350-001

Exam: 350-001

Title: CCIE Routing and Switching Written Exam

Version: Demo

350-001

QUESTION NO: 1
What does the root guard feature provide in a bridged network?

A. it ensures that the bridge is elected as Root Bridge in the network.
B. it enforces the root bridge placement in the network
C. It ensures that BPDUs sent by the root bridge are forwarded in a timely manner.
D. It ensures that all ports receiving BPDUs from the root bridge are in the forwarding state.

Answer: B
Explanation:
Root Guard—Enabled per port; ignores any received superior BPDUs to prevent a switch connected to this port from becoming root. Upon receipt of superior BPDUs, this switch puts the port in a loop-inconsistent state, ceasing forwarding and receiving frames until the superior BPDUs cease.
The STP topology can be changed based on one of these unexpected and undesired switches being added to the network. For instance, this newly added and unexpected switch might have the lowest bridge ID and become the root. To prevent such problems, BPDU Guard and Root Guard can be enabled on these access ports to monitor for incoming BPDUs.

QUESTION NO: 2
Which two of these statements about WCCP version 2 are false? (Choose two.)

A. It allows for the redirection of traffic other than HTTP, including a variety of UDP and TCP traffic.
B. Only one router can redirect content requests.
C. Multiple routers can redirect content requests.
D. It works only with IP networks.
E. The Cache Engine defines one central "home router" and stores it in its memory.
F. The Cache Engine defines one central "home router," and stores it in its memory.

Answer: B, F
Explanation:
WCCP transparently redirects Hypertext Transfer Protocol (HTTP) requests going to the intended server to a Cache Engine. End users do not know that the page came from the Cache Engine rather than the originally requested web server.
WCCP Version 2 now contains the following new features:
- Multiple router support
- Improved security
- Faster throughput
- Redirection of multiple TCP port-destined traffic
- Load distributing applications capability
- Client IP addressing transparency

Multirouter Support:
WCCP Version 2 enables a series of Cache Engines, called a *Cache Engine cluster*, to connect to multiple routers. This feature provides redundancy and a more distributed architecture for instances when a Cache Engine needs to connect to a large number of interfaces. This strategy

also has the benefit of keeping all the Cache Engines in a single cluster, avoiding unnecessary duplication of web pages across several clusters.
Reference:
http://www.cisco.com/en/US/products/sw/conntsw/ps547/products_user_guide_chapter09186a008009f1ae.html

QUESTION NO: 3
According to the exhibit provided, what will be the purpose of this route map when applied to traffic passing through a router?

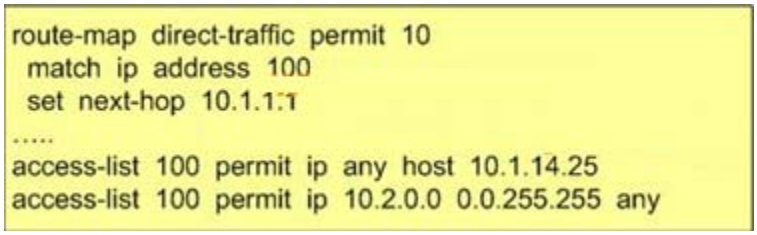

A. take any packet sourced from any address in the 10.2.0.0/16 network or destined to 10.1.14.25 and set the next hop to 10.1.1.1
B. nothing; extended access lists are not allowed in route maps used for policy-based routing
C. take any packet sourced from any address in the 10.2.0.0/16 network and destined to 10.1.14.25 and set the next hop to 10.1.1.1
D. drop any packet sourced from 10.2.0.0/16

Answer: A
Explanation:
In this configuration example, any traffic matching access list 100 will have their next hop set to 10.1.1.1 overriding the normal behavior of the routing table. Access list 100 has two entries, so any traffic matching either will be policy routed.

QUESTION NO: 4
Which two benefits are of applying WRED? (Choose two.)

A. helps to avoid TCP synchronization
B. allows a different drop profile to be manually enabled for each IP precedence or DSCP
C. provides minimal bandwidth guarantees
D. provides bounded low latency

Answer: A,B

Explanation:
WRED and distributed WRED (DWRED)—both of which are the Cisco implementations of RED—combine the capabilities of the RED algorithm with the IP Precedence feature. Within the section on WRED, the following related features are discussed:

– Flow-based WRED. Flow-based WRED extends WRED to provide greater fairness to all flows on an interface in regard to how packets are dropped.
– DiffServ Compliant WRED-DiffServ Compliant WRED extends WRED to support Differentiated Services (DiffServ) and Assured Forwarding (AF) Per Hop Behavior (PHB). This feature enables customers to implement AF PHB by coloring packets according to differentiated services code point (DSCP) values and then assigning preferential drop probabilities to those packets.
WRED avoids the globalization problems that occur when tail drop is used as the congestion avoidance mechanism on the router. Global TCP synchronization occurs as waves of congestion crest only to be followed by troughs during which the transmission link is not fully utilized. Global synchronization of TCP hosts, for example, can occur because packets are dropped all at once. Global synchronization manifests when multiple TCP hosts reduce their transmission rates in response to packet dropping, then increase their transmission rates once again when the congestion is reduced.

