

New ENCOR Questions – Part 3

Question 1

Which deployment option of Cisco NGFW provides scalability?

- A. tap
- B. clustering
- C. inline tap
- D. high availability


Answer: B

Explanation

Clustering lets you group multiple Firepower Threat Defense (FTD) units together as a single logical device. Clustering is only supported for the FTD device on the Firepower 9300 and the Firepower 4100 series. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the **increased throughput and redundancy** of multiple devices.

Question 2

Refer to the exhibit.



```

R1#traceroute
Protocol [ip]:
Target IP address: 3.3.3.3
Source address: 1.1.1.1
Numeric display [n]:
Timeout in seconds [3]:
Probe count [3]:
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]: Record
Number of hops [9]:
Loose, Strict, Record, Timestamp, Verbose [RV]:
Type escape sequence to abort.

Continued --->

```

```

Tracing the route to 3.3.3.3
 1 10.99.69.2 36 msec
Received packet has options
Total option bytes = 40, padded length=40
Record route:
(10.99.69.1) <*>
(0.0.0.0)
(0.0.0.0)
End of list
---output omitted---

 2 10.99.69.6 1A
Received packet has options
Total option bytes = 40, padded length=40
Record route:
(10.99.69.1)
(10.99.69.5) <*>
(0.0.0.0)
(0.0.0.0)
End of list
1A
---output omitted---

```

The traceroute fails from R1 to R3. What is the cause of the failure?

- A. An ACL applied inbound on fa0/1 of R3 is dropping the traffic
- B. An ACL applied inbound on loopback0 of R2 is dropping the traffic
- C. The loopback on R3 is in a shutdown state
- D. Redistribution of connected routes into OSPF is not configured

Answer: A

Explanation

We see in the traceroute result the packet could reach 10.99.69.5 (on R2) but it could not go any further so we can deduce an ACL on R3 was blocking it.

Note: Record option displays the address(es) of the hops (up to nine) the packet goes through.

Question 3

Refer to the exhibit.

```

flow record v4_r1
 match ipv4 tos
 match ipv4 protocol
 match ipv4 source address
 match ipv4 destination address
 match transport source-port
 match transport destination-port
 collect counter bytes long
 collect counter packets long
!
flow monitor FLOW-MONITOR-1
 record v_r1
 exit
!
sampler SAMPLER-1
 mode random 1 out-of 2
 exit
!
ip cef
!
interface GigabitEthernet0/0/0
 ip address 172.16.6.2 255.255.255.0

```

Option A sampler SAMPLER-1 mode random 1-out-of 2 flow FLOW-MONITOR-1 interface GigabitEthernet0/0/0 ip flow monitor SAMPLER-1 input	Option B sampler SAMPLER-1 no mode random 1-out-of 2 mode percent 50 interface GigabitEthernet0/0/0 ip flow monitor FLOW_MONITOR-1 sampler S.
Option C interface GigabitEthernet0/0/0 ip flow monitor FLOW-MONITOR-1 sampler S.	Option D flow monitor FLOW-MONITOR-1 record v4_r1 sampler SAMPLER-1 interface GigabitEthernet0/0/0 ip flow monitor FLOW-MONITOR-1 sampler S.

Which command set must be added to the configuration to analyze 50 packets out of every 100?

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

Question 4


An engineer must configure a ACL that permits packets which include an ACK in the TCP header. Which entry must be included in the ACL?

- A. access-list 110 permit tcp any any eq 21 tcp-ack
- B. access-list 10 permit ip any any eq 21 tcp-ack
- C. access-list 10 permit tcp any any eq 21 established
- D. access-list 110 permit tcp any any eq 21 established

Answer: D

Explanation

The **established** keyword is only applicable to TCP access list entries to match TCP segments that have the ACK and/or RST control bit set (regardless of the source and destination ports), which assumes that a TCP connection has already been established in one direction only. Let's see an example below:




Suppose you only want to allow the hosts inside your company to telnet to an outside server but not vice versa, you can simply use an “established” access-list like this:

```
access-list 100 permit tcp any any established
access-list 101 permit tcp any any eq telnet
!
interface S0/0
ip access-group 100 in
ip access-group 101 out
```

Note:

Suppose host A wants to start communicating with host B using TCP. Before they can send real data, a three-way handshake must be established first. Let's see how this process takes place:



1. First host A will send a **SYN message** (a TCP segment with SYN flag set to 1, SYN is short for SYNchronize) to indicate it wants to setup a connection with host B. This message includes a sequence (SEQ) number for tracking purpose. This sequence number can be any 32-bit number (range from 0 to 2^{32}) so we use “x” to represent it.

2. After receiving SYN message from host A, host B replies with **SYN-ACK message** (some books may call it “SYN/ACK” or “SYN, ACK” message. ACK is short for ACKnowledge). This message includes a SYN sequence number and an ACK number:

- + SYN sequence number (let’s called it “y”) is a random number and does not have any relationship with Host A’s SYN SEQ number.

- + ACK number is the next number of Host A’s SYN sequence number it received, so we represent it with “x+1”. It means “I received your part. Now send me the next part (x + 1)”.

The SYN-ACK message indicates host B accepts to talk to host A (via ACK part). And ask if host A still wants to talk to it as well (via SYN part).

3. After Host A received the SYN-ACK message from host B, it sends an **ACK message** with ACK number “y+1” to host B. This confirms host A still wants to talk to host B.

Question 5

Which two sources cause interference for Wi-Fi networks? (Choose two)

- A. mirrored wall
- B. fish tank
- C. 900MHz baby monitor
- D. DECT 6.0 cordless
- E. incandescent lights

Answer: A B

Explanation

Windows can actually block your WiFi signal. How? Because the signals will be reflected by the glass.

Some new windows have transparent films that can block certain wave types, and this can make it harder for your WiFi signal to pass through.

Tinted glass is another problem for the same reasons. They sometimes contain metallic films that can completely block out your signal.

Mirrors, like windows, can reflect your signal. They're also a source of electromagnetic interference because of their metal backings.

Reference: <https://dis-dot-dat.net/what-materials-can-block-a-wifi-signal/>

An incandescent light bulb, incandescent lamp or incandescent light globe is an electric light with a wire filament heated until it glows. WiFi operates in the gigahertz microwave band. The FCC has strict regulations on RFI (radio frequency interference) from all sorts of things, including light bulbs -> Incandescent lights do not interfere Wi-Fi networks.

Note:

- + Many baby monitors operate at 900MHz and won't interfere with Wi-Fi, which uses the 2.4GHz band.
- + DECT cordless phone 6.0 is designed to eliminate wifi interference by operating on a different frequency. There is essentially no such thing as DECT wifi interference.

Question 6

What are two considerations when using SSO as a network redundancy feature? (Choose two)

- A. must be combined with NSF to support uninterrupted Layer 2 operations
- B. must be combined with NSF to support uninterrupted Layer 3 operations
- C. both supervisors must be configured separately
- D. the multicast state is preserved during switchover
- E. requires synchronization between supervisors in order to guarantee continuous connectivity

Answer: B E

Explanation

Cisco IOS Nonstop Forwarding(NSF) always runs with stateful switchover (SSO) and provides redundancy for Layer 3 traffic.

Reference:

https://www.cisco.com/en/US/docs/switches/lan/catalyst3850/software/release/3se/consolidated_guide/b_consolidated_3850_3se_cg_chapter_01101110.pdf

Question 7

What is the responsibility of a secondary WLC?

- A. It shares the traffic load of the LAPs with the primary controller.
- B. It avoids congestion on the primary controller by sharing the registration load on the LAPs.
- C. It registers the LAPs if the primary controller fails.
- D. It enables Layer 2 and Layer 3 roaming between itself and the primary controller.

Answer: C

Explanation

When the primary controller (WLC-1) goes down, the APs automatically get registered with the secondary controller (WLC-2). The APs register back to the primary controller when the primary controller comes back on line.

Reference: <https://www.cisco.com/c/en/us/support/docs/wireless-mobility/wireless-lan-wlan/69639-wlc-failover.html>

Question 8

What is the purpose of the LISP routing and addressing architecture?

- A. It creates head-end replication used to deliver broadcast and multicast frames to the entire network.
- B. It allows LISP to be applied as a network visualization overlay through encapsulation.
- C. It allows multiple instances of a routing table to co-exist within the same router.
- D. It creates two entries for each network node, one for its identity and another for its location on the network.

Answer: D

Explanation

Locator ID Separation Protocol (LISP) solves this issue by separating the location and identity of a device through the Routing locator (RLOC) and Endpoint identifier (EID):

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

Question 9

How does the EIGRP metric differ from the OSPF metric?

- A. The EIGRP metric is calculated based on bandwidth only. The OSPF metric is calculated on delay only.
- B. The EIGRP metric is calculated based on delay only. The OSPF metric is calculated on

bandwidth and delay.

C. The EIGRP metric is calculated based on hop count and bandwidth. The OSPF metric is calculated on bandwidth and delay.

D. The EIGRP metric is calculated based on bandwidth and delay. The OSPF metric is calculated on bandwidth only.

Answer: D

Explanation

By default, EIGRP metric is calculated:

metric = bandwidth + delay

While OSPF is calculated by:

OSPF metric = Reference bandwidth / Interface bandwidth in bps

(Or Cisco uses 100Mbps (10^8) bandwidth as reference bandwidth. With this bandwidth, our equation would be:

Cost = 10^8 /interface bandwidth in bps)

Question 10

What is one fact about Cisco SD-Access wireless network deployments?

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

D. The wireless client is part of the fabric overlay

Answer: C

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 11

What are two differences between the RIB and the FIB? (Choose two)

- A. The FIB is derived from the data plane, and the RIB is derived from the FIB.
- B. The RIB is a database of routing prefixes, and the FIB is the information used to choose the egress interface for each packet.
- C. FIB is a database of routing prefixes, and the RIB is the information used to choose the egress interface for each packet.
- D. The FIB is derived from the control plane, and the RIB is derived from the FIB.
- E. The RIB is derived from the control plane, and the FIB is derived from the RIB.

Answer: B E

Explanation

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 FIB maintains next-hop address information based on the information in the IP routing table (RIB).

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 12

What is the function of the fabric control plane node in a Cisco SD-Access deployment?

- A. It is responsible for policy application and network segmentation in the fabric.
- B. It performs traffic encapsulation and security profiles enforcement in the fabric.
- C. It holds a comprehensive database that tracks endpoints and networks in the fabric.
- D. It provides integration with legacy nonfabric-enabled environments.

Answer: C


Explanation

Fabric control plane node (C): One or more network elements that implement the LISP Map-Server (MS) and Map-Resolver (MR) functionality. The control plane node's host tracking database keep track of all endpoints in a fabric site and associates the endpoints to fabric nodes in what is known as an EID-to-RLOC binding in LISP.

Reference: <https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-sda-macro-segmentation-deploy-guide.html>

Question 13

Refer to the exhibit.



An engineer must allow all users in the 10.2.2.0/24 subnet to access the Internet. To conserve address space, the public interface address of 209.165.201.1 must be used for all external communication. Which command set accomplishes these requirements?

Option A access-list 10 permit 10.2.2.0 0.0.0.255 interface G0/3 ip nat outside interface G0/2 ip nat inside ip nat inside source list 10 interface G0/2 overload	Option B access-list 10 permit 10.2.2.0 0.0.0.255 interface G0/3 ip nat outside interface G0/2 ip nat inside ip nat inside source list 10 209.165.201.1
Option C access-list 10 permit 10.2.2.0 0.0.0.255	Option D access-list 10 permit 10.2.2.0 0.0.0.255

<pre> interface G0/3 ip nat outside interface G0/2 ip nat inside ip nat inside source list 10 interface G0/3 </pre>	<pre> interface G0/3 ip nat outside interface G0/2 ip nat inside ip nat inside source list 10 interface G0/3 overload </pre>
---	--


Answer: D

Explanation

The command “ip nat inside source list 10 interface G0/3 overload” configures NAT to overload (PAT) on the address that is assigned to the G0/3 interface.

Question 14

Refer to the exhibit.



```

SW2#
08:33:23: %PM-4-ERR_DISABLE: channel-misconfig error detection on Gi0/0, putting
Gi0/0 in err-disable state
08:33:23: %PM-4-ERR_DISABLE: channel-misconfig error detection on Gi0/1, putting
Gi0/1 in err-disable state

```

After an engineer configures an EtherChannel between switch SW1 and switch SW2, this error message is logged on switch SW2. Based on the output from SW1 and the log message received on Switch SW2, what action should the engineer take to resolve this issue?

- A. Configure the same protocol on the EtherChannel on switch SW1 and SW2.
- B. Connect the configuration error on interface Gi0/1 on switch SW1.

- C. Define the correct port members on the EtherChannel on switch SW1.
- D. Correct the configuration error on interface Gi0/0 switch SW1.

Answer: A

Explanation

In this case, we are using your EtherChannel without a negotiation protocol. As a result, if the opposite switch is not also configured for EtherChannel operation on the respective ports, there is a danger of a switching loop. The EtherChannel Misconfiguration Guard tries to prevent that loop from occurring by disabling all the ports bundled in the EtherChannel.

