP heartbeat timeout
- D. enable coverage hole detection

Answer: A

Explanation

These message logs inform that the radio channel has been reset (and the AP must be down briefly). With dynamic channel assignment (DCA), the radios can frequently switch from one channel to another but it also makes disruption. The default DCA interval is 10 minutes, which is matched with the time of the message logs. By increasing the DCA interval, we can reduce the number of times our users are disconnected for changing radio channels.

Question 36

Which DHCP option helps lightweight APs find the IP address of a wireless LAN controller?

- A. Option 43
- B. Option 60
- C. Option 67
- D. Option 150

Answer: A

Question 37

A network administrator applies the following configuration to an IOS device.

```
aaa new-model
aaa authentication login default local group tacacs+
```

What is the process of password checks when a login attempt is made to the device?

- A. A TACACS+ server is checked first. If that check fail, a database is checked
- B. A TACACS+ server is checked first. If that check fail, a RADIUS server is checked. If that check fail, a local database is checked
- C. A local database is checked first. If that fails, a TACACS+server is checked, if that check fails, a RADIUS server is checked
- D. A local database is checked first. If that check fails, a TACACS+server is checked**

Answer: D

Explanation

The "aaa authentication login default local group tacacs+" command is broken down as follows:

- + The 'aaa authentication' part is simply saying we want to configure authentication settings.
- + The 'login' is stating that we want to prompt for a username/password when a connection is made to the device.
- + The 'default' means we want to apply for all login connections (such as tty, vty, console and aux). If we use this keyword, we don't need to configure anything else under tty, vty and aux lines. If we don't use this keyword then we have to specify which line(s) we want to apply the authentication feature.
- + The 'local group tacacs+' means all users are authenticated using router's local database (the first method). If the credentials are not found on the local database, then the TACACS+ server is used (the second method).

Question 38

What is the role of the vsmart controller in a Cisco SD-WAN environment?

- A. IT performs authentication and authorization
- B. It manages the control plane.**
- C. It is the centralized network management system.
- D. It manages the data plane.

Answer: B

Explanation

- + Control plane (vSmart) builds and maintains the network topology and make decisions on the traffic flows. The vSmart controller disseminates control plane information between WAN Edge devices, implements control plane policies and distributes data plane policies to network devices for enforcement.

Question 39

Why is an AP joining a different WLC than the one specified through option 43?

- A. The WLC is running a different software version
- B. The API is joining a primed WLC**
- C. The AP multicast traffic unable to reach the WLC through Layer 3
- D. The APs broadcast traffic is unable to reach the WLC through Layer 2

Answer: B

Question 40

Which devices does Cisco Center configure when deploying an IP-based access control policy?

- A. All devices integrating with ISE**
- B. selected individual devices
- C. all devices in selected sites
- D. all wired devices

Answer: A

Explanation

When you click Deploy, Cisco DNA Center requests the Cisco Identity Services Engine (Cisco ISE) to send notifications about the policy changes to the network devices.

Reference: https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-3-1-0/user_guide/b_cisco_dna_center_ug_1_3_1_0/b_cisco_dna_center_ug_1_3_1_0_chapter_01011.html

Question 41

Which method of account authentication does OAuth 2.0 within REST APIs?

- A. username/role combination
- B. access tokens**
- C. cookie authentication
- D. basic signature workflow

Answer: B

Explanation

The most common implementations of OAuth (OAuth 2.0) use one or both of these tokens:

- + access token: sent like an API key, it allows the application to access a user's data; optionally, access tokens can expire.
- + refresh token: optionally part of an OAuth flow, refresh tokens retrieve a new access token if they have expired. OAuth2 combines Authentication and Authorization to allow more sophisticated scope and validity control.

Question 42

What does the Cisco DNA Center use to enable the delivery of applications through a network and to yield analytics for innovation?

- A. process adapters
- B. Command Runner
- C. intent-based APIs**
- D. domain adapters

Answer: C

Explanation

The Cisco DNA Center open platform for intent-based networking provides 360-degree extensibility across multiple components, including:

+ Intent-based APIs leverage the controller to enable business and IT applications to deliver intent to the network and to reap network analytics and insights for IT and business innovation. These enable APIs that allow Cisco DNA Center to receive input from a variety of sources, both internal to IT and from line-of-business applications, related to application policy, provisioning, software image management, and assurance.

...

Reference: <https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/dna-center/nb-06-dna-cent-plat-sol-over-cte-en.html>

Question 43

Which action is a function of VTEP in VXLAN?


- A. tunneling traffic from IPv6 to IPv4 VXLANs
- B. allowing encrypted communication on the local VXLAN Ethernet segment
- C. encapsulating and de-encapsulating VXLAN Ethernet frames**
- D. tunneling traffic from IPv4 to IPv6 VXLANs

Answer: C

Explanation


VTEPs connect between Overlay and Underlay network and they are responsible for encapsulating frame into VXLAN packets to send across IP network (Underlay) then decapsulating when the packets leaves the VXLAN tunnel.

VTEPs connect between Overlay and Underlay network and they are responsible for encapsulating frame into VXLAN packets to send across IP network (Underlay) then decapsulating when the packets leaves the VXLAN tunnel.



Question 44

Which type of antenna does the radiation pattern represent?



- A. Yagi
- B. multidirectional
- C. directional patch
- D. omnidirectional

Answer: A

Explanation

A Yagi antenna is formed by driving a simple antenna, typically a dipole or dipole-like antenna, and shaping the beam using a well-chosen series of non-driven elements whose length and spacing are tightly controlled.



Reference: https://www.cisco.com/c/en/us/products/collateral/wireless/aironet-antennas-accessories/prod_white_paper0900aecd806a1a3e.htm

Etherchannel Quiz

Question 1

Which PAgP mode combination prevents an Etherchannel from forming?

- A. desirable
- B. auto/desirable
- C. auto/auto correct
- D. desirable/desirable

Explanation

There are two PAgP modes:


Auto	Responds to PAgP messages but does not aggressively negotiate a PAgP EtherChannel. Answer 'auto/auto' channel is formed only if the port on the other end is set to Desirable. This is the default mode.
Desirable	Port actively negotiates channeling status with the interface on the other end of the link. Answer 'auto/auto' channel is formed if the other side is Auto or Desirable.

The table below lists if an EtherChannel will be formed or not for PAgP:

PAgP	Desirable	Auto
Desirable	Yes	Yes
Auto	Yes	No

Question 2

Refer to the exhibit. A port channel is configured between SW2 and SW3. SW2 is not running Cisco operating system. When all physical connections are made, the port channel does not establish. Based on the configuration excerpt of SW3, what is the cause of the problem?



```
interface gi1/2
channel-group 30 mode desirable
port-channel load-balance src-ip
```

```
interface gi1/3
channel-group 30 mode desirable
port-channel load-balance src-ip
```

```
interface PortChannel 30
switchport mode trunk
switchport encapsulation dot1q
switchport trunk allowed vlan 10-100
```

- A. The port-channel interface load balance should be set to src-mac
- B. The port channel on SW2 is using an incompatible protocol** correct
- C. The port-channel trunk is not allowing the native VLAN
- D. The port-channel should be set to auto

Explanation

The Cisco switch was configured with PAgP, which is a Cisco proprietary protocol so non-Cisco switch could not communicate.

Trunking Quiz

Question 1

Refer to exhibit. VLANs 50 and 60 exist on the trunk links between all switches. All access ports on SW3 are configured for VLAN 50 and SW1 is the VTP server. Which command ensures that SW3 receives frames only from VLAN 50?



- A. SW1 (config)#vtp pruning** correct
- B. SW3(config)#vtp mode transparent
- C. SW2(config)#vtp pruning
- D. SW1(config)>vtp mode transparent

Explanation

SW3 does not have VLAN 60 so it should not receive traffic for this VLAN (sent from SW2). Therefore we should configure VTP Pruning on SW3 so that SW2 does not forward VLAN 60 traffic to SW3. Also notice that we need to configure pruning on SW1 (the VTP Server), not SW2.

Question 2

Refer to the exhibit. SwitchC connects HR and Sales to the Core switch. However, business needs require that no traffic from the Finance VLAN traverse this switch. Which command meets this requirement?

```
SwitchC#show vtp status
```

```
VTP Version                : 2
Configuration Revision      : 0
Maximum VLANs supported locally : 255
Number of existing VLANs    : 8
VTP Operating Mode          : Transparent
VTP Domain Name              : MyDomain.com
VTP Pruning Mode             : Disabled
VTP V2 Mode                  : Disabled
VTP Traps Generation        : Disabled
MD5 digest                   : 0xCC 0x77 0x02 0x40 0x93 0xB5 0xC1 0xA2
Configuration last modified by 0.0.0.0 at 3-1-93 00:00:00
```

```
SwitchC#show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Fa0/3, Fa0/4, Fa0/5, Fa0/6

Fa0/7, Fa0/8, Fa0/9,
Fa0/10
Fa0/11, Fa0/12, Fa0/13,
Fa0/14
Fa0/15, Fa0/16, Fa0/17,
Fa0/18
Fa0/19, Fa0/20, Fa0/21,
Fa0/22
Fa0/23, Fa0/24, Po1

110	Finance	active	
210	HR	active	Fa0/1
310	Sales	active	Fa0/2

```
SwitchC#show int trunk
```

Port	Mode	Encapsulation	Status	Native vlan
Gig1/1	on	802.1q	trunking	1
Gig1/2	on	802.1q	trunking	1

Port	Vlans allowed on trunk
Gig1/1	1-1005
Gig1/2	1-1005

Port	Vlans allowed and active in management domain
Gig1/1	1,110,210,310
Gig1/2	1,110,210,310