Reference:
http://www.google.com/search?hl=en&q=WRED%2C+A+different+drop+profile+can+be+manually+enabled+per+IP+Precedence+or+DSCP

QUESTION NO: 5
Refer to the exhibit. The output of the show interface command for the link between R2 and R5 in this EIGRP network shows that the link load varies between 10 and 35. What K value setting could you use to ensure that this link is not used by EIGRP when the link load reaches 35, but can be used again when the link load drops below 20?

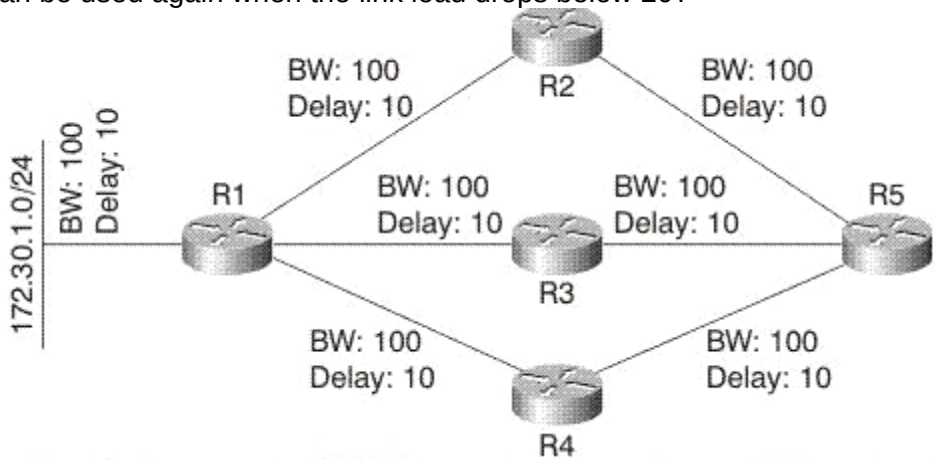

A. Link load is not read in real-time, so there is no way to set the K values to make EIGRP choose to use or not use a link based on the link load.
B. Use the K5 setting to include load in EIGRP's metric calculations.
C. Use the K2 setting to include load in EIGRP's metric calculations.
D. There is not enough information in the question to determine the correct answer.

Answer: A
Explanation:
EIGRP computes its composite metric from five parameters, one of them being interface load, therefore raising the theoretical possibility of having route metrics that include interface load.

However, tweaking EIGRP K-values with the “metric weights” command to include interface load in metric calculations is highly discouraged - every change in interface load could lead to network instability. Even worse, whenever an interface load would increase, the increased composite metric of the affected routes in EIGRP topology table would cause them to enter *active* state (and the router to start the DUAL algorithm trying to find more optimum paths toward the destination).
To make the whole idea even more impractical, EIGRP does not scan the interface load (and other parameters influencing the metric) on periodical basis, but only when triggered by a change in network topology (for example, interface or neighbor up/down even).

QUESTION NO: 6
NBAR is used to provide which QoS function?

A. classification
B. policing
C. CBWFQ bandwidth guarantees
D. shaping

Answer: A

Explanation:
NBAR addresses IP QoS classification requirements by classifying application-level protocols so that QoS policies can be applied to the classified traffic. NBAR addresses the ongoing need to extend the classification engine for the many existing and emerging application protocols by providing an extensible Packet Description Language (PDL). NBAR can determine which protocols and applications are currently running on a network so that an appropriate QoS policy can be created based upon the current traffic mix and application requirements.
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800c75d1.html#54116

QUESTION NO: 7
You work as a network technician at Pas4sure.com, study the exhibit carefully. Upon examining the EIGRP topology table, you see that ROUTER1 has routes to 10.1.3.0/24 and 10.1.4.0/24, but not to 192.168.2.0/24. ROUTER3 has routes to 10.1.1.0/24 and 10.1.2.0/24, but not to 192.168.1.0/24. Which would most likely cause this problem?

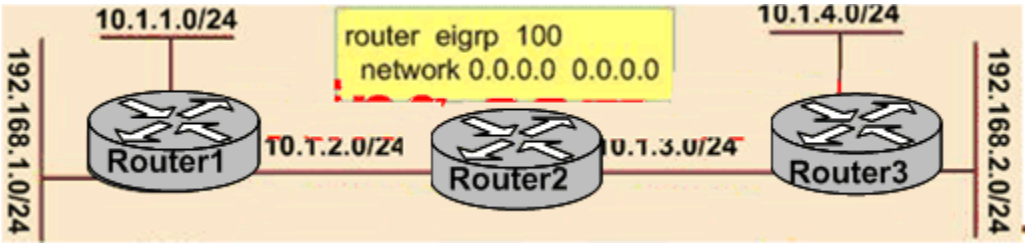

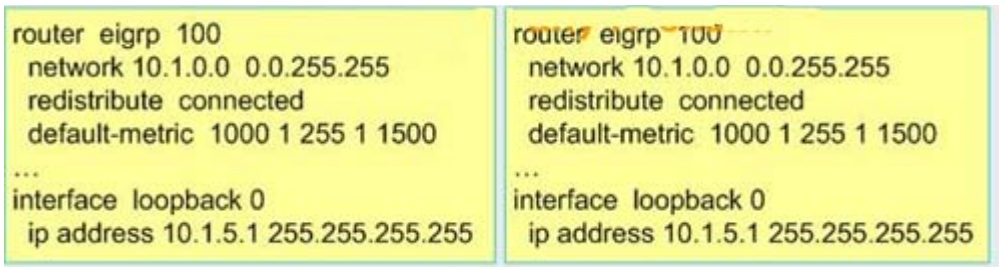

A. ROUTER2 is most likely filtering EIGRP externals, but you cannot be certain without examining its configuration
B. ROUTER1 and ROUTER3 have the same router ID, so they will reject each other's redistributed (external) EIGRP routes.
C. Autosummarization is removing the routes to 192.168.1.0/24 and 192.168.2.0/24. ROUTER1 and ROUTER3 should have routes to 192.168.0.0/16 instead.
D. The redistribution at ROUTER1 and ROUTER3 is configured incorrectly.