Question 15


Which element enables communication between guest VMs within a virtualized environment?

- A. vSwitch
- B. virtual router
- C. hypervisor
- D. pNIC

Answer: A

Explanation

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.



Question 16

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
```

These commands have been added to the configuration of a switch. Which command flags an error if it is added to this configuration?

- A. monitor session 1 source interface FastEthernet0/1 rx
- B. monitor session 1 source interface port-channel 6
- C. monitor session 1 source vlan 10
- D. monitor session 1 source interface port-channel 7, port-channel 8

Answer: C

Explanation

RSPAN consists of at least one RSPAN source session, an RSPAN VLAN, and at least one RSPAN destination session. You separately configure RSPAN source sessions and RSPAN destination sessions on different network devices. To configure an RSPAN source session on a device, **you associate a set of source ports or source VLANs** with an RSPAN VLAN.

Traffic monitoring in a SPAN session has these restrictions:

+ Sources can be ports or VLANs, but you cannot mix source ports and source VLANs in the same session.

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

Therefore in this question, we cannot configure a source VLAN because we configured source ports for RSPAN session 1 already.

Question 17

Which entity is responsible for maintaining Layer 2 isolation between segments in a VXLAN environment?

- A. switch fabric
- B. host switch

- C. VTEP
- D. VNID


Answer: D

Explanation

VXLAN uses an 8-byte VXLAN header that consists of a 24-bit VNID and a few reserved bits. The VXLAN header together with the original Ethernet frame goes in the UDP payload. The 24-bit VNID is used to identify Layer 2 segments and to maintain Layer 2 isolation between the segments.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/vxlan/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_VXLAN_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_VXLAN_Configuration_Guide_7x_chapter_010.html

Let's see the structure of a VXLAN packet to understand how (note: VNI = VNID)



The key fields for the VXLAN packet in each of the protocol headers are:

- + **Outer MAC header** (14 bytes with 4 bytes optional) – Contains the MAC address of the source VTEP and the MAC address of the next-hop router. Each router along the packet's path rewrites this header so that the source address is the router's MAC address and the destination address is the next-hop router's MAC address.
- + **Outer IP header** (20 bytes)- Contains the IP addresses of the source and destination VTEPs.
- + **(Outer) UDP header** (8 bytes)- Contains source and destination UDP ports:
 - Source UDP port: The VXLAN protocol repurposes this standard field in a UDP packet

header. Instead of using this field for the source UDP port, the protocol uses it as a numeric identifier for the particular flow between VTEPs. The VXLAN standard does not define how this number is derived, but the source VTEP usually calculates it from a hash of some combination of fields from the inner Layer 2 packet and the Layer 3 or Layer 4 headers of the original frame.

– Destination UDP port: The VXLAN UDP port. The Internet Assigned Numbers Authority (IANA) allocates port 4789 to VXLAN.

+ **VXLAN header** (8 bytes)- Contains the 24-bit VNI (or VNID)

+ **Original Ethernet/L2 Frame** – Contains the original Layer 2 Ethernet frame.

Question 18

Which method does Cisco DNA Center use to allow management of non-Cisco devices through southbound protocols?

- A. It creates device packs through the use of an SDK
- B. It obtains MIBs from each vendor that details the APIs available.
- C. It uses an API call to interrogate the devices and register the returned data.
- D. It imports available APIs for the non-Cisco device in a CSV format.

Answer: A

Explanation

Cisco DNA Center allows customers to manage their non-Cisco devices through the use of a Software Development Kit (SDK) that can be used to create Device Packages for third-party devices.

Reference: <https://developer.cisco.com/docs/dna-center/#!/cisco-dna-center-platform-overview/multivendor-support-southbound>

Question 19

Refer to the exhibit.

R1 key chain cisco123 key 1 key-string Cisco123!	R2 key chain cisco123 key 1 key-string cisco123!
Ethernet0/0 - Group 10 State is Active 8 state changes, last state change 00:03:33 Virtual IP address is 192.168.0.1 Active virtual MAC address is 0000.0c07.ac0a Local virtual MAC address is 0000.0c07.ac0a (v1 default) Hello time 5 sec, hold time 15 sec Next hello sent in 2.704 secs Authentication MD5, key-chain "cisco123" Preemption enabled Active router is local Standby router is unknown Priority 255 (configured 255) Group name is "workstation-group" (cfgd)	Ethernet0/0 - Group 10 State is Active 17 state changes, last state change 00:03:33 Virtual IP address is 192.168.0.1 Active virtual MAC address is 0000.0c07.ac0a Local virtual MAC address is 0000.0c07.ac0a (v1 default) Hello time 10 sec, hold time 30 sec Next hello sent in 6.704 secs Authentication MD5, key-chain "cisco123" Preemption disabled Active router is local Standby router is unknown Priority 200 (configured 200) Group name is "workstation-group" (cfgd)

An engineer is installing a new pair of routers in a redundant configuration. When checking on the standby status of each router the engineer notices that the routers are not functioning as expected. Which action will resolve the configuration error?

- A. configure matching hold and delay timers
- B. configure matching key-strings
- C. configure matching priority values
- D. configure unique virtual IP addresses

Answer: B

Explanation

From the output exhibit, we notice that the key-string of R1 is "Cisco123!" (letter "C" is in capital) while that of R2 is "cisco123!". This causes a mismatch in the authentication so we have to fix their key-strings.

Note:

key-string [encryption-type] *text-string*: Configures the text string for the key. The text-string argument is alphanumeric, case-sensitive, and supports special characters.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_chapter_01111.pdf

Question 20

Refer to the exhibit.

```
line vty 0 4
  session-timeout 30
  exec-timeout 120 0
  session-limit 30
  login local
line vty 5 15
  session-timeout 30
  exec-timeout 30 0
  session-limit 30
  login local
```

Only administrators from the subnet 10.10.10.0/24 are permitted to have access to the router. A secure protocol must be used for the remote access and management of the router instead of clear-text protocols. Which configuration achieves this goal?

Option A access-list 23 permit 10.10.10.0 0.0.0.255 line vty 0 4 access-class 23 in transport input ssh	Option B access-list 23 permit 10.10.10.0 0.0.0.255 line vty 0 15 access-class 23 in transport input ssh
Option C access-list 23 permit 10.10.10.0 0.0.0.255 line vty 0 15 access-class 23 out transport input all	Option D access-list 23 permit 10.10.10.0 255.0.0.0 line vty 0 15 access-class 23 in transport input ssh

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

Question 21

What is used to validate the authenticity of the client and is sent in HTTP requests as a JSON object?

- A. SSH
- B. HTTPS
- C. JVVVT
- D. TLS

Answer: B

Question 22

Refer to the exhibit.

```
monitor session 1 source vlan 10 -12 rx  
monitor session 1 destination interface gigabitethernet0/1
```

An engineer must configure a SPAN session. What is the effect of the configuration?

- A. Traffic sent on VLANs 10, 11, and 12 is copied and sent to interface g0/1.
- B. Traffic sent on VLANs 10 and 12 only is copied and sent to interface g0/1.
- C. Traffic received on VLANs 10 and 12 only is copied and sent to interface g0/1.
- D. Traffic received on VLANs 10, 11, and 12 is copied and sent to interface g0/1.

Answer: D

Question 23

In a Cisco SD-Access wireless architecture, which device manages endpoint ID to Edge Node bindings?

- A. fabric control plane node
- B. fabric wireless controller
- C. fabric border node
- D. fabric edge node

Answer: A

Explanation

SD-Access Wireless Architecture Control Plane Node –A Closer Look

Fabric Control-Plane Node is based on a LISP Map Server / Resolver

Runs the LISP Endpoint ID Database to provide overlay reachability information

- + A simple Host Database, that tracks Endpoint ID to Edge Node bindings (RLOCs)
- + Host Database supports multiple types of Endpoint ID (EID), such as IPv4 /32, IPv6 /128* or MAC/48
- + Receives prefix registrations from Edge Nodes for wired clients, and from Fabric mode WLCs for wireless clients
- + Resolves lookup requests from FE to locate Endpoints
- + Updates Fabric Edge nodes, Border nodes with wireless client mobility and RLOC information

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

===== New Questions (added on 24th-Dec-2020)
=====

Question 24


Which control plane protocol is used between Cisco SD-WAN routers and vSmart controllers?

- A. BGP
- B. OMP
- C. TCP
- D. UDP

Answer: B

Explanation

Cisco SD-WAN uses Overlay Management Protocol (OMP) which manages the overlay network. OMP runs between the vSmart controllers and WAN Edge routers (and among vSmarts themselves) where control plane information, such as the routing, policy, and management information, is exchanged over a secure connection.



Question 25

In a Cisco Catalyst switch equipped with two supervisor modules an administrator must temporally remove the active supervisor from the chassis to perform hardware maintenance on it. Which mechanism ensure that the active supervisor removal is not disruptive to the network operation?

- A. NSF/NSR
- B. SSO

- C. HSRP
- D. VRRP

Answer: B

Explanation

Stateful Switchover (SSO) provides protection for network edge devices with dual Route Processors (RPs) that represent a single point of failure in the network design, and where an outage might result in loss of service for customers.

Reference: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SY/configuration/guide/sy_swcg/stateful_switchover.html

Question 26

Which router is elected the IGMP Querier when more than one router is in the same LAN segment?

- A. The router with the shortest uptime
- B. The router with the lowest IP address
- C. The router with the highest IP address
- D. The router with the longest uptime

Answer: B

Explanation

Query messages are used to elect the IGMP querier as follows:

1. When IGMPv2 devices start, they each multicast a general query message to the all-systems group address of 224.0.0.1 with their interface address in the source IP address field of the message.
2. When an IGMPv2 device receives a general query message, the device compares the source IP address in the message with its own interface address. **The device with the lowest IP address on the subnet is elected the IGMP querier.**
3. All devices (excluding the querier) start the query timer, which is reset whenever a general query message is received from the IGMP querier. If the query timer expires, it is assumed that the IGMP querier has gone down, and the election process is performed again to elect a new IGMP querier.

Reference:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750x_3560x/software/release/15-2-2_e/multicast/configuration_guide/b_mc_1522e_3750x_3560x_cg/b_ipmc_3750x_3560x_chapter_01000.html

Question 27

Refer to the exhibit.

```
ip sla 10
  icmp-echo 192.168.10.20
  timeout 500
  frequency 3
ip sla schedule 10 life forever start-time now
track 10 ip sla 10 reachability
```

The IP SLA is configured in a router. An engineer must configure an EEM applet to shut down the interface and bring it back up when there is a problem with the IP SLA. Which configuration should the engineer use?

- A. event manager applet EEM_IP_SLA
event track 10 state down
- B. event manager applet EEM_IP_SLA
event track 10 state unreachable
- C. event manager applet EEM_IP_SLA
event sla 10 state unreachable
- D. event manager applet EEM_IP_SLA
event sla 10 state down

Answer: A

Explanation

The “ip sla 10” will ping the IP 192.168.10.20 every 3 seconds to make sure the connection is still up. We can configure an EEM applet if there is any problem with this IP SLA via the command “event track 10 state down”.

Reference: <https://www.theroutingtable.com/ip-sla-and-cisco-eem/>

Question 28

A network engineer is configuring Flexible NetFlow and enters these commands:

```
Sampler Netflow1
mode random one-out-of 100
interface fastethernet 1/0
flow-sampler netflow1
```

Which are two results of implementing this feature instead of traditional NetFlow? (Choose two)

- A. Only the flows of top 100 talkers are exported
- B. CPU and memory utilization are reduced
- C. The data export flow is more secure
- D. The accuracy of the data to be analyzed is improved
- E. The number of packets to be analyzed are reduced

Answer: B E

New ENCOR Questions – Part 2

Question 1

Which two LISP infrastructure elements are needed to support LISP to non-LISP internetworking? (Choose two)


- A. PETR
- B. PITR
- C. MR
- D. MS
- E. ALT

Answer: A C

Explanation

In this question we suppose that we only need to send packets from LISP site to non-LISP site successfully. We don't care about the way back (if we care about the way back then all PETR, PITR, MS & MR are needed).

Proxy Egress Tunnel Router (PETR): A LISP device that de-encapsulates packets from LISP sites to deliver them to non-LISP sites.



When the xTR in LISP Site 1 want to sends traffic to Non-LISP site, the ITR (not PETR) needs a Map Resolver (MR) to send Map Request to. When the ITR (the xTR in LISP Site 1 in the figure above) receives negative MAP-Reply packet from MR, it caches that prefix and map it to the PETR.

Good reference: <https://netmindblog.com/2019/12/04/lisp-locator-id-separation-protocol-part-ii-pxtr/>

Question 2

Which statement about dynamic GRE between a headend router and a remote router is true?

- A. The headend router learns the IP address of the remote end router statically
- B. A GRE tunnel without an IP address has a status of administratively down
- C. GRE tunnels can be established when the remote router has a dynamic IP address
- D. The remote router initiates the tunnel connection

Answer: D

Question 3

Which two statements about AAA authentication are true? (Choose two)

- A. RADIUS authentication queries the router's local username database
- B. TACACS+ authentication uses an RSA server to authenticate users
- C. Local user names are case-insensitive

- D. Local authentication is maintained on the router
- E. KRB5 authentication disables user access when an incorrect password is entered

Answer: D E

Question 4

Which action is performed by Link Management Protocol in a Cisco stackwise virtual domain?

- A. It discovers the stackwise domain and brings up SVL interfaces
- B. It rejects any unidirectional link traffic forwarding
- C. It determines if the hardware is compatible to form the stackwise virtual domain
- D. It determines which switch becomes active or standby

Answer: B

Explanation

The Link Management Protocol (LMP) performs the following functions:

- + Verifies link integrity by establishing bidirectional traffic forwarding, and rejects any unidirectional links
- + Exchanges periodic hellos to monitor and maintain the health of the links
- + Negotiates the version of StackWise Virtual header between the switches StackWise Virtual link role resolution

Reference: <https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9000/nb-06-cat-9k-stack-wp-cte-en.html>

Question 5

Which two actions provide controlled Layer 2 network connectivity between virtual machines running on the same hypervisor? (Choose two)

- A. Use a single trunk link to an external Layer2 switch
- B. Use a virtual switch provided by the hypervisor
- C. Use VXLAN fabric after installing VXLAN tunnelling drivers on the virtual machines
- D. Use a single routed link to an external router on stick
- E. Use a virtual switch running as a separate virtual machine

Answer: B E

Question 6

How does SSO work with HSRP to minimize network disruptions?

- A. It enables HSRP to elect another switch in the group as the active HSRP switch
- B. It ensures fast failover in the case of link failure
- C. It enables data forwarding along known routes following a switchover, while the routing protocol reconverges
- D. It enables HSRP to failover to the standby RP on the same device

Answer: D

Explanation


SSO HSRP alters the behavior of HSRP when a device with redundant Route Processors (RPs) is configured for stateful switchover (SSO) redundancy mode. When an RP is active and the other RP is standby, SSO enables the standby RP to take over if the active RP fails.

The SSO HSRP feature enables the Cisco IOS HSRP subsystem software to detect that a standby RP is installed and the system is configured in SSO redundancy mode. Further, if the active RP fails, no change occurs to the HSRP group itself and traffic continues to be forwarded through the current active gateway device.

Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp_fhrp/configuration/15-s/fhp-15-s-book/fhp-hsrp-ss0.html

Question 7

Refer to the exhibit.



Dallas#show ip route ospf

3.0.0.0/32 is subnetted, 1 subnets

- O 3.3.3.3 [110/40001] via 192.168.4.3, 00:33:32, FastEthernet0/0
- O IA 192.168.0.0/24 [110/145535] via 192.168.4.3, 00:33:32, FastEthernet0/0
- O IA 192.168.1.0/24 [110/80000] via 192.168.4.3, 00:33:32, FastEthernet0/0
- O IA 192.168.2.0/24 [110/80000] via 192.168.4.3, 00:33:32, FastEthernet0/0
- O IA 192.168.3.0/24 [110/44000] via 192.168.4.3, 00:33:32, FastEthernet0/0


Which command when applied to the Atlanta router reduces type 3 LSA flooding into the backbone area and summarizes the inter-area routes on the Dallas router?

- A. Atlanta(config-router)#area 0 range 192.168.0.0 255.255.252.0
- B. Atlanta(config-router)#area 1 range 192.168.0.0 255.255.252.0
- C. Atlanta(config-router)#area 0 range 192.168.0.0 255.255.248.0
- D. Atlanta(config-router)#area 1 range 192.168.0.0 255.255.248.0

Answer: B

Question 8

Refer the exhibit.



```
Chicago#show ip ospf nei
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
3.3.3.3	1	FULL/BDR	00:00:35	192.168.2.3	FastEthernet0/0
2.2.2.2	0	FULL/-	00:00:35	192.168.0.2	Serial0/0

```
Chicago#show ip ospf int bri
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Fa0/0	1	1	192.168.2.1/24	40444	DR	1/1	
Se0/0	1	1	192.168.0.1/24	65535	P2P	1/1	

```
Chicago#
```

Which router is the designated router on the segment 192.168.0.0/24?

- A. Router Chicago because it has a lower router ID
- B. Router NewYork because it has a higher router ID
- C. This segment has no designated router because it is a nonbroadcast network type.
- D. This segment has no designated router because it is a p2p network type.

Answer: D

Question 9

An engineer must configure interface GigabitEthernet0/0 for VRRP group 10. When the router has the highest priority in the group, it must assume the master role. Which command set must be added to the initial configuration to accomplish this task?

Initial Configuration

```
interface GigabitEthernet0/0
description to IDF
ip address 172.16.13.2 255.255.255.0
```

- A.
vrrp 10 ip 172.16.13.254
vrrp 10 preempt
- B.
standby 10 ip 172.16.13.254
standby 10 priority 120
- C.
vrrp group 10 ip 172.16.13.254 255.255.255.0
vrrp group 10 priority 120
- D.
standby 10 ip 172.16.13.254 255.255.255.0
standby 10 preempt

Answer: A

Explanation

In fact, VRRP has the preemption enabled by default so we don't need the "vrrp 10 preempt" command. The default priority is 100 so we don't need to configure it either. But notice that the correct command to configure the virtual IP address for the group is "vrrp 10 ip {ip-address}" (not "vrrp group 10 ip ...") and this command does not include a subnet mask.

Question 10

Drag and drop the characteristics from the left onto the infrastructure types on the right.

slow upgrade lifecycle	On-Premises Infrastructure
low capital expenditure	
provider maintains the infrastructure	
high capital expenditure	
enterprise owns the hardware	Cloud-Hosted Infrastructure
fast upgrade lifecycle	

Answer:

On-Premises Infrastructure:

- + slow upgrade lifecycle
- + high capital expenditure
- + enterprise owns the hardware

Cloud-Hosted Infrastructure:

- + low capital expenditure
- + provider maintains the infrastructure
- + fast upgrade lifecycle

Question 11

Drag and drop the threat defense solutions from the left onto their descriptions on the right.

StealWatch	provides IPS/IDS capabilities
ESA	provides malware protection on endpoints
AMP4E	protects against email threat vector
Umbrella	performs security analytics by collecting network flows
FTD	provides DNS protection

Answer:

- + StealWatch: performs security analytics by collecting network flows
- + ESA: protects against email threat vector
- + AMP4E: provides malware protection on endpoints
- + Umbrella: provides DNS protection
- + FTD: provides IPS/IDS capabilities

Question 12

Refer to the exhibit. An engineer configures CoPP and enters the show command to verify the implementation. What is the result of the configuration?

```
Router2#show policy-map control-plane

Control Plane
Service-policy input:CISCO
Class-map:CISCO (match-all)
 20 packets, 11280 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
Match:access-group 120
police:
 8000 bps, 1500 limit, 1500 extended limit
conformed 15 packets, 6210 bytes; action:transmit
exceeded 5 packets, 5070 bytes; action:drop
violated 0 packets, 0 bytes; action:drop
conformed 0 bps, exceed 0 bps, violate 0 bps
Class-map:class-default (match-any)
105325 packets, 11415151 bytes
 5 minute offered rate 0 bps, drop rate 0 bps
Match:any
```

- A. All traffic will be policed based on access-list 120
- B. If traffic exceeds the specified rate, it will be transmitted and remarked

- C. Class-default traffic will be dropped
- D. ICMP will be denied based on this configuration

Answer: A

Question 13

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

Answer: A B D

Question 14

Which technology is used to provide Layer 2 and Layer 3 logical networks in the Cisco SD-Access architecture?

- A. underlay network
- B. overlay network
- C. VPN routing/forwarding
- D. easy virtual network

Answer: B

Explanation

An overlay network creates a logical topology used to virtually connect devices that are built over an arbitrary physical underlay topology.

An overlay network is created on top of the underlay network through virtualization (virtual networks). The data plane traffic and control plane signaling are contained within each virtualized network, maintaining isolation among the networks and an independence from the underlay network.

SD-Access allows for the extension of Layer 2 and Layer 3 connectivity across the overlay through the services provided by through LISP.

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

Question 15

An engineer uses the Design workflow to create a new network infrastructure in Cisco DNA Center. How is the physical network device hierarchy structured?


- A. by location
- B. by role
- C. by organization
- D. by hostname naming convention

Answer: A

Explanation

You can create a network hierarchy that represents your network's geographical locations. Your network hierarchy can contain sites, which in turn contain buildings and areas. You can create site and building IDs to easily identify where to apply design settings or configurations later.

Reference: https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/network-automation-and-management/dna-center/1-2-5/user_guide/b_dnac_ug_1_2_5/b_dnac_ug_1_2_4_chapter_0110.html



Question 16

Refer to the exhibit.

```
(WLC) >show interface summary
Interface Name          Vlan Id
-----
deadnet                 999
users1                  14
users2                  15
users3                  16

(WLC) >show wlan 1
WLAN Identifier . . . . . 1
Network Name (SSID) . . . . . wlan1
AAA Policy Override . . . . . Enabled
Interface . . . . . deadnet
FlexConnect Local Switching . . . . . Enabled
FlexConnect Central Association . . . . . Disabled
flexconnect Central Dhcp Flag . . . . . Disabled
flexconnect nat-pat Flag . . . . . Disabled
flexconnect DNS Override Flag . . . . . Disabled
flexconnect PPPoE pass-through . . . . . Disabled
flexconnect local-switching IP-source-guar . . . . . Disabled
FlexConnect Vlan based Central Switching . . . . . Enabled
FlexConnect Local Authentication . . . . . Disabled
FlexConnect Learn IP Address . . . . . Enabled

(WLC) >show ap config general FlexAP1
AP Mode . . . . . FlexConnect
FlexConnect Vlan mode : . . . . . Enabled
Native ID : . . . . . 1
WLAN 1 : . . . . . 10 (AP-Specific)
FlexConnect VLAN ACL Mappings
Vlan : . . . . . 10
Ingress ACL : . . . . . None
Egress ACL : . . . . . None
VLAN with least priority : . . . . . 13
FlexConnect Group . . . . . flexgroup1
Group VLAN ACL Mappings
Vlan : . . . . . 11
Ingress ACL : . . . . . None
Egress ACL : . . . . . None
Vlan : . . . . . 12
```

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.

Answer: B C G

Explanation

- + From the output of WLC “show interface summary”, we learned that the WLC has four VLANs: 999, 14, 15 and 16.
- + From the “show ap config general FlexAP1” output, we learned that FlexConnect AP has four VLANs: 10, 11, 12 and 13. Also the WLAN of FlexConnect AP is mapped to VLAN 10 (from the line “WLAN 1: 10 (AP-Specific)).

From the reference at: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-1/Enterprise-Mobility-8-1-Design-Guide/Enterprise_Mobility_8-1_Deployment_Guide/ch7_HREA.html

FlexConnect VLAN Central Switching Summary

Traffic flow on WLANs configured for Local Switching when FlexConnect APs are in connected mode are as follows:

- + If the VLAN is returned as one of the AAA attributes and that VLAN is not present in the FlexConnect AP database, traffic will switch centrally and the client is assigned this VLAN/Interface returned from the AAA server provided that the VLAN exists on the WLC. (-> as VLAN 15 exists on the WLC so the client in connected mode would be assigned this VLAN -> Answer G is correct)
- + If the VLAN is returned as one of the AAA attributes and that VLAN is not present in the FlexConnect AP database, traffic will switch centrally. If that VLAN is also not present on the WLC, the client will be assigned a VLAN/Interface mapped to a WLAN on the WLC.
- + If the VLAN is returned as one of the AAA attributes and that VLAN is present in the FlexConnect AP database, traffic will switch locally.
- + If the VLAN is not returned from the AAA server, the client is assigned a WLAN mapped VLAN on that FlexConnect AP and traffic is switched locally.

Traffic flow on WLANs configured for Local Switching when FlexConnect APs are in standalone mode are as follows:

- + If the VLAN returned by the AAA server is not present in the FlexConnect AP database, the client will be put on a default VLAN (that is, a WLAN mapped VLAN on a FlexConnect

AP) (-> Therefore answer B is correct). When the AP connects back, this client is de-authenticated (-> Therefore answer C is correct) and will switch traffic centrally.

Question 17

Which three methods does Cisco DNA Center use to discover devices? (Choose three)

- A. CDP
- B. LLDP
- C. SNMP
- D. ping
- E. NETCONF
- F. a specified range of IP addresses

Answer: A B F

Question 18

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

Answer: B

Explanation

Where to Use MST

This diagram shows a common design that features access Switch A with 1000 VLANs redundantly connected to two distribution Switches, D1 and D2. In this setup, users connect to Switch A, and the network administrator typically seeks to achieve load balancing on the access switch Uplinks based on even or odd VLANs, or any other scheme deemed appropriate.



Reference:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/24248-147.html>

Question 19

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

Link State Protocol	OSPF
selects routes using the DUAL algorithm	
maintains alternative loop-free backup path if available	
supports only equal multipath load balancing	EIGRP
Advanced Distance Vector Protocol	
quickly computes new path upon link failure	

Answer:

OSPF

- + Link State Protocol
- + supports only equal multipath load balancing
- + quickly computes new path upon link failure

EIGRP

- + selects routes using the DUAL algorithm
- + maintains alternative loop-free backup path if available
- + Advanced Distance Vector Protocol

Explanation

EIGRP maintains alternative loop-free backup via the feasible successors. To qualify as a feasible successor, a router must have an Advertised Distance (AD) less than the Feasible distance (FD) of the current successor route.

Advertised distance (AD): the cost from the neighbor to the destination.

Feasible distance (FD): The sum of the AD plus the cost between the local router and the next-hop router

Question 20

How does the RIB differ from the FIB?

- A. The RIB includes many routes to the same destination prefix. The FIB contains only the best route.
- B. The FIB maintains network topologies and routing tables. The RIB is a list of routes to particular network destinations.
- C. The RIB is used to create network topologies and routing tables. The FIB is a list of routes to particular network destinations.
- D. The FIB includes many routes a single destination. The RIB is the best route to a single destination.

Answer: C

Question 21


What is the purpose of an RP in PIM?

- A. secure the communication channel between the multicast sender and receiver.
- B. ensure the shortest path from the multicast source to the receiver.
- C. receive IGMP joins from multicast receivers.
- D. send join messages toward a multicast source SPT

Answer: C

Question 22

Refer to the exhibit.