```
SwitchC#show run interface port-channel 1
```

```
interface Port-channel 1
```

```
description Uplink_to_Core
```

```
switchport mode trunk
```

A. SwitchC(config)#vtp pruning vlan 110

B. SwitchC(config)#vtp pruning

C. SwitchC(config)#interface port-channel 1

SwitchC(config-if)#switchport trunk allowed vlan add 210,310

D. SwitchC(config)#interface port-channel 1

SwitchC(config-if)#switchport trunk allowed vlan remove 110 correct

Explanation

From the "show vlan brief" we learn that Finance belongs to VLAN 110 and all VLANs (from 1 to 1005) are allowed to traverse the trunk (port-channel 1). Therefore we have to remove VLAN 110 from the allowed VLAN list with the "switchport trunk allowed vlan remove" command. The pruning feature cannot do this job as Finance VLAN is active.

SD-WAN & SD-Access Solutions

Your answers are shown below:

Question 1

Which function does a fabric edge node perform in an SD-Access deployment?

A. Encapsulates end-user data traffic into LISP.

B. Provides reachability border nodes in the fabric underlay


C. Connects endpoints to the fabric and forwards their traffic correct

D. Connects the SD-Access fabric to another fabric or external Layer 3 networks

Explanation

There are five basic device roles in the fabric overlay:

- + Control plane node: This node contains the settings, protocols, and mapping tables to provide the endpoint-to-location (EID-to-RLOC) mapping system for the fabric overlay.
- + Fabric border node: This fabric device (for example, core layer device) connects external Layer 3 networks to the SDA fabric.
- + Fabric edge node: This fabric device (for example, access or distribution layer device) connects wired endpoints to the SDA fabric.
- + Fabric WLAN controller (WLC): This fabric device connects APs and wireless endpoints to the SDA fabric.
- + Intermediate nodes: These are intermediate routers or extended switches that do not provide any sort of SD-Access fabric role other than underlay services.



Reference: CCNP and CCIE Enterprise Core ENCOR 350-401 Official Cert Guide

Question 2

On which protocol or technology is the fabric data plane based in Cisco SD-Access fabric?

- A. LISP
- B. Cisco TrustSec
- C. VXLAN correct**
- D. IS-IS

Explanation

The tunneling technology used for the fabric data plane is based on Virtual Extensible LAN (VXLAN). VXLAN encapsulation is UDP based, meaning that it can be forwarded by any IP-based

network (legacy or third party) and creates the overlay network for the SD-Access fabric. Although LISP is the control plane for the SD-Access fabric, it does not use LISP data encapsulation for the data plane; instead, it uses VXLAN encapsulation because it is capable of encapsulating the original Ethernet header to perform MAC-in-IP encapsulation, while LISP does not. Using VXLAN allows the SD-Access fabric to support Layer 2 and Layer 3 virtual topologies (overlays) and the ability to operate over any IP-based network with built-in network segmentation (VRF instance/VN) and built-in group-based policy.

Reference: CCNP and CCIE Enterprise Core ENCOR 350-401 Official Cert Guide

Question 3

Which statement about a fabric access point is true?

- A. It is in local mode and must be connected directly to the fabric edge switch** correct
- B. It is in FlexConnect mode and must be connected directly to the fabric edge switch
- C. It is in local mode and must be connected directly to the fabric border node
- D. It is in FlexConnect mode and must be connected directly to the fabric border node

Explanation

Fabric mode APs continue to support the same wireless media services that traditional APs support; apply AVC, quality of service (QoS), and other wireless policies; and establish the CAPWAP control plane to the fabric WLC. **Fabric APs join as local-mode APs and must be directly connected to the fabric edge node switch** to enable fabric registration events, including RLOC assignment via the fabric WLC. The fabric edge nodes use CDP to recognize APs as special wired hosts, applying special port configurations and assigning the APs to a unique overlay network within a common EID space across a fabric. The assignment allows management simplification by using a single subnet to cover the AP infrastructure at a fabric site.

Reference: <https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/sda-sdg-2019oct.html>

Question 4

Which controller is the single plane of management for Cisco SD-WAN?

- A. vBond
- B. vSmart
- C. vEdge
- D. vManage** correct

Explanation

The primary components for the Cisco SD-WAN solution consist of the **vManage network management system (management plane)**, the vSmart controller (control plane), the vBond orchestrator (orchestration plane), and the vEdge router (data plane).

+ vManage – This centralized network management system provides a GUI interface to easily monitor, configure, and maintain all Cisco SD-WAN devices and links in the underlay and overlay network.

+ vSmart controller – This software-based component is responsible for the centralized control plane of the SD-WAN network. It establishes a secure connection to each vEdge router and distributes routes and policy information via the Overlay Management Protocol (OMP), acting as a route reflector. It also orchestrates the secure data plane connectivity between the vEdge routers by distributing crypto key information, allowing for a very scalable, IKE-less architecture.

+ vBond orchestrator – This software-based component performs the initial authentication of vEdge devices and orchestrates vSmart and vEdge connectivity. It also has an important role in enabling the communication of devices that sit behind Network Address Translation (NAT).

+ vEdge router – This device, available as either a hardware appliance or software-based router, sits at a physical site or in the cloud and provides secure data plane connectivity among the sites

over one or more WAN transports. It is responsible for traffic forwarding, security, encryption, Quality of Service (QoS), routing protocols such as Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF), and more.

Reference: <https://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/SDWAN/CVD-SD-WAN-Design-2018OCT.pdf>

Question 5

Which action is the vSmart controller responsible for in an SD-WAN deployment?

- A. gather telemetry data from vEdge routers
- B. onboard vEdge nodes into the SD-WAN fabric
- C. manage, maintain, and gather configuration and status for nodes within the SD-WAN fabric
- D. distribute security information for tunnel establishment between vEdge routers** correct

Explanation

+ **Orchestration plane** (vBond) assists in securely onboarding the SD-WAN WAN Edge routers into the SD-WAN overlay (-> Therefore answer "onboard vEdge nodes into the SD-WAN fabric" mentioned about vBond). The vBond controller, or orchestrator, authenticates and authorizes the SD-WAN components onto the network. The vBond orchestrator takes an added responsibility to distribute the list of vSmart and vManage controller information to the WAN Edge routers. vBond is the only device in SD-WAN that requires a public IP address as it is the first point of contact and authentication for all SD-WAN components to join the SD-WAN fabric. All other components need to know the vBond IP or DNS information.

+ **Management plane** (vManage) is responsible for central configuration and monitoring. The vManage controller is the centralized network management system that provides a single pane of glass GUI interface to easily deploy, configure, monitor and troubleshoot all Cisco SD-WAN components in the network. (-> Answer "manage, maintain, and gather configuration and status for nodes within the SD-WAN fabric" and answer "gather telemetry data from vEdge routers" are about vManage)

+ **Control plane** (vSmart) builds and maintains the network topology and make decisions on the traffic flows. The vSmart controller disseminates control plane information between WAN Edge devices, implements control plane policies and distributes data plane policies to network devices for enforcement (-> Answer "distribute security information for tunnel establishment between vEdge routers" is about vSmart)

Question 6

Which statement about a Cisco APIC controller versus a more traditional SDN controller is true?

- A. APIC does support a Southbound REST API
- B. APIC supports OpFlex as a Northbound protocol
- C. APIC uses a policy agent to translate policies into instructions** correct
- D. APIC uses an imperative model

Explanation

The southbound protocol used by APIC is OpFlex that is pushed by Cisco as the protocol for policy enablement across physical and virtual switches.

Southbound interfaces are implemented with some called Service Abstraction Layer (SAL), which talks to the network elements via SNMP and CLI.

Note: Cisco OpFlex is a southbound protocol in a software-defined network (SDN).

Question 7

Which description of an SD-Access wireless network infrastructure deployment is true?

- A. The access point is part of the fabric underlay
- B. The WLC is part of the fabric underlay
- C. The access point is part the fabric overlay** correct
- D. The wireless client is part of the fabric overlay

Explanation

Access Points

- + AP is directly connected to FE (or to an extended node switch)
- + AP is part of Fabric overlay

Reference: <https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2018/pdf/BRKEWN-2020.pdf>

Question 8

What the role of a fusion in an SD-Access solution?

A. performs route leaking between user-defined virtual networks and shared services correct

- B. provides additional forwarding capacity to the fabric
- C. provides connectivity to external networks
- D. acts as a DNS server

Explanation

Today the Dynamic Network Architecture Software Defined Access (DNA-SDA) solution requires a fusion router to perform VRF route leaking between user VRFs and Shared-Services, which may be in the Global routing table (GRT) or another VRF. Shared Services may consist of DHCP, Domain Name System (DNS), Network Time Protocol (NTP), Wireless LAN Controller (WLC), Identity Services Engine (ISE), DNAC components which must be made available to other virtual networks (VN's) in the Campus.

Reference: <https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/dna-center/213525-sda-steps-to-configure-fusion-router.html>

QoS Quiz

Question 1

Which QoS mechanism will prevent a decrease in TCP performance?

- A. Shaper
- B. Rate-Limit
- C. Policer
- D. Fair-Queue
- E. WRED** correct
- F. LLQ

Explanation

Weighted Random Early Detection (WRED) is just a congestion avoidance mechanism. WRED drops packets selectively based on IP precedence. Edge routers assign IP precedences to packets as they enter the network. When a packet arrives, the following events occur:

The average queue size is calculated.

2. If the average is less than the minimum queue threshold, the arriving packet is queued.
3. If the average is between the minimum queue threshold for that type of traffic and the maximum threshold for the interface, the packet is either dropped or queued, depending on the packet drop probability for that type of traffic.
4. If the average queue size is greater than the maximum threshold, the packet is dropped.

WRED reduces the chances of tail drop (when the queue is full, the packet is dropped) by selectively dropping packets when the output interface begins to show signs of congestion (thus it can mitigate congestion by preventing the queue from filling up). By dropping some packets early rather than waiting until the queue is full, WRED avoids dropping large numbers of packets at once and minimizes the chances of global synchronization. Thus, WRED allows the transmission line to be used fully at all times.


WRED generally drops packets selectively based on IP precedence. Packets with a higher IP precedence are less likely to be dropped than packets with a lower precedence. Thus, the higher the priority of a packet, the higher the probability that the packet will be delivered.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_conavd/configuration/15-mt/qos-conavd-15-mt-book/qos-conavd-cfg-wred.html

WRED is only useful when the bulk of the traffic is TCP/IP traffic. With TCP, dropped packets indicate congestion, so the packet source will reduce its transmission rate. With other protocols, packet sources may not respond or may resend dropped packets at the same rate. Thus, dropping packets does not decrease congestion.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_conavd/configuration/xe-16/qos-conavd-xe-16-book/qos-conavd-oview.html

Note: Global synchronization occurs when multiple TCP hosts reduce their transmission rates in response to congestion. But when congestion is reduced, TCP hosts try to increase their transmission rates again simultaneously (known as slow-start algorithm), which causes another congestion. Global synchronization produces this graph:



Question 2

Which statement about the default QoS configuration on a Cisco switch is true?

- A. The Cos value of each tagged packet is modified
- B. Port trust is enabled

C. The Port Cos value is 0 correct

D. All traffic is sent through four egress queues

Question 3

Which QoS component alters a packet to change the way that traffic is treated in the network?

- A. Shaping
- B. Classification
- C. Marking correct**
- D. Policing

Explanation

QoS Packet Marking refers to changing a field within a packet either at Layer 2 (802.1Q/p CoS, MPLS EXP) or Layer 3 (IP Precedence, DSCP and/or IP ECN).

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_mqc/configuration/xr-16/qos-mqc-xr-16-book/qos-mrkg.html

Question 4

Which marking field is used only as an internal marking within a router?

- A. Discard Eligibility
- B. QoS Group correct**
- C. IP Precedence
- D. MPLS Experimental

Explanation

Cisco routers allow you to mark two internal values (qos-group and discard-class) that travel with the packet within the router but do not modify the packet's contents.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/qos_mqc/configuration/xr-16-6/qos-mqc-xr-16-6-book/qos-mrkg.html

Switching Mechanism Quiz

Question 1

How are the Cisco Express Forwarding table and the FIB related to each other?

- A. Cisco Express Forwarding uses a FIB to make IP destination prefix-based switching decisions correct**
- B. The FIB is used to populate the Cisco Express Forwarding table
- C. There can be only one FIB but multiple Cisco Express Forwarding tables on IOS devices
- D. The Cisco Express Forwarding table allows route lookups to be forwarded to the route processor for processing before they are sent to the FIB

Explanation

The Forwarding Information Base (FIB) table – CEF uses a FIB to make IP destination prefix-based switching decisions. The FIB is conceptually similar to a routing table or information base. It maintains a mirror image of the forwarding information contained in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and these changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table.

Reference: <https://www.cisco.com/c/en/us/support/docs/routers/12000-series-routers/47321-ciscoef.html>

Question 2

What is the difference between a RIB and a FIB?

- A. The FIB is populated based on RIB content correct**
- B. The FIB is where all IP routing information is stored

- C. The RIB is used to make IP source prefix-based switching decisions
- D. The RIB maintains a mirror image of the FIB

Explanation

CEF uses a Forwarding Information Base (FIB) to make IP destination prefix-based switching decisions. The FIB is conceptually similar to a routing table or information base. It maintains a mirror image of the forwarding information contained in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table. Because there is a one-to-one correlation between FIB entries and routing table entries, the FIB contains all known routes and eliminates the need for route cache maintenance that is associated with earlier switching paths such as fast switching and optimum switching.

Note: In order to view the Routing information base (RIB) table, use the "show ip route" command. To view the Forwarding Information Base (FIB), use the "show ip cef" command. RIB is in Control plane while FIB is in Data plane.

Question 3

Which statement about Cisco Express Forwarding is true?

- A. It maintains two tables in the data plane the FIB and adjacency table **correct**
- B. The CPU of a router becomes directly involved with packet-switching decisions
- C. It makes forwarding decisions by a process that is scheduled through the IOS scheduler
- D. It uses a fast cache that is maintained in a router data plane

Explanation

Cisco Express Forwarding (CEF) provides the ability to switch packets through a device in a very quick and efficient way while also keeping the load on the router's processor low. CEF is made up of two different main components: the **Forwarding Information Base (FIB)** and the **Adjacency Table**. These are automatically updated at the same time as the routing table.

The Forwarding Information Base (FIB) contains destination reachability information as well as next hop information. This information is then used by the router to make forwarding decisions. The FIB allows for very efficient and easy lookups. Below is an example of the FIB table:

```
R2#show ip cef
```

Prefix	Next Hop	Interface
0.0.0.0/0	192.168.201.1	FastEthernet0/0
0.0.0.0/32	receive	
192.168.201.0/27	attached	FastEthernet0/0
192.168.201.0/32	receive	
192.168.201.1/32	192.168.201.1	FastEthernet0/0
192.168.201.2/32	receive	
192.168.201.31/32	receive	
224.0.0.0/4	drop	
224.0.0.0/24	receive	
255.255.255.255/32	receive	

The adjacency table is tasked with maintaining the layer 2 next-hop information for the FIB. An example of the adjacency table is shown below:

Router#show adjacency		
Protocol	Interface	Address

IP	Serial0	192.168.209.130 (2) (incomplete)
IP	Serial0	192.168.209.131 (7)
IP	Ethernet0	192.168.201.1 (7)

Note: answer 'It uses a fast cache that is maintained in a router data plane' fast cache is only used when fast switching is enabled while CEF is disabled.

Question 4

Which two statements about Cisco Express Forwarding load balancing are true? (Choose two)

- A. Each hash maps directly to a single entry in the RIB
- B. It combines the source IP address subnet mask to create a hash for each destination
- C. Cisco Express Forwarding can load-balance over a maximum of two destinations
- D. It combines the source and destination IP addresses to create a hash for each destination** correct
- E. Each hash maps directly to a single entry in the adjacency table** correct

Explanation

Cisco IOS software basically supports two modes of CEF load balancing: On per-destination or per-packet basis.

For per destination load balancing a hash is computed out of the source and destination IP address (-> Answer 'It combines the source and destination IP addresses to create a hash for each destination' is correct). **This hash points to exactly one of the adjacency entries in the adjacency table** (-> Answer 'Each hash maps directly to a single entry in the adjacency table' is correct), providing that the same path is used for all packets with this source/destination address pair. If per packet load balancing is used the packets are distributed round robin over the available paths. In either case the information in the FIB and adjacency tables provide all the necessary forwarding information, just like for non-load balancing operation.

The number of paths used is limited by the number of entries the routing protocol puts in the routing table, the default in IOS is 4 entries for most IP routing protocols with the exception of BGP, where it is one entry. **The maximum number that can be configured is 6 different paths** -> Answer 'Cisco Express Forwarding can load-balance over a maximum of two destinations' is not correct.

Reference: https://www.cisco.com/en/US/products/hw/modules/ps2033/prod_technical_reference09186a00800afeb7.html

Virtualization Quiz

Result of Virtualization Quiz:

Total Questions	Full Score	Passing Rate	Your Score	Correct Answer Percentage	Elapsed
8	110	90%	110	100%	00:02:02

Congratulations!

You passed this test!

If you want to retake this quiz, please press Ctrl + F5 on Windows or press CMD + R on Mac.

Your answers are shown below:

Question 1

Which two statements about VRF-lite are true? (Choose two)

- A. It can support multiple customers on a single switch** correct
- B. It supports most routing protocols, including EIGRP, ISIS, and OSPF** correct
- C. It should be used when a customer's router is connected to an ISP over OSPF
- D. It can increase the packet switching rate
- E. It supports MPLS-VRF label exchange and labeled packets

Explanation

In VRF-Lite, Route distinguisher (RD) identifies the customer routing table and allows customers to be assigned overlapping addresses. Therefore it can support multiple customers with overlapping addresses -> Answer 'It can support multiple customers on a single switch' is correct.

VRFs are commonly used for MPLS deployments, when we use VRFs without MPLS then we call it VRF lite -> Answer 'It supports MPLS-VRF label exchange and labeled packets' is not correct.

VRF-Lite supports most popular routing protocols: BGP, OSPF, EIGRP, RIP, and static routing -> Answer 'It supports most routing protocols, including EIGRP, ISIS, and OSPF' is correct.

Question 2

Which statement describes the IP and MAC allocation requirements for virtual machines on type 1 hypervisors?

- A. Each virtual machine requires a unique IP and MAC addresses to be able to reach to other nodes** correct
- B. Each virtual machine requires a unique MAC address but shares the IP address with the physical server
- C. Each virtual machines requires a unique IP address but shares the MAC address with the address of the physical server
- D. Each virtual machine requires a unique IP address but shares the MAC address with the physical server

Explanation


A virtual machine (VM) is a software emulation of a physical server with an operating system. From an application's point of view, the VM provides the look and feel of a real physical server, including all its components, such as CPU, memory, and network interface cards (NICs).

The virtualization software that creates VMs and performs the hardware abstraction that allows multiple VMs to run concurrently is known as a hypervisor.

There are two types of hypervisors: type 1 and type 2 hypervisor.

In type 1 hypervisor (or native hypervisor), the hypervisor is installed directly on the physical server. Then instances of an operating system (OS) are installed on the hypervisor. Type 1 hypervisor has direct access to the hardware resources. Therefore they are more efficient than hosted architectures. Some examples of type 1 hypervisor are VMware vSphere/ESXi, Oracle VM Server, KVM and Microsoft Hyper-V.

In contrast to type 1 hypervisor, a type 2 hypervisor (or hosted hypervisor) runs on top of an operating system and not the physical hardware directly. answer 'Each virtual machine requires a unique IP and MAC addresses to be able to reach to other nodes' big advantage of Type 2 hypervisors is that management console software is not required. Examples of type 2 hypervisor are VMware Workstation (which can run on Windows, Mac and Linux) or Microsoft Virtual PC (only runs on Windows).



Question 3

Which statement about route targets is true when using VRF-Lite?

A. Route targets control the import and export of routes into a customer routing table correct

B. Route targets allow customers to be assigned overlapping addresses

C. When BGP is configured, route targets are transmitted as BGP standard communities

D. Route targets uniquely identify the customer routing table


Explanation

Answer 'Route targets allow customers to be assigned overlapping addresses' and answer 'Route targets uniquely identify the customer routing table' are not correct as only route distinguisher (RD) identifies the customer routing table and "allows customers to be assigned overlapping addresses".

Answer 'When BGP is configured, route targets are transmitted as BGP standard communities' is not correct as "When BGP is configured, route targets are transmitted as BGP **extended** communities"

Question 4

Refer to the exhibit. Assuming that R1 is a CE router, which VRF is assigned to Gi0/0 on R1?



A. Default VRF correct

B. VRF VPN_A

C. Management VRF

D. VRF VPN_B

Explanation

There is nothing special with the configuration of Gi0/0 on R1. Only Gi0/0 interface on R2 is assigned to VRF VPN_A. The default VRF here is similar to the global routing table concept in Cisco IOS

Question 5

Which statement explains why Type 1 hypervisor is considered more efficient than Type 2 hypervisor?


- A. Type 1 hypervisor relies on the existing OS of the host machine to access CPU, memory, storage, and network resources
- B. Type 1 hypervisor enables other operating systems to run on it
- C. Type 1 hypervisor runs directly on the physical hardware of the host machine without relying on the underlying OS correct**
- D. Type 1 hypervisor is the only type of hypervisor that supports hardware acceleration techniques

Explanation

There are two types of hypervisors: type 1 and type 2 hypervisor.

In type 1 hypervisor (or native hypervisor), the hypervisor is installed directly on the physical server. Then instances of an operating system (OS) are installed on the hypervisor. Type 1 hypervisor has direct access to the hardware resources. Therefore they are more efficient than hosted architectures. Some examples of type 1 hypervisor are VMware vSphere/ESXi, Oracle VM Server, KVM and Microsoft Hyper-V.

In contrast to type 1 hypervisor, a type 2 hypervisor (or hosted hypervisor) runs on top of an operating system and not the physical hardware directly. answer 'Type 1 hypervisor runs directly on the physical hardware of the host machine without relying on the underlying OS' big advantage of Type 2 hypervisors is that management console software is not required. Examples of type 2 hypervisor are VMware Workstation (which can run on Windows, Mac and Linux) or Microsoft Virtual PC (only runs on Windows).



Question 6

What are two benefits of virtualizing the server with the use of VMs in data center environment? (Choose two)

- A. reduced rack space, power, and cooling requirements correct**
- B. smaller Layer 2 domain
- C. speedy deployment correct**
- D. increased security
- E. reduced IP and MAC address requirements

Explanation

Server virtualization and the use of virtual machines is profoundly changing data center dynamics. Most organizations are struggling with the cost and complexity of hosting multiple physical servers in their data centers. The expansion of the data center, a result of both scale-out server architectures and traditional "one application, one server" sprawl, has created problems in housing, powering, and cooling large numbers of underutilized servers. In addition, IT organizations continue to deal with the traditional cost and operational challenges of matching server resources to organizational needs that seem fickle and ever changing.

Virtual machines can significantly mitigate many of these challenges by enabling multiple application and operating system environments to be hosted on a single physical server while maintaining complete isolation between the guest operating systems and their respective applications. Hence, server virtualization facilitates server consolidation by enabling organizations to exchange a number of underutilized servers for a single highly utilized server running multiple virtual machines.

By consolidating multiple physical servers, organizations can gain several benefits:

- + Underutilized servers can be retired or redeployed.
- + Rack space can be reclaimed.
- + Power and cooling loads can be reduced.
- + New virtual servers can be rapidly deployed.
- + CapEx (higher utilization means fewer servers need to be purchased) and OpEx (few servers means a simpler environment and lower maintenance costs) can be reduced.

Reference: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/net_implementation_white_paper0900aecd806a9c05.html

Question 7

Refer to the exhibit. You have just created a new VRF on PE3. You have enabled debug ip bgp vpnv4 unicast updates on PE1, and you can see the route in the debug, but not in the BGP VPNv4 table. Which two statements are true? (Choose two)