Answer: B
Explanation:
Many times, EIGRP will not install routes because of a duplicate router ID problem. EIGRP does not use router ID as extensively as OSPF. EIGRP uses the notion of router ID only on external routes to prevent loops. EIGRP chooses the router ID based on the highest IP address of the loopback interfaces on the router. If the router doesn't have any loopback interfaces, the highest active IP address of all the interfaces is chosen as the router ID for EIGRP. In this case, the loopback addresses are both 10.1.5.1 so the redistributed routes will be rejected as Router1 and Router3 will assume that there is a routing loop.
Reference: Troubleshooting EIGRP by Zaheer Aziz, Johnson Lui, Abe Martey, Faraz Shamim, Cisco Press.

QUESTION NO: 8
Which of these potential issues is eliminated by the use of split horizon?

A. asymmetric routing throughout the network
B. packet forwarding loops
C. joined horizons
D. Cisco Express Forwarding load-balancing inconsistency

Answer: B
Explanation:
Distance-vector routing protocols employ the split horizon rule which prohibits a router from advertising a route back out the interface from which it was learned. Split horizon is one of the methods used to prevent routing loops due to the slow convergence times of distance-vector routing protocols.

QUESTION NO: 9
The 802.1w protocol is seen as the next evolution beyond the 802.1 D standard protocol. Which of these statements regarding port states is true of both 802.1 D and 802.1w?

A. All 802.1 D port states (Disabled, Blocking, Listening, Learning, and Forwarding) are identical in 802.1w.
B. The 802.1 D port states Disabled and Blocking have become the 802.1w port state Discarding, and all other 802.1D port states remain the same in 802.1w.
C. The 802.1 D port states Disabled, Blocking, and Listening have become the 802.1w port state Discarding, and all other 802.1D port states remain the same in 802.1w.
D. The 802.1 D port states Disabled, Blocking, and Listening have been removed completely from 802.1w (there is no corresponding port state), and all other 802.1 D port states remain the same in 802.1w.) E. The 802.1 D port state Disabled has been removed from 802.1w, and the 802.1 D port states Blocking and Listening have become the 802.1w port state Discarding; all other 802.1D port states remain the same in 802.1w.

Answer: C

Explanation:

Table 7-3 Port State Comparison

Operational Status	STP Port State	RSTP Port State	Is Port Included in the Active Topology?
Enabled	Blocking	Discarding	No
Enabled	Listening	Discarding	No
Enabled	Learning	Learning	Yes
Enabled	Forwarding	Forwarding	Yes
Disabled	Disabled	Discarding	No

QUESTION NO: 10
When a router makes a forwarding decision, which of these routes in the routing table always wins?

A. administrative distance
B. router ID
C. longest prefix match
D. routing process ID

Answer: C

Explanation:
Making a forwarding decision actually consists of three sets of processes: the routing protocols, the routing table, and the actual process, which makes a forwarding decision and switches packets. These three sets of processes are illustrated, along with their relationship, below:

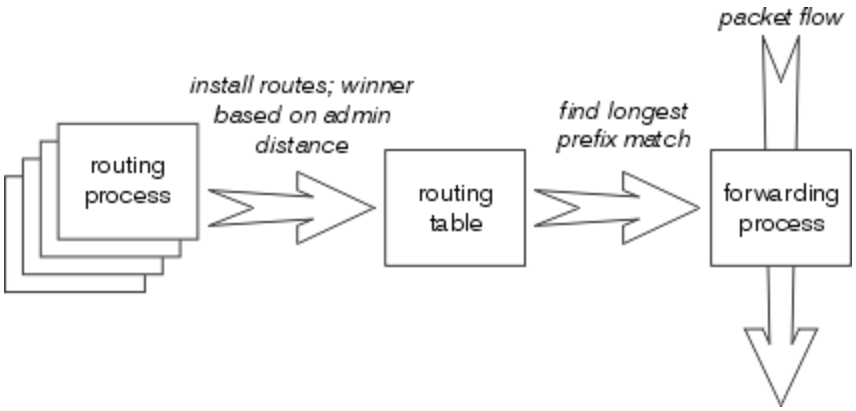

The longest prefix match always wins among the routes actually installed in the routing table, while the routing protocol with the lowest administrative distance always wins when installing routes into the routing table.
Reference:
www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094823.shtml

QUESTION NO: 11
On what type of ports would STP Port Fast BPDU guard be most appropriate?

A. root ports
B. Designated ports
C. Host ports
D. alternate ports

Answer: C

QUESTION NO: 12
Which of these statements best describes how neighbor adjacencies are formed in a multi-access OSPF network?

