351-001 Q&As

CCIE Routing and Switching Written

Pass Cisco 351-001 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/351-001.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Cisco Official Exam Center

QUESTION 1

Refer to the exhibit.

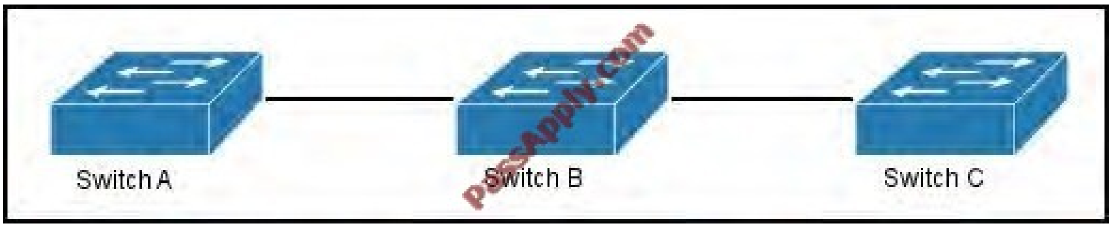

All switches are Cisco switches. Assume that Cisco Discovery Protocol is enabled only on switches A and C. Which information is returned when you issue the command show cdp neighbors on switch C?

A. a limited amount of information about switch B

B. no neighbor details will be returned

C. neighbor details for switch B

D. neighbor details for switch A

E. neighbor details for switch C

Correct Answer: B

CDP is used to discover information on directly connected neighbors only, so in this case SwitchC would only be able to obtain CDP information from SwitchB. However, since SwitchB is not running CDP then no neighbor information will be seen on SwitchC. Same goes for Switch A also in this topology.

QUESTION 2

Which three components comprise the structure of a pseudowire FEC element? (Choose three.)

A. pseudowire ID

B. pseudowire type

C. control word

D. Layer 3 PDU

E. header checksum

F. type of service

Correct Answer: ABC

The Pseudowire ID FEC element has the following components: Pseudowire ID FEC -- The first octet has a value of 128 that identifies it as a Pseudowire ID FEC element. Control Word Bit (C-Bit) -- The C-bit indicates whether the advertising PE expects the control word to be present for pseudowire packets. A control word is an optional 4-byte field located between the MPLS label stack and the Layer 2 payload in the pseudowire packet. The control word carries generic and

Layer 2 payload-specific information. If the C-bit is set to 1, the advertising PE expects the control word to be present in every pseudowire packet on the pseudowire that is being signaled. If the C-bit is set to 0, no control word is expected to be present. Pseudowire Type -- PW Type is a 15-bit field that represents the type of pseudowire. Examples of pseudowire types are shown in Table 6-1. Pseudowire Information Length -- Pseudowire Information Length is the length of the Pseudowire ID field and the interface parameters in octets. When the length is set to 0, this FEC element stands for all pseudowires using the specified Group ID. The Pseudowire ID and Interface Parameters fields are not present. Group ID -- The Group ID field is a 32-bit arbitrary value that is assigned to a group of pseudowires. Pseudowire ID -- The Pseudowire ID, also known as VC ID, is a non-zero, 32-bit identifier that distinguishes one pseudowire from another. To connect two attachment circuits through a pseudowire, you need to associate each one with the same Pseudowire ID. Interface Parameters -- The variable-length Interface Parameters field provides attachment circuit-specific information, such as interface MTU, maximum number of concatenated ATM cells, interface description, and so on.

Reference: http://www.ciscopress.com/articles/article.asp?p=386788andseqNum=2

QUESTION 3

Refer to the exhibit.

```
class-map match-any CM-EXAMPLE-1
   match dscp AF11
   match dscp AF21
class-map match-any CM-EXAMPLE-2
   match access-group 100
   match dscp EF
class-map match-all CM-EXAMPLE-3
   match dscp [illegible]
   match dscp [illegible]
class-map match-all CM-EXAMPLE-4
   match ip rtp 16384 16383
   match precedence 5
class-map match-any CM-EXAMPLE-5
   match dscp AF41 EF
```

If you apply this configuration to a device on your network, which class map cannot match traffic?

A. CM-EXAMPLE-3

B. CM- EXAMPLE-1

C. CM-EXAMPLE-4

D. CM-EXAMPLE-5

E. CM-EXAMPLE-2

Correct Answer: A

QUESTION 4

Which two statements are true about a 6to4 tunnel connecting two IPv6 islands over the IPv4 Internet? (Choose two.)

