- **Vendor:** Cisco
- **Exam Code:** 350-401
- **Exam Name:** Implementing and Operating Cisco Enterprise Network Core Technologies (ENCOR)
- **New Updated Questions from Braindump2go (Updated in August/2020)**

QUESTION 84
Which feature must be configured to allow packet capture over Layer 3 infrastructure'?

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

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

QUESTION 85
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
Explanation:
Cisco Express Forwarding (CEF) provides the ability to switch packets through a device in a very quick and efficient way while also keeping the load on the router’s processor low. CEF is made up of two different main components: the Forwarding Information Base (FIB) and the Adjacency Table. These are automatically updated at the same time as the routing table.
The Forwarding Information Base (FIB) contains destination reachability information as well as next hop information. This information is then used by the router to make forwarding decisions. The FIB allows for very efficient and easy lookups.
The adjacency table is tasked with maintaining the layer 2 next-hop information for the FIB.
Note: A fast cache is only used when fast switching is enabled while CEF is disabled.

QUESTION 86
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
Explanation:
Answer C and answer D are not correct as only route distinguisher (RD) identifies the customer routing table and "allows customers to be assigned overlapping addresses".
Answer A is not correct as "When BGP is configured, route targets are transmitted as BGP extended communities"

QUESTION 87
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: BF
Explanation:
The IP protocol was designed for use on a wide variety of transmission links. Although the maximum length of an IP datagram is 65535, most transmission links enforce a smaller maximum packet length limit, called an MTU. The value of the MTU depends on the type of the transmission link. The design of IP accommodates MTU differences since it allows routers to fragment IP datagrams as necessary. The receiving station is responsible for the reassembly of the fragments back into the original full size IP datagram.
Fragmentation and Path Maximum Transmission Unit Discovery (PMTUD) is a standardized technique to determine the maximum transmission unit (MTU) size on the network path between two hosts, usually with the goal of avoiding IP fragmentation. PMTUD was originally intended for routers in IPv4. However, all modern operating systems use it on endpoints.
The TCP Maximum Segment Size (TCP MSS) defines the maximum amount of data that a host is willing to accept in a single TCP/IP datagram. This TCP/IP datagram might be fragmented at the IP layer. The MSS value is sent as a TCP header option only in TCP SYN segments. Each side of a TCP connection reports its MSS value to the other side. Contrary to popular belief, the MSS value is not negotiated between hosts. The sending host is required to limit the size of data in a single TCP segment to a value less than or equal to the MSS reported by the receiving host.
TCP MSS takes care of fragmentation at the two endpoints of a TCP connection, but it does not handle the case where there is a smaller MTU link in the middle between these two endpoints. PMTUD was developed in order to avoid fragmentation in the path between the endpoints. It is used to dynamically determine the lowest MTU along the path from a packet's source to its destination.
Reference: http://www.cisco.com/c/en/us/support/docs/ip/generic-routing-encapsulation-gre/25885-pmtud-ipfrag.html (there is some examples of how TCP MSS avoids IP Fragmentation in this link but it is too long so if you want to read please visit this link)
Note: IP fragmentation involves breaking a datagram into a number of pieces that can be reassembled later.
If the DF bit is set to clear, routers can fragment packets regardless of the original DF bit setting -> Answer D is not correct.

QUESTION 88
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 GigabitEthernetO/O 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?

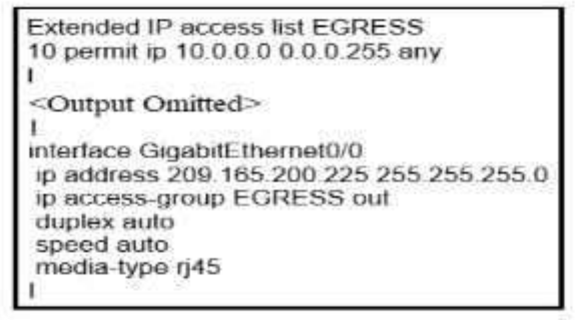

```
Extended IP access list EGRESS
10 permit ip 10.0.0.0 0.0.0.255 any
!
<Output Omitted>
!
interface GigabitEthernet0/0
 ip address 209.165.200.225 255.255.255.0
 ip access-group EGRESS out
 duplex auto
 speed auto
 media-type rj45
!
```

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 89
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
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 90
Which LISP device is responsible for publishing EID-to-RLOC mappings for a site?