```
hostname R1
router ospf 1
network 0.0.0.0 255.255.255.255 area 0
auto-cost reference-bandwidth 1000
!
hostname R2
router ospf 2
network 20.0.0.0 0.0.0.255 area 0
```

Which command must be applied to R2 for an OSPF neighborship to form?

- A. network 20.1.1.2 0.0.255.255 area 0
- B. network 20.1.1.2 255.255.255.255 area 0
- C. network 20.1.1.2 0.0.0.0 area 0
- D. network 20.1.1.2 255.255.0.0 area 0

Answer: C

Explanation

The “network 20.0.0.0 0.0.0.255 area 0” command on R2 did not cover the IP address of Fa1/1 interface of R2 so OSPF did not run on this interface. Therefore we have to use the command “network 20.1.1.2 0.0.255.255 area 0” to turn on OSPF on this interface.

Note: The command “network 20.1.1.2 0.0.255.255 area 0” can be used too so this answer is also correct but answer C is the best answer here.

The “network 0.0.0.0 255.255.255.255 area 0” command on R1 will run OSPF on all active interfaces of R1.

Question 23

Which antenna type should be used for a site-to-site wireless connection?

- A. Omnidirectional
- B. Yagi
- C. dipole
- D. patch

Answer: B

Question 24

Refer to the exhibit. An engineer is using XML in an application to send information to a RESTCONF-enabled device. After sending the request, the engineer gets this response message and a HTTP response code of 400. What do these responses tell the engineer?

```
<errors xmlns="urn:ietf:params:xml:ns:yang:ietf-restconf">
  <error>
    <error-message>End-of-file reached in XML
stream</error-message>
    <error-path>/ietf-interfaces:interfaces/interface=Giga
bitEthernet2</error-path>
    <error-tag>malformed-message</error-tag>
    <error-type>application</error-type>
  </error>
</errors>
```

- A. POST was used instead of PUT to update
- B. The Accept header sent was application/xml
- C. The Content-Type header sent was application/xml.
- D. JSON body was used

Answer: B

Explanation

Accept and Content-type are both headers sent from a client (a browser) to a service. Accept header is a way for a client to specify the media type of the response content it is expecting and Content-type is a way to specify the media type of request being sent from the client to the server.

The response was sent in XML so we can say the Accept header sent was application/xml.

Question 25

Refer to the exhibit. Which two commands ensure that DSW1 becomes root bridge for VLAN 10 and 20? (Choose two)

DSW1#show spanning-tree

MST1

Spanning tree enabled protocol mstp

Root ID Priority 32769
 Address 0018.7363.4300
 Cost 2
 Port 13 (FastEthernet1/0/11)
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 32769 (priority 32768 sys-id- ext 1)
 Address 001b.0d8e.e080
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa1/0/7	Desg FWD	2	128.1	P2p Bound (PVST)	
Fa1/0/10	Desg FWD	2	128.12	P2p Bound (PVST)	
Fa1/0/11	Root FWD	2	128.13	P2p	
Fa1/0/12	Altn BLK	2	128.14	P2p	

DSW1#show spanning-tree mst

MST1 vlans mapped: 10,20
Bridge address 001b.0d0e.e000 priority 32769 (32768 sysid 1)
Root address 0018.7363.4300 priority 32769 (32768 sysid 1)
 port Fa1/0/11 cost 2 (rem hops 19)

----- output omitted -----

- A. spanning-tree mstp 1 priority 0
- B. spanning-tree mst 1 root primary
- C. spanning-tree mst vlan 10,20 priority root
- D. spanning-tree mst 1 priority 4096
- E. spanning-tree mst 1 priority 1
- F. spanning-tree mstp vlan 10,20 root primary

Answer: B D

Explanation

From the second command output (show spanning-tree mst) we learn that MST1 includes VLANs 10 & 20. Therefore if we want DSW1 to become root bridge for these VLANs we need to set the MST 1 region to root -> The command “spanning-tree mst 1 root primary” can do the trick. In fact, this command runs a macro and sets the priority lower than the current root.

Also we can see the current root bridge for these VLANs has the priority of 32769 (default value + sysid) so we can set the priority of DSW1 to a specific lower value. But notice that the priority must be a multiple of 4096. Therefore D is a correct answer.

Question 26

Which feature of EIGRP is not supported in OSPF?

- A. load balancing of unequal-cost paths
- B. load balance over four equal-costs paths
- C. uses interface bandwidth to determine best path
- D. per-packet load balancing over multiple paths

Answer: A

Question 27

Which two characteristics define the Intent API provided by Cisco DNA Center? (Choose two)

- A. northbound API
- B. southbound API
- C. device-oriented
- D. business outcome oriented
- E. procedural

Answer: A D

Explanation

The Intent API is a **Northbound REST API** that exposes specific capabilities of the Cisco DNA Center platform.

The Intent API provides policy-based abstraction of **business intent, allowing focus on an outcome** rather than struggling with individual mechanisms steps.

Reference: <https://developer.cisco.com/docs/dna-center/#!/cisco-dna-center-platform-overview/intent-api-northbound>

Question 28

What is the difference between CEF and process switching?

- A. CEF processes packets that are too complex for process switching to manage.
- B. CEF is more CPU-intensive than process switching.
- C. CEF uses the FIB and the adjacency table to make forwarding decisions, whereas process

switching punts each packet.
D. Process switching is faster than CEF.

Answer: C

Explanation

“Punt” is often used to describe the action of moving a packet from the fast path (CEF) to the route processor for handling.

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**.

Process switching is the slowest switching methods (compared to fast switching and Cisco Express Forwarding) because it must find a destination in the routing table. Process switching must also construct a new Layer 2 frame header for every packet. With process switching, when a packet comes in, the scheduler calls a process that examines the routing table, determines which interface the packet should be switched to and then switches the packet. The problem is, this happens for the every packet.

Reference: <http://www.cisco.com/web/about/security/intelligence/acl-logging.html>

Question 29

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

Answer: A

Explanation


CoS value 5 is commonly used for VOIP and CoS value 5 should be mapped to DSCP 46. DSCP 46 is defined as being for EF (Expedited Forwarding) traffic flows and is the value usually assigned to all interactive voice and video traffic. This is to keep the uniformity from end-to-end that DSCP EF (mostly for VOICE RTP) is mapped to COS 5.

Note:

- + CoS is a L2 marking contained within an 802.1q tag,. The values for CoS are 0 – 7
- + DSCP is a L3 marking and has values 0 – 63
- + The default DSCP-to-CoS mapping for CoS 5 is DSCP 40

Question 30

Refer to the exhibit. Edge-01 is currently operational as the HSRP primary with priority 110. Which command on Edge-02 causes it to take over the forwarding role when Edge-01 is down?



- A. standby 10 priority
- B. standby 10 timers
- C. standby 10 track
- D. standby 10 preempt

Answer: D

Explanation

The “preempt” command enables the HSRP router with the highest priority to immediately become the active router.

Question 31

What is a Type 1 hypervisor?

- A. runs directly on a physical server and depends on a previously installed operating system
- B. runs directly on a physical server and includes its own operating system
- C. runs on a virtual server and depends on an already installed operating system
- D. run on a virtual server and includes its own operating system


Answer: B

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. A 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 32

An engineer reviews a router's logs and discovers the following entry. What is the event's logging severity level?

```
Router# *Feb 03 11:13:44 334: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up
```

- A. error
- B. notification
- C. informational
- D. warning

Answer: A

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 “3” in “%LINK-3-UPDOWN” is the severity level of this message so in this case it is “errors”.

Question 33

Refer to the exhibit. An engineer attempts to configure a router on a stick to route packets between Clients, Servers, and Printers; however, initial tests show that this configuration is not working. Which command set resolves this issue?

```

interface Vlan10
ip vrf forwarding Clients
ip address 192.168.1.1 255.255.255.0
!
interface Vlan20
ip vrf forwarding Servers
ip address 172.16.1.1 255.255.255.0
!
interface Vlan30
ip vrf forwarding Printers
ip address 10.1.1.1 255.255.255.0
<output omitted>
router eigrp 1
network 10.0.0.0
network 172.16.0.0
network 192.168.1.0

```

Option A router eigrp 1 network 10.0.0.0 255.0.0.0 network 172.16.0.0 255.255.0.0 network 192.168.1.0 255.255.0.0	Option B router eigrp 1 network 10.0.0.0 255.255.255.0 network 172.16.0.0 255.255.255.0 network 192.168.1.0 255.255.255.0
Option C interface Vlan10 no ip vrf forwarding Clients ! interface Vlan20	Option D interface Vlan10 no ip vrf forwarding Clients ip address 192.168.1.2 255.255.255.0 !

no ip vrf forwarding Servers ! interface Vlan30 no ip vrf forwarding Printers	interface Vlan20 no ip vrf forwarding Servers ip address 172.16.1.2 255.255.255.0 ! interface Vlan30 no ip vrf forwarding Printers ip address 10.1.1.2 255.255.255.0
--	--

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

Explanation

We must reconfigure the IP address after assigning or removing an interface to a VRF. Otherwise that interface does not have an IP address.

Question 34

How is a data modeling language used?

- A. To enable data to be easily structured, grouped validated, and replicated
- B. To represent finite and well-defined network elements that cannot be changed
- C. To model the flows of unstructured data within the infrastructure
- D. To provide human readability to scripting languages

Answer: A

Explanation

Customer needs are fast evolving. Typically, a network center is a heterogenous mix of various devices at multiple layers of the network. Bulk and automatic configurations need to be accomplished. CLI scraping is not flexible and optimal. Re-writing scripts many times, even for small configuration changes is cumbersome. Bulk configuration changes through CLIs are error-prone and may cause system issues. The solution lies in using data models-a programmatic and standards-based way of writing configurations to any network device, replacing the process of manual configuration. Data models are written in a standard, industry-defined language. Although configurations using CLIs are easier (more human-friendly), automating the configuration using data models results in scalability.

Reference:

https://www.cisco.com/c/en/us/td/docs/optical/ncs1000/60x/b_Datamodels_cg_ncs1000/b_Datamodels_cg_ncs1000_chapter_00.pdf

Question 35

Refer to the exhibit.

```
aaa new-model
aaa authentication login authorizationlist tacacs+
tacacs-server host 192.168.0.202
tacacs-server key ciscotestkey
line vty 0 4
login authentication authorizationlist
```


What is the effect of the configuration?

- A. The device will allow users at 192.168.0.202 to connect to vty lines 0 through 4 using the password ciscotestkey
- B. The device will allow only users at 192 168.0.202 to connect to vty lines 0 through 4
- C. When users attempt to connect to vty lines 0 through 4, the device will authenticate them against TACACS+ if local authentication fails
- D. The device will authenticate all users connecting to vty lines 0 through 4 against TACACS+

Answer: D

Question 36

Refer to the exhibit. R1 is able to ping the R3 fa0/1 interface. Why do the extended pings fail?



```

R1#ping
Protocol [ip]:
Target IP address: 3.3.3.3
Repeat count [5]: 3
Datagram size [100]: 1500
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 1.1.1.1
Type of service [0]:
Set DF bit in IP header? [no]: yes
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]: Record
Number of hops [9]:
Loose, Strict, Record, Timestamp, Verbose[RV]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 3, 1500-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
Packet sent with the DF bit set
Packet has IP options: Total option bytes= 39, padded length=40
Record route: <*>
(0.0.0.0)
(0.0.0.0)

Unreachable from 10.99.69.2, maximum MTU 1492, Received packet has options
Total option bytes= 39, padded length=40
Record route: <*>
(0.0.0.0)
(0.0.0.0)
<output omitted>

```

- A. R2 and R3 do not have an OSPF adjacency
- B. R3 is missing a return route to 10.99.69.0/30
- C. The maximum packet size accepted by the command is 1476 bytes
- D. The DF bit has been set

Answer: D

Explanation

If the DF bit is set, routers cannot fragment packets. From the output below, we learn that the maximum MTU of R2 is 1492 bytes while we sent ping with 1500 bytes. Therefore these ICMP packets were dropped.

Note: Record option displays the address(es) of the hops (up to nine) the packet goes through.

Question 37

Refer to the exhibit. A network engineer configures a GRE tunnel and enters the show interface tunnel command. What does the output confirm about the configuration?

