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?

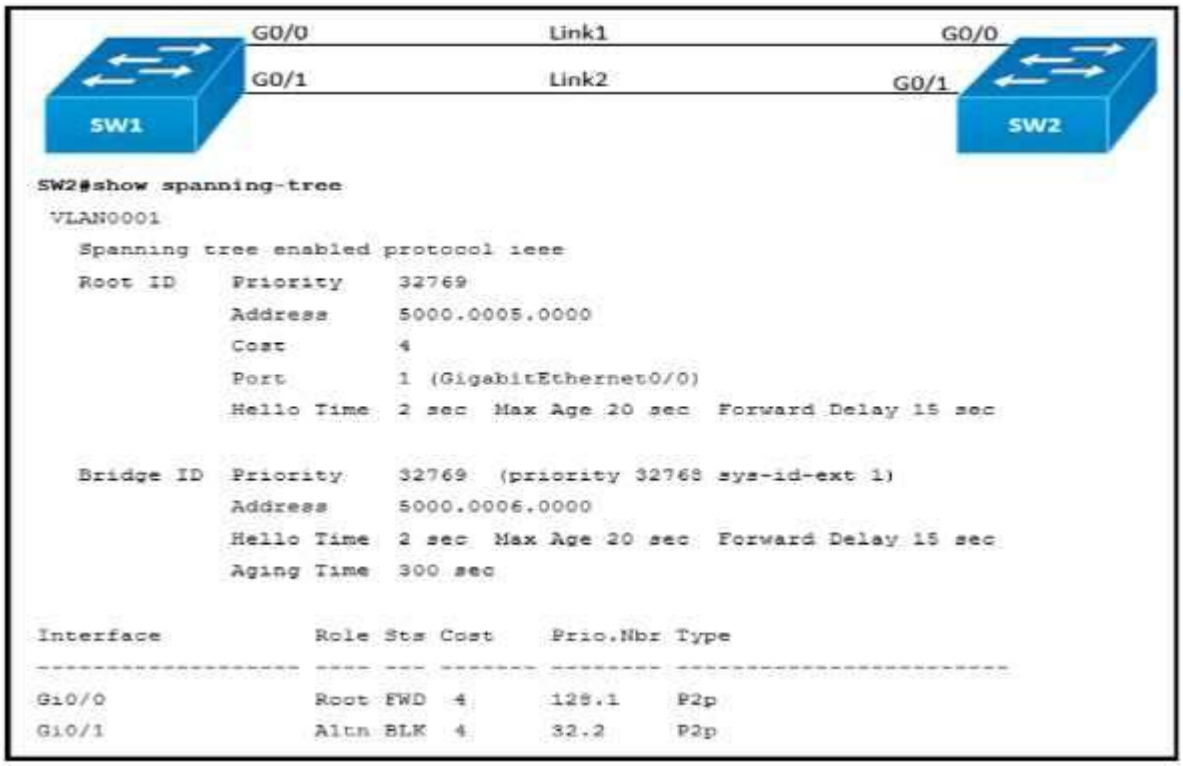

A. Enter spanning-tree port-priority 32 on SW1.
B. Enter spanning-tree port-priority 224 on SW1.
C. Enter spanning-tree port-priority 4 on SW2.
D. Enter spanning-tree port-priority 64 on SW2.

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.

QUESTION 96
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
Explanation:
Syslog levels are listed below:

Level	Keyword	Description
0	emergencies	System is unusable
1	alerts	Immediate action is needed
2	critical	Critical conditions exist
3	errors	Error conditions exist
4	warnings	Warning conditions exist
5	notification	Normal, but significant, conditions exist
6	informational	Informational messages
7	debugging	Debugging messages

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

QUESTION 97
Which DNS lookup does an access point perform when attempting CAPWAP discovery?

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

Answer: D
Explanation:
The Lightweight AP (LAP) can discover controllers through your domain name server (DNS). For the access point (AP) to do so, you must configure your DNS to return controller IP addresses in response to CISCO-LWAPP-

CONTROLLER.localdomain, where localdomain is the AP domain name. When an AP receives an IP address and DNS information from a DHCP server, it contacts the DNS to resolve CISCO-CAPWAP-CONTROLLER.localdomain. When the DNS sends a list of controller IP addresses, the AP sends discovery requests to the controllers.
The AP will attempt to resolve the DNS name CISCO-CAPWAP-CONTROLLER.localdomain. When the AP is able to resolve this name to one or more IP addresses, the AP sends a unicast CAPWAP Discovery Message to the resolved IP address(es). Each WLC that receives the CAPWAP Discovery Request Message replies with a unicast CAPWAP Discovery Response to the AP.
Reference: https://www.cisco.com/c/en/us/support/docs/wireless/4400-series-wireless-lan-controllers/107606-dns-wlc-config.html

QUESTION 98
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 99
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: BC
Explanation:
IP SLAs allows Cisco customers to analyze IP service levels for IP applications and services, to increase productivity, to lower operational costs, and to reduce the frequency of network outages. IP SLAs uses active traffic monitoring–the generation of traffic in a continuous, reliable, and predictable manner–for measuring network performance.
Being Layer-2 transport independent, IP SLAs can be configured end-to-end over disparate networks to best reflect the metrics that an end-user is likely to experience.
Reference: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/configuration/15-mt/sla-15-mt-book/sla_overview.html

QUESTION 100
Which two statements about Cisco Express Forwarding load balancing are true?

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: DE
Explanation:
Cisco IOS software basically supports two modes of CEF load balancing: On per-destination or per-packet basis. For per destination load balancing a hash is computed out of the source and destination IP address (-> Answer E is correct). This hash points to exactly one of the adjacency entries in the adjacency table (-> Answer D is correct), providing that the same path is used for all packets with this source/destination address pair. If per packet load

balancing is used the packets are distributed round robin over the available paths. In either case the information in the FIB and adjacency tables provide all the necessary forwarding information, just like for non-load balancing operation. The number of paths used is limited by the number of entries the routing protocol puts in the routing table, the default in IOS is 4 entries for most IP routing protocols with the exception of BGP, where it is one entry. The maximum number that can be configured is 6 different paths -> Answer A is not correct.
Reference:
https://www.cisco.com/en/US/products/hw/modules/ps2033/prod_technical_reference09186a00800afeb7.html

QUESTION 101
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 102
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: DE

QUESTION 103
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
Explanation:
On Ethernet interfaces the OSPF hello intervl is 10 second by default so in this case there would be a Hello interval mismatch -> the OSPF adjacency would not be established.