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

Answer: A
Explanation:
An Egress Tunnel Router (ETR) connects a site to the LISP-capable part of a core network (such as the Internet), publishes EID-to-RLOC mappings for the site, responds to Map-Request messages, and decapsulates and delivers LISP-encapsulated user data to end systems at the site.
Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_lisp/configuration/xe-3s/irl-xe-3s-book/irl-overview.html

QUESTION 91
Which access controls list allows only TCP traffic with a destination port range of 22-433, excluding port 80?

A. Deny tcp any any eq 80
   Permit tco any any gt 21 it 444
B. Permit tcp any any ne 80
C. Permit tco any any range 22 443
   Deny tcp any any eq 80
D. Deny tcp any any ne 80
   Permit tcp any any range 22 443

Answer: A

QUESTION 92
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
Explanation:
Cisco TrustSec uses tags to represent logical group privilege. This tag, called a Security Group Tag (SGT), is used in access policies. The SGT is understood and is used to enforce traffic by Cisco switches, routers and firewalls . Cisco TrustSec is defined in three phases: classification, propagation and enforcement.
When users and devices connect to a network, the network assigns a specific security group. This process is called classification. Classification can be based on the results of the authentication or by associating the SGT with an IP, VLAN, or port-profile (-> Answer A and answer C are not correct as they say “assigned … on a switch” only. Answer D is not correct either as it says “assigned to each router”).

QUESTION 93
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: A
Explanation:
The major components of the vSmart controller are:
+ Control plane connections - Each vSmart controller establishes and maintains a control plane connection with each vEdge router in the overlay network. (In a network with multiple vSmart controllers, a single vSmart controller may have connections only to a subset of the vEdge routers, for load-balancing purposes.) Each connection, which runs as a DTLS tunnel, is established after device authentication succeeds, and it carries the encrypted payload between the vSmart controller and the vEdge router. This payload consists of route information necessary for the vSmart controller to determine the network topology, and then to calculate the best routes to network destinations and distribute this route information to the vEdge routers. The DTLS connection between a vSmart controller and a vEdge router is a permanent connection. The vSmart controller has no direct peering relationships with any devices that a vEdge router is connected to on the service side (so answer C is not correct as vSmart only manages vEdge routers only, not the whole nodes within SD-WAN fabric).
+ OMP (Overlay Management Protocol) - The OMP protocol is a routing protocol similar to BGP that manages the Cisco SD-WAN overlay network. OMP runs inside DTLS control plane connections and carries the routes, next hops, keys, and policy information needed to establish and maintain the overlay network. OMP runs between the vSmart controller and the vEdge routers and carries only control plane information. The vSmart controller processes the routes and advertises reachability information learned from these routes to other vEdge routers in the overlay network.
+ Authentication - The vSmart controller has pre-installed credentials that allow it to authenticate every new vEdge router that comes online (-> Answer A is correct). These credentials ensure that only authenticated devices are allowed access to the network.

+ Key reflection and rekeying - The vSmart controller receives data plane keys from a vEdge router and reflects them to other relevant vEdge routers that need to send data plane traffic.
+ Policy engine - The vSmart controller provides rich inbound and outbound policy constructs to manipulate routing information, access control, segmentation, extranets, and other network needs.
+ Netconf and CLI - Netconf is a standards-based protocol used by the vManage NMS to provision a vSmart controller. In addition, each vSmart controller provides local CLI access and AAA.
Reference: https://www.cisco.com/c/en/us/td/docs/routers/sdwan/configuration/sdwan-xe-gs-book/system-overview.html

QUESTION 94
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 GO/1 on SW2. but the port remains blocked.
Which command should be entered on the ports that are connected to Lmk2 to resolve the issue?