A. The router with the highest priority will become the DR
B. Only those routers with the Cisco default priority of 0 are eligible to become the DR or BDR.
C. The router with the highest loop-back address will become the DR if two or more routers have the same priority.
D. The router with the lowest Router ID will become the DR and the router with the next lowest Router ID will become the BDR.
E. Election of the DR and BDR begins only after a router that wants to become either the DR or BDR enters the ExStart state.

Answer: A
Explanation:
The router with the highest priority is elected the DR on a multiaccess network. A router with a priority of 0 is ineligible to become a DR or BDR. In the event of a tie in priority, the router with the highest router ID is elected the DR. If no router ID has been manually configured on a router, the router uses its numerically highest loopback address as its router ID. If no loopback

interfaces have been configured, the router uses its numerically highest IP address of any physical interface.

QUESTION NO: 13
Which two fundamental modifications, related to traffic forwarding, does MPLS introduce? (Choose two.)

A. IP lookup is performed on every hop within the MPLS core.
B. IP destination routing is reduced to label lookup within the MPLS network.
C. For unicast routing, labels are assigned to FECs (in other words, IP prefixes).
D. For multicast routing, labels are assigned to IP multicast groups.

Answer: B, C

Explanation:
MPLS works by tagging packets with an identifier (a label) to distinguish the LSPs. When a packet is received, the router uses this label (and sometimes also the link over which it was received) to identify the LSP. It then looks up the LSP in its own forwarding table to determine the best link over which to forward the packet, and the label to use on this next hop.
A different label is used for each hop, and it is chosen by the router or switch performing the forwarding operation. This allows the use of very fast and simple forwarding engines, as the router can select the label to minimize processing.
Ingress routers at the edge of the MPLS network use the packet's destination address to determine which LSP to use. Inside the network, the MPLS routers use only the LSP labels to forward the packet to the egress router.

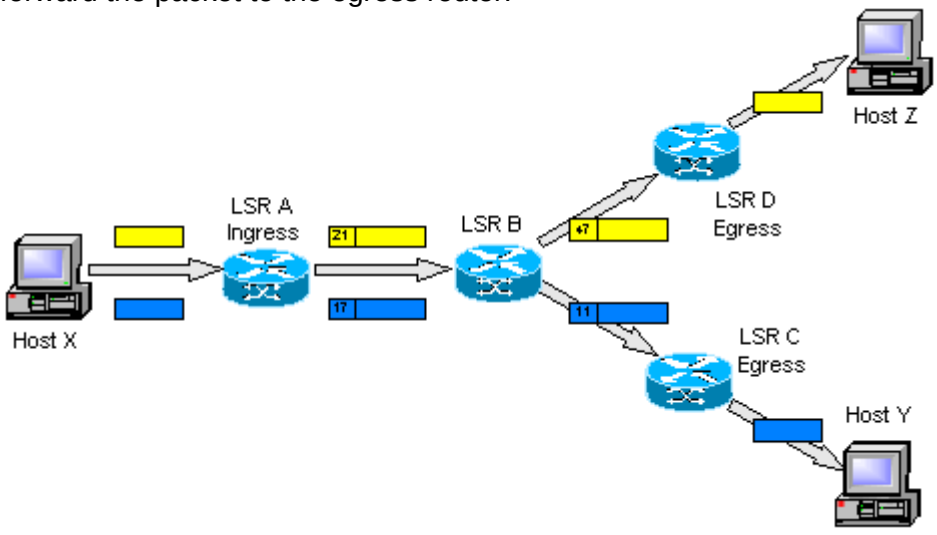

In the diagram above, LSR (Label Switched Router) A uses the destination IP address on each packet to select the LSP, which determines the next hop and initial label for each packet (21 and 17). When LSR B receives the packets, it uses these labels to identify the LSPs, from which it determines the next hops (LSRs D and C) and labels (47 and 11). The egress routers (LSRs D and C) strip off the final label and route the packet out of the network.
As MPLS uses only the label to forward packets, it is protocol-independent, hence the term "Multi-Protocol" in MPLS. Packet forwarding has been defined for all types of layer-2 link technologies, with a different label encoding used in each case.

QUESTION NO: 14
You are using IPv6, and would like to configure EIGRPv3. Which three of these correctly describe how you can perform this configuration? (Choose three.)

A. EIGRP for IPv6 is directly configured on the interfaces over which it runs.
B. EIGRP for IPv6 is not configured on the interfaces over which it runs, but if a user uses passive-interface configuration, EIGRP for IPv6 needs to be configured on the interface that is made passive.
C. There is a network statement configuration in EIGRP for IPv6, the same as for IPv4.
D. There is no network statement configuration in EIGRP for IPv6.
E. When a user uses a passive-interface configuration, EIGRP for IPv6 does not need to be configured on the interface that is made passive.
F. When a user uses a non-passive-interface configuration, EIGRP for IPv6 does not need to be configured on the interface that is made passive

Answer: A, D, E

Explanation:
Restrictions for Implementing EIGRP for IPv6:
This section lists ways in which EIGRP for IPv6 differs from EIGRP IPv4 as well as EIGRP for IPv6 restrictions.
- EIGRP for IPv6 is directly configured on the interfaces over which it runs. This feature allows EIGRP for IPv6 to be configured without the use of a global IPv6 address. There is no network statement in EIGRP for IPv6.