```
Tunnel100 is up, line protocol is up
Hardware is Tunnel
Internet address is 192.168.200.1/24
MTU 17912 bytes, BW 100 Kbit/sec, DLY 50000 usec,
  reliability 255/255, txload 1/255, rxload 1/255
Encapsulation TUNNEL, loopback not set
Keepalive set (10 sec), retries 3
Tunnel source 209.165.202.129 (GigabitEthernet0/1)
Tunnel Subblocks:
  src-track:
    Tunnel100 source tracking subblock associated with GigabitEthernet0/1
    Set of tunnels with source GigabitEthernet0/1, 1 members (includes iterators),
    on interface <OK>
Tunnel protocol/transport GRE/IP
Key disabled, sequencing disabled
Checksumming of packets disabled
Tunnel TTL 255, Fast tunneling enabled
Tunnel transport MTU 1476 bytes
```

- A. The keepalive value is modified from the default value.
- B. Interface tracking is configured.
- C. The tunnel mode is set to the default.
- D. The physical interface MTU is 1476 bytes.


Answer: C

Explanation

From the “Tunnel protocol/transport GRE/IP” line, we can deduce this tunnel is using the default IPv4 Layer-3 tunnel mode. We can return to this default mode with the “tunnel mode gre ip” command.

Question 38

Refer to the exhibit. Which two commands are needed to allow for full reachability between AS 1000 and AS 2000? (Choose two)



- A. R2#no network 10.0.0.0 255.255.255.0
- B. R1#network 19.168.0.0 mask 255.255.0.0
- C. R1#no network 10.0.0.0 255.255.255.0
- D. R2#network 209.165.201.0 mask 255.255.192.0
- E. R2#network 192.168.0.0 mask 255.255.0.0

Answer: A E

Question 39

Refer to the exhibit.

```
SW1#show monitor session all
Session 1
-----
Type                : Remote Destination Session
Source RSPAN VLAN   : 50

Session 2
-----
Type                : Local Session
Source Ports        :
Both                : Fa0/14
Destination Ports   : Fa0/15
Encapsulation       : Native
Ingress             : Disabled
```

An engineer configures monitoring on SW1 and enters the show command to verify operation. What does the output confirm?

- A. SPAN session 1 monitors activity on VLAN 50 of a remote switch
- B. SPAN session 2 only monitors egress traffic exiting port FastEthernet 0/14.
- C. SPAN session 2 monitors all traffic entering and exiting port FastEthernet 0/15.
- D. RSPAN session 1 is incompletely configured for monitoring

Answer: D

Explanation

SW1 has been configured with the following commands:


```
SW1(config)#monitor session 1 source remote vlan 50
SW1(config)#monitor session 2 source interface fa0/14
SW1(config)#monitor session 2 destination interface fa0/15
```

The session 1 on SW1 was configured for Remote SPAN (RSPAN) while session 2 was configured for local SPAN. For RSPAN we need to configure the destination port to complete the configuration.

Note: In fact we cannot create such a session like session 1 because if we only configure “Source RSPAN VLAN 50” (with the command “monitor session 1 source remote vlan 50”) then we will receive a “Type: Remote Source Session” (not “Remote Destination Session”).

Question 40

Refer to the exhibit. PC-1 must access the web server on port 8080. To allow this traffic, which statement must be added to an access control list that is applied on SW2 port G0/0 in the inbound direction?



- A. permit host 172.16.0.2 host 192.168.0.5 eq 8080
- B. permit host 192.168.0.5 host 172.16.0.2 eq 8080
- C. permit host 192.168.0.5 eq 8080 host 172.16.0.2
- D. permit host 192.168.0.5 it 8080 host 172.16.0.2

Answer: C

Explanation

The inbound direction of G0/0 of SW2 only filter traffic from Web Server to PC-1 so the source IP address and port is of the Web Server.

Question 41

Refer to the exhibit.

R1 key chain cisco123 key 1 key-string Cisco123! Ethernet0/0 - Group 10 State is Active 8 state changes, last state change 00:03:33 Virtual IP address is 192.168.0.1 Active virtual MAC address is 0000.0c07.ac0a	R2 key chain cisco123 key 1 key-string Cisco123! Ethernet0/0 - Group 10 State is Active 17 state changes, last state change 00:03:33 Virtual IP address is 192.168.0.1 Active virtual MAC address is 0000.0c07.ac0a
---	--

An engineer is installing a new pair of routers in a redundant configuration. Which protocol ensures that traffic is not disrupted in the event of a hardware failure?

- A. HSRPv2
- B. VRRP
- C. GLBP
- D. HSRPv1

Answer: D

Explanation

The “virtual MAC address” is 0000.0c07.acXX (XX is the hexadecimal group number) so it is using HSRPv1.

Note: HSRP Version 2 uses a new MAC address which ranges from 0000.0C9F.F000 to 0000.0C9F.FFFF.

Question 42

Refer to the exhibit.

```
aaa new-model
aaa authentication login default local-case enable
aaa authentication login ADMIN local-case
username CCNP secret Str0ngP@ssw0rd!
line 0 4
login authentication ADMIN
```

How can you change this configuration so that when user CCNP logs in, the show run command is executed and the session is terminated?

- A. Add the autocmd keyword to the aaa authentication command
- B. Assign privilege level 15 to the CCNP username

- C. Add the access-class keyword to the aaa authentication command
- D. Assign privilege level 14 to the CCNP username
- E. Add the access-class keyword to the username command
- F. Add the autocommand keyword to the username command

Answer: F

Explanation

The “autocommand” causes the specified command to be issued automatically after the user logs in. When the command is complete, the session is terminated. Because the command can be any length and can contain embedded spaces, commands using the autocommand keyword must be the last option on the line. In this specific question, we have to enter this line “username CCNP autocommand show running-config”.

Question 43

Refer to the exhibit. What does the error message relay to the administrator who is trying to configure a Cisco IOS device?

```
<?xml version="1.0" encoding="utf-8"?>
<data xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"/>
```

- A. A NETCONF request was made for a data model that does not exist.
- B. The device received a valid NETCONF request and serviced it without error.
- C. A NETCONF message with valid content based on the YANG data models was made, but the request failed.
- D. The NETCONF running datastore is currently locked.

Answer: A

Explanation

Missing Data Model RPC Error Reply Message

If a request is made for a data model that doesn't exist on the Catalyst 3850 or a request is made for a leaf that is not implemented in a data model, the Server (Catalyst 3850) responds with an empty data response. This is expected behavior.

Reference: <https://www.cisco.com/c/en/us/support/docs/storage-networking/management/200933-YANG-NETCONF-Configuration-Validation.html>

Question 44

In an SD-WAN deployment, which action in the vSmart controller responsible for?

- A. handle, maintain, and gather configuration and status for nodes within the SD-WAN fabric
- B. onboard vEdge nodes into the SD-WAN fabric
- C. gather telemetry data from vEdge routers
- D. distribute policies that govern data forwarding performed within the SD-WAN fabric

Answer: D

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 45

What does Call Admission Control require the client to send in order to reserve the bandwidth?

- A. SIP flow information
- B. Wi-Fi multimedia
- C. traffic specification
- D. VoIP media session awareness

Answer: D

===== New Questions (added on 10th-Oct-2020)
=====

Question 46

Refer to the exhibit.

```
R1
interface GigabitEthernet0/0
 ip address 192.168.250.2 255.255.255.0
 standby 20 ip 192.168.250.1
 standby 20 priority 120
```

```
R2
interface GigabitEthernet0/0
 ip address 192.168.250.3 255.255.255.0
 standby 20 ip 192.168.250.1
 standby 20 priority 110
```

What are two effects of this configuration? (Choose two)

- A. R1 becomes the active router
- B. R1 becomes the standby router
- C. If R2 goes down, R1 becomes active but reverts to standby when R2 comes back online
- D. If R1 goes down, R2 becomes active but reverts to standby when R1 comes back online
- E. If R1 goes down, R2 becomes active and remains the active device when R1 comes back online

Answer: A E

Question 47

Using the EIRP formula, what parameter is subtracted to determine the EIRP value?


- A. antenna cable loss
- B. antenna gain
- C. transmitter power
- D. signal-to-noise ratio

Answer: A

Explanation

Once you know the complete combination of transmitter power level, the length of cable, and the antenna gain, you can figure out the actual power level that will be radiated from the antenna. This is known as the effective isotropic radiated power (EIRP), measured in dBm.

EIRP is a very important parameter because it is regulated by governmental agencies in most countries. In those cases, a system cannot radiate signals higher than a maximum allowable EIRP. To find the EIRP of a system, simply add the transmitter power level to the antenna gain and subtract the cable loss.



$$\text{EIRP} = \text{Tx Power} - \text{Tx Cable} + \text{Tx Antenna}$$


Suppose a transmitter is configured for a power level of 10 dBm (10 mW). A cable with 5-dB loss connects the transmitter to an antenna with an 8-dBi gain. The resulting EIRP of the system is 10 dBm – 5 dB + 8 dBi, or 13 dBm.

You might notice that the EIRP is made up of decibel-milliwatt (dBm), dB relative to an isotropic antenna (dBi), and decibel (dB) values. Even though the units appear to be different, you can safely combine them because they are all in the dB “domain”.

Reference: CCNA Wireless 640-722 Official Cert Guide

Question 48

Refer to the exhibit.



An engineer must ensure that all traffic entering AS 200 will choose Link 2 as an entry 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?

Option A

```
R3(config)#route-map PREPEND permit 10
R3(config-route-map)#set as-path prepend
200 200 200
```

```
R3(config)# router bgp 200
R3(config-router)#neighbor 10.1.1.1 route-
map PREPEND out
```

Option C

```
R3(config)#route-map PREPEND permit 10
R3(config-route-map)#set as-path prepend
```

Option B

```
R3(config)#route-map PREPEND permit 10
R3(config-route-map)#set as-path prepend
100 100 100
```

```
R3(config)# router bgp 200
R3(config-router)#neighbor 10.2.2.2 route-
map PREPEND in
```

Option D

```
R3(config)#route-map PREPEND permit 10
R3(config-route-map)#set as-path prepend
```

100 100 100	200 200 200
R3(config)# router bgp 200 R3(config-router)#neighbor 10.1.1.1 route-map PREPEND in	R3(config)# router bgp 200 R3(config-router)#neighbor 10.2.2.2 route-map PREPEND out

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

Explanation

R3 advertises BGP updates to R1 with multiple AS 100 so R3 believes the path to reach AS 200 via R3 is farther than R2 so R3 will choose R2 to forward traffic to AS 200.

Question 49

Drag and drop the DHCP messages that are exchanged between a client and an AP into the order they are exchanged on the right.

DHCP Request	Step 1
DHCP Offer	Step 2
DHCP Discover	Step 3
DHCP ACK	Step 4

Answer:

- + Step 1: DHCP Discover
- + Step 2: DHCP Offer
- + Step 3: DHCP Request
- + Step 4: DHCP ACK

Explanation

There are four messages sent between the DHCP Client and DHCP Server: DHCPDISCOVER, DHCPOFFER, DHCPREQUEST and DHCPACKNOWLEDGEMENT.

This process is often abbreviated as **DORA** (for Discover, Offer, Request, Acknowledgement).

Question 50

Refer to the exhibit.

```
interface FastEthernet0/1
ip address 209.165.200.225 255.255.255.224
ip nat outside
!
interface FastEthernet0/2
ip address 10.10.10.1 255.255.255.0
ip nat inside
!
access-list 10 permit 10.10.10.0 0.0.0.255
!
```

Which command allows hosts that are connected to FastEthernet0/2 to access the Internet?

- A. ip nat inside source list 10 interface FastEthernet0/1 overload
- B. ip nat outside source static 209.165.200.225 10.10.10.0 overload
- C. ip nat inside source list 10 interface FastEthernet0/2 overload
- D. ip nat outside source list 10 interface FastEthernet0/2 overload

Answer: A

Explanation

The command “ip nat inside source list 10 interface FastEthernet0/1 overload” configures NAT to overload on the address that is assigned to the Fa0/1 interface.

===== New Questions (added on 12th-Oct-2020)
=====

Question 51

In a Cisco SD-Access fabric, which control plane protocol is used for mapping and resolving endpoints?

- A. LISP
- B. DHCP
- C. SXP
- D. VXLAN

Answer: A

Question 52

Refer to the exhibit. What does the snippet of code achieve?