```
*Jun19 11:12: BGP(4):10.1.1.2 rcvd UPDATE w/ attr:nexthop 10.1.1.2, origin ?,  
localpref 100,metric 0,extended community RT:999:999  
*Jun19 11:12: BGP(4):10.1.1.2 rcvd 999:999:192.168.1.99/32,label 29-DENIED due to:  
extended community not supported
```

A. After you configure route-target import 999:999 for a VRF on PE1, the route will be accepted correct

B. VPNv4 is not configured between PE1 and PE3

C. address-family ipv4 vrf is not configured on PE3

D. PE1 will reject the route due to automatic route filtering correct

E. After you configure route-target import 999:999 for a VRF on PE3, the route will be accepted

Explanation

Because some PE routers might receive routing information they do not require, a basic requirement is to be able to filter the MP-iBGP updates at the ingress to the PE router so that the router does not need to keep this information in memory.

The **Automatic Route Filtering** feature fulfills this filtering requirement. This feature is available by default on all PE routers, and no additional configuration is necessary to enable it. Its function is to filter automatically VPN-IPv4 routes that contain a route target extended community that does not match any of the PE's configured VRFs. This effectively discards any unwanted VPN-IPv4 routes silently, thus reducing the amount of information that the PE has to store in memory -> Answer 'PE1 will reject the route due to automatic route filtering' is correct.

Reference: MPLS and VPN Architectures Book, Volume 1

The reason that PE1 dropped the route is there is no "route-target import 999:999" command on PE1 (so we see the "DENIED due to: extended community not supported" in the debug) so we need to type this command to accept this route -> Answer 'After you configure route-target import 999:999 for a VRF on PE1, the route will be accepted' is correct.

Question 8

What is the main function of VRF-lite?

- A. To connect different autonomous systems together to share routes
- B. To allow devices to use labels to make Layer 2 Path decisions
- C. To route IPv6 traffic across an IPv4 backbone
- D. To segregate multiple routing tables on a single device**

LISP & VXLAN Quiz

Result of LISP & VXLAN Quiz:

Total Questions	Full Score	Passing Rate	Your Score	Correct Answer Percentage	Elapsed
5	60	90%	60	100%	00:00:46

Congratulations!

You passed this test!

If you want to retake this quiz, please press Ctrl + F5 on Windows or press CMD + R on Mac.

Your answers are shown below:

Question 1

Refer to the exhibit. Which LISP component do routers in the public IP network use to forward traffic between the two networks?



A. RLOC correct

- B. map resolver
- C. EID
- D. map server

Explanation

Locator ID Separation Protocol (LISP) is a network architecture and protocol that implements the use of two namespaces instead of a single IP address:

- + Endpoint identifiers (EIDs) – assigned to end hosts.
- + Routing locators (RLOCs) – assigned to devices (primarily routers) that make up the global routing system.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_lisp/configuration/xen-3s/irl-xe-3s-book/irl-overview.html

Question 2

Which LISP device is responsible for publishing EID-to-RLOC mappings for a site?

- A. MS
- B. MR
- C. ITR
- D. ETR** correct

Explanation

An Egress Tunnel Router (ETR) connects a site to the LISP-capable part of a core network (such as the Internet), publishes EID-to-RLOC mappings for the site, responds to Map-Request messages, and decapsulates and delivers LISP-encapsulated user data to end systems at the site.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_lisp/configuration/xen-3s/irl-xe-3s-book/irl-overview.html

Question 3

Which LISP infrastructure device provides connectivity between non-sites and LISP sites by receiving non-LISP traffic with a LISP site destination?

- A. map server
- B. PETR
- C. map resolver
- D. PITR correct**

Explanation

Proxy ingress tunnel router (PITR): answer 'PETR' PITR is an infrastructure LISP network entity that receives packets from non-LISP sites and encapsulates the packets to LISP sites or natively forwards them to non-LISP sites.

Reference: <https://www.ciscopress.com/articles/article.asp?p=2992605>

Question 4

Which statement about VXLAN is true?

- A. VXLAN uses TCP 35 the transport protocol over the physical data center network
- B. VXLAN uses the Spanning Tree Protocol for loop prevention
- C. VXLAN encapsulates a Layer 2 frame in an IP-UDP header, which allows Layer 2 adjacency across router boundaries correct**
- D. VXLAN extends the Layer 2 Segment ID field to 24-bits, which allows up to 4094 unique Layer 2 segments over the same network

Explanation

802.1Q VLAN identifier space is only 12 bits. The VXLAN identifier space is 24 bits. This doubling in size allows the VXLAN ID space to support 16 million Layer 2 segments -> Answer 'VXLAN extends the Layer 2 Segment ID field to 24-bits, which allows up to 4094 unique Layer 2 segments over the same network' is not correct.

VXLAN is a MAC-in-UDP encapsulation method that is used in order to extend a Layer 2 or Layer 3 overlay network over a Layer 3 infrastructure that already exists.

Reference: <https://www.cisco.com/c/en/us/support/docs/lan-switching/vlan/212682-virtual-extensible-lan-and-ethernet-virt.html>

Question 5

Into which two pieces of information does the LISP protocol split the device identity? (Choose two)

- A. Device ID
- B. Enterprise Identifier
- C. LISP ID
- D. Routing Locator correct**
- E. Resource Location
- F. Endpoint Identifier correct**

Explanation

Locator ID Separation Protocol (LISP) is a network architecture and protocol that implements the use of two namespaces instead of a single IP address:

- + Endpoint identifiers (EIDs)—assigned to end hosts.
- + Routing locators (RLOCs)—assigned to devices (primarily routers) that make up the global routing system.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_lisp/configuration/xen3s/irl-xe-3s-book/irl-overview.html

EIGRP & OSPF Quiz

Result of EIGRP & OSPF Quiz:

Total Questions	Full Score	Passing Rate	Your Score	Correct Answer Percentage	Elapsed
10	110	90%	110	100%	00:01:08

Congratulations!

You passed this test!

If you want to retake this quiz, please press Ctrl + F5 on Windows or press CMD + R on Mac.

Your answers are shown below:

Question 1

Based on this interface configuration, what is the expected state of OSPF adjacency?

```
R1
interface GigabitEthernet0/1
ip address 192.0.2.1 255.255.255.252
ip ospf 1 area 0
ip ospf hello-interval 2
ip ospf cost 1
```

```
R2
interface GigabitEthernet0/1
ip address 192.0.2.2 255.255.255.252
ip ospf 1 area 0
ip ospf cost 500
```

- A. Full on both routers
- B. 2WAY/DROTHER on both routers
- C. FULL/BDR on R1 and FULL/BDR on R2
- D. not established correct**

Explanation

On Ethernet interfaces the OSPF hello interval is 10 second by default so in this case there would be a Hello interval mismatch -> the OSPF adjacency would not be established.

Question 2

Refer to the exhibit. Which statement about the OPSF debug output is true?

```
R1#debug ip ospf hello
R1#debug condition interface fa0/1
Condition 1 set
```

- A. The output displays all OSPF messages which router R1 has sent or received on interface Fa0/1
B. The output displays OSPF hello messages which router R1 has sent or received on interface Fa0/1 correct
 C. The output displays OSPF hello and LSACK messages which router R1 has sent or received
 D. The output displays all OSPF messages which router R1 has sent or received on all interfaces


Explanation

This combination of commands is known as "Conditional debug" and will filter the debug output based on your conditions. Each condition added, will behave like an 'And' operator in Boolean logic. Some examples of the "debug ip ospf hello" are shown below:

```
*Oct 12 14:03:32.595: OSPF: Send hello to 224.0.0.5 area 0 on FastEthernet1/0 from 192.168.12.2
*Oct 12 14:03:33.227: OSPF: Rcv hello from 1.1.1.1 area 0 on FastEthernet1/0 from 192.168.12.1
*Oct 12 14:03:33.227: OSPF: Mismatched hello parameters from 192.168.12.1
```

Question 3

Refer to the exhibit. Which IPv6 OSPF network type is applied to interface Fa0/0 of R2 by default?



- A. multipoint
B. broadcast correct
 C. Ethernet
 D. point-to-point

Explanation

The Broadcast network type is the default for an OSPF enabled ethernet interface (while Point-to-Point is the default OSPF network type for Serial interface with HDLC and PPP encapsulation).

Reference: <https://www.oreilly.com/library/view/cisco-ios-cookbook/0596527225/ch08s15.html>

Question 4

Which OSPF networks types are compatible and allow communication through the two peering devices?

- A. point-to-multipoint to nonbroadcast
 B. point-to-multipoint to broadcast
 C. broadcast to point-to-point
D. broadcast to nonbroadcast correct

Explanation

The following different OSPF types are compatible with each other:

- + Broadcast and Non-Broadcast (adjust hello/dead timers)
- + Point-to-Point and Point-to-Multipoint (adjust hello/dead timers)

Broadcast and Non-Broadcast networks elect DR/BDR so they are compatible. Point-to-point/multipoint do not elect DR/BDR so they are compatible.

Question 5

Which reason could cause an OSPF neighborship to be in the EXSTART/EXCHANGE state?

- A. Mismatched areas
- B. Mismatched OSPF network type
- C. Mismatched OSPF link costs
- D. Mismatched MTU size correct**

Explanation

When OSPF adjacency is formed, a router goes through several state changes before it becomes fully adjacent with its neighbor. The states are Down -> Attempt (optional) -> Init -> 2-Way -> Exstart -> Exchange -> Loading -> Full. Short descriptions about these states are listed below:

Down: no information (hellos) has been received from this neighbor.

Attempt: only valid for manually configured neighbors in an NBMA environment. In Attempt state, the router sends unicast hello packets every poll interval to the neighbor, from which hellos have not been received within the dead interval.

Init: specifies that the router has received a hello packet from its neighbor, but the receiving router's ID was not included in the hello packet

2-Way: indicates bi-directional communication has been established between two routers.

Exstart: Once the DR and BDR are elected, the actual process of exchanging link state information can start between the routers and their DR and BDR.

Exchange: OSPF routers exchange database descriptor (DBD) packets

Loading: In this state, the actual exchange of link state information occurs

Full: routers are fully adjacent with each other

(Reference: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0e.shtml)

Neighbors Stuck in Exstart/Exchange State

the problem occurs most frequently when attempting to run OSPF between a Cisco router and another vendor's router. The problem occurs when the maximum transmission unit (MTU) settings for neighboring router interfaces don't match. If the router with the higher MTU sends a packet larger than the MTU set on the neighboring router, the neighboring router ignores the packet.

Question 6

Which EIGRP feature allows the use of leak maps?

- A. neighbor
- B. stub correct**
- C. offset-list
- D. address-family

Explanation

If we configured an EIGRP stub router so that it only advertises connected and summary routes. But we also want to have an exception to this rule then we can configure a leak-map. For example:

```
R4(config-if)#router eigrp 1
R4(config-router)#eigrp stub
R4(config)#ip access-list standard R4_L0opback0
R4(config-std-nacl)#permit host 4.4.4.4
R4(config)#route-map R4_L0opback0_LEAKMAP
R4(config-route-map)#match ip address R4_L0opback0
```

```
R4(config)#router eigrp 1
R4(config-router)#eigrp stub leak-map R4_L0opback0_LEAKMAP
```

As we can see the leak-map feature goes along with 'eigrp stub' command.

Question 7

Which two statements about EIGRP load balancing are true? (Choose two)

- A. Cisco Express Forwarding is required to load-balance across interfaces
- B. A path can be used for load balancing only if it is a feasible successor correct**
- C. EIGRP supports unequal-cost paths by default
- D. Any path in the EIGRP topology table can be used for unequal-cost load balancing
- E. EIGRP supports unequal-cost paths correct**

Explanation

EIGRP provides a mechanism to load balance over unequal cost paths (or called unequal cost load balancing) through the "variance" command. In other words, EIGRP will install all paths with **metric < variance * best metric** into the local routing table, provided that it meets the feasibility condition to prevent routing loop. The path that meets this requirement is called a feasible successor. If a path is not a feasible successor, it is not used in load balancing.

Note: The feasibility condition states that, the Advertised Distance (AD) of a route must be lower than the feasible distance of the current successor route.

Question 8

Which statement about LISP encapsulation in an EIGRP OTP implementation is true?


- A. LISP learns the next hop
- B. OTP uses LISP encapsulation to obtain routes from neighbors
- C. OTP uses LISP encapsulation for dynamic multipoint tunneling
- D. OTP maintains the LISP control plane correct**

Explanation

The EIGRP Over the Top solution can be used to ensure connectivity between disparate EIGRP sites. This feature uses EIGRP on the control plane and Locator ID Separation Protocol (LISP) encapsulation on the data plane to route traffic across the underlying WAN architecture. EIGRP is used to distribute routes between customer edge (CE) devices within the network, and the traffic forwarded across the WAN architecture is LISP encapsulated.

EIGRP OTP only uses LISP for the data plane, EIGRP is still used for the control plane. Therefore we cannot say OTP uses LISP encapsulation for dynamic multipoint tunneling as this requires encapsulating both data and control plane traffic -> Answer 'OTP uses LISP encapsulation for dynamic multipoint tunneling' is not correct.

In OTP, EIGRP serves as the replacement for LISP control plane protocols (therefore EIGRP will learn the next hop, not LISP -> Answer 'LISP learns the next hop' is not correct). Instead of doing dynamic EID-to-RLOC mappings in native LISP-mapping services, EIGRP routers running OTP over a service provider cloud create targeted sessions, use the IP addresses provided by the service provider as RLOCs, and exchange routes as EIDs. Let's take an example:



If R1 and R2 ran OTP to each other, R1 would learn about the network 10.0.2.0/24 from R2 through EIGRP, treat the prefix 10.0.2.0/24 as an EID prefix, and take the advertising next hop 198.51.100.62 as the RLOC for this EID prefix. Similarly, R2 would learn from R1 about the network 10.0.1.0/24 through EIGRP, treat the prefix 10.0.1.0/24 as an EID prefix, and take the advertising next hop 192.0.2.31 as the RLOC for this EID prefix. On both routers, this information would be used to populate the LISP mapping tables. Whenever a packet from 10.0.1.0/24 to 10.0.2.0/24 would arrive at R1, it would use its LISP mapping tables just like in ordinary LISP to discover that the packet has to be LISP encapsulated and tunneled toward 198.51.100.62, and vice versa. The LISP data plane is reused in OTP and does not change; however, the native LISP mapping and resolving mechanisms are replaced by EIGRP.

Reference: CCIE Routing and Switching V5.0 Official Cert Guide, Volume 1, Fifth Edition

Question 9

Which feature is supported by EIGRP but is not supported by OSPF?

- A. equal-cost load balancing
- B. route filtering
- C. unequal-cost load balancing correct**
- D. route summarization

Explanation

EIGRP support unequal-cost load balancing via the "variance ..." while OSPF only supports equal-cost load balancing.

Question 10


In OSPF, which LSA type is responsible for pointing to the ASBR router?

- A. type 1
- B. type 3
- C. type 2
- D. type 4 correct**

Explanation

Summary ASBR LSA (Type 4) – Generated by the ABR to describe an ASBR to routers in other areas so that routers in other areas know how to get to external routes through that ASBR. For example, suppose R8 is redistributing external route (EIGRP, RIP...) to R3. This makes R3 an Autonomous System Boundary Router (ASBR). When R2 (which is an ABR) receives this LSA Type 1 update, R2 will create LSA Type 4 and flood into Area 0 to inform them how to reach R3. When R5 receives this LSA it also floods into Area 2.

In the above example, the only ASBR belongs to area 1 so the two ABRs (R2 & R5) send LSA Type 4 to area 0 & area 2 (not vice versa). This is an indication of the existence of the ASBR in area 1.



Note:

+ Type 4 LSAs contain the router ID of the ASBR.

+ There are no LSA Type 4 injected into Area 1 because every router inside area 1 knows how to reach R3. R3 only uses LSA Type 1 to inform R2 about R8 and inform R2 that R3 is an ASBR.

BGP Quiz

Result of BGP Quiz:

Total Questions	Full Score	Passing Rate	Your Score	Correct Answer Percentage	Elapsed
4	40	90%	40	100%	00:00:40

Congratulations!

You passed this test!

If you want to retake this quiz, please press Ctrl + F5 on Windows or press CMD + R on Mac.

Your answers are shown below:

Question 1

Refer to the exhibit. Which IP address becomes the next active next hop for 192.168.102.0/24 when 192.168.101.2 fails?