A. It embeds the IPv6 packet into the IPv4 payload with the protocol type set to 51.

B. It works by appending the private IPv4 address (converted into hexadecimal format) to the 2002::/16 prefix.

C. It embeds the IPv6 packet into the IPv4 payload with the protocol type set to 41.

D. It works by appending the public IPv4 address (converted into hexadecimal format) to the 2002::/16 prefix.

Correct Answer: CD

6to4 embeds an IPv6 packet in the payload portion of an IPv4 packet with protocol type 41. To send an IPv6 packet over an IPv4 network to a 6to4 destination address, an IPv4 header with protocol type 41 is prepended to the IPv6 packet. The IPv4 destination address for the prepended packet header is derived from the IPv6 destination address of the inner packet (which is in the format of a 6to4 address), by extracting the 32 bits immediately following the IPv6 destination address\\'s 2002::/16 prefix. The IPv4 source address in the prepended packet header is the IPv4 address of the host or router which is sending the packet over IPv4. The resulting IPv4 packet is then routed to its IPv4 destination address just like any other IPv4 packet.

Reference: http://en.wikipedia.org/wiki/6to4

QUESTION 5

Which statement about shaped round robin queuing is true?

A. Queues with higher configured weights are serviced first.

B. The device waits a period of time, set by the configured weight, before servicing the next queue.

C. The device services a single queue completely before moving on to the next queue.

D. Shaped mode is available on both the ingress and egress queues.

Correct Answer: A

SRR is scheduling service for specifying the rate at which packets are dequeued. With SRR there are two modes, shaped and shared. Shaped mode is only available on the egress queues SRR differs from typical WRR. With WRR queues are serviced based on the weight. Q1 is serviced for weight 1 period of time, Q2 is served for weight 2 period of time, and so forth. The servicing mechanism works by moving from queue to queue and services them for the weighted amount of time. With SRR weights are still followed; however, SRR services Q1, moves to Q2, then Q3 and Q4 in a different way. It does not wait at and service each queue for a weighted amount of time before moving on to the next queue. Instead, SRR makes several rapid passes at the queues; in each pass, each queue might or might not be serviced. For each given pass, the more highly weighted queues are more likely to be serviced than the lower priority queues.

Reference: http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-3560-e- series-switches/prod_qas0900aecd805bacc7.html

---

QUESTION 6

Refer to the exhibit.

```
ip route 10.0.0.0 255.255.255.0 192.168.192.9
ip default-network 172.16.199.9
```

The device with this configuration is unable to reach network 172.31.31.0/24. The next hop router has been verified to have full connectivity to the network. Which two actions can you take to establish connectivity to the network? (Choose two.)

A. Create a static route to 172.16.199.0 using the address of the next hop router.

B. Create a default route to the link address of the next hop router.

C. Create a static route to the loopback address of the next hop router.

D. Create a default route to 172.16.199.9.

E. Modify the existing static route so that the next hop is 0.0.0.0.

F. Replace the ip default-network command with the ip default-gateway command.

Correct Answer: AB

Unlike the ip default-gateway command, you can use ip default-network when ip routing is enabled on the Cisco router. When you configure ip default-network the router considers routes to that network for installation as the gateway of last resort on the router.

For every network configured with ip default-network, if a router has a route to that network, that route is flagged as a candidate default route. However, in this case if the router does not a route to the drfault network of 172.16.199.9, then you would need to ensure that this route exisits by creating a static route to 172.16.199.0 using the address of the next hop router, or simply create a default route using the address of the next hop router.

QUESTION 7

What is a reason for 6PE to use two MPLS labels in the data plane instead of one?

A. 6PE allows penultimate hop popping and has a requirement that all P routers do not have to be IPv6 aware.

B. 6PE does not allow penultimate hop popping.

C. It allows MPLS traffic engineering to work in a 6PE network.

D. It allows 6PE to work in an MPLS network where 6VPE is also deployed.

Correct Answer: A

Q. Why does 6PE use two MPLS labels in the data plane? A. 6PE uses two labels: The top label is the transport label, which is assigned hop-by-hop by the Label Distribution Protocol (LDP) or by MPLS traffic engineering (TE). The bottom label is the label assigned by the Border Gateway Protocol (BGP) and advertised by the internal BGP (iBGP) between the Provider Edge (PE) routers. When the 6PE was released, a main requirement was that none of the MPLS core routers (the P routers) had to be IPv6-aware. That requirement drove the need for two labels in the data plane. There are two reasons why the 6PE needs both labels. PHP Functionality If only the transport label were used, and if penultimate hop popping (PHP) were used, the penultimate hop router (the P router) would need to understand IPv6. With PHP, this penultimate hop router would need to remove the MPLS label and forward the packet as an IPv6 packet.

