QUESTION 104
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 no supported on Cisco devices.
C. It required certificates for authentication.
D. It is used for HTTP and HTTPs requests.

Answer: C
Explanation:
The https-based protocol-RESTCONF (RFC 8040), which is a stateless protocol, uses secure HTTP methods to provide CREATE, READ, UPDATE and DELETE (CRUD) operations on a conceptual datastore containing YANG-defined data -> RESTCONF only uses HTTPs.
RESTCONF servers MUST present an X.509v3-based certificate when establishing a TLS connection with a RESTCONF client. The use of X.509v3-based certificates is consistent with NETCONF over TLS -> Answer C is correct.
Reference: https://tools.ietf.org/html/rfc8040

QUESTION 105
Which controller is the single plane of management for Cisco SD-WAN?

A. vBond
B. vEdge
C. vSmart
D. vManange

Answer: D
Explanation:
The primary components for the Cisco SD-WAN solution consist of the vManage network management system (management plane), the vSmart controller (control plane), the vBond orchestrator (orchestration plane), and the vEdge router (data plane).
+ vManage - This centralized network management system provides a GUI interface to easily monitor, configure, and maintain all Cisco SD-WAN devices and links in the underlay and overlay network.
+ vSmart controller - This software-based component is responsible for the centralized control plane of the SD-WAN network. It establishes a secure connection to each vEdge router and distributes routes and policy information via the Overlay Management Protocol (OMP), acting as a route reflector. It also orchestrates the secure data plane connectivity between the vEdge routers by distributing crypto key information, allowing for a very scalable, IKE-less architecture.

+ vBond orchestrator - This software-based component performs the initial authentication of vEdge devices and orchestrates vSmart and vEdge connectivity. It also has an important role in enabling the communication of devices that sit behind Network Address Translation (NAT).
+ vEdge router - This device, available as either a hardware appliance or software-based router, sits at a physical site or in the cloud and provides secure data plane connectivity among the sites over one or more WAN transports. It is responsible for traffic forwarding, security, encryption, Quality of Service (QoS), routing protocols such as Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF), and more.
Reference: https://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/SDWAN/CVD-SD-WAN-Design-2018OCT.pdf

QUESTION 106
Drag and Drop Question
Drag and drop the characteristics from the left onto the correct infrastructure deployment types on the right.

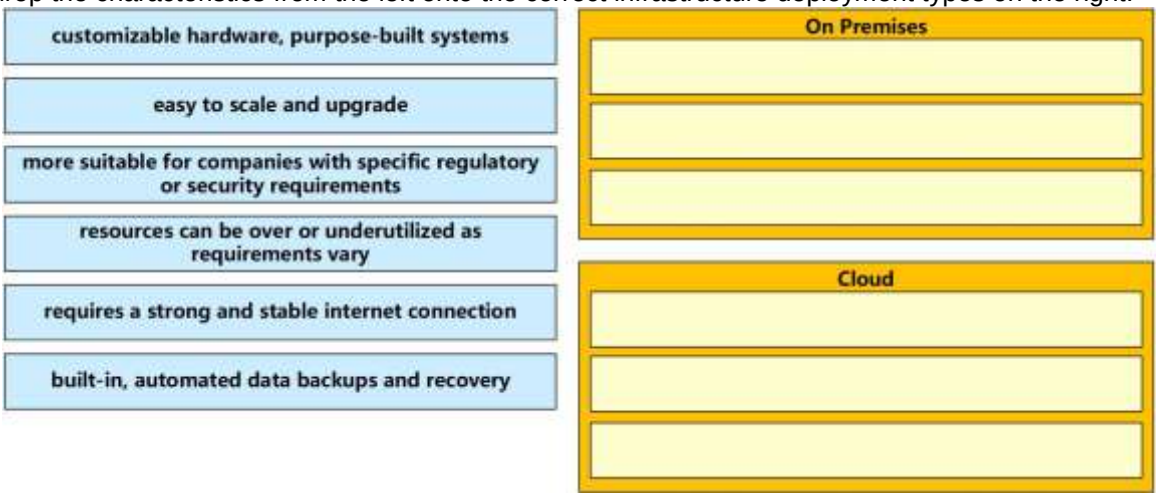

Answer:

QUESTION 107
Drag and Drop Question
Drag and drop the description from the left onto the correct QoS components on the right.