```
R1#show ip bgp
```

```
BGP table version is 32, local router ID is 192.168.101.5
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
                r RIB-failure, S Stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```

Network          Next Hop           Metric LocPrf Weight Path
* 192.168.102.0   192.168.101.18      80             0 64517 i
*                  192.168.101.14      80             80 0 64516 i
*                  192.168.101.10              0 64515 64515
i
*>                192.168.101.2              0 64513 i

```

*	192.168.101.6	80	0 64514 64514
i			

- A. 192.168.101.6
- B. 192.168.101.10
- C. 192.168.101.18 correct**
- D. 192.168.101.14

Explanation

The '>' shown in the output above indicates that the path with a next hop of 192.168.101.2 is the current best path.

Path Selection Attributes: Weight > Local Preference > Originate > AS Path > Origin > MED > External > IGP Cost > eBGP Peering > Router ID

BGP prefers the path with highest weight but the weights here are all 0 (which indicate all routes that are not originated by the local router) so we need to check the Local Preference. Answer '192.168.101.18' path without LOCAL_PREF (LocPrf column) means it has the default value of 100. Therefore we can find the two next best paths with the next hop of 192.168.101.18 and 192.168.101.10.

We have to move to the next path selection attribute: Originate. BGP prefers the path that the local router originated (which is indicated with the "next hop 0.0.0.0"). But none of the two best paths is self-originated.

The AS Path of the next hop 192.168.101.18 is shorter than the AS Path of the next hop 192.168.101.10 then the next hop 192.168.101.18 will be chosen as the next best path.

Question 2

A local router shows an EBGP neighbor in the Active state. Which statement is true about the local router?

- A. The local router has active prefix in the forwarding table from the neighboring router
- B. The local router is attempting to open a TCP session with the neighboring router. correct**
- C. The local router has BGP passive mode configured for the neighboring router
- D. The local router is receiving prefixes from the neighboring router and adding them in RIB-IN

Explanation

The BGP session may report in the following states

1 – Idle: the initial state of a BGP connection. In this state, the BGP speaker is waiting for a BGP start event, generally either the establishment of a TCP connection or the re-establishment of a previous connection. Once the connection is established, BGP moves to the next state.

2 – Connect: In this state, BGP is waiting for the TCP connection to be formed. If the TCP connection completes, BGP will move to the Open Sent stage; if the connection cannot complete, BGP goes to Active

3 – Active: In the Active state, the BGP speaker is attempting to initiate a TCP session with the BGP speaker it wants to peer with. If this can be done, the BGP state goes to Open Sent state.

4 – Open Sent: the BGP speaker is waiting to receive an OPEN message from the remote BGP speaker

5 – Open Confirm: Once the BGP speaker receives the OPEN message and no error is detected, the BGP speaker sends a KEEPALIVE message to the remote BGP speaker

6 – Established: All of the neighbor negotiations are complete. You will see a number, which tells us the number of prefixes the router has received from a neighbor or peer group.

Question 3

What is the correct EBGP path attribute list, ordered from most preferred to the least preferred, that the BGP best-path algorithm uses?


- A. weight, AS path, local preference, MED
- B. weight, local preference, AS path, MED correct**
- C. local preference, weight, AS path, MED
- D. local preference, weight, MED, AS path

Explanation

Path Selection Attributes: Weight > Local Preference > Originate > AS Path > Origin > MED > External > IGP Cost > eBGP Peering > Router ID

Question 4

Refer to the exhibit. Which configuration establishes EBGP neighborship between these two directly connected neighbors and exchanges the loopback network of the two routers through BGP?



- A. R1 (config) #router bgp 1
R1 (config-router) #neighbor 10.2.2.2 remote-as 2
R1 (config-router) #neighbor 10.2.2.2 update-source lo0
R1 (config-router) #network 10.1.1.0 mask 255.255.255.0
R2 (config) #router bgp 2
R2 (config-router) #neighbor 10.1.1.1 remote-as 1
R2 (config-router) #neighbor 10.1.1.1 update-source lo0
R2 (config-router) #network 10.2.2.0 mask 255.255.255.0

- B. R1(config)#router bgp 1
R1(config-router)#neighbor 192.168.10.2 remote-as 2
R1(config-router)#network 10.1.1.0 mask 255.255.255.0
R2(config)#router bgp 2
R2(config-router)#neighbor 192.168.10.1 remote-as 1
R2(config-router)#network 10.2.2.0 mask 255.255.255.0**

- C. R1 (config) #router bgp 1
R1 (config-router) #neighbor 192.168.10.2 remote-as 2
R1 (config-router) #network 10.0.0.0 mask 255.0.0.0
R2 (config) #router bgp 2
R2 (config-router) #neighbor 192.168.10.1 remote-as 1
R2 (config-router) #network 10.0.0.0 mask 255.0.0.0

- D. R1 (config) #router bgp 1
R1 (config-router) #neighbor 10.2.2.2 remote-as 2
R1 (config-router) #network 10.1.1.0 mask 255.255.255.0
R2 (config) #router bgp 2
R2 (config-router) #neighbor 10.1.1.1 remote-as 1
R2 (config-router) #network 10.2.2.0 mask 255.255.255.0

Explanation

With BGP, we must advertise the correct network and subnet mask in the "network" command (in this case network 10.1.1.0/24 on R1 and network 10.2.2.0/24 on R2). BGP is very strict in the routing advertisements. In other words, BGP only advertises the network which exists exactly in the routing table. In this case, if you put the command "network x.x.0.0 mask 255.255.0.0" or "network x.0.0.0 mask 255.0.0.0" or "network x.x.x.x mask 255.255.255.255" then BGP will not advertise anything.

It is easy to establish eBGP neighborship via the direct link. But let's see what are required when we want to establish eBGP neighborship via their loopback interfaces. We will need two commands:

- + the command "neighbor 10.1.1.1 ebgp-multihop 2" on R1 and "neighbor 10.2.2.2 ebgp-multihop 2" on R1. This command increases the TTL value to 2 so that BGP updates can reach the BGP neighbor which is two hops away.

- + Answer 'R1 (config) #router bgp 1

R1 (config-router) #neighbor 192.168.10.2 remote-as 2

R1 (config-router) #network 10.1.1.0 mask 255.255.255.0

R2 (config) #router bgp 2

R2 (config-router) #neighbor 192.168.10.1 remote-as 1

R2 (config-router) #network 10.2.2.0 mask 255.255.255.0

Quick Wireless Summary

Cisco Access Points (APs) can operate in one of two modes: autonomous or lightweight

- + **Autonomous**: self-sufficient and standalone. Used for small wireless networks.

- + **Lightweight**: A Cisco lightweight AP (LAP) has to join a Wireless LAN Controller (WLC) to function. LAP and WLC communicate with each other via a logical pair of CAPWAP tunnels.

- **Control and Provisioning for Wireless Access Point (CAPWAP)** is an IETF standard for control messaging for setup, authentication and operations between APs and WLCs. CAPWAP is similar to LWAPP except the following differences:

- + CAPWAP uses Datagram Transport Layer Security (DTLS) for authentication and encryption to protect traffic between APs and controllers. LWAPP uses AES.

- + CAPWAP has a dynamic maximum transmission unit (MTU) discovery mechanism.

- + CAPWAP runs on UDP ports 5246 (control messages) and 5247 (data messages)

An LAP operates in one of six different modes:

- + **Local mode** (default mode): measures noise floor and interference, and scans for intrusion detection (IDS) events every 180 seconds on unused channels

- + **FlexConnect**, formerly known as **Hybrid Remote Edge AP (H-REAP)**, mode: allows data traffic to be switched locally and not go back to the controller. The FlexConnect AP can perform standalone client authentication and switch VLAN traffic locally even when it's disconnected to the WLC (Local Switched). FlexConnect AP can also tunnel (via CAPWAP) both user wireless data and control traffic to a centralized WLC (Central Switched).

- + **Monitor mode**: does not handle data traffic between clients and the infrastructure. It acts like a sensor for location-based services (LBS), rogue AP detection, and IDS

- + **Rogue detector mode**: monitor for rogue APs. It does not handle data at all.


- + **Sniffer mode**: run as a sniffer and captures and forwards all the packets on a particular channel to a remote machine where you can use protocol analysis tool (Wireshark, Airopeek, etc) to review the packets and diagnose issues. Strictly used for troubleshooting purposes.

- + **Bridge mode**: bridge together the WLAN and the wired infrastructure together.

Mobility Express is the ability to use an access point (AP) as a controller instead of a real WLAN controller. But this solution is only suitable for small to midsize, or multi-site branch locations where you might not want to invest in a dedicated WLC. A Mobility Express WLC can support up to 100 APs

Question 1

Refer to the exhibit. Which type of antenna do the radiation patterns present?



- A. ☐ Dipole
- B. ☐ Yagi
- C. ☒ Patch
- D. ☐ Omnidirectional

Question 2

An engineer is configuring local web authentication on a WLAN. The engineer chooses the Authentication radio button under the Layer 3 Security options for Web Policy. Which device presents the web authentication for the WLAN?

- A. ☐ RADIUS server
- B. ☒ local WLC
- C. ☐ ISE server
- D. ☐ anchor WLC

Question 3

A client device fails to see the enterprise SSID, but other devices are connected to it. What is the cause of this issue?

- A. ☒ The hidden SSID was not manually configured on the client.
- B. ☐ The client has incorrect credentials stored for the configured broadcast SSID.
- C. ☐ The broadcast SSID was not manually configured on the client.
- D. ☐ The client has incorrect credentials stored for the configured hidden SSID.

Question 4

Which two methods are used by an AP that is trying to discover a wireless LAN controller? (Choose two)

- A. ☐ DNS lookup cisco-DNA-PRIMARY.local domain
- B. ☐ querying other APs
- C. ☒ broadcasting on the local subnet
- D. ☒ DHCP Option 43
- E. ☐ Cisco Discovery Protocol neighbor

Question 5

Which statement about Cisco EAP-FAST is true?

- A. ☐ It is an IETF standard.
- B. ☐ It requires a client certificate
- C. ☒ It does not require a RADIUS server certificate
- D. ☐ It operates in transparent mode

Question 6

What are two common sources of interference for WI-FI networks? (Choose two)

- A. ☐ conventional oven
- B. ☒ rogue AP
- C. ☐ LED lights
- D. ☐ fire alarm
- E. ☒ radar

Question 7

Refer to the exhibit. The WLC administrator sees that the controller to which a roaming client associates has Mobility Role Anchor configured under Clients > Detail. Which type of roaming is supported?

Clients > Detail

Client Properties

MAC Address	00:09:ee:12:34:d2
IP Address	192.168.100.199
Client Type	Regular
User Name	
Port Number	20
Interface	00:09:ee:12:34:d2
VLAN ID	3602
CCX Version	Not Supported
E2E Version	E2Ev1
Mobility Role	Anchor
Mobility Peer IP Address	172.22.253.20
Policy Manager State	RUN
Management Frame Protection	No
UpTime (Sec)	944581
Power Save Mode	OFF
Current TxRateSet	48.0
Data RateSet	6.0,9.0,12.0,18.0,24.0,36.0,48.0,54.0

AP Properties

AP Address	
AP Name	172.22.253.20
AP Type	Mobile
WLAN Profile	
Status	Associated
Association ID	16
802.11 Authentication	Open System
Reason Code	1
Status Code	0
CF Pollable	Not Implemented
CF Poll Request	Not Implemented
Short Preamble	Not Implemented
PBCC	Not Implemented
Channel Agility	Not Implemented
Timeout	0
WEP State	WEP Enable

- A. ☒ Layer 3 intercontroller
- B. ☐ Intercontroller
- C. ☐ Layer 2 intercontroller
- D. ☐ Indirect

Question 8


Which DNS lookup does an access point perform when attempting CAPWAP discovery?

- A. ☐ CISCO-CONTROLLER.local

- B. ☐ CAPWAP-CONTROLLER.local
- C. ☒ CISCO-CAPWAP-CONTROLLER.local
- D. ☐ CISCO-DNA-CONTROILLER.local

Question 9

Refer to the exhibit. Based on the configuration in this WLAN security setting. Which method can a client use to authenticate to the network?



The screenshot shows the 'Security' tab of a WLAN configuration interface. Under the 'Layer 3' sub-tab, the 'AAA Servers' section is visible. The 'Fast Transition' checkbox is unchecked. The 'Protected Management Frame' (PMF) is set to 'Disabled'. Under 'WPA+WPA2 Parameters', 'WPA Policy' is unchecked and 'WPA2 Policy-AES' is checked. In the 'Authentication Key Management' section, '802.1X', 'CCKM', 'FT 802.1X', and 'FT PSK' are all unchecked, while 'PSK' is checked. The 'PSK Format' is set to 'ASCII'.

- A. ☐ RADIUS token
- B. ☒ text string
- C. ☐ username and password
- D. ☐ certificate

Question 10

An engineer configures a WLAN with fast transition enabled. Some legacy clients fail to connect to this WLAN. Which feature allows the legacy clients to connect while still allowing other clients to use fast transition based on their OLTIs?

- A. ☐ 802.11V
- B. ☐ over the DS
- C. ☐ 802.11k
- D. ☒ adaptive R

Question 11

Which two pieces of information are necessary to compute SNR? (Choose two)

- A. ☐ transmit power
- B. ☐ EIRP
- C. ☒ RSSI
- D. ☐ antenna gain
- E. ☒ noise floor

Question 12

A customer has several small branches and wants to deploy a WI-FI solution with local management using CAPWAP. Which deployment model meets this requirement?

- A. ☒ Mobility express
- B. ☐ Local mode
- C. ☐ SD-Access wireless
- D. ☐ Autonomous

Question 13

To increase total throughput and redundancy on the links between the wireless controller and switch, the customer enabled LAG on the wireless controller. Which EtherChannel mode must be configured on the switch to allow the WLC to connect?

- A. ☐ Auto
- B. ☐ Passive
- C. ☒ On
- D. ☐ Active

Question 14

When configuration WPA2 Enterprise on a WLAN, which additional security component configuration is required?

- A. ☐ TACACS server
- B. ☐ NTP server
- C. ☒ RADIUS server
- D. ☐ PKI server

Question 15

Which two descriptions of FlexConnect mode for Cisco APs are true? (Choose two)

- A. ☐ APs that operate in FlexConnect mode cannot detect rogue APs
- B. ☒ FlexConnect mode is a wireless solution for branch office and remote office deployments
- C. ☐ FlexConnect mode is a feature that is designed to allow specified CAPWAP-enabled APs to exclude themselves from managing data traffic between clients and infrastructure
- D. ☐ FlexConnect mode is used when the APs are set up in a mesh environment and used to bridge between each other
- E. ☒ When connected to the controller, FlexConnect APs can tunnel traffic back to the controller

HSRP & VRRP Quiz

Question 1

Which First Hop Redundancy Protocol maximizes uplink utilization and minimizes the amount of configuration that is necessary?

- A. ☐ VRRP
- B. ☒ GLBP
- C. ☐ HSRP v2
- D. ☐ HSRP v1

Question 2

If a VRRP master router fails, which router is selected as the new master router?

- A. ☐ router with the highest loopback address
- B. ☐ router with the lowest loopback address
- C. ☒ router with the highest priority
- D. ☐ router with the lowest priority

Question 3

Which two statements about VRRP are true? (Choose two)

- A. ☐ It supports both MD5 and SHA1 authentication.
- B. ☐ It is assigned multicast address 224.0.0.9.
- C. ☒ Three versions of the VRRP protocol have been defined.
- D. ☐ It is assigned multicast address 224.0.0.8.
- E. ☒ The TTL for VRRP packets must be 255.
- F. ☐ Its IP address number is 115.

Question 4

Which two statements about HSRP are true? (Choose two)

- A. ☐ It supports unique virtual MAC addresses
- B. ☒ Its virtual MAC is 0000.0C07.ACxx

- C. ☐ Its default configuration allows for pre-emption
- D. ☒ It supports tracking
- E. ☐ Its multicast virtual MAC is 0000.5E00.01xx

Question 5

Which behavior can be expected when the HSRP versions is changed from 1 to 2?

- A. ☐ Each HSRP group reinitializes because the multicast address has changed
- B. ☒ Each HSRP group reinitializes because the virtual MAC address has changed
- C. ☐ No changes occur because version 1 and 2 use the same virtual MAC OUI
- D. ☐ No changes occur because the standby router is upgraded before the active router

Question 6

What are three valid HSRP states? (Choose three)

- A. ☐ INIT
- B. ☒ listen
- C. ☐ full
- D. ☒ learning
- E. ☒ speak
- F. ☐ established

Network Assurance Questions

Result of Network Assurance Quiz:

Total Questions	Full Score	Passing Rate	Your Score	Correct Answer Percentage	Elapsed
9	110	90%	110	100%	00:01:31

Congratulations!

You passed this test!

If you want to retake this quiz, please press Ctrl + F5 on Windows or press CMD + R on Mac.

Your answers are shown below:

Question 1

Which two steps are required for a complete Cisco DNA Center upgrade? (Choose two)

- A. application updates correct
- B. proxy configuration
- C. golden image selection
- D. automation backup
- E. system update correct

Explanation

A complete Cisco DNA Center upgrade includes "System Update" and "Appplication Updates"

System Update

System 1.3.0.109

✔ Your system package is up to date. Proceed with Application updates.

Application Updates

Update All

Cisco DNA Center Core

	Size	Version	
Automation - Base	493.25 MB	2.1.78.60109	Update failed
NCP - Base	167.84 MB	2.1.78.60109	Update failed
NCP - Services	326.84 MB	2.1.78.60109	Update failed
Network Controller Platform	3.65 GB	2.1.78.60109	Update failed

Automation

	Size	Version	
Command Runner	55.20 MB	2.1.78.60109	Update failed
Device Onboarding	162.41 MB	2.1.78.60109	Update failed
Image Management	362.85 MB	2.1.78.60109	Update failed

Question 2

Which two statements about IP SLA are true? (Choose two)

- A. It uses NetFlow for passive traffic monitoring
- B. It can measure MOS
- C. The IP SLA responder is a component in the source Cisco device
- D. It is Layer 2 transport-independent correct**
- E. It uses active traffic monitoring correct**
- F. SNMP access is not supported

Explanation

IP SLAs allows Cisco customers to analyze IP service levels for IP applications and services, to increase productivity, to lower operational costs, and to reduce the frequency of network outages. **IP SLAs uses active traffic monitoring**—the generation of traffic in a continuous, reliable, and predictable manner—for measuring network performance.

Being **Layer-2 transport independent**, IP SLAs can be configured end-to-end over disparate networks to best reflect the metrics that an end-user is likely to experience.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/configuration/15-mt/sla-15-mt-book/sla_overview.html

Question 3

Which IP SLA operation requires the IP SLA responder to be configured on the remote end?

- A. UDP jitter correct**
- B. CMP jitter
- C. ICMP echo
- D. TCP connect

Explanation

Cisco IOS IP SLA Responder is a Cisco IOS Software component whose functionality is to respond to Cisco IOS IP SLA request packets. The IP SLA source sends control packets before the operation starts to establish a connection to the responder. Once the control packet is acknowledged, test packets are sent to the responder. **The responder inserts a time-stamp when it receives a packet** and factors out the destination processing time **and adds time-stamps to the sent**

packets. This feature allows the calculation of unidirectional packet loss, latency, and jitter measurements with the kind of accuracy that is not possible with ping or other dedicated probe testing.

Reference: https://www.cisco.com/en/US/technologies/tk869/tk769/technologies_white_paper0900aecd806bfb52.html

The IP SLAs responder is a component embedded in the destination Cisco device that allows the system to anticipate and respond to IP SLAs request packets. The responder provides accurate measurements without the need for dedicated probes.

Reference: <https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/46sg/configuration/guide/Wrapper-46SG/swipsla.html>

UDP Jitter measures the delay, delay variation(jitter), corruption, misordering and packet loss by generating periodic UDP traffic. This operation always requires IP SLA responder.

Reference: <https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2017/pdf/BRKNMS-3043.pdf>

Question 4

Which statement about an RSPAN session configuration is true?

A. A special VLAN type must be used as the RSPAN destination. Correct

- B. A filter must be configured for RSPAN Regions
- C. Only one session can be configured at a time
- D. Only incoming traffic can be monitored

Explanation

The traffic for each RSPAN session is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches -> This VLAN can be considered a special VLAN type -> Answer 'A special VLAN type must be used as the RSPAN destination' is correct.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750x_3560x/software/release/12-2_55_se/configuration/guide/3750xscg/swspan.html

We can configure multiple RSPAN sessions on a switch at a time, then continue configuring multiple RSPAN sessions on the other switch without any problem -> Answer 'Only one session can be configured at a time' is not correct.

This is how to configure Remote SPAN (RSPAN) feature on two switches. Traffic on FastEthernet0/1 of Switch 1 will be sent to Fa0/10 of Switch2 via VLAN 40.

+ Configure on both switches

Switch1,2(config)#vlan 40

Switch1,2(config-vlan)#remote-span

+ Configure on Switch1

Switch1(config)# monitor session 1 source interface FastEthernet 0/1

Switch1(config)# monitor session 1 destination remote vlan 40

+ Configure on Switch2

Switch2(config)#monitor session 5 source remote vlan 40

Switch2(config)# monitor session 5 destination interface FastEthernet 0/10

Question 5

Which feature must be configured to allow packet capture over Layer 3 infrastructure?

- A. IPSPAN
- B. RSPAN
- C. VSPAN
- D. ERSPAN** correct

Explanation

Encapsulated remote SPAN (ERSPAN): encapsulated Remote SPAN (ERSPAN), as the name says, brings generic routing encapsulation (GRE) for all captured traffic and allows it to be extended across Layer 3 domains.

Question 6

A network is being migrated from IPv4 to IPv6 using a dual-stack approach. Network management is already 100% IPv6 enabled. In a dual-stack network with two dual-stack NetFlow collections, how many flow exporters are needed per network device in the flexible NetFlow configuration?

- A. 1
- B. 8
- C. 4
- D. 2 correct**

Question 7

When using TLS for syslog, which configuration allows for secure and reliable transportation of messages to its default port?

- A. logging host 10.2.3.4 vrf mgmt transport udp port 6514
- B. logging host 10.2.3.4 vrf mgmt transport tcp port 514
- C. logging host 10.2.3.4 vrf mgmt transport udp port 514
- D. logging host 10.2.3.4 vrf mgmt transport tcp port 6514 correct**

Explanation

The TCP port 6514 has been allocated as the default port for syslog over Transport Layer Security (TLS).

Reference: <https://tools.ietf.org/html/rfc5425>

Question 8

At which layer does Cisco DNA Center support REST controls?

- A. EEM applets or scripts
- B. Northbound APIs correct**
- C. Session layer
- D. YMAL output from responses to API calls

Question 9

Refer to this output What is the logging severity level?

R1#Feb 14 37:15:12:429: %LINEPROTO-5-UPDOWN Line protocol on interface GigabitEthernet0/1. Change state to up

A. Notification correct

- B. Alert
- C. Critical
- D. Emergency

Explanation

Syslog levels are listed below:

Level	Keyword	Description
0	emergencies	System is unusable
1	alerts	Immediate action is needed

2	critical	Critical conditions exist
3	errors	Error conditions exist
4	warnings	Warning conditions exist
5	notification	Normal, but significant, conditions exist
6	informational	Informational messages
7	debugging	Debugging messages

Number "5" in "%LINEPROTO-5- UPDOWN" is the severity level of this message so in this case it is "notification".

Security Questions

Result of Security Quiz:

Total Questions	Full Score	Passing Rate	Your Score	Correct Answer Percentage	Elapsed
8	80	90%	80	100%	00:01:51

Congratulations!

You passed this test!

If you want to retake this quiz, please press Ctrl + F5 on Windows or press CMD + R on Mac.

Your answers are shown below:

Question 1

Which method does the enable secret password option use to encrypt device passwords?

- A. AES
- B. PAP
- C. MD5 correct**
- D. CHAP

Question 2

Which feature does Cisco TrustSec use to provide scalable, secure communication throughout a network?

- A. security group tag number assigned to each port on a network correct**
- B. security group tag number assigned to each user on a switch
- C. security group tag ACL assigned to each router on a network
- D. security group tag ACL assigned to each port on a switch

Explanation

Cisco TrustSec uses tags to represent logical group privilege. This tag, called a Security Group Tag (SGT), is used in access policies. The SGT is understood and is used to enforce traffic by Cisco switches, routers and firewalls . Cisco TrustSec is defined in three phases: classification, propagation and enforcement.

When users and devices connect to a network, the network assigns a specific security group. This

process is called classification. Classification can be based on the **results of the authentication or by associating the SGT with an IP, VLAN, or port-profile** (-> Answer 'security group tag ACL assigned to each port on a switch' and answer 'security group tag number assigned to each user on a switch' are not correct as they say "assigned ... on a switch" only. Answer 'security group tag ACL assigned to each router on a network' is not correct either as it says "assigned to each router").

Question 3

Refer to the exhibit. Which privilege level is assigned to VTY users?