This P router would need to know that the packet is IPv6 because the P router would need to use the correct Layer 2 encapsulation type for IPv6. (The encapsulation type is different for IPv6 and IPv4; for example, for Ethernet, the encapsulation type is 0x86DD for IPv6, while it is 0x0800 for IPv4.) If the penultimate hop router is not IPv6-capable, it would likely put the Layer 2 encapsulation type for IPv4 for the IPv6 packet. The egress PE router would then believe that the packet was IPv4. There is time-to-live (TTL) processing in both the IPv4 and IPv6 headers. In IPv6, the field is called Hop Limit. The IPv4 and IPv6 fields are at different locations in the headers. Also, the Header Checksum in the IPv4 header would also need to be changed; there is no Header Checksum field in IPv6. If the penultimate hop router is not IPv6- capable, it would cause the IPv6 packet to be malformed since the router expects to find the TTL field and Header Checksum field in the header. Because of these differences, the penultimate hop router would need to know it is an IPv6 packet. How would this router know that the packet is an IPv6 packet, since it did not assign a label to the IPv6 Forwarding Equivalence Class (FEC), and there is no encapsulation field in the MPLS header? It could scan for the first nibble after the label stack and determine that the packet is IPv6 if the value is 6. However, that implies that the penultimate hop router needs to be IPv6-capable. This scenario could work if the explicit null label is used (hence no PHP). However, the decision was to require PHP. Load Balancing Typical load balancing on a P router follows this process. The P router goes to the end of the label stack and determines if it is an IPv4 packet by looking at the first nibble after the label stack. If the nibble has a value of 4, the MPLS payload is an IPv4 packet, and the P router load balances by hashing the source and destination IPv4 addresses. If the P router is IPv6-capable and the value of the nibble is 6, the P router load balances by hashing the source and destination IPv6 addresses. If the P router is not IPv6-capable and the value of the nibble is not 4 (it could be 6 if the packet is an IPv6 packet), the P router determines it is not an IPv4 packet and makes the load balancing decision based on the bottom label. In the 6PE scenario, imagine there are two egress PE routers advertising one IPv6 prefix in BGP towards the ingress PE router. This IPv6 prefix would be advertised with two different labels in BGP. Hence, in the data plane, the bottom label would be either of the two labels. This would allow a P router to load balance on the bottom label on a per-flow basis. If 6PE used only the transport label to transport the 6PE packets through the MPLS core, the P routers would not be able to load balance these packets on a per-flow basis unless the P routers were IPv6-capable. If the P routers were IPv6-capable, they could use the source and destination IPv6 addresses in order to make a load balancing decision.

Reference: http://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching- mpls/mpls/116061-qa-6pe-00.html

QUESTION 8

Which statement is true about trunking?

A. Cisco switches that run PVST+ do not transmit BPDUs on nonnative VLANs when using a dot1q trunk.

B. When removing VLAN 1 from a trunk, management traffic such as CDP is no longer passed in that VLAN.

C. DTP only supports autonegotiation on 802.1q and does not support autonegotiation for ISL.

D. DTP is a point-to-point protocol.

Correct Answer: D

Ethernet trunk interfaces support different trunking modes. You can set an interface as trunking or nontrunking or to negotiate trunking with the neighboring interface. To autonegotiate trunking, the interfaces must be in the same VTP domain.

Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which is a Point-to- Point Protocol. However, some internetworking devices might forward DTP frames improperly, which could cause misconfigurations.

Reference:

http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/12-2_55_se/configuration/guide/scg3750/swvlan.html

QUESTION 9

Refer to the exhibit.

```
ip route vrf red 0.0.0.0 0.0.0.0 192.168.1.1 global
```

Which three statements about this configuration are true? (Choose three.)

A. The default route appears in the global routing table.

B. The static route appears in the VRF red routing table.

C. The subnet 192.168.1.0 is unique to the VRF red routing table.

D. The static route is added to the global routing table and leaked from the VRF red.

E. The subnet 192.168.1.0 is unique to the global routing table.

F. 192.168.1.1 is reachable using any of the addresses on the router where the static route is configured.

Correct Answer: ABE

This is an example of the route leaking feature. Here, this static route is created for the red VRF so it will be installed into the red VRF routing table, but the use of the global keyword will cause this default route to appear in the global routing table.