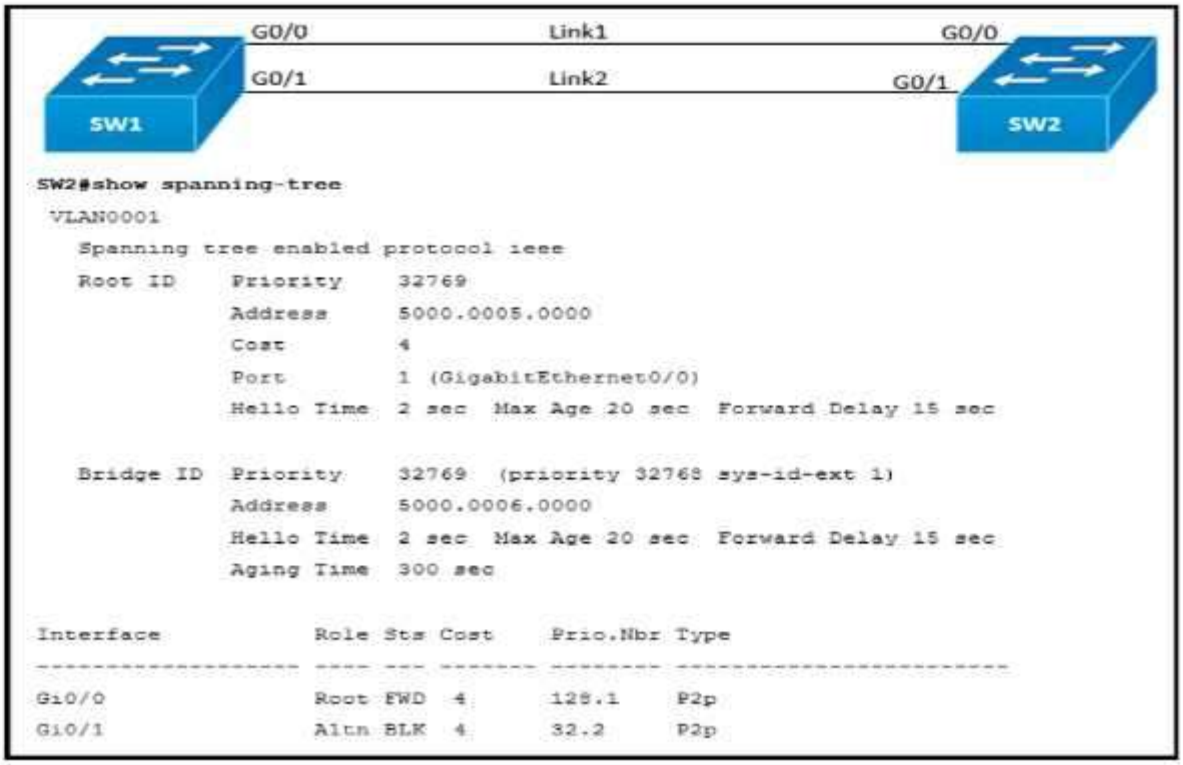

```
SW2#show spanning-tree
 VLAN0001
   Spanning tree enabled protocol ieee
   Root ID     Priority    32769
               Address     5000.0005.0000
               Cost        4
               Port        1 (GigabitEthernet0/0)
               Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec

   Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
               Address     5000.0006.0000
               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               Altn BLK  4       32.2     P2p
```

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
Explanation:
SW1 needs to block one of its ports to SW2 to avoid a bridging loop between the two switches. Unfortunately, it blocked the fiber port Link2. But how does SW2 select its blocked port? Well, the answer is based on the BPDUs it receives from SW1. A BPDU is superior than another if it has:
1. A lower Root Bridge ID
2. A lower path cost to the Root
3. A lower Sending Bridge ID
4. A lower Sending Port ID
These four parameters are examined in order. In this specific case, all the BPDUs sent by SW1 have the same Root Bridge ID, the same path cost to the Root and the same Sending Bridge ID. The only parameter left to select the best

one is the Sending Port ID (Port ID = port priority + port index). And the port index of Gi0/0 is lower than the port index of Gi0/1 so Link 1 has been chosen as the primary link.
Therefore we must change the port priority to change the primary link. The lower numerical value of port priority, the higher priority that port has. In other words, we must change the port-priority on Gi0/1 of SW1 (not on Gi0/1 of SW2) to a lower value than that of Gi0/0.

QUESTION 95
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
Explanation:
Ansible can communicate with modern Cisco devices via SSH or HTTPS so it does not require an SSH server -> Answer B is not correct.
An Ansible ad-hoc command uses the /usr/bin/ansible command-line tool to automate a single task on one or more managed nodes. Ad-hoc commands are quick and easy, but they are not reusable -> It is not a requirement either -> Answer C is not correct.
Ansible Tower is a web-based solution that makes Ansible even more easy to use for IT teams of all kinds. But it is not a requirement to run Ansible -> Answer D is not correct.
Therefore only answer A is the best choice left. An Ansible controller (the main component that manages the nodes), is supported on multiple flavors of Linux, but it cannot be installed on Windows.