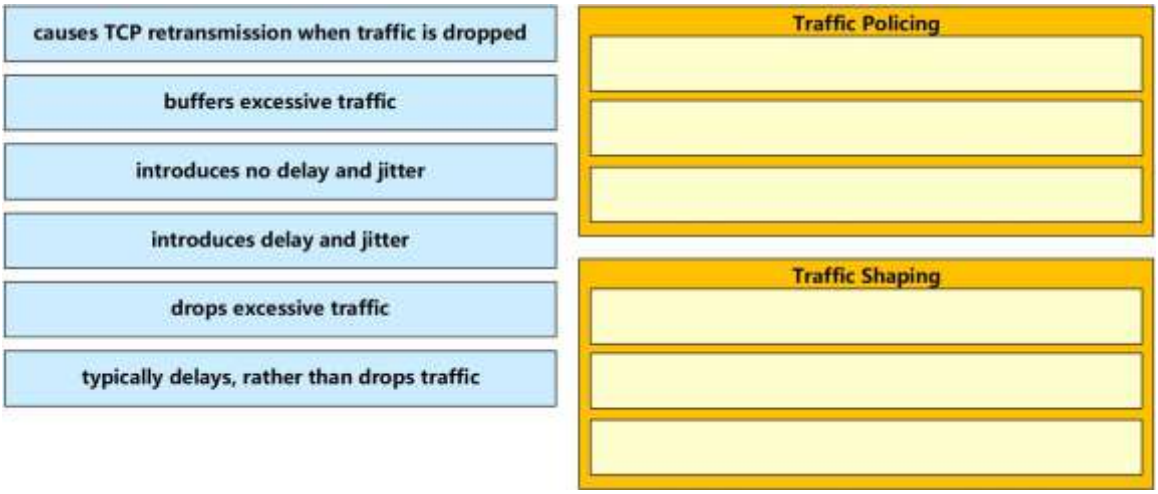

Answer:

Explanation:
The following diagram illustrates the key difference between traffic policing and traffic shaping. 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. In contrast to policing, 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.
Note: Committed information rate (CIR): The minimum guaranteed data transfer rate agreed to by the routing device.

QUESTION 108
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 g 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 GO.
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
Explanation:
EEM offers the ability to monitor events and take informational or corrective action when the monitored events occur or reach a threshold. An EEM policy is an entity that defines an event and the actions to be taken when that event occurs. There are two types of EEM policies: an applet or a script. An applet is a simple form of policy that is defined within the CLI configuration.

To specify the event criteria for an Embedded Event Manager (EEM) applet that is run by sampling Simple Network Management Protocol (SNMP) object identifier values, use the event snmp command in applet configuration mode.
event snmp oid oid-value get-type {exact | next} entry-op operator entry-val entry-value [exit-comb {or | and}] [exit-op operator] [exit-val exit-value] [exit-time exit-time-value] poll-interval poll-int-value
+ oid: Specifies the SNMP object identifier (object ID)
+ get-type: Specifies the type of SNMP get operation to be applied to the object ID specified by the oid-value argument.
- next - Retrieves the object ID that is the alphanumeric successor to the object ID specified by the oid-value argument.
+ entry-op: Compares the contents of the current object ID with the entry value using the specified operator. If there is a match, an event is triggered and event monitoring is disabled until the exit criteria are met.
+ entry-val: Specifies the value with which the contents of the current object ID are compared to decide if an SNMP event should be raised.
+ exit-op: Compares the contents of the current object ID with the exit value using the specified operator. If there is a match, an event is triggered and event monitoring is reenabled.
+ poll-interval: Specifies the time interval between consecutive polls (in seconds)
Reference: https://www.cisco.com/en/US/docs/ios/12_3t/12_3t4/feature/guide/gtioseem.html