QUESTION 10

When BGP route reflectors are used, which attribute ensures that a routing loop is not created?

A. weight

B. local preference

C. multiexit discriminator

D. originator ID

Correct Answer: D

As the iBGP learned routes are reflected, routing information may loop. The route reflector model has the following mechanisms to avoid routing loops:

Originator ID is an optional, nontransitive BGP attribute. It is a 4-byte attributed created by a route reflector. The attribute carries the router ID of the originator of the route in the local autonomous system. Therefore, if a misconfiguration causes routing information to come

back to the originator, the information is ignored. ? Cluster-list is an optional, nontransitive BGP attribute. It is a sequence of cluster IDs that the route has passed. When a route reflector reflects a route from its clients to nonclient peers, and vice versa, it appends the local cluster ID to the cluster-list. If the cluster-list is empty, a new cluster-list is created. Using this attribute, a route reflector can identify if routing information is looped back to the same cluster due to misconfiguration. If the local cluster ID is found in the cluster-list, the advertisement is ignored.

Reference: http://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfbgp.html

QUESTION 11

Which three conditions can cause excessive unicast flooding? (Choose three.)

A. Asymmetric routing

B. Repeated TCNs

C. The use of HSRP

D. Frames sent to FFFF.FFFF.FFFF

E. MAC forwarding table overflow

F. The use of Unicast Reverse Path Forwarding

Correct Answer: ABE

Causes of Flooding The very cause of flooding is that destination MAC address of the packet is not in the L2 forwarding table of the switch. In this case the packet will be flooded out of all forwarding ports in its VLAN (except the port it was received on). Below case studies display most common reasons for destination MAC address not being known to the switch. Cause 1: Asymmetric Routing Large amounts of flooded traffic might saturate low-bandwidth links causing network performance issues or complete connectivity outage to devices connected across such low-bandwidth links. Cause 2: Spanning-Tree Protocol Topology Changes Another common issue caused by flooding is Spanning-Tree Protocol (STP) Topology Change Notification (TCN). TCN is designed to correct forwarding tables after the forwarding topology has changed. This is necessary to avoid a connectivity outage, as after a topology change some destinations previously accessible via particular ports might become accessible via different ports. TCN operates by shortening the forwarding table aging time, such that if the address is not relearned, it will age out and flooding will occur. TCNs are triggered by a port that is transitioning to or from the forwarding state. After the TCN, even if the particular destination MAC address has aged out, flooding should not happen for long in most cases since the address will be relearned. The issue might arise when TCNs are occurring repeatedly with short intervals. The switches will constantly be fast-aging their forwarding tables so flooding will be nearly constant. Normally, a TCN is rare in a well-configured network. When the port on a switch goes up or down, there is eventually a TCN once the STP state of the port is changing to or from forwarding. When the port is flapping, repetitive TCNs and flooding occurs. Cause 3: Forwarding Table Overflow Another possible cause of flooding can be overflow of the switch forwarding table. In this case, new addresses cannot be learned and packets destined to such addresses are flooded until some space becomes available in the forwarding table. New addresses will then be learned. This is possible but rare, since most modern switches have large enough forwarding tables to accommodate MAC addresses for most designs. Forwarding table exhaustion can also be caused by an attack on the network where one host starts generating frames each sourced with different MAC address. This will tie up all the forwarding table resources. Once the forwarding tables become saturated, other traffic will be flooded because new learning cannot occur. This kind of attack can be detected by examining the switch forwarding table. Most of the MAC addresses will point to the same port or group of ports. Such attacks can be prevented by limiting the number of MAC addresses learned on untrusted ports by using the port security feature.

Reference: http://www.cisco.com/c/en/us/support/docs/switches/catalyst-6000-series- switches/23563-143.html#causes

QUESTION 12

Which two statements about IBGP multipath are true? (Choose two.)

A. The IGP metric of the BGP next hop can be different from the best-path IGP metric if you configure the router for unequal-cost IBGP multipath.

B. The IGP metric of the BGP next hop must be the same as the best-path IGP metric.

C. The equivalent next-hop-self is performed on the best path from among the IBGP multipaths before it is forwarded to external peers.

D. The path should be learned from an external neighbor.

E. The router BGP process must learn the path from a confederation-external or external neighbor.

F. The router BGP process must learn the path from an internal neighbor.

Correct Answer: AF