In per-interface configuration at system startup, if EIGRP has been configured on an interface, then the EIGRP protocol may start running before any EIGRP router mode commands have been executed.
- An EIGRP for IPv6 protocol instance requires a router ID before it can start running.
- EIGRP for IPv6 has a shutdown feature. The routing process should be in "no shutdown" mode in order to start running.
- When a user uses passive-interface configuration, EIGRP for IPv6 does not need to be configured on the interface that is made passive.
- EIGRP for IPv6 provides route filtering using the distribute-list prefix-list command. Use of the routE.map command is not supported for route filtering with a distribute list.

Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00805fc867.html

QUESTION NO: 15
Refer to the exhibit. Which switching feature is being tested?

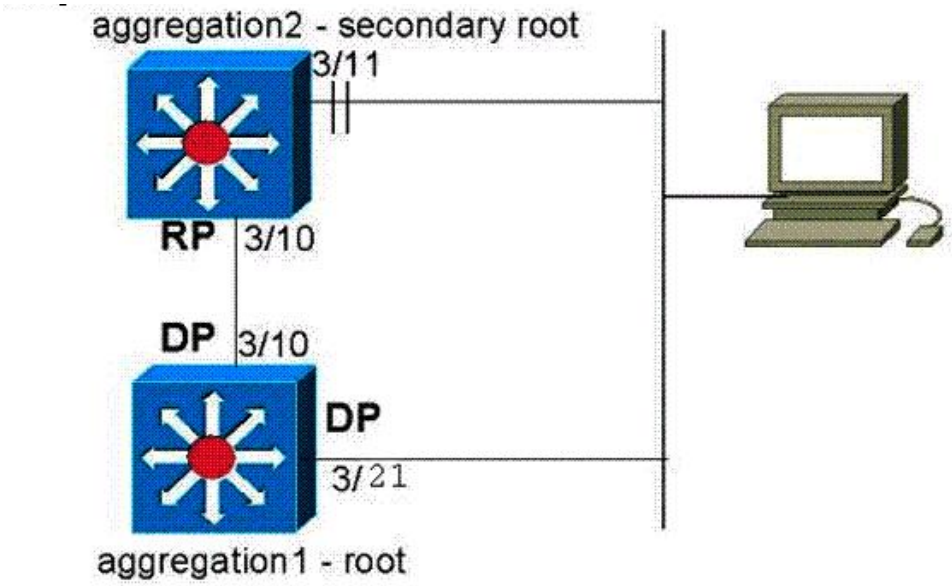

aggregation-2 (enable) set spantree portfast 3/11 ena
Warning: Spantree port fast start should only be enabled on ports connected to a single host.
Connecting hubs, concentrators, switches, bridges, etc. to a fast start port can cause temporary spanning tree loops. Use with caution.
Spantree port 3/11 fast start enabled.
aggregation-2 (enable) set spantree portfast bpdu-filter ena
Spantree portfast bpdu-filter enabled on this switch.
2001 Feb 06 13:32:14 %SPANTREE-4-LOOPGUARDBLOCK: No BPDUs were received on port 3/21 in VLAN 99. Moved to loop-inconsistent state

A. Loop guard
B. Port Fast
C. root guard
D. BDPU guard

Answer: A

Explanation:
Loop guard checks if a root port or an alternate/backup root port receives BPDUs. If the port does not receive BPDUs, loop guard puts the port into an inconsistent state (blocking) until it starts to receive BPDUs again. A port in the inconsistent state does not transmit BPDUs. If such a port receives BPDUs again, the port (and link) is deemed viable again. The loop-inconsistent condition is removed from the port, and STP determines the port state. In this way, recovery is automatic.
Loop guard isolates the failure and lets spanning tree converge to a stable topology without the failed link or bridge. Loop guard prevents STP loops with the speed of the STP version that is in use. There is no dependency on STP itself (802.1D or 802.1w) or when tuning the STP timers. For these reasons, Cisco recommends that you implement loop guard in conjunction with UDLD in topologies that rely on STP and where the software supports the features.
When loop guard blocks an inconsistent port, this message is logged:

%SPANTREE-SP-2-LOOPGUARD_BLOCK: Loop guard blocking port GigabitEthernet2/1 on VLAN0010

Reference:
http://www.cisco.com/en/US/products/hw/switches/ps700/products_white_paper09186a00801b49a4.shtml

QUESTION NO: 17
A new router has been allocated a single /24 subnet (172.16.123.0/24). The interface between this new router and the upstream router has already been configured from a different IP subnet. The four other interfaces on this router require 56,10, 72, and 24 IP addresses, respectively. The router always uses the first IP address on any subnet.
Which one of these combinations of IP addresses allow the router to meet the interface requirements?

Which of these combinations of IP addresses allow the router to meet the interface requirements?

A. 172.16.123.1 255.255.255.128
172.16.123.129 255.255.255.192
172.16.123.193 255.255.255.224
172.16.123.225 255.255.255.240

B. 172.16.123.1 255.255.255.192
172.16.123.65 255.255.255.192
172.16.123.129 255.255.255.192
172.16.123.193 255.255.255.192

C. 172.16.123.1 255.255.255.128
172.16.123.129 255.255.255.192
172.16.123.193 255.255.255.224
172.16.123.225 255.255.255.248

D. 172.16.123.1 255.255.255.128
172.16.123.129 255.255.255.224
172.16.123.161 255.255.255.224
172.16.123.193 255.255.255.224

E. 172.16.123.1 255.255.255.128
172.16.123.129 255.255.255.192
172.16.123.193 255.255.255.240
172.16.123.209 255.255.255.240

Answer: A

Explanation:
The subnet sizes needed to meet the address requirements are:
56 = /26 = 255.255.255.192
10 = /28 = 255.255.255.240

72 = /25 = 255.255.255.128
24 = /27 = 255.255.255.224
Answer A most efficiently meets these requirements.