```
R1# sh run | begin line con
line con 0
exec-timeout 0 0
privilege level 15
logging synchronous
stopbits 1
line aux 0
exec-timeout 0 0
privilege level 15
logging synchronous
stopbits 1
line vty 0 4
password 7 03384737389E938
login
line vty 5 15
password 7 03384737389E938
login
!
end
```

```
R1#sh run | include aaa | enable
no aaa new-model
R1#
```

- A. 13
- B. 15
- C. 7
- D. 1** correct

Explanation

Lines (CON, AUX, VTY) default to level 1 privileges.

Question 4

Which technology provides a secure communication channel for all traffic at Layer 2 of the OSI model?

- A. SSL
- B. IPsec
- C. Cisco Trustsec

D. MACsec correct

Explanation

MACsec, defined in 802.1AE, provides MAC-layer encryption over wired networks by using out-of-band methods for encryption keying. The MACsec Key Agreement (MKA) Protocol provides the required session keys and manages the required encryption keys. MKA and MACsec are implemented after successful authentication using the 802.1x Extensible Authentication Protocol (EAP-TLS) or Pre Shared Key (PSK) framework.

A switch using MACsec accepts either MACsec or non-MACsec frames, depending on the policy associated with the MKA peer. MACsec frames are encrypted and protected with an integrity check value (ICV). When the switch receives frames from the MKA peer, it decrypts them and calculates the correct ICV by using session keys provided by MKA. The switch compares that ICV to the ICV within the frame. If they are not identical, the frame is dropped. The switch also encrypts and adds an ICV to any frames sent over the secured port (the access point used to provide the secure MAC service to a MKA peer) using the current session key.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/16-9/configuration_guide/sec/b_169_sec_9300_cg/macsec_encryption.html

Note: Cisco Trustsec is the solution which includes MACsec.

Question 5

What is the difference between the enable password and the enable secret password when password encryption is enable on an IOS device?

A. The enable secret password is protected via stronger cryptography mechanisms correct

B. The enable password cannot be decrypted

C. The enable password is encrypted with a stronger encryption method

D. There is no difference and both passwords are encrypted identically

Explanation

The "enable secret" password is always encrypted (independent of the "service password-encryption" command) using MD5 hash algorithm. The "enable password" does not encrypt the password and can be view in clear text in the running-config. In order to encrypt the "enable password", use the "service password-encryption" command. This command will encrypt the passwords by using the Vigenere encryption algorithm. Unfortunately, the Vigenere encryption method is cryptographically weak and trivial to reverse.

The MD5 hash is a stronger algorithm than Vigenere so answer 'The enable secret password is protected via stronger cryptography mechanisms' is correct.

Question 6

The login method is configured on the VTY lines of a router with these parameters.

– The first method for authentication is TACACS

– If TACACS is unavailable, login is allowed without any provided credentials

Which configuration accomplishes this task?

```
R1#sh run | include aaa
aaa new-model
aaa authentication login VTY group tacacs+ none
aaa session-id common
R1#sh run | section vty
line vty 0 4
password 7 0202039485748
```

```
R1#sh run | include username
R1#
```

```
B. R1#sh run | include aaa
aaa new-model
```

```
aaa authentication login telnet group tacacs+ none
aaa session-id common
R1#sh run | section vty
line vty 0 4
```

```
R1#sh run | include username
R1#
```

```
C. R1#sh run | include aaa
aaa new-model
aaa authentication login default group tacacs+ none
aaa session-id common
R1#sh run | section vty
line vty 0 4
password 7 0202039485748
```

```
D. R1#sh run | include aaa
aaa new-model
aaa authentication login default group tacacs+
aaa session-id common
R1#sh run | section vty
line vty 0 4
transport input none
R1#
```

Explanation

According to the requirements (first use TACACS+, then allow login with no authentication), we have to use "aaa authentication login ... **group tacacs+ none**" for AAA command.

The next thing to check is the if the "aaa authentication login **default**" or "aaa authentication login *list-name*" is used. The '**default**' keyword means we want to apply for all login connections (such as tty, vty, console and aux). If we use this keyword, we don't need to configure anything else under tty, vty and aux lines. If we don't use this keyword then we have to specify which line(s) we want to apply the authentication feature.

From above information, we can find out answer 'R1#sh run | include aaa
aaa new-model
aaa authentication login default group tacacs+ none
aaa session-id common

```
R1#sh run | section vty
line vty 0 4
password 7 0202039485748
```

If you want to learn more about AAA configuration, please read our [AAA TACACS+ and RADIUS Tutorial – Part 2](#).

For your information, answer 'R1#sh run | include aaa
aaa new-model
aaa authentication login telnet group tacacs+ none
aaa session-id common

```
R1#sh run | section vty
line vty 0 4
```

```
R1#sh run | include username
R1#' would be correct if we add the following command under vty line ("line vty 0 4"): "login authentication telnet" ("telnet" is the name of the AAA list above)
```

Question 7

Which NGFW mode block flows crossing the firewall?

A. Inline tap

- B. Tap
C. Inline correct
D. Passive

Explanation

Firepower Threat Defense (FTD) provides six interface modes which are: Routed, Switched, Inline Pair, Inline Pair with Tap, Passive, Passive (ERSPAN).

When Inline Pair Mode is in use, packets can be blocked since they are processed inline

When you use Inline Pair mode, the packet goes mainly through the FTD Snort engine
When Tap Mode is enabled, a copy of the packet is inspected and dropped internally while the actual traffic goes through FTD unmodified

<https://www.cisco.com/c/en/us/support/docs/security/firepower-ngfw/200924-configuring-firepower-threat-defense-int.html>

Question 8

How does Cisco Trustsec enable more access controls for dynamic networking environments and data centers?

- A. classifies traffic based on advanced application recognition
B. uses flexible NetFlow
C. classifies traffic based on the contextual identity of the endpoint rather than its IP address correct
D. assigns a VLAN to the endpoint

Explanation

The Cisco TrustSec solution simplifies the provisioning and management of network access control through the use of software-defined segmentation to classify network traffic and enforce policies for more flexible access controls. **Traffic classification is based on endpoint identity, not IP address**, enabling policy change without network redesign.

Reference: https://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/Apr2016/User-to-DC_Access_Control_Using_TrustSec_Deployment_April2016.pdf

Access-list Quiz

Result of Access-list Quiz:

Total Questions	Full Score	Passing Rate	Your Score	Correct Answer Percentage	Elapsed
5	50	90%	50	100%	00:01:46

Congratulations!

You passed this test!

If you want to retake this quiz, please press Ctrl + F5 on Windows or press CMD + R on Mac.

Your answers are shown below:

Question 1

Refer to the exhibit. An engineer must modify the access control list EGRESS to allow all IP traffic from subnet 10.1.10.0/24 to 10.1.2.0/24. The access control list is applied in the outbound direction on router interface GigabitEthernet 0/1.