QUESTION 109
What are three valid HSRP states? (Choose three)

A. listen
B. learning
C. full
D. established
E. speak
F. IN IT

Answer: ABE
Explanation:
HSRP consists of 6 states:

State	Description
Initial	This is the beginning state. It indicates HSRP is not running. It happens when the configuration changes or the interface is first turned on
Learn	The router has not determined the virtual IP address and has not yet seen an authenticated hello message from the active router. In this state, the router still waits to hear from the active router.
Listen	The router knows both IP and MAC address of the virtual router but it is not the active or standby router. For example, if there are 3 routers in HSRP group, the router which is not in active or standby state will remain in listen state.
Speak	The router sends periodic HSRP hellos and participates in the election of the active or standby router.
Standby	In this state, the router monitors hellos from the active router and it will take the active state when the current active router fails (no packets heard from active router)
Active	The router forwards packets that are sent to the HSRP group. The router also sends periodic hello messages

Please notice that not all routers in a HSRP group go through all states above. In a HSRP group, only one router reaches active state and one router reaches standby state. Other routers will stop at listen state.

QUESTION 110
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: AD
Explanation:

When you change the HSRP version, Cisco NX-OS reinitializes the group because it now has a new virtual MAC address. HSRP version 1 uses the MAC address range 0000.0C07.ACxx while HSRP version 2 uses the MAC address range 0000.0C9F.F0xx.
HSRP supports interface tracking which allows to specify another interface on the router for the HSRP process to monitor in order to alter the HSRP priority for a given group.

QUESTION 111
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?

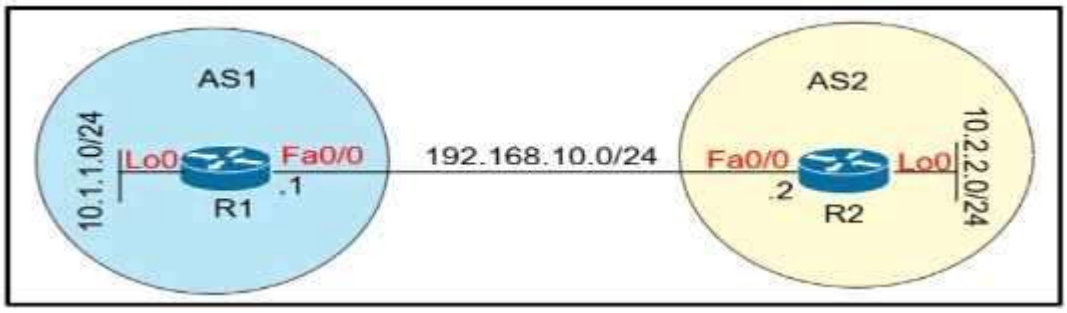

A. R1(config)#router bgp 1
R1(config-router)#neighbor 192.168.10.2 remote-as 2
R1(config-router)#network 10.1.1.0 mask 255.255.255.0
R2(config)#router bgp 2
R2(config-router)#neighbor 192.168.10.1 remote-as 1
R2(config-router)#network 10.2.2.0 mask 255.255.255.0

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
Explanation:
With BGP, we must advertise the correct network and subnet mask in the “network” command (in this case network 10.1.1.0/24 on R1 and network 10.2.2.0/24 on R2). BGP is very strict in the routing advertisements. In other words, BGP only advertises the network which exists exactly in the routing table. In this case, if you put the command “network x.x.0.0 mask 255.255.0.0” or “network x.0.0.0 mask 255.0.0.0” or “network x.x.x.x mask 255.255.255.255” then BGP will not advertise anything.
It is easy to establish eBGP neighborship via the direct link. But let’s see what are required when we want to establish eBGP neighborship via their loopback interfaces. We will need two commands:
+ The command “neighbor 10.1.1.1 ebgp-multihop 2” on R1 and “neighbor 10.2.2.2 ebgp-multihop 2” on R1. This command increases the TTL value to 2 so that BGP updates can reach the BGP neighbor which is two hops away.