QUESTION NO: 18
Exhibit:

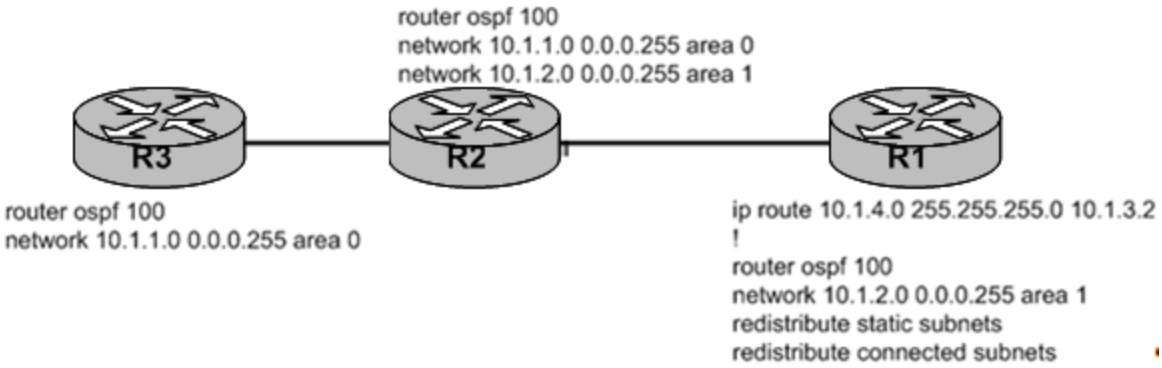

Refer to the exhibits. At R1 in this network, there is no route to 10.1.4.0/24 in the local routing table. Based on the output for R1 in the exhibit, what is the most likely reason 10.1.4.0/24 is not in R1's routing table?

network 10.1.2.0 0.0.0.255 area 1
redistribute static subnets
redistribute connected subnets

```
R1#show ip ospf data external 10.1.4.0/24

        OSPF Router with ID (10.1.3.1) (Process ID 1)
        Type- 5 AS External Link States
        LS age: 72
        Options: (No TOS- capability, DC)
        LS Type: AS External Link
        Link State ID: 10.1.4.0 (External Network Number )
        Advertising Router: 10.1.3.1
        LS Seq Number: 80000001
        Checksum: 0xF161
        Length: 36
```

A. The forwarding address, 10.1.3.2, is also redistributed into OSPF, and an OSPF external route cannot use another OSPF external as its next hop.
B. R2 is not properly configured as an Area Border Router.
C. Area 1 is a stub area, and external routes cannot be originated in a stub area.
D. R3 is not redistributing 10.1.4.0/24 properly.

Answer: A

Explanation:
The forwarding address, 10.1.3.2, is also redistributed into OSPF, and an OSPF external route cannot use another OSPF external as its next hop.

QUESTION NO: 19
Which bits are copied to the EXP bits in an MPLS label by default?

A. TOS
B. CoS
C. IP precedence
D. DSCP

Answer: C
Explanation:
MPLS has 3 EXP bits in the label header that are used in much the same way as IP Precedence bits or the DSCP CS bits. By default, when Cisco IOS Software pushes labels onto an IP packet, the most significant bits in the DiffServ field (the IP Precedence bits) are copied to the EXP field of all imposed labels.
Reference: "Traffic Engineering with MPLS" By Eric Osborne, Ajay Simha, Cisco Press.
http://www.ciscopress.com/articles/article.asp?p=28688&seqNum=5

QUESTION NO: 20
You work as a network technician. Study the exhibit carefully. ROUTER1 is the root bridge for both VLAN 1 and VLAN 2. Which way is the easiest to load-share traffic across both trunks and maintain redundancy in case a link fails, without using any type of EtherChannel link-bundling?

A. Increase the root bridge priority (increasing the numerical priority number) for VLAN 2 on ROUTER2 so that port B2 becomes the root port on ROUTER2 for VLAN 2.
B. Increase the root bridge priority (decreasing the numerical priority number) for VLAN 2 on ROUTER1 so that A2 becomes the root port on ROUTER2 for VLAN 2.
C. Decrease the path cost on A2 on ROUTER1 for VLAN 2 so that port B1 will be blocked for VLAN 2 and port B2 will remain blocked for VLAN 1.
D. Decrease the port priority on A2 for VLAN 2 on ROUTER1 so that port B1 will be blocked for VLAN 2 and port B2 will remain blocked for VLAN 1.

Answer: D

Explanation:
To achieve VLAN load sharing you will need to decrease the port priority value for VLAN 2 on port A2. This way, the corresponding port B2 on Router2 receives better BPDUs than the ones that are sent on port A2 (that still has a port priority default value of 32).

Example:
Router1> (enable) set spantree portvlanpri 2/2 16 1
Port 3/2 vlans 1 using portpri 16.
Port 3/2 vlans 2-1004 using portpri 32.
Port 3/2 vlans 1005 using portpri 4.
Router1> (enable)

The exact same scenario as this question is documented on the Cisco site at the reference link listed below.