```
with manager.connect(host=192.168.0.1, port=22,  
    username='admin', password='password1', hostkey_verify=True,  
    device_params={'name':'nexus'}) as m:
```

- A. It creates an SSH connection using the SSH key that is stored and the password is ignored
- B. It creates a temporary connection to a Cisco Nexus device and retrieves a token to be used for API calls
- C. It opens an ncclient connection to a Cisco Nexus device and maintains it for the duration of the context
- D. It opens a tunnel and encapsulates the login information, if the host key is correct

Answer: C

Explanation

ncclient is a Python library that facilitates client-side scripting and application development around the NETCONF protocol.

The above Python snippet uses the ncclient to connect and establish a NETCONF session to a Nexus device (which is also a NETCONF server).

===== New Questions (added on 6th-Nov-2020)
=====

Question 53

What are two reasons a company would choose a cloud deployment over an on-prem deployment? (Choose two)

- A. Cloud deployments require long implementation times due to capital expenditure processes. OnPrem deployments can be accomplished quickly using operational expenditure processes
- B. Cloud costs adjust up or down depending on the amount of resources consumed. On- Prem costs for hardware, power, and space are ongoing regardless of usage
- C. In a cloud environment, the company controls technical issues. On-prem environments rely on the service provider to resolve technical issue
- D. Cloud resources scale automatically to an increase in demand. On-prem requires additional capital expenditure
- E. In a cloud environment, the company is in full control of access to their data. On-prem risks access to data due to service provider outages

Answer: B D

Question 54


Which outbound access list, applied to the WAN interface of a router, permits all traffic except for http traffic sourced from the workstation with IP address 10.10.10.1?

- A. ip access-list extended 200
deny tcp host 10.10.10.1 eq 80 any
permit ip any any
- B. ip access-list extended 10
deny tcp host 10.10.10.1 any eq 80
permit ip any any
- C. ip access-list extended NO_HTTP
deny tcp host 10.10.10.1 any eq 80
- D. ip access-list extended 100
deny tcp host 10.10.10.1 any eq 80
permit ip any any

Answer: D

Question 55

Refer to the exhibit. Assuming all links are functional, which path does PC1 take to reach DSW1?



- A. PC1 goes from ALSW1 to DSW1
- B. PC1 goes from ALSW1 to DSW2 to ALSW2 to DSW1
- C. PC1 goes from ALSW1 to DSW2 to Core to DSW1
- D. PC1 goes from ALSW1 to DSW2 to DSW1

Answer: D

Explanation

In the topology above, we see DSW2 has lowest priority 24576 so it is the root bridge for VLAN 10 so surely all traffic for this VLAN must go through it. All of DSW2 ports must be in forwarding state. And:


- + The direct link between DSW1 and ALSW1 is blocked by STP.
- + The direct link between DSW1 and ALSW2 is also blocked by STP.

Therefore PC1 must go via this path: PC1 -> ALSW1 -> DSW2 -> DSW1.

===== New Questions (added on 11th-Nov-2020)
=====

Question 56

Refer to the exhibit.



Both controllers are in the same mobility group. Which result occurs when Client 1 roams between APs that are registered to different controllers in the same WLAN?

- A. Client 1 contact controller B by using an EoIP tunnel
- B. CAPWAP tunnel is created between controller A and controller B
- C. Client 1 users an EoIP tunnel to contact controller A
- D. The client database entry moves from controller A to controller B

Answer: D

Explanation

This is called Inter Controller-L2 Roaming. Inter-Controller (normally layer 2) roaming occurs when a client roam between two APs registered to two different controllers, where each controller has an interface in the client subnet. In this instance, controllers exchange mobility control messages (over UDP port 16666) and the client database entry is moved from the original controller to the new controller.

Question 57

Drag and drop the LIPS components on the left to the correct description on the right.

map server	IPv4 or IPv6 address of an endpoint within a LISP site
ETR	network infrastructure component that learns of EID-prefix mapping entries from an ETR
EID	de-encapsulates LISP packets coming from outside of the LISP site to destinations inside of the site

Answer:

- + IPv4 or IPv6 address of an endpoint within a LISP site: EID
- + network infrastructure component that learns of EID-prefix mapping entries from an ETR: map server
- + de-encapsulates LISP packets coming from outside of the LISP site to destinations inside of the site: ETR

New ENCOR Questions

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.

(Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800c95bb.shtml#localpref)

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. But we cannot name the control-plane (like “transit”).

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/bsm/16-6-1/b-sm-xe-16-6-1-asr920/bsm-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 sparse 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/xs-3s/irl-xe-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: A B

Explanation

Broadcast radiation refers to the processing that is required every time a broadcast is received on a host. Although IP is very efficient from a broadcast perspective when compared to traditional protocols such as Novell Internetwork Packet Exchange (IPX) Service Advertising Protocol (SAP), virtual machines and the vswitch implementation require special consideration. **Because the vswitch is software based, as broadcasts are received the vswitch must interrupt the server CPU** to change contexts to enable the vswitch to process the packet. After the vswitch has determined that the packet is a broadcast, it copies the packet to all the VMNICs, which then pass the broadcast packet up the stack to process. This processing overhead can have a tangible effect on overall server performance if a single domain is hosting a large number of virtual machines.


Note: This overhead effect is not a limitation of the vswitch implementation. It is a result of the software-based nature of the vswitch embedded in the ESX hypervisor.

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

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 have 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.



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

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 nonroutable 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
- + summaries 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 VirtualBox
- 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
- + **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.

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: D

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

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 Stealthwatch 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 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.



Question 44

Which type of antenna does the radiation pattern represent?



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

Answer: A

Question 45

Drag and drop the REST API authentication method from the left to the description on the right.

HTTP basic authentication	public API resource
token-based authentication	username and password in an encoded string
secure vault	API-dependent secret
OAuth	authorization through identity provider

Answer:

- + public API resource: secure vault
- + username and password in an encoded string: HTTP basic authentication
- + API-dependent secret: OAuth
- + authorization through identity provider: token-based authentication

Explanation

When **Secure Vault** is not in use, all information stored in its container is encrypted. When a user wants to use the files and notes stored within the app, they have to first decrypt the database. This happens by filling in a previously determined Security Lock – which could be a PIN or a password of the user's choosing.

When a user leaves the app, it automatically encrypts everything again. This way all data stored in Secure Vault is decrypted only while a user is actively using the app. In all other instances, it remains locked to any attacker, malware or spyware trying to access the data.

How **token-based authentication** works: Users log in to a system and – once authenticated – are provided with a token to access other services without having to enter their username and password multiple times. In short, token-based authentication adds a second layer of security to application, network, or service access.

OAuth is an open standard for authorization used by many APIs and modern applications. The simplest example of OAuth is when you go to log onto a website and it offers one or more opportunities to log on using another website's/service's login. You then click on the button linked to the other website, the other website authenticates you, and the website you were originally connecting to logs you on itself afterward using permission gained from the second website.

Question 46

Which characteristic distinguishes Ansible from Chef?

- A. Ansible lacks redundancy support for the master server. Chef runs two masters in an active/active mode
- B. Ansible uses Ruby to manage configurations. Chef uses YAML to manage configurations
- C. Ansible pushes the configuration to the client. Chef client pulls the configuration from the server
- D. The Ansible server can run on Linux, Unix or Windows. The Chef server must run on Linux or Unix

Answer: C

Explanation

Ansible works by connecting to your nodes and pushing out small programs, called “Ansible modules” to them. These programs are written to be resource models of the desired state of the system. Ansible then executes these modules (over SSH by default), and removes them when finished.

Chef is a much older, mature solution to configure management. Unlike Ansible, it does require an installation of an agent on each server, named chef-client. Also, unlike Ansible, it has a Chef server that each client pulls configuration from.

Question 47

Drag and drop the QoS mechanisms from the left to the correct descriptions on the right.

DSCP	bandwidth management technique which delays datagrams
policy map	mechanism to create a scheduler for packets prior to forwarding
shaping	portion of the IP header used to classify packets
service policy	mechanism to apply a QoS policy to an interface
policing	tool to enforce rate-limiting on ingress/egress
CoS	portion of the 802.1Q header used to classify packets

Answer:

- + bandwidth management technique which delays datagrams: shaping
- + mechanism to create a scheduler for packets prior to forwarding: policy map
- + portion of the IP header used to classify packets: DSCP
- + mechanism to apply a QoS policy to an interface: service policy

- + tool to enforce rate-limiting on ingress/egress: policing
- + portion of the 802.1Q header used to classify packets: CoS

Explanation

To attach a policy map to an input interface, a virtual circuit (VC), an output interface, or a VC that will be used as the service policy for the interface or VC, use the **service-policy** command in the appropriate configuration mode.

Class of Service (CoS) is a 3 bit field within an Ethernet frame header when we use 802.1q which supports virtual LANs on an Ethernet network. This field specifies a priority value which is between 0 and 63 inclusive which can be used in the Quality of Service (QoS) to differentiate traffic.

The **Differentiated Services Code Point (DSCP)** is a 6-bit field in the IP header for the classification of packets. Differentiated Services is a technique which is used to classify and manage network traffic and it helps to provide QoS for modern Internet networks. It can provide services to all kinds of networks.

Traffic policing is also known as rate limiting as it 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.

Traffic shaping retains excess packets in a queue and then schedules the excess for later transmission over increments of time -> It causes delay.

Question 48

In a Cisco SD-WAN solution, how is the health of a data plane tunnel monitored?

- A. with IP SLA
- B. ARP probing
- C. using BFD
- D. with OMP

Answer: C

Explanation

The BFD (Bidirectional Forwarding Detection) is a protocol that detects link failures as part of the Cisco SD-WAN (Viptela) high availability solution, is enabled by default on all vEdge routers, and you cannot disable it.

Question 49

What function does VXLAN perform in an SD-Access deployment?

- A. policy plane forwarding
- B. control plane forwarding
- C. data plane forwarding
- D. systems management and orchestration

Answer: C

Question 50


A server running Linux is providing support for virtual machines along with DNS and DHCP services for a small business. Which technology does this represent?

- A. container
- B. Type 1 hypervisor
- C. hardware pass-through
- D. Type 2 hypervisor

Answer: D

Explanation

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. A 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 51

What is the primary effect of the spanning-tree portfast command?

- A. It enables BPDU messages
- B. It minimizes spanning-tree convergence time

- C. It immediately puts the port into the forwarding state when the switch is reloaded
- D. It immediately enables the port in the listening state

Answer: C

Explanation

Portfast feature should only be used on edge ports (ports directly connected to end stations). Neither edge ports or PortFast enabled ports generate topology changes when the link toggles so we cannot say Portfast reduces the STP convergence time.

PortFast causes a switch or trunk port to enter the spanning tree forwarding state immediately, bypassing the listening and learning states so answer C is the best choice.

===== New Questions (added on 5th-July-2020) =====

Question 52

What is calculated using the numerical values of the transmitter power level, cable loss and antenna gain?


- A. SNR
- B. RSSI
- C. dBi
- D. EIRP

Answer: D

Explanation

Once you know the complete combination of transmitter power level, the length of cable, and the antenna gain, you can figure out the actual power level that will be radiated from the antenna. This is known as the effective isotropic radiated power (EIRP), measured in dBm.

EIRP is a very important parameter because it is regulated by governmental agencies in most countries. In those cases, a system cannot radiate signals higher than a maximum allowable EIRP. To find the EIRP of a system, simply add the transmitter power level to the antenna gain and subtract the cable loss.



$$\text{EIRP} = \text{Tx Power} - \text{Tx Cable} + \text{Tx Antenna}$$

Suppose a transmitter is configured for a power level of 10 dBm (10 mW). A cable with 5-dB loss connects the transmitter to an antenna with an 8-dBi gain. The resulting EIRP of the system is 10 dBm – 5 dB + 8 dBi, or 13 dBm.

You might notice that the EIRP is made up of decibel-milliwatt (dBm), dB relative to an isotropic antenna (dBi), and decibel (dB) values. Even though the units appear to be different, you can safely combine them because they are all in the dB “domain”.

Reference: CCNA Wireless 640-722 Official Cert Guide

Question 53

Which two security features are available when implementing NTP? (Choose two)

- A. encrypted authentication mechanism
- B. dock offset authentication
- C. broadcast association mode
- D. access list based restriction scheme
- E. symmetric server passwords

Answer: A D

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 54

Refer to the exhibit.



An engineer reconfigures the port-channel between SW1 and SW2 from an access port to a trunk and immediately notices this error in SW1's log.

%PM-SP-4-ERR_DISABLE: bpduguard error detected on Gi0/0, putting Gi0/0 in err-disable state.

Which command set resolves this error?

A.

```
Sw1(config)# interface G0/0
Sw1(config-if)# no spanning-tree bpduguard enable
Sw1(config-if)# shut
Sw1(config-if)# no shut
```

B.

```
Sw1(config)# interface G0/0
Sw1(config-if)# spanning-tree bpduguard enable
Sw1(config-if)# shut
Sw1(config-if)# no shut
```

C.

```
Sw1(config)# interface G0/1
Sw1(config-if)# spanning-tree bpduguard enable
Sw1(config-if)# shut
Sw1(config-if)# no shut
```

D.