```
Extended IP access list EGRESS
10 permit ip 10.1.100.0 0.0.0.255 10.1.2.0 0.0.0.255
20 deny ip any any
```

Which configuration commands can the engineer use to allow this traffic without disrupting existing traffic flows?

A. config t

ip access-list extended EGRESS

5 permit ip 10.1.10.0 0.0.0.255 10.1.2.0 0.0.0.255

<p< span="" style="border: 1px solid black; padding: 2px 5px; display: inline-block;">correct</p>

B. config t

ip access-list extended EGRESS

permit ip 10.1.10.0 255.255.255.0 10.1.2.0 255.255.255.0

<p< span="" style="border: 1px solid black; padding: 2px 5px; display: inline-block;"></p>

C. config t

ip access-list extended EGRESS2

permit ip 10.1.10.0 0.0.0.255 10.1.2.0 0.0.0.255

permit ip 10.1.100.0 0.0.0.255 10.1.2.0 0.0.0.255

deny ip any any

!

interface g0/1

no ip access-group EGRESS out

ip access-group EGRESS2 out

<p< span="" style="border: 1px solid black; padding: 2px 5px; display: inline-block;"></p>

D. config t

ip access-list extended EGRESS

permit ip 10.1.10.0 0.0.0.255 10.1.2.0 0.0.0.255

Explanation

We can insert a line (statement) between entries into an existing ACL by a number in between.

```
R1#sh access-list
Extended IP access list EGRESS
 10 permit ip 10.1.100.0 0.0.0.255 10.1.2.0 0.0.0.255
 20 deny ip any any
R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#ip access-list extended EGRESS
R1(config-ext-nacl)#5 permit ip 10.1.10.0 0.0.0.255 10.1.2.0 0.0.0.255
R1(config-ext-nacl)#do sh access-list
Extended IP access list EGRESS
 5 permit ip 10.1.10.0 0.0.0.255 10.1.2.0 0.0.0.255
10 permit ip 10.1.100.0 0.0.0.255 10.1.2.0 0.0.0.255
20 deny ip any any
```

So what will happen if we just enter a statement without the number? Well, that statement would be added at the bottom of an ACL. But in this case we already had an explicit "deny ip any any" statement so we cannot put another line under it.

Question 2

Refer to the exhibit. An engineer must block all traffic from a router to its directly connected subnet 209.165.200.0/24. The engineer applies access control list EGRESS in the outbound direction on the GigabitEthernet0/0 interface of the router. However, the

router can still ping hosts on the 209.165.200.0/24 subnet. Which explanation of this behavior is true?

```
Extended IP access list EGRESS
10 permit ip 10.0.0.0 0.0.0.255 any
!
---output omitted---
!
interface GigabitEthernet0/0
ip address 209.165.200.255 255.255.255.0
ip access-group EGRESS out
duplex auto
speed auto
media-type rj45
!
```

- A. Only standard access control lists can block traffic from a source IP address
- B. The access control list must contain an explicit deny to block traffic from the router
- C. After an access control list is applied to an interface, that interface must be shut and no shut for the access control list to take effect
- D. Access control lists that are applied outbound to a router interface do not affect traffic that is sourced from the router**correct

Question 3

Which standard access control entry permits from odd-numbered hosts in the 10.0.0.0/24 subnet?


- A. Permit 10.0.0.0 255.255.255.254
- B. Permit 10.0.0.0 0.0.0.1
- C. Permit 10.0.0.1 0.0.0.254**correct
- D. Permit 10.0.0.1 0.0.0.0

Explanation

Remember, for the wildcard mask, **1's are I DON'T CARE**, and **0's are I CARE**. So now let's analyze a simple ACL:

access-list 1 permit 172.23.16.0 0.0.15.255

Two first octets are all 0's meaning that we care about the network **172.23.x.x**. The third octet of the wildcard mask, 15 (0000 1111 in binary), means that we care about first 4 bits but don't care about last 4 bits so we allow the third octet in the form of 0001xxxx (minimum: 0001**0000** = 16; maximum: 000**1111** = 31).



The fourth octet is 255 (all 1 bits) that means I don't care.

Therefore **network 172.23.16.0 0.0.15.255** ranges from **172.23.16.0** to **172.23.31.255**.

Now let's consider the wildcard mask of 0.0.0.254 (four octet: 254 = 1111 1110) which means we only care the last bit. Therefore if the last bit of the IP address is a "1" (0000 0001) then only odd numbers are allowed. If the last bit of the IP address is a "0" (0000 0000) then only even numbers are allowed.

Note: In binary, odd numbers are always end with a "1" while even numbers are always end with a "0".

Therefore in this question, only the statement "permit 10.0.0.1 0.0.0.254" will allow all odd-numbered hosts in the 10.0.0.0/24 subnet.

Question 4

Which access controls list allows only TCP traffic with a destination port range of 22-443, excluding port 80?

A. Permit tcp any any range 22 443

Deny tcp any any eq 80

<p< span="" style="border: 1px solid black; padding: 2px; display: inline-block;">correct</p>

B. Deny tcp any any eq 80

Permit tcp any any gt 21 lt 444

<p< span="" style="border: 1px solid black; padding: 2px; display: inline-block;"></p>

C. Permit tcp any any neq 80

<p< span="" style="border: 1px solid black; padding: 2px; display: inline-block;"></p>

D. Deny tcp any any neq 80

Permit tcp any any range 22 443

Explanation

Although the statement "permit tcp any any gt ... lt ..." seems to be correct but in fact it is not. Each ACL statement only supports either "gt" or "lt" but not both:


```

R1(config)#ip access-list extended test2
R1(config-ext-nacl)#deny tcp any any eq 80
R1(config-ext-nacl)#permit tcp any any gt 21 lt 444
                                     ^
% Invalid input detected at '^' marker.

R1(config-ext-nacl)#permit tcp any any lt 444 gt 21
                                     ^
% Invalid input detected at '^' marker.

R1(config-ext-nacl)#permit tcp any any gt 21
R1(config-ext-nacl)#permit tcp any any lt 444

```

In fact answer 'Permit tcp any any range 22 443
Deny tcp any any eq 80

eq 80".

Question 5

A client with IP address 209.165.201.25 must access a web server on port 80 at 209.165.200.225. To allow this traffic, an engineer must add a statement to an access control list that is applied in the inbound direction on the port connecting to the web server. Which statement allows this traffic?


- A. permit tcp host 209.165.200.225 host 209.165.201.25 eq 80
- B. permit tcp host 209.165.201.25 host 209.165.200.225 eq 80
- C. permit tcp host 209.165.201.25 eq 80 host 209.165.200.225

D. permit tcp host 209.165.200.225 eq 80 host 209.165.201.25 correct

Explanation

The syntax of an extended ACL is shown below:

access-list *access-list-number* {permit | deny} *protocol* source {source-mask} destination {destination-mask} [eq destination-port]



According to the request in this question, we must apply the ACL on the port connected to the Web Server and with inbound direction. So it can only filter traffic sent from the Web Server to the Client. Please notice that the Client communicate to the Web Server with destination port of 80 but with random source port. So the Web Server must answer the Client with this random port (as the destination port) -> Therefore the destination port in the required ACL must be ignored. Also the Web Server must use port 80 as its source port.

So the structure of the ACL should be: permit tcp host <IP-address-of-Web-Server> eq 80 host <IP-address-of-Client>

Answer "permit tcp host 209.165.200.225 eq 80 host 209.165.201.25" is correct.

Automation Quiz

Result of Automation Quiz:

Total Questions	Full Score	Passing Rate	Your Score	Correct Answer Percentage	Elapsed
12	150	90%	150	100%	00:01:52

Congratulations!

You passed this test!

If you want to retake this quiz, please press Ctrl + F5 on Windows or press CMD + R on Mac.

Your answers are shown below:

Question 1

Which two protocols are used with YANG data models? (Choose two)

- A. NETCONFcorrect
- B. RESTCONFcorrect
- C. SSH
- D. HTTPS
- E. TLS

Explanation

YANG (Yet Another Next Generation) is a data modeling language for the definition of data sent over network management protocols such as the NETCONF and RESTCONF.

Question 2

Which JSON syntax is valid?

- A. {"switch":{"name":"dist1","interfaces":["gig1","gig2","gig3"]}}correct
- B. {/"switch/":{/"name/":"dist1"/,"interfaces/":["gig1","gig2","gig3"]}}
- C. {"switch":"name":"dist1","interfaces":["gig1","gig2","gig3"]}
- D. {'switch':('name':'dist1','interfaces':['gig1','gig2','gig3'])}

Explanation

This JSON can be written as follows:

```
{
  'switch': {
    'name': 'dist1',
    'interfaces': ['gig1', 'gig2', 'gig3']
  }
}
```

Question 3

Which requirement for an Ansible-managed node is true?

- A. It must have an SSH server running
- B. It must support ad hoc commands.

C. It must have an Ansible Tower installed

D. It must be a Linux server or a Cisco devicecorrect

Explanation

Ansible can communicate with modern Cisco devices via SSH or HTTPS so it does not require an SSH server -> Answer 'It must have an SSH server running' is not correct.

An Ansible ad-hoc command uses the /usr/bin/ansible command-line tool to automate a single task on one or more managed nodes. Ad-hoc commands are quick and easy, but they are not reusable -> It is not a requirement either -> Answer 'It must support ad hoc commands' is not correct.

Ansible Tower is a web-based solution that makes Ansible even more easy to use for IT teams of all kinds. But it is not a requirement to run Ansible -> Answer 'It must have an Ansible Tower installed' is not correct.

Therefore only answer 'It must be a Linux server or a Cisco device' is the best choice left. An Ansible controller (the main component that manages the nodes), is supported on multiple flavors of Linux, but it cannot be installed on Windows.

Question 4

Which method creates an EEM applet policy that is registered with EEM and runs on demand or manually?

A. event manager applet ondemand
event register

action 1.0 syslog priority critical msg 'This is a message from ondemand'
<p< span="" style="border: 1px solid black; padding: 2px; display: inline-block;"></p></p>

B. event manager applet ondemand

action 1.0 syslog priority critical msg 'This is a message from ondemand'

C. event manager applet ondemand

event none

action 1.0 syslog priority critical msg 'This is a message from ondemand'

<p< span="" style="border: 1px solid black; padding: 2px; display: inline-block;">correct</p></p>

D. event manager applet ondemand

event manual

action 1.0 syslog priority critical msg 'This is a message from ondemand'

<p< span="" style="border: 1px solid black; padding: 2px; display: inline-block;"></p></p>

Explanation

An EEM policy is an entity that defines an event and the actions to be taken when that event occurs. There are two types of EEM policies: an applet or a script. An applet is a simple form of policy that is defined within the CLI configuration. answer 'event manager applet ondemand event register

action 1.0 syslog priority critical msg 'This is a message from ondemand'

<= "" p="" style="border: 1px solid black; padding: 2px; display: inline-block;">

There are two ways to manually run an EEM policy. EEM usually schedules and runs policies on the basis of an event specification that is contained within the policy itself. The **event none** command allows EEM to identify an EEM policy that can be manually triggered. To run the policy, use either the **action policy** command in applet configuration mode or the **event manager run** command in privileged EXEC mode.

Reference: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/eem/configuration/xr-3s/eem-xr-3s-book/eem-policy-cli.html>

Question 5

Which exhibit displays a valid JSON file?

A. {

"hostname": "edge_router_1",

"interfaces": [

```

"GigabitEthernet1/1",
"GigabitEthernet1/2",
"GigabitEthernet1/3"
]
}correct
B. {
  "hostname": "edge_router_1"
  "interfaces": [
    "GigabitEthernet1/1"
    "GigabitEthernet1/2"
    "GigabitEthernet1/3"
  ]
}
<p< span="" style="box-sizing: border-box;"></p><
C. {
  "hostname": "edge_router_1"
  "interfaces": {
    "GigabitEthernet1/1",
    "GigabitEthernet1/2",
    "GigabitEthernet1/3",
  },
}
<p< span="" style="box-sizing: border-box;"></p><
D. {
  "hostname": "edge_router_1"
  "interfaces": {
    "GigabitEthernet1/1"
    "GigabitEthernet1/2"
    "GigabitEthernet1/3"
  }
}
<p< span="" style="box-sizing: border-box;"></p><

```

Question 6

Which statement about TLS is true when using RESTCONF to write configurations on network devices?

- A. It is used for HTTP and HTTPs requests
- B. It is no supported on Cisco devices
- C. It is provided using NGINX acting as a proxy web server correct**
- D. It required certificates for authentication

Explanation

When a device boots up with the startup configuration, the *nginx* process will be running. **NGINX is an internal webserver that acts as a proxy webserver.** It provides Transport Layer Security (TLS)-based HTTPS. RESTCONF request sent via HTTPS is first received by the NGINX proxy web server, and the request is transferred to the confd web server for further syntax/semantics check.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/prog/configuration/168/b_168_programmability_cg/RESTCONF.html

The https-based protocol-RESTCONF (RFC 8040), which is a stateless protocol, uses secure HTTP methods to provide CREATE, READ, UPDATE and DELETE (CRUD) operations on a conceptual datastore containing YANG-defined data -> RESTCONF only uses HTTPs.

Note: In fact answer 'It required certificates for authentication' is also correct:

RESTCONF servers MUST present an X.509v3-based certificate when establishing a TLS connection with a RESTCONF client. The use of X.509v3-based certificates is consistent with NETCONF over TLS.

Reference: <https://tools.ietf.org/html/rfc8040>

But answer 'It is provided using NGINX acting as a proxy web server' is still a better choice.

Question 7

Refer to the exhibit. Which two statements about the EEM applet configuration are true? (Choose two)

```
event manager applet LARGECONFIG
event cli pattern 'show running-config' sync yes
action 1.0 puts 'Warning! This device has a VERY LARGE configuration
```