Reference:
VLAN Load Balancing Between Trunks Using the Spanning-Tree Protocol Port Priority
http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a00800ae96a.shtml

QUESTION NO: 21
On the basis of the network provided in the exhibit, R3 and R4 are configured to run all connected links in OSPF Area 1. The network administrator is complaining that traffic destined to 192.168.1.0/24 is being routed to R2, even if R2 is not running OSPF. Which would be the cause of this problem?

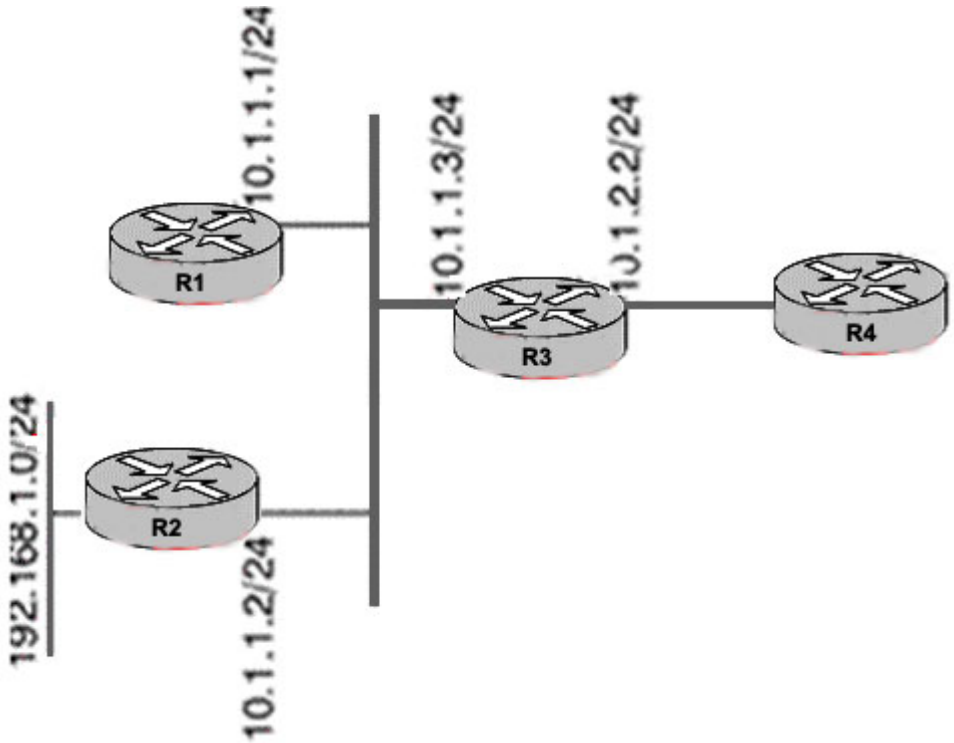

R1 configuration:

```
router eigrp 100
 network 0.0.0.0 0.0.0.0

router eigrp 100
 network 0.0.0.0 0.0.0.0 area 1
 redistribute eigrp 100 cost 10 subnets
```

R2 configuration:

```
router eigrp 100
 network 0.0.0.0 0.0.0.0
```

A. The next hop towards 192.168.1.0/24 at R4 should be 10.1.1.1, since R1 is redistributing the route from EIGRP into OSPF. R3 is forwarding traffic incorrectly
B. R4 would not have a route towards 192.168.1.0/24, so the network administrator is wrong in thinking any traffic is being forwarded there
C. The next hop towards 192.168.1.0/24 at R4 should be 10.1.1.2 which is R2
D. The next hop towards 192.168.1.0/24 at R4 would be 10.1.2.2, which is R3. R3 should be load sharing between R1 and R2 for its next hop

Answer: A
Explanation:
Since OSPF and EIGRP are being redistributed between R 1 and R 2, the route will appear to R4 as an external route, with the next hop being the IP address at R2.

QUESTION NO: 22
Refer to the exhibit. R12 (in Area 4) receives a packet destined for a network in Area 1. What routing table entry will R12 have that will enable it to forward the packet?