```
Sw1(config)# interface G0/0
Sw1(config-if)# no spanning-tree bpdufilter
Sw1(config-if)# shut
Sw1(config-if)# no shut
```

Answer: A

Question 55




Company policy restricts VLAN 10 to be allowed only on SW1 and SW2. All other VLANs can be on all three switches. An administrator has noticed that VLAN 10 has propagated to SW3. Which configuration corrects the issue?

- A.
SW2(config)#interface gi1/2
SW2(config)#switchport trunk allowed vlan 10
- B.
SW1(config)#interface gi1/1
SW1(config)#switchport trunk allowed vlan 1-9,11-4094
- C.
SW2(config)#interface gi1/1
SW2(config)#switchport trunk allowed vlan 10
- D.
SW2(config)#interface gi1/2
SW2(config)#switchport trunk allowed vlan 1-9,11-4094

Answer: D

Question 56

Refer to the exhibit.



An engineer must establish eBGP peering between router R3 and router R4. Both routers should use their loopback interfaces as the BGP router ID. Which configuration set accomplishes this task?

- A.
R3(config)#router bgp 200
R3(config-router)#neighbor 10.24.24.4 remote-as 100
R3(config-router)#bgp router-id 10.3.3.3
- R4(config)#router bgp 100
R4(config-router)#neighbor 10.24.24.3 remote-as 200
R4(config-router)#bgp router-id 10.4.4.4
- B.
R3(config)#router bgp 200

```
R3(config-router)#neighbor 10.4.4.4 remote-as 100
R3(config-router)#neighbor 10.4.4.4 update-source loopback0
```

```
R4(config)#router bgp 100
R4(config-router)#neighbor 10.3.3.3 remote-as 200
R4(config-router)#neighbor 10.3.3.3 update-source loopback0
```

C.

```
R3(config)#router bgp 200
R3(config-router)#neighbor 10.24.24.4 remote-as 100
R3(config-router)#neighbor 10.24.24.4 update-source loopback0
```

```
R4(config)#router bgp 100
R4(config-router)#neighbor 10.24.24.3 remote-as 200
R4(config-router)#neighbor 10.24.24.3 update-source loopback0
```

Answer: A

Question 57

Refer to the exhibit.



```
R1(config)# ip nat inside source static 10.70.5.1 10.45.1.7
```

A network architect has partially configured static NAT. which commands should be asked to complete the configuration?

A.

```
R1(config)#interface GigabitEthernet0/0
R1(config)#ip nat outside
```

```
R1(config)#interface GigabitEthernet0/1
R1(config)#ip nat inside
```

B.

```
R1(config)#interface GigabitEthernet0/0
R1(config)#ip nat outside
```

```
R1(config)#interface GigabitEthernet0/1
R1(config)#ip nat inside
```

C.

```
R1(config)#interface GigabitEthernet0/0  
R1(config)#ip nat inside
```

```
R1(config)#interface GigabitEthernet0/1  
R1(config)#ip nat outside
```

D.

```
R1(config)#interface GigabitEthernet0/0  
R1(config)#ip nat inside
```

```
R1(config)#interface GigabitEthernet0/1  
R1(config)#ip nat outside
```

Answer: B

Question 58

What is the result of applying this access control list?


```
ip access-list extended STATEFUL  
10 permit tcp any any established  
20 deny ip any any
```

- A. TCP traffic with the DF bit set is allowed
- B. TCP traffic with the SYN bit set is allowed
- C. TCP traffic with the ACK bit set is allowed
- D. TCP traffic with the URG bit set is allowed

Answer: C

Explanation

The **established** keyword is only applicable to TCP access list entries to match TCP segments that have the ACK and/or RST control bit set (regardless of the source and destination ports), which assumes that a TCP connection has already been established in one direction only. Let's see an example below:




Suppose you only want to allow the hosts inside your company to telnet to an outside server but not vice versa, you can simply use an “established” access-list like this:

```
access-list 100 permit tcp any any established
access-list 101 permit tcp any any eq telnet
!
interface S0/0
ip access-group 100 in
ip access-group 101 out
```

Note:

Suppose host A wants to start communicating with host B using TCP. Before they can send real data, a three-way handshake must be established first. Let’s see how this process takes place:



1. First host A will send a **SYN message** (a TCP segment with SYN flag set to 1, SYN is short for SYNchronize) to indicate it wants to setup a connection with host B. This message includes a sequence (SEQ) number for tracking purpose. This sequence number can be any 32-bit number (range from 0 to 2^{32}) so we use “x” to represent it.

2. After receiving SYN message from host A, host B replies with **SYN-ACK message** (some books may call it “SYN/ACK” or “SYN, ACK” message. ACK is short for ACKnowledge). This message includes a SYN sequence number and an ACK number:

+ SYN sequence number (let’s called it “y”) is a random number and does not have any relationship with Host A’s SYN SEQ number.


+ ACK number is the next number of Host A’s SYN sequence number it received, so we represent it with “x+1”. It means “I received your part. Now send me the next part (x + 1)”.

The SYN-ACK message indicates host B accepts to talk to host A (via ACK part). And ask if host A still wants to talk to it as well (via SYN part).

3. After Host A received the SYN-ACK message from host B, it sends an **ACK message** with ACK number “y+1” to host B. This confirms host A still wants to talk to host B.

Question 59

Refer to exhibit.



MTU has been configured on the underlying physical topology, and no MTU command has been configured on the tunnel interfaces. What happens when a 1500-byte IPv4 packet traverses the GRE tunnel from host X to host Y, assuming the DF bit is cleared?

- A. The packet arrives on router C without fragmentation.
- B. The packet is discarded on router A
- C. The packet is discarded on router B
- D. The packet arrives on router C fragmented.

Answer: D

Explanation

If the DF bit is set to clear (not set), routers can fragment packets regardless of the original DF bit setting.

Whenever we create tunnel interfaces, the GRE IP MTU is automatically configured 24 bytes less than the outbound physical interface MTU. Ethernet interfaces have an MTU value of 1500 bytes so tunnel interfaces by default will have 1476 bytes MTU, which is 24 bytes less than the physical interface. The process of sending a 1500-byte IPv4 packet (with DF bit set to clear) is shown below:

1. The sender sends a 1500-byte packet (20 byte IPv4 header + 1480 bytes of TCP payload).
2. Since the MTU of the GRE tunnel is 1476, the 1500-byte packet is broken into two IPv4 fragments of 1476 and 44 bytes, each in anticipation of the additional 24 bytes of GRE header.
3. The 24 bytes of GRE header is added to each IPv4 fragment. Now the fragments are 1500 (1476 + 24) and 68 (44 + 24) bytes each.
4. The GRE + IPv4 packets that contain the two IPv4 fragments are forwarded to the GRE tunnel peer router.
5. The GRE tunnel peer router removes the GRE headers from the two packets.
6. This router forwards the two packets to the destination host.
7. The destination host reassembles the IPv4 fragments back into the original IPv4 datagram.

Reference: <https://www.cisco.com/c/en/us/support/docs/ip/generic-routing-encapsulation-gre/25885-pmtud-ipfrag.html> (Scenario 5)

Question 60

What is used to measure the total output energy of a Wi-Fi device?

- A. dBi
- B. EIRP
- C. mW
- D. dBm

Answer: C

Explanation

Output power is measured in mW (milliwatts). A milliwatt is equal to one thousandth (10^{-3}) of a watt.

Question 61

Drag and drop the characteristics from the left onto the correct infrastructure deployment types on the right.

significant initial investment but lower reoccurring costs	On Premises
pay-as-you-go model	
physical location of data can be defined in contract with provider	
very scalable and fast delivery of changes in scale	Cloud
company has control over the physical security of equipment	

Answer:

On Premises:

- + significant initial investment but lower reoccurring costs
- + company has control over the physical security of equipment

Cloud:

- + pay-as-you-go model
- + very scalable and fast delivery of changes in scale
- + physical location of data can be defined in contract with provider

Etherchannel Questions

Question 1


Which PAgP mode combination prevents an Etherchannel from forming?

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

Answer: A

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

Answer: A

Trunking Questions

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

Answer: A

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
```

Fa0/18			Fa0/15, Fa0/16, Fa0/17,
			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
- B. SwitchC(config)#vtp pruning vlan 110
- 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

Answer: D

SD-WAN & SD-Access Solutions

Question 1

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

- A. Connects the SD-Access fabric to another fabric or external Layer 3 networks
- B. Connects endpoints to the fabric and forwards their traffic
- C. Provides reachability border nodes in the fabric underlay
- D. Encapsulates end-user data traffic into LISP.

Answer: B

Question 2

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

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

Answer: B

Question 3

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

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

Answer: A

Question 4

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

- A. provides connectivity to external networks
- B. acts as a DNS server
- C. performs route leaking between user-defined virtual networks and shared services
- D. provides additional forwarding capacity to the fabric

Answer: C

Question 5

Which statement about a fabric access point is true?

- A. It is in local mode and must be connected directly to the fabric border node
- B. It is in FlexConnect mode and must be connected directly to the fabric border node

- C. It is in local mode and must be connected directly to the fabric edge switch
- D. It is in FlexConnect mode and must be connected directly to the fabric edge switch

Answer: C

Question 6

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

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

Answer: D

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 of the fabric overlay
- D. The wireless client is part of the fabric overlay

Answer: C

Question 8

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

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

Answer: D

QoS Questions

Question 1

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

- A. All traffic is sent through four egress queues
- B. Port trust is enabled
- C. The Port Cos value is 0
- D. The Cos value of each tagged packet is modified

Answer: C

Question 2

Which QoS mechanism will prevent a decrease in TCP performance?

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

Answer: C

Question 3

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

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

Answer: A

Question 4

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

- A. QOS Group
- B. Discard Eligibility
- C. IP Precedence
- D. MPLS Experimental

Answer: A

Switching Mechanism Questions

Question 1

Which statement about Cisco Express Forwarding is true?

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

Answer: B

Question 2

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

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

Answer: D E

Question 3

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

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

Answer: D

Question 4

What is the difference between a RIB and a FIB?

- A. The RIB is used to make IP source prefix-based switching decisions
- B. The FIB is where all IP routing information is stored
- C. The RIB maintains a mirror image of the FIB
- D. The FIB is populated based on RIB content

Answer: D

Virtualization Questions

Question 1

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



- A. VRF VPN_B
- B. Default VRF
- C. Management VRF
- D. VRF VPN_A

Answer: B

Question 2

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

- A. When BGP is configured, route targets are transmitted as BGP standard communities
- B. Route targets control the import and export of routes into a customer routing table
- C. Route targets allow customers to be assigned overlapping addresses
- D. Route targets uniquely identify the customer routing table

Answer: B

Question 3

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

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

Answer: D E

Question 4

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

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

Answer: A

Question 5

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

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

Answer: B D

Question 6

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
- B. Each virtual machine requires a unique IP address but shares the MAC 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 MAC address but shares the IP address with the physical server

Answer: A

Question 7

What is the main function of VRF-lite?

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

Answer: B

Question 8

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. VPNv4 is not configured between PE1 and PE3
- B. address-family ipv4 vrf is not configured on PE3
- C. After you configure route-target import 999:999 for a VRF on PE3, the route will be accepted
- D. PE1 will reject the route due to automatic route filtering
- E. After you configure route-target import 999:999 for a VRF on PE1, the route will be accepted

Answer: D E

LISP & VXLAN Questions

Question 1

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

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

Answer: A

Question 2

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

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

Answer: B

Question 3

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

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

Answer: A B

Question 4

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



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

Answer: B

Question 5

Which statement about VXLAN is true?

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

Answer: C

EIGRP & OSPF Questions

Question 1

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

- A. broadcast to nonbroadcast
- B. point-to-multipoint to nonbroadcast

- C. broadcast to point-to-point
- D. point-to-multipoint to broadcast

Answer: A

Question 2

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. not established
- C. 2WAY/DROTHER on both routers
- D. FULL/BDR on R1 and FULL/BDR on R2

Answer: B

Question 3

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

Answer: C

Question 4

Which EIGRP feature allows the use of leak maps?

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

Answer: D

Question 5

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

- A. EIGRP supports 6 unequal-cost paths
- B. A path can be used for load balancing only if it is a feasible successor
- 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. Cisco Express Forwarding is required to load-balance across interfaces

Answer: A B

Question 6

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

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

Answer: A

Question 7

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

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

Answer: C

Question 8


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

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

Answer: C

Question 9

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



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

Answer: A

Question 10

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

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

Answer: D

BGP Questions

Question 1


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

Answer: C

Question 2

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 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

B. 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

```
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)#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
```

Answer: A

Question 3

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          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. 192.168.101.18
- B. 192.168.101.6
- C. 192.168.101.10
- D. 192.168.101.14

Answer: A

Question 4

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
- C. local preference, weight, AS path, MED
- D. local preference, weight, MED, AS path

Answer: B

Wireless Questions

Question 1

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

- A. CISCO-DNA-CONTROILLER.local
- B. CAPWAP-CONTROLLER.local
- C. CISCO-CONTROLLER.local
- D. CISCO-CAPWAP-CONTROLLER.local

Answer: D

Question 2

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

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

Answer: B D

Question 3

Which statement about Cisco EAP-FAST is true?

- A. It does not require a RADIUS server certificate
- B. It requires a client certificate

- C. It is an IETF standard.
- D. It operates in transparent mode

Answer: A

Question 4

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. Indirect
- B. Layer 3 intercontroller
- C. Layer 2 intercontroller
- D. Intercontroller

Answer: B

Question 5

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 'Fast Transition' checkbox is unchecked. The 'Protected Management Frame' section shows 'PMF' set to 'Disabled'. The 'WPA+WPA2 Parameters' section shows 'WPA Policy' unchecked and 'WPA2 Policy-AES' checked. The 'Authentication Key Management' section lists several options: '802.1X' (unchecked), 'CCKM' (unchecked), 'PSK' (checked), 'FT 802.1X' (unchecked), and 'FT PSK' (unchecked). The 'PSK Format' is set to 'ASCII' with a dropdown arrow, and a password field with dots is visible below it.