```
and may take some time to process'
```

```
action 1.1 puts newline 'Do you wish to continue [Y/N] '
action 1.2 gets response
action 1.3 string toupper '$response'
action 1.4 string match '$_string_result' 'Y'
action 2.0 if $_string_result eq 1
action 2.1 cli command 'enable'
action 2.2 cli command 'show running-config'
action 2.3 puts $_cli_result
action 2.4 cli command 'exit'
action 2.9 end
```

- A. The EEM applet runs after the CLI command is executed
- B. The running configuration is displayed only if the letter Y is entered at the CLI correct**
- C. The EEM applet runs before the CLI command is executed correct**
- D. The EEM applet requires a case-insensitive response

Explanation

When you use the **sync yes** option in the event cli command, the EEM applet runs before the CLI command is executed. The EEM applet should set the `_exit_status` variable to indicate whether the CLI command should be executed (`_exit_status` set to one) or not (`_exit_status` set to zero).

With the **sync no** option, the EEM applet is executed in background in parallel with the CLI command.

Reference: <https://blog.ipSPACE.net/2011/01/eem-event-cli-command-options-and.html>

Question 8

Which protocol does REST API rely on to secure the communication channel?

- A. HTTPS correct**
- B. TCP
- C. SSH
- D. HTTP

Explanation

The REST API accepts and returns HTTP (not enabled by default) or HTTPS messages that contain JavaScript Object Notation (JSON) or Extensible Markup Language (XML) documents. You can use any programming language to generate the messages and the JSON or XML documents that contain the API methods or Managed Object (MO) descriptions.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/2-x/rest_cfg/2_1_x/b_Cisco_APIC_REST_API_Configuration_Guide/b_Cisco_APIC_REST_API_Configuration_Guide_chapter_01.html

Question 9

Which two operations are valid for RESTCONF? (Choose two)

- A. ADD
- B. HEAD correct**
- C. PULL
- D. PUSH
- E. REMOVE
- F. PATCH correct**

Explanation

RESTCONF operations include OPTIONS, HEAD, GET, POST, PATCH, DELETE.

Question 10

What does this EEM applet event accomplish?

“event snmp oid 1.3.6.1.3.7.1.5.1.2.4.2.9 get-type next entry-op go entry-val 75 poll-interval 5”

A. It reads an SNMP variable, and when the value exceeds 75%, it triggers an action GO correct

- B. Upon the value reaching 75%, a SNMP event is generated and sent to the trap server
- C. It issues email when the value is greater than 75% for five polling cycles
- D. It presents a SNMP variable that can be interrogated

Explanation

EEM offers the ability to monitor events and take informational or corrective action when the monitored events occur or reach a threshold. An EEM policy is an entity that defines an event and the actions to be taken when that event occurs. There are two types of EEM policies: an applet or a script. An applet is a simple form of policy that is defined within the CLI configuration.

To specify the event criteria for an Embedded Event Manager (EEM) applet that is run by sampling Simple Network Management Protocol (SNMP) object identifier values, use the event snmp command in applet configuration mode.

event snmp oid *oid-value* **get-type** {exact | next} **entry-op** *operator* **entry-val** *entry-value* [**exit-comb** {or | and}] [**exit-op** *operator*] [**exit-val** *exit-value*] [**exit-time** *exit-time-value*] **poll-interval** *poll-int-value*

- + oid: Specifies the SNMP object identifier (object ID)
- + get-type: Specifies the type of SNMP get operation to be applied to the object ID specified by the oid-value argument.
- next – Retrieves the object ID that is the alphanumeric successor to the object ID specified by the oid-value argument.
- + entry-op: Compares the contents of the current object ID with the entry value using the specified operator. **If there is a match, an event is triggered** and event monitoring is disabled until the exit criteria are met.
- + entry-val: Specifies the value with which the contents of the current object ID are compared to decide if an SNMP event should be raised.
- + exit-op: Compares the contents of the current object ID with the exit value using the specified operator. If there is a match, an event is triggered and event monitoring is reenabled.
- + poll-interval: Specifies the time interval between consecutive polls (in seconds)

Reference: https://www.cisco.com/en/US/docs/ios/12_3t/12_3t4/feature/guide/gtioseem.html

Question 11

What is the structure of a JSON web token?

A. header and payload

B. three parts separated by dots header payload, and signature correct

C. three parts separated by dots version header and signature

D. payload and signature

Explanation

JSON Web Token (JWT) is an open standard (RFC 7519) that defines a compact and self-contained way for securely transmitting information between parties as a JSON object. This information can be verified and trusted because it is digitally signed. JWTs can be signed using a secret (with the HMAC algorithm) or a public/private key pair using RSA or ECDSA.

JSON Web Tokens are composed of three parts, separated by a dot (.): Header, Payload, Signature. Therefore, a JWT typically looks like the following:

xxxxx.yyyyy.zzzzz

The header typically consists of two parts: the type of the token, which is JWT, and the signing algorithm being used, such as HMAC SHA256 or RSA.

The second part of the token is the payload, which contains the claims. Claims are statements about an entity (typically, the user) and additional data.

To create the signature part you have to take the encoded header, the encoded payload, a secret, the algorithm specified in the header, and sign that.

Reference: <https://jwt.io/introduction/>

Question 12

Refer to the exhibit. Which network script automation option or tool is used in the exhibit?

<https://mydevice.mycompany.com/getstuff?queryName=errors&queryResults=yes>

A. EEM

B. Bash script

C. REST correct

D. NETCONF

E. Python

Automation Quiz 2

Result of Automation Quiz 2:

Total Questions	Full Score	Passing Rate	Your Score	Correct Answer Percentage	Elapsed
12	120	90%	120	100%	00:01:11

Congratulations!

You passed this test!

If you want to retake this quiz, please press Ctrl + F5 on Windows or press CMD + R on Mac.

Your answers are shown below:

Question 1

What do Cisco DNA southbound APIs provide?


A. Interface between the controller and the network devices correct

B. Interface between the controller and the consumer

- C. RESTful API interface for orchestrator communication
- D. NETCONF API interface for orchestration communication

Explanation

The Southbound API is used to communicate with network devices.



Question 2

Which statement about agent-based versus agentless configuration management tools is true?

- A. Agentless tools use proxy nodes to interface with slave nodes.
- B. Agentless tools require no messaging systems between master and slaves.
- C. Agent-based tools do not require installation of additional software packages on the slave nodes.

D. Agent-based tools do not require a high-level language interpreter such as Python or Ruby on slave nodes.correct

Explanation

Agentless tool means that no software or agent needs to be installed on the client machines that are to be managed. Ansible is such an agentless tool. In contrast to agentless tool, agent-based tool requires software or agent to be installed on the client. Therefore the master and slave nodes can communicate directly without the need of high-level language interpreter.

An agentless tool uses standard protocols, such as SSH, to push configurations down to a device (and it can be considered a "messaging system").

Question 3

What is a benefit of data modeling languages like YANG?

- A. They enable programmers to change or write their own application within the device operating system.

B. They provide a standardized data structure, which results in configuration scalability and consistency.correct

- C. They create more secure and efficient SNMP OIDs.
- D. They make the CLI simpler and more efficient.

Explanation

Yet Another Next Generation (YANG) is a language which is only used to describe data models (structure). It is not XML or JSON.

Question 4

Which method displays text directly into the active console with a synchronous EEM applet policy?

A. event manager applet boom
 event syslog pattern 'UP'
 action 1.0 syslog priority direct msg 'log directly to console'
 <p< span="" style="box-sizing: border-box;"></p><>

B. event manager applet boom
 event syslog pattern 'UP'
 action 1.0 gets 'logging directly to console'
 <p< span="" style="box-sizing: border-box;"></p><>

C. event manager applet boom
 event syslog pattern 'UP'
 action 1.0 string 'logging directly to console'

D. event manager applet boom
event syslog pattern 'UP'
action 1.0 puts 'logging directly to console'
 <p< span="" style="box-sizing: border-box;">correct</p><>

Explanation

To enable the action of printing data directly to the local tty when an Embedded Event Manager (EEM) applet is triggered, use the **action puts** command in applet configuration mode.

The following example shows how to print data directly to the local tty:

```
Router(config-applet)# event manager applet puts
Router(config-applet)# event none
Router(config-applet)# action 1 regexp "(.*) (.*) (.*)" "one two three" _match _sub1
Router(config-applet)# action 2 puts "match is $_match"
Router(config-applet)# action 3 puts "submatch 1 is $_sub1"
Router# event manager run puts
match is one two three
submatch 1 is one
Router#
```

The **action puts** command applies to synchronous events. The output of this command for a synchronous applet is directly displayed to the tty, bypassing the syslog.

Reference: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/eem/command/eem-cr-book/eem-cr-a1.html>

Question 5

Which data modeling language is commonly used by NETCONF?

- A. REST
- B. YANG**correct
- C. HTML
- D. XML

Explanation

Cisco IOS XE supports the Yet Another Next Generation (YANG) data modeling language. YANG can be used with the Network Configuration Protocol (NETCONF) to provide the desired solution of automated and programmable network operations. NETCONF(RFC6241) is an XML-based protocol that client applications use to request information from and make configuration changes to the device. YANG is primarily used to model the configuration and state data used by NETCONF operations.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9500/software/release/16-5/configuration_guide/prog/b_165_prog_9500_cg/data_models.pdf

Note: Although NETCONF also uses XML but XML is not a data modeling language.

Question 6

Which variable in an EEM applet is set when you use the sync yes option?

- A. \$_cli_result
- B. \$_exit_statuscorrect
- C. \$_string_result
- D. \$_result

Explanation

With Synchronous (sync yes), the CLI command in question is not executed until the policy exits. Whether or not the command runs depends on the value for the variable _exit_status. If _exit_status is 1, the command runs, if it is 0, the command is skipped.

Question 7

Refer to the exhibit. Which HTTP JSON response does the python code output give?

PYTHON CODE

```
import requests
import json
```

```
url='http://YOURIP/ins'
switchuser='USERID'
switchpassword='PASSWORD'
```

```
myheaders={'content-type':'application/json'}
payload={
  'ins_api': {
    'version':'1.0',
    'type':'cli_show',
    'chunk':'0',
    'sid':'1',
    'input':'show version',
    'output_format':'json'
  }
}
response = requests.post(url,data=json.dumps(payload) ,
headers=myheaders,auth=(switchuser,switchpassword)).json()
```

```
print(response['ins_api']['outputs'][output]['body']['kickstart_ver_str']
)
```

HTTP JSON Response:

```
{
  'ins_api': {
    'type': 'cli_show',
    'version': '1.0',
    'sid': 'eoc',
    'outputs': {
      'output': {
        'input': 'show version',
        'msg': 'Success',
        'code': '200',
        'body': {
          'bios_ver_str': '07.61',
          'kickstart_ver_str': '7.0(3)I7(4)',
          'bios_cmpl_time': '04/08/2017',
          'kick_file_name': 'bootflash:///nxos.7.0.3.I7.4.bin',
          'kick_cmpl_time': '6/14/1970 09:49:04',
          'chassis_id': 'Nexus9000 93180YC-EX chassis',
          'cpu_name': 'Intel(R) Xeon(R) CPU @1.80GHz',
          'memory': 24633488,
          'mem_type': 'kB',
          'rr_usecs': 134703,
          'rr_ctime': 'Sun Mar 10 15:41:46 2019',
          'rr_reason': 'Reset Requested by CLI command reload',
          'rr_sys_ver': '7.0(3)I7(4)',
          'rr_service': '',
          'manufacturer': 'Cisco Systems, Inc',
          'TABLE_package_list': {
```

```
    'ROW_package_list': {
      'package_id': {}
    }
  }
}
```

```
}
}
}
}
}
}
```

- A. KeyError 'kickstart_ver_str'
- B. NameError: name 'json' is not defined
- C. 7.61
- D. 7.0(3)I7(4)correct**

Explanation

If we have a JSON string, we can parse it by using the `json.loads()` method so we don't need to have a response server to test this question. Therefore, in order to test the result above, you can try this Python code:

```
import json
json_string = '''
{
  'ins_api': { <!!!Please copy the code above and put here. We omitted it to
save some space!!!>
}
}
'''
response = json.loads(json_string)
```

```
print(response['ins_api']['outputs']['output']['body']['kickstart_ver_str'
])
```

And this is the result:

```

C:\WINDOWS\system32>python
Python 3.7.0 (v3.7.0:1bf9cc5093, Jun 27 2018, 04:06:47) [MSC v.1914 32 bit (I
Type "help", "copyright", "credits" or "license" for more information.
>>> import json
>>> json_string = """
... {
...   "ins_api": {
...     "type": "cli_show",
...     "version": "1.0",
...     "sid": "eoc",
...     "outputs": {
...       "output": {
...         "input": "show version",
...         "msg": "Success",
...         "code": "200",
...         "body": {
...           "bios_ver_str": "07.61",
...           "kickstart_ver_str": "7.0(3)I7(4)",
...           "bios_cmpl_time": "04/08/2017",
...           "kick_file_name": "bootflash:///nxos.7.0.3.I7.4.bin",
...           "kick_cmpl_time": "6/14/1970 09:49:04",
...           "chassis_id": "Nexus9000 93180YC-EX chassis",
...           "cpu_name": "Intel(R) Xeon(R) CPU @1.80GHz",
...           "memory": 24633488,
...           "mem_type": "kB",
...           "rr_usecs": 134703,
...           "rr_ctime": "Mon Jun 10 12:34:46 2019",
...           "rr_reason": "Reset Requested by CLI command reload",
...           "rr_sys_ver": "7.0(3)I7(4)",
...           "rr_service": "",
...           "manufacturer": "Cisco Systems, Inc",
...           "TABLE_package_list": {
...             "ROW_package_list": {
...               "package_id": {}
...             }
...           }
...         }
...       }
...     }
...   }
... }
... """
>>> response = json.loads(json_string)
>>> print(response['ins_api']['outputs']['output']['body']['kickstart_ver_str
7.0(3)I7(4)

```

Note:

- + If you want to run the full code in this question in Python (with a real HTTP JSON response), you must first install "requests" package before "import requests".
- + The error "NameError: name 'json' is not defined" is only shown if we forgot the line "import json" in Python code -> Answer 'NameError: name 'json' is not defined' is not correct.
- + We only see the "KeyError" message if we try to print out an unknown attribute (key). For example:

```
print(response['ins_api']['outputs']['output']['body']['unknown_attribute'])
```

```
>>> print(response['ins_api']['outputs']['output']['body']['unknown_attribute'])
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
KeyError: 'unknown_attribute'
>>>
```

- + Triple quotes (""") in Python allows strings to span multiple lines, including verbatim NEWLINES, TABs, and any other special characters.

Question 8

In which part of the HTTP message is the content type specified?

- A. HTTP method
- B. URI
- C. body
- D. header** correct

Question 9

Refer to the exhibit. What is the JSON syntax that is formed the data?

Name is Bob Johnson
Age is 76
Is alive

Favorite foods are:
+ Cereal
+ Mustard
+ Onions

- A. Name: Bob, Johnson, Age: 76, Alive: true, Favourite Foods. [Cereal, "Mustard", "Onions"]}
- B. Name', 'Bob Johnson', 'Age', 76, 'Alive', true, 'favourite Foods' 'Cereal Mustard', 'Onions'}
- C. {"Name": "Bob Johnson", "age": 76, "alive": true, "favorite foods": ["Cereal", "Mustard", "Onions"]}** correct
- D. Name", "Bob Johnson", "Age", 76, "Alive", true, "favourite Foods", ["Cereal, "Mustard", Onions"]}
- E. Name", "Bob Johnson", "Age": Seventysix, "Alive" true, "favourite Foods", [Cereal" "Mustard" "Onions"]}

Explanation

JSON data is written as name/value pairs.

A name/value pair consists of a field name (in double quotes), followed by a colon, followed by a value:

"name": "Mark"

JSON can use arrays. Array values must be of type string, number, object, array, boolean or null. For example:

```
{
  "name": "John",
  "age": 30,
  "alive": true,
  "cars": [ "Ford", "BMW", "Fiat" ]
}
```

Question 10

Which statements are used for error handling in Python?

- A. try/catch
- B. block/rescue
- C. try/except** correct
- D. catch/release

Explanation

The words "try" and "except" are Python keywords and are used to catch exceptions. For example:

```
try:
    print 1/0
except ZeroDivisionError:
    print 'Error! We cannot divide by zero!!!'
```

Question 11

Which HTTP status code is the correct response for a request with an incorrect password applied to a REST API session?

- A. HTTP Status Code 401** correct
- B. HTTP Status Code 504
- C. HTTP Status Code 302
- D. HTTP Status Code 200

Explanation

A 401 error response indicates that the client tried to operate on a protected resource without providing the proper authorization. It may have provided the wrong credentials or none at all.

Note: answer 'HTTP Status Code 200' 4xx code indicates a "client error" while a 5xx code indicates a "server error".

Reference: <https://restfulapi.net/http-status-codes/>

Question 12

A response code of 404 is received while using the REST API on Cisco UNA Center to POST to this URL

/dna/intent/api/v1 /template-programmer/project

What does the code mean?

- A. The POST/PUT request was fulfilled and a new resource was created, information about the resource is in the response body
- B. The request accepted for processing, but the processing was not completed
- C. The server has not implemented the functionality that is needed to fulfill the request
- D. The client made a request a resource that does not exist** correct

Explanation

The 404 (Not Found) error status code indicates that the REST API can't map the client's URI to a resource but may be available in the future. Subsequent requests by the client are permissible.

Reference: <https://restfulapi.net/http-status-codes/>

Miscellaneous Questions

Result of Miscellaneous Quiz:

Total Questions	Full Score	Passing Rate	Your Score	Correct Answer Percentage	Elapsed
10	150	90%	140	90%	00:01:18

Congratulations!

You passed this test!

If you want to retake this quiz, please press Ctrl + F5 on Windows or press CMD + R on Mac.

Your answers are shown below:

Question 1

Which two GRE features are configured to prevent fragmentation? (Choose two)

- A. TCP MSScorrect
- B. PMTUDcorrect
- C. DF bit Clear
- D. MTU ignore
- E. IP MTU
- F. TCP window size

Explanation

The IP protocol was designed for use on a wide variety of transmission links. Although the maximum length of an IP datagram is 65535, most transmission links enforce a smaller maximum packet length limit, called an MTU. The value of the MTU depends on the type of the transmission link. The design of IP accommodates MTU differences since it allows routers to fragment IP datagrams as necessary. The receiving station is responsible for the reassembly of the fragments back into the original full size IP datagram.

Fragmentation and Path Maximum Transmission Unit Discovery (PMTUD) is a standardized technique to determine the maximum transmission unit (MTU) size on the network path between two hosts, usually with the goal of **avoiding IP fragmentation**. PMTUD was originally intended for routers in IPv4. However, all modern operating systems use it on endpoints.

The TCP Maximum Segment Size (TCP MSS) defines the maximum amount of data that a host is willing to accept in a single TCP/IP datagram. This TCP/IP datagram might be fragmented at the IP layer. The MSS value is sent as a TCP header option only in TCP SYN segments. Each side of a TCP connection reports its MSS value to the other side. Contrary to popular belief, the MSS value is not negotiated between hosts. The sending host is required to limit the size of data in a single TCP segment to a value less than or equal to the MSS reported by the receiving host.

TCP MSS takes care of fragmentation at the two endpoints of a TCP connection, but it does not handle the case where there is a smaller MTU link in the middle between these two endpoints. PMTUD was developed in order to avoid fragmentation in the path between the endpoints. It is used to dynamically determine the lowest MTU along the path from a packet's source to its destination.

Reference: <http://www.cisco.com/c/en/us/support/docs/ip/generic-routing-encapsulation-gre/25885-pmtud-ipfrag.html> (there is some examples of how TCP MSS avoids IP Fragmentation in this link but it is too long so if you want to read please visit this link)

Note: IP fragmentation involves breaking a datagram into a number of pieces that can be reassembled later.

If the DF bit is set to clear, routers can fragment packets regardless of the original DF bit setting -
> Answer 'DF bit Clear' is not correct.

Question 2

Which IPv6 migration method relies on dynamic tunnels that use the 2002::/16 reserved address space?

- A. GRE
- B. 6RD
- C. 6to4**correct
- D. ISATAP

Explanation

6to4 tunnel is a technique which relies on reserved address space 2002::/16 (you must remember this range). These tunnels determine the appropriate destination address by combining the IPv6 prefix with the globally unique destination 6to4 border router's IPv4 address, beginning with the 2002::/16 prefix, in this format:

2002:border-router-IPv4-address::/48

For example, if the border-router-IPv4-address is 64.101.64.1, the tunnel interface will have an IPv6 prefix of 2002:4065:4001:1::/64, where 4065:4001 is the hexadecimal equivalent of 64.101.64.1. This technique allows IPv6 sites to communicate with each other over the IPv4 network without explicit tunnel setup but we have to implement it on all routers on the path.

Question 3

Which two mechanisms are available to secure NTP? (Choose two)

- A. TACACS-based authentication
- B. IP access list-based**correct
- C. IP prefix list-based
- D. IPsec
- E. Encrypted authentication**correct

Explanation

The time kept on a machine is a critical resource and it is strongly recommend that you use the security features of NTP to avoid the accidental or malicious setting of incorrect time. **The two security features available are an access list-based restriction scheme and an encrypted authentication mechanism.**

Reference: <https://www.cisco.com/c/en/us/support/docs/availability/high-availability/19643-ntp.html>

Question 4

A GRE tunnel is down with the error message %TUN-5-RECUR DOWN:

Tunnel0 temporarily disabled due to recursive routing error.

Which two options describe possible causes of the error? (Choose two)

- A. There is link flapping on the tunnel
- B. Incorrect destination IP addresses are configured on the tunnel
- C. The tunnel mode and tunnel IP address are misconfigured
- D. There is instability in the network due to route flapping**correct
- E. The tunnel destination is being routed out of the tunnel interface**correct

Explanation

The **%TUN-5-RECURDOWN: Tunnel0 temporarily disabled due to recursive routing error** message means that the generic routing encapsulation (GRE) tunnel router has discovered a recursive routing problem. This condition is usually due to one of these causes:

+ A misconfiguration that causes the router to try to route to the tunnel destination address using

the tunnel interface itself (recursive routing)

+ A temporary instability caused by route flapping elsewhere in the network

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/22327-gre-flap.html>

Question 5

A network administrator is implementing a routing configuration change and enables routing debugs to track routing behavior during the change. The logging output on the terminal is interrupting the command typing process. Which two actions can the network administrator take to minimize the possibility of typing commands incorrectly? (Choose two)

A. Configure the logging synchronous global configuration command

B. Configure the logging delimiter feature

C. Press the TAB key to reprint the command in a new line correct

D. Configure the logging synchronous command under the vty correct

E. Increase the number of lines on the screen using the terminal length command

Explanation

Although some Cisco webpages (like [this one](#)) mentioned about "logging synchronous" command in global configuration mode, which means "Router(config)#logging synchronous", but in fact we cannot use it under global configuration mode. We can only use this command in line mode. Therefore answer 'Configure the logging synchronous command under the vty' is better than answer A.

Let's see how the "logging synchronous" command affect the typing command:

Without this command, a message may pop up and you may not know what you typed if that message is too long. When trying to erase (backspace) your command, you realize you are erasing the message instead.

```
NVbos2811-1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
NVbos2811-1(config)#^Z
NVbos2811-1#sh
Jan 18 16:38:02: %SYS-5-CONFIG_I: Configured from console by admin on vty0 (10.0.1.111)
```

With this command enabled, when a message pops up you will be put to a new line with your typing command which is very nice:

```
NVbos2811-1(config)#line con 0
NVbos2811-1(config-line)#logging synch
NVbos2811-1(config-line)#line vty 0 4
NVbos2811-1(config-line)#logging synchr
NVbos2811-1(config-line)#logging synchronous
NVbos2811-1(config-line)#^Z
NVbos2811-1#sh ip
Jan 18 16:39:33: %SYS-5-CONFIG_I: Configured from console by admin
NVbos2811-1#sh ip
```

Question 6

Which TCP setting is tuned to minimize the risk of fragmentation on a GRE/IP tunnel?

A. MTU

B. MRU

C. Window size

D. MSS correct

Explanation

The TCP Maximum Segment Size (TCP MSS) defines the maximum amount of data that a host is willing to accept in a single TCP/IP datagram. This TCP/IP datagram might be fragmented at the IP layer. The MSS value is sent as a TCP header option only in TCP SYN segments. Each side of a TCP connection reports its MSS value to the other side. Contrary to popular belief, the MSS value is not negotiated between hosts. The sending host is required to limit the size of data in a single TCP segment to a value less than or equal to the MSS reported by the receiving host.

TCP MSS takes care of fragmentation at the two endpoints of a TCP connection, but it does not handle the case where there is a smaller MTU link in the middle between these two endpoints. PMTUD was developed in order to avoid fragmentation in the path between the endpoints. It is used to dynamically determine the lowest MTU along the path from a packet's source to its destination.

Reference: <http://www.cisco.com/c/en/us/support/docs/ip/generic-routing-encapsulation-gre/25885-pmtud-ipfrag.html> (there is some examples of how TCP MSS avoids IP Fragmentation in this link but it is too long so if you want to read please visit this link)

Note: IP fragmentation involves breaking a datagram into a number of pieces that can be reassembled later.

Question 7

Refer to the exhibit. What are two effect of this configuration? (Choose two)

```
access-list 1 permit 10.1.1.0 0.0.0.31
ip nat pool CISCO 209.165.201.1 209.165.201.30 netmask 255.255.255.224
ip nat inside source list 1 pool CISCO
```

- A. It establishes a one-to-one NAT translation
- B. The 10.1.1.0/27 subnet is assigned as the inside global address range
- C. The 209.165.201.0/27 subnet is assigned as the outside local address range
- D. Inside source addresses are translated to the 209.165.201.0/27 subnet**correct
- E. The 10.1.1.0/27 subnet is assigned as the inside local addresses**correct

Explanation

In this question, the inside local addresses of the 10.1.1.0/27 subnet are translated into 209.165.201.0/27 subnet. This is one-to-one NAT translation as the keyword "overload" is missing so in fact answer 'It establishes a one-to-one NAT translation' is also correct.

Question 8

What is a benefit of deploying an on-premises infrastructure versus a cloud infrastructure deployment?

- A. less power and cooling resources needed to run infrastructure on-premises
 - B. faster deployment times because additional infrastructure does not need to be purchased
 - C. lower latency between systems that are physically located near each othercorrect
 - D. ability to quickly increase compute power without the need to install additional hardware
- Question was not answered


Explanation

The difference between on-premise and cloud is essentially where this hardware and software resides. On-premise means that a company keeps all of this IT environment onsite either managed by themselves or a third-party. Cloud means that it is housed offsite with someone else responsible for monitoring and maintaining it.

Question 9

Refer to the exhibit. Link1 is a copper connection and Link2 is a fiber connection. The fiber port must be the primary port for all forwarding. The output of the show spanning-tree command on SW2 shows that the fiber port is blocked by spanning tree. An

engineer enters the spanning-tree port-priority 32 command on G0/1 on SW2, but the port remains blocked.



```
SW2#show spanning-tree
VLAN0010
Spanning tree enabled protocol ieee
Root ID      Priority  24596
```

```
Address      0018.7363.4300
Cost         2
Port         13 (GigabitEthernet0/0)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID      Priority  28692 (priority 28672 sys-id-ext 20)
```

```
Address      001b.0d8e.e080
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Aging Time 300 sec
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Gi0/0	Root	FWD	4	128.1	P2p
Gi0/1	Atln	BLK	4	32.2	P2p

Which command should be entered on the ports that are connected to Link2 to resolve the issue?

- A. Enter spanning-tree port-priority 32 on SW1 correct
- B. Enter spanning-tree port-priority 64 on SW2
- C. Enter spanning-tree port-priority 224 on SW1
- D. Enter spanning-tree port-priority 4 on SW2

Explanation

SW1 needs to block one of its ports to SW2 to avoid a bridging loop between the two switches. Unfortunately, it blocked the fiber port Link2. But how does SW2 select its blocked port? Well, the answer is based on the BPDUs it receives from SW1. answer 'Enter spanning-tree port-priority 32 on SW1' BPDUs are superior than another if it has:

1. answer 'Enter spanning-tree port-priority 32 on SW1' lower Root Bridge ID
2. answer 'Enter spanning-tree port-priority 32 on SW1' lower path cost to the Root
3. answer 'Enter spanning-tree port-priority 32 on SW1' lower Sending Bridge ID
4. answer 'Enter spanning-tree port-priority 32 on SW1' lower Sending Port ID

These four parameters are examined in order. In this specific case, all the BPDUs sent by SW1 have the same Root Bridge ID, the same path cost to the Root and the same Sending Bridge ID. The only parameter left to select the best one is the Sending Port ID (Port ID = port priority +

port index). And the port index of Gi0/0 is lower than the port index of Gi0/1 so Link 1 has been chosen as the primary link.

Therefore we must change the port priority to change the primary link. The lower numerical value of port priority, the higher priority that port has. In other words, we must change the port-priority on Gi0/1 of SW1 (not on Gi0/1 of SW2) to a lower value than that of Gi0/0.

Question 10

Which statement about multicast RPs is true?

- A. By default, the RP is needed periodically to maintain sessions with sources and receivers
- B. RPs are required for protocol independent multicast sparse mode and dense mode
- C. RPs are required only when using protocol independent multicast dense mode
- D. By default, the RP is needed only to start new sessions with sources and receivers** correct

Explanation

A rendezvous point (RP) is required only in networks running Protocol Independent Multicast sparse mode (PIM-SM).

By default, the RP is needed only to start new sessions with sources and receivers.

Reference: https://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/ip_multicast/White_papers/rps.html

For your information, in PIM-SM, only network segments with active receivers that have explicitly requested multicast data will be forwarded the traffic. This method of delivering multicast data is in contrast to the PIM dense mode (PIM-DM) model. In PIM-DM, multicast traffic is initially flooded to all segments of the network. Routers that have no downstream neighbors or directly connected receivers prune back the unwanted traffic.

You are configuring a controller that runs Cisco IOS XE by using the CLI. Which three configuration options are used for 802.11w Protected Management Frames? (Choose three.)

- A. mandatory**
- B. association-comeback**
- C. SA teardown protection
- D. saquery-retry-time**
- E. enable
- F. comeback-time

A,B,D

During deployment, a network engineer notices that voice traffic is not being tagged correctly as it traverses the network. Which COS to DSCP map must be modified to ensure that voice traffic is treated properly?

- A. COS of 5 to DSCP 46**
- B. COS of 7 to DSCP 48
- C. COS of 6 to DSCP 46
- D. COS of 3 to DSCP of 26

Correct Answer: A

Refer to the exhibit. A wireless client is connecting to FlexAP1 which is currently working standalone mode. The AAA authentication process is returning the following AVPs:

Tunnel-Private-Group-Id(81): 15 Tunnel-Medium-Type(65): IEEE-802(6) Tunnel-Type(64): VLAN(13)

Which three behaviors will the client experience? (Choose three.)

- A. While the AP is in standalone mode, the client will be placed in VLAN 15.**
- B. While the AP is in standalone mode, the client will be placed in VLAN 10.

- C. When the AP transitions to connected mode, the client will be de-authenticated.
 - D. While the AP is in standalone mode, the client will be placed in VLAN 13.**
 - E. When the AP is in connected mode, the client will be placed in VLAN 13.**
 - F. When the AP transitions to connected mode, the client will remain associated.
 - G. When the AP is in connected mode, the client will be placed in VLAN 15.
 - H. When the AP is in connected mode, the client will be placed in VLAN 10.
- Correct Answer: ADE

Drag and Drop

Drag and drop the LISP devices from the left onto the correct descriptions on the right.
Select and Place:

ETR – receives packets from site facing interfaces
 ITR – receives packets from core-facing interfaces
 PETR – provides connectivity from non-LISP sites and LISP sites
 PITR – allows IPV6 LISP sites without native IPV6 RLOC connectivity

Refer to the exhibit. An engineer is designing a guest portal on Cisco ISE using the default configuration. During the testing phase, the engineer receives a warning when displaying the guest portal.(YOUR CONNECTION IS NOT PRIVATE WARNING) Which issue is occurring?

- A. The server that is providing the portal has an expired certificate
- B. The server that is providing the portal has a self-signed certificate**
- C. The connection is using an unsupported protocol
- D. The connection is using an unsupported browser

What would be the preferred way to implement a loopless switch network where there are 1500 defined VLANs and it is necessary to load the shared traffic through two main aggregation points based on the VLAN identifier?

- a. 802.1D
- b. 802.1s**
- c. 802.1W
- d. 802.1AE

Not A

Not D: providing secure MAC service to the client.

802.1 w : Rapid Spanning Tree Protocol (RSTP) as 802.1w. RSTP provides significantly faster recovery

802.1 s : This Supplement to IEEE Std 802.1Q adds the facility for VLAN bridges to use multiple spanning trees, providing for traffic belonging to different VLANs to flow over potentially different paths within the virtual bridged LAN.

So I think the best answer is : 802.1 s