+ A route to the neighbor loopback interface. For example: “ip route 10.2.2.0 255.255.255.0 192.168.10.2” on R1 and “ip route 10.1.1.0 255.255.255.0 192.168.10.1” on R2

QUESTION 112
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: DE
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-ntpm.html

QUESTION 113
Which standard access control entry permits from odd-numbered hosts in the 10.0.0.0/24 subnet?

A. Permit 10.0.0.0.0.0.1
B. Permit 10.0.0.1.0.0.0
C. Permit 10.0.0.1.0.0.254
D. Permit 10.0.0.0.255.255.255.254

Answer: C
Explanation:
Remember, for the wildcard mask, 1′s are I DON’T CARE, and 0′s are I CARE. So now let’s analyze a simple ACL:
access-list 1 permit 172.23.16.0 0.0.15.255
Two first octets are all 0’s meaning that we care about the network 172.23.x.x. The third octet of the wildcard mask, 15 (0000 1111 in binary), means that we care about first 4 bits but don’t care about last 4 bits so we allow the third octet in the form of 0001xxxx (minimum:00010000 = 16; maximum: 0001111 = 31).
The fourth octet is 255 (all 1 bits) that means I don’t care.
Therefore network 172.23.16.0 0.0.15.255 ranges from 172.23.16.0 to 172.23.31.255.
Now let’s consider the wildcard mask of 0.0.0.254 (four octet: 254 = 1111 1110) which means we only care the last bit.
Therefore if the last bit of the IP address is a “1” (0000 0001) then only odd numbers are allowed. If the last bit of the IP address is a “0” (0000 0000) then only even numbers are allowed.
Note: In binary, odd numbers are always end with a “1” while even numbers are always end with a “0”.
Therefore in this question, only the statement “permit 10.0.0.1 0.0.0.254” will allow all odd-numbered hosts in the 10.0.0.0/24 subnet.

QUESTION 114
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: AE
Explanation:
In this question, the inside local addresses of the 10.1.1.0/27 subnet are translated into 209.165.201.0/27 subnet. This is one-to-one NAT translation as the keyword “overload” is missing so in fact answer B is also correct.

QUESTION 115
Which statement about a fabric access point is true?

A. It is in local mode an 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 an must 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
Explanation:
Fabric mode APs continue to support the same wireless media services that traditional APs support; apply AVC, quality of service (QoS), and other wireless policies; and establish the CAPWAP control plane to the fabric WLC. Fabric APs join as local-mode APs and must be directly connected to the fabric edge node switch to enable fabric registration events, including RLOC assignment via the fabric WLC. The fabric edge nodes use CDP to recognize APs as special wired hosts, applying special port configurations and assigning the APs to a unique overlay network within a common EID space across a fabric. The assignment allows management simplification by using a single subnet to cover the AP infrastructure at a fabric site.
Reference: https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/sda-sdg-2019oct.html

QUESTION 116
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
Explanation:
The BGP session may report in the following states
1 - Idle: the initial state of a BGP connection. In this state, the BGP speaker is waiting for a BGP start event, generally either the establishment of a TCP connection or the re-establishment of a previous connection. Once the connection is established, BGP moves to the next state.
2 - Connect: In this state, BGP is waiting for the TCP connection to be formed. If the TCP connection completes, BGP will move to the OpenSent stage; if the connection cannot complete, BGP goes to Active
3 - Active: In the Active state, the BGP speaker is attempting to initiate a TCP session with the BGP speaker it wants to peer with. If this can be done, the BGP state goes to OpenSent state.
4 - OpenSent: the BGP speaker is waiting to receive an OPEN message from the remote BGP speaker
5 - OpenConfirm: Once the BGP speaker receives the OPEN message and no error is detected, the BGP speaker sends a KEEPALIVE message to the remote BGP speaker
6 - Established: All of the neighbor negotiations are complete. You will see a number, which tells us the number of prefixes the router has received from a neighbor or peer group.