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

Answer: A

Question 6

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

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

Answer: A C

Question 7

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. ISE server
- B. local WLC
- C. RADIUS server
- D. anchor WLC

Answer: B

Question 8

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 used when the APs are set up in a mesh environment and used to bridge between each other
- 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. When connected to the controller, FlexConnect APs can tunnel traffic back to the controller
- E. FlexConnect mode is a wireless solution for branch office and remote office deployments

Answer: D E

Question 9

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

- A. NTP server
- B. PKI server

- C. RADIUS server
- D. TACACS server

Answer: C

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

Answer: B

Question 11

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. Active
- C. On
- D. Passive

Answer: C

Question 12

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

Answer: A

Question 13

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

Answer: B

Question 14


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

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

Answer: B D

Question 15

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



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

Answer: A

HSRP & VRRP Questions

Question 1

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

- A. Its virtual MAC is 0000.0C07.ACxx
- B. Its multicast virtual MAC is 0000.5E00.01xx
- C. Its default configuration allows for pre-emption
- D. It supports tracking
- E. It supports unique virtual MAC addresses

Answer: A D

Question 2

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

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

Answer: A

Question 3

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

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

Answer: A

Question 4

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

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

Answer: A

Question 5

What are three valid HSRP states? (Choose three)

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

Answer: A B E

Question 6

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

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

Answer: B E

Network Assurance Questions

Question 1

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

Answer: A

Question 2

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

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

Answer: D

Question 3

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

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

Answer: B C

Question 4

At which layer does Cisco DNA Center support REST controls?

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

Answer: D

Question 5

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

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

Answer: D E

Question 6

Which statement about an RSPAN session configuration is true?

- A. A filter must be configured for RSPAN Regions
- B. Only one session can be configured at a time
- C. A special VLAN type must be used as the RSPAN destination.
- D. Only incoming traffic can be monitored

Answer: C

Question 7

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

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

Answer: B

Question 8

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. 2
- C. 4
- D. 8

Answer: B

Question 9

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

Answer: A

Security Questions

Question 1

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
  stoppbits 1
line aux 0
  exec-timeout 0 0
  privilege level 15
  logging synchronous
  stoppbits 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. 1
- B. 7
- C. 13
- D. 15

Answer: A

Question 2

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

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

Answer: A

Question 3

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

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

Answer: B

Question 4

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

- A. uses flexible NetFlow
- B. assigns a VLAN to the endpoint

- C. classifies traffic based on the contextual identity of the endpoint rather than its IP address
- D. classifies traffic based on advanced application recognition

Answer: C

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 password is encrypted with a stronger encryption method
- B. There is no difference and both passwords are encrypted identically
- C. The enable password cannot be decrypted
- D. The enable secret password is protected via stronger cryptography mechanisms

Answer: D

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?

A. 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 default group tacacs+
aaa session-id common

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

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

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

Answer: C

Question 7

Which NGFW mode block flows crossing the firewall?

- A. Passive
- B. Tap
- C. Inline tap
- D. Inline

Answer: D

Question 8

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

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

Answer: D

Access-list Questions

Question 1

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

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

Answer: C

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. Access control lists that are applied outbound to a router interface do not affect traffic that is sourced from the router
- B. Only standard access control lists can block traffic from a source IP address
- 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. The access control list must contain an explicit deny to block traffic from the router

Answer: A

Question 3

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

Answer: C

Question 4

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

- A. Deny tcp any any eq 80
Permit tcp any any gt 21 lt 444
- B. Permit tcp any any neq 80
- C. Permit tcp any any range 22 443
Deny tcp any any eq 80
- D. Deny tcp any any neq 80
Permit tcp any any range 22 443

Answer: A (?)

Question 5

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
permit ip 10.1.10.0 255.255.255.0 10.1.2.0 255.255.255.0

```
B. 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
```

```
C. config t
ip access-list extended EGRESS2
permit ip 10.1.10.0 0.0.0.295 10.1.2.0 0.0.0.299
permit ip 10.1.100.0 0.0.0.299 10.1.2.0 0.0.0.299
deny ip any any
!
interface g0/1
no ip access-group EGRESS out
ip access-group EGRESS2 out
```

```
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
```

Answer: B

Automation Questions

Question 1

Which requirement for an Ansible-managed node is true?

- A. It must be a Linux server or a Cisco device
- B. It must have an SSH server running
- C. It must support ad hoc commands.
- D. It must have an Ansible Tower installed

Answer: A

Question 2

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

- A. It is provided using NGINX acting as a proxy web server
- B. It is not supported on Cisco devices
- C. It requires certificates for authentication
- D. It is used for HTTP and HTTPS requests

Answer: A

Question 3

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

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

Answer: A D

Question 4

Which exhibit displays a valid JSON file?

A. {
 "hostname": "edge_router_1"
 "interfaces": {
 "GigabitEthernet1/1"
 "GigabitEthernet1/2"
 "GigabitEthernet1/3"
 }
}

B. {
 "hostname": "edge_router_1"
 "interfaces": {
 "GigabitEthernet1/1",
 "GigabitEthernet1/2",
 "GigabitEthernet1/3",
 },
}

C. {
 "hostname": "edge_router_1"
 "interfaces": [
 "GigabitEthernet1/1"
 "GigabitEthernet1/2"
 "GigabitEthernet1/3"
]
}

```
D. {  
  "hostname": "edge_router_1",  
  "interfaces": [  
    "GigabitEthernet1/1",  
    "GigabitEthernet1/2",  
    "GigabitEthernet1/3"  
  ]  
}
```

Answer: D

Question 5

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'

B. event manager applet ondemand
event manual
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'

D. event manager applet ondemand
action 1.0 syslog priority critical msg 'This is a message from ondemand'

Answer: C

Question 6

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

Answer: B

Question 7

What is the structure of a JSON web token?

- A. three parts separated by dots header payload, and signature
- B. header and payload
- C. three parts separated by dots version header and signature
- D. payload and signature

Answer: A

Question 8

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

Answer: A D

Question 9

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. Python
- C. Bash script

- D. NETCONF
- E. REST

Answer: E

Question 10

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

- A. HTTPS
- B. SSH
- C. RESTCONF
- D. TLS
- E. NETCONF

Answer: C E

Question 11

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

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

Answer: B

Question 12

Which JSON syntax is valid?

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

Answer: C

Automation Questions 2

Question 1

Which statements are used for error handling in Python?

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

Answer: B

Question 2

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",
```


- A. The client made a request a resource that does not exist
- B. The server has not implemented the functionality that is needed to fulfill the request
- C. The request accepted for processing, but the processing was not completed
- D. The POST/PUT request was fulfilled and a new resource was created, information about the resource is in the response body

Answer: A

Question 5

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

Answer: C

Question 6

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

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

Answer: C

Question 7

What do Cisco DNA southbound APIs provide?

- A. Interface between the controller and the network devices
- B. NETCONF API interface for orchestration communication
- C. RESTful API interface for orchestrator communication
- D. Interface between the controller and the consumer

Answer: A

Question 8

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 gets 'logging directly to console'
- B. event manager applet boom
event syslog pattern 'UP'
action 1.0 syslog priority direct msg 'log directly to console'
- C. event manager applet boom
event syslog pattern 'UP'
action 1.0 puts 'logging directly to console'
- D. event manager applet boom
event syslog pattern 'UP'
action 1.0 string 'logging directly to console'

Answer: C

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

Answer: E

Question 10

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

- A. Agentless tools require no messaging systems between master and slaves.
- B. Agentless tools use proxy nodes to interface with slave nodes.
- C. Agent-based tools do not require a high-level language interpreter such as Python or Ruby on slave nodes.
- D. Agent-based tools do not require installation of additional software packages on the slave nodes.

Answer: C

Question 11

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 create more secure and efficient SNMP OIDs.
- C. They make the CLI simpler and more efficient.
- D. They provide a standardized data structure, which results in configuration scalability and consistency.

Answer: D

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

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

Answer: D

Miscellaneous Questions

Question 1

Which two mechanisms are available to secure NTP? (Choose two)

- A. IP prefix list-based
- B. IPsec
- C. TACACS-based authentication
- D. IP access list-based
- E. Encrypted authentication

Answer: D E

Question 2

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. Inside source addresses are translated to the 209.165.201.0/27 subnet
- B. It establishes a one-to-one NAT translation
- C. The 10.1.1.0/27 subnet is assigned as the inside global address range
- D. The 209.165.201.0/27 subnet is assigned as the outside local address range
- E. The 10.1.1.0/27 subnet is assigned as the inside local addresses

Answer: A E

Question 3

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
- B. Enter spanning-tree port-priority 224 on SW1
- C. Enter spanning-tree port-priority 4 on SW2
- D. Enter spanning-tree port-priority 64 on SW2

Answer: A

Question 4

What is a benefit of deploying an on-premises infrastructure versus a cloud infrastructure deployment?

- A. faster deployment times because additional infrastructure does not need to be purchased
- B. lower latency between systems that are physically located near each other
- C. less power and cooling resources needed to run infrastructure on-premises
- D. ability to quickly increase compute power without the need to install additional hardware

Answer: B

Question 5

Which two GRE features are configured to prevent fragmentation? (Choose two)

- A. TCP window size
- B. TCP MSS
- C. IP MTU
- D. DF bit Clear
- E. MTU ignore
- F. PMTUD

Answer: B F

Question 6

Which TCP setting is tuned to minimize the risk of fragmentation on a GRE/IP tunnel?

- A. MTU
- B. Window size
- C. MRU
- D. MSS

Answer: D

Question 7

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. Configure the logging synchronous command under the vty
- D. Press the TAB key to reprint the command in a new line
- E. Increase the number of lines on the screen using the terminal length command

Answer: C D

Question 8

Which statement about multicast RPs is true?

- A. RPs are required only when using protocol independent multicast dense mode
- B. RPs are required for protocol independent multicast sparse mode and dense mode
- C. By default, the RP is needed periodically to maintain sessions with sources and receivers
- D. By default, the RP is needed only to start new sessions with sources and receivers

Answer: D

Question 9

Which IPv6 migration method relies on dynamic tunnels that use the 2002::/16 reserved address space?

- A. 6RD
- B. 6to4
- C. ISATAP
- D. GRE

Answer: B

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. Incorrect destination IP addresses are configured on the tunnel
- B. There is link flapping on the tunnel
- C. There is instability in the network due to route flapping
- D. The tunnel mode and tunnel IP address are misconfigured
- E. The tunnel destination is being routed out of the tunnel interface

Answer: C E

Drag Drop Questions

Question 1

Drag and drop the characteristics from the left onto the correct routing protocol types on the right.

	OSPF
supports unequal path load balancing	
link state routing protocol	
distance vector routing protocol	
metric is based on delay and reliability by default	EIGRP
makes it easy to segment the network logically	
constructs three tables as part of its operation: neighbor table, topology table and routing table	

Answer:

OSPF:

- + link state routing protocol
- + makes it easy to segment the network logically
- + constructs three tables as part of its operation: neighbor table, topology table and routing table

EIGRP:

- + supports unequal path load balancing
- + distance vector routing protocol
- + metric is based on delay and reliability by default (?)

Explanation

Maybe there is something wrong with the answer "metric is based on delay and reliability by default" as OSPF metric is only dependent on the interface bandwidth & reference bandwidth while EIGRP metric is dependent on bandwidth and delay by default. But only EIGRP metric is based on delay so "EIGRP" is a better answer.

Both OSPF and EIGRP have three tables to operate: neighbor table (store information about OSPF/EIGRP neighbors), topology table (store topology structure of the network) and routing table (store the best routes).

Question 2

Drag and drop the characteristics from the left onto the correct infrastructure deployment types on the right.

	On Premises
customizable hardware, purpose-built systems	
easy to scale and upgrade	
more suitable for companies with specific regulatory or security requirements	
resources can be over or underutilized as requirements vary	
requires a strong and stable internet connection	
built-in, automated data backups and recovery	

	Cloud

Answer:

On Premises:

- + resources can be over or underutilized as requirements vary
- + customizable hardware, purpose-built systems
- + more suitable for companies with specific regulatory or security requirements

Cloud:

- + easy to scale and upgrade
- + requires a strong and stable internet connection
- + built-in, automated data backups and recovery

Question 3

Drag and drop the description from the left onto the correct QoS components on the right.

	Traffic Policing
causes TCP retransmission when traffic is dropped	
buffers excessive traffic	
introduces no delay and jitter	
introduces delay and jitter	Traffic Shaping
drops excessive traffic	
typically delays, rather than drops traffic	

Answer:**Traffic Policing:**

- + introduces no delay and jitter
- + drops excessive traffic
- + causes TCP retransmission when traffic is dropped

Traffic Shaping:

- + buffers excessive traffic
- + introduces delay and jitter
- + typically delays, rather than drops traffic