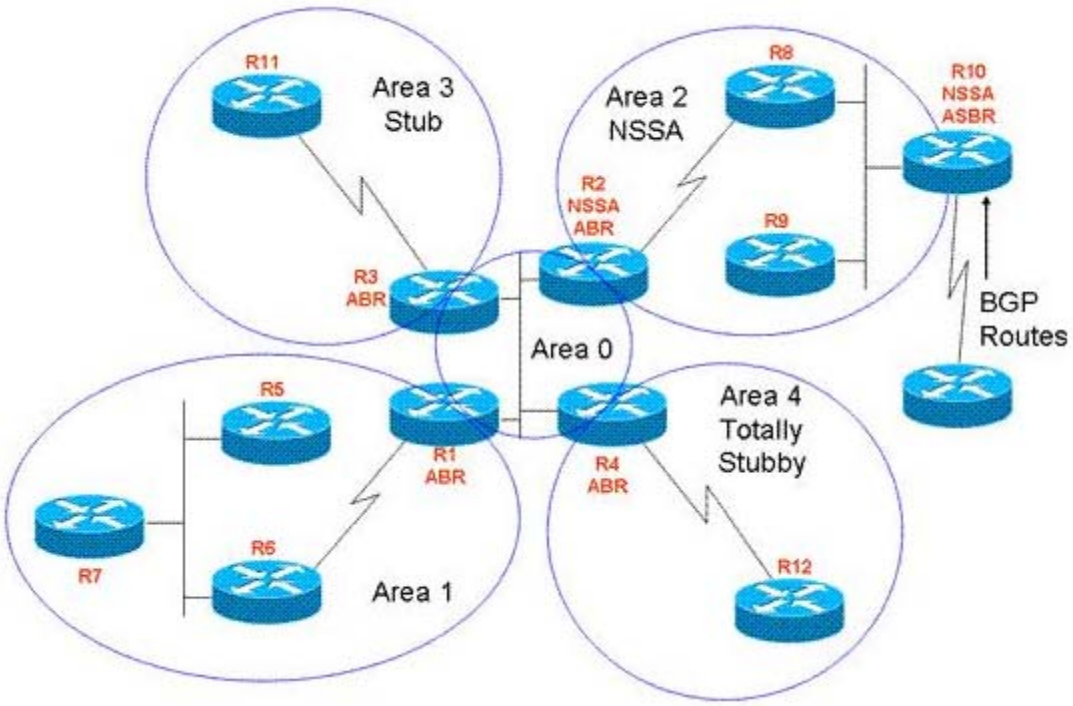

A. a summary route generated by R1 and propagated through the OSPF domain
B. a default route generated by R1 and propagated through the OSPF domain
C. a summary route generated by R4 and propagated to R12
D. a default route generated by R4 and propagated to R12

Answer: D

QUESTION NO: 23
When the NTP peer statement is used in a Cisco IOS router, what functionality does this imply is also being used on the router?

A. static client
B. symmetric active mode
C. static server
D. NTP broadcast client

Answer: B

Explanation:
When a networking device is operating in the *symmetric active mode*, it polls its assigned time-serving hosts for the current time and it responds to polls by its hosts. Because this is a peer-to-peer relationship, the host also retains time-related information about the local networking device that it is communicating with. This mode should be used when there is a number of mutually redundant servers that are interconnected via diverse network paths. Most stratum 1 and stratum 2 servers on the Internet today adopt this form of network setup. Use the **peer** command to individually specify the time-serving hosts that you want your networking device to

consider synchronizing with and to set your networking device to operate in the *symmetric active mode*.

Reference:
http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.3/system_management/configuration/guide/yc33ntp.html

QUESTION NO: 24
In Frame Relay, BECN messages indicating congestion are sent or received by which of these?

A. received by the sender
B. sent by the sender
C. received by the destination
D. sent by the destination

Answer: A
Explanation:
Backward Explicit Congestion Notification (BECN) - The router receiving the frame with BECN set knows that a frame it sent experienced congestion. A frame relay switch sends the BECN to the original sender of the frame to indicate congestion in the network.

QUESTION NO: 25
Refer to the exhibit. In this network, R1 has been configured to advertise a summary route, 192.168.0.0/22, to R2. R2 has been configured to advertise a summary route, 192.168.0.0/21, to R1. Both routers have been configured to remove the discard route (the route to null created when a summary route is configured) by setting the administrative distance of the discard route to 255.
What will happen if R1 receives a packet destined to 192.168.3.1?

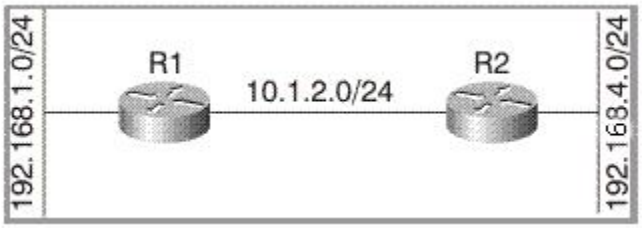

A. The packet will loop between R1 and R2.
B. It is not possible to set the administrative distance on a summary to 255.
C. The packet will be forwarded to R2, where it will be routed to null0.
D. The packet will be dropped by R1, since there is no route to 192.168.3.1

Answer: A

QUESTION NO: 26
Which two options help minimize router resource requirements and improve manageability? (Choose two.)

A. autosummarization
B. Simple Network Management Protocol
C. CPU optimization
D. prefix aggregation

Answer: A,D

Explanation:
Automatic route summarization and prefix aggregation is always a recommended best design practice whenever possible, as it means less routing table entries for the router to store. For example, many subnets can be hidden behind a single routing table entry, making these entries smaller, and routing more efficient).

QUESTION NO: 27
Which of these best identifies the types of prefixes a router running BGP will advertise to an EBGP peer?

A. prefixes received from any other BGP peer and prefixes locally originated via network statements or redistributed to BGP
B. all prefixes in its IP routing table
C. only prefixes received from EBGP peers and prefixes locally originated via network statements or redistributed
D. only prefixes received from EBGP peers and prefixes received from route reflectors
E. all prefixes in its routing table except the prefixes received from other EBGP peers
F. a prefixes in its routing table except the prefixes received from other IBGP peers

Answer: A

Explanation:
By default, a BGP router will advertise routes that were received form other BGP peers (both IBGP and EBGP peers) as well as any locally generated routes via the network command or via redistribution. The default configuration of BGP on a circuit does not advertise any routes or allow any learned routes into the IGP routing table, these have to be manually entered as Network statements or be redistributed into the IGP.
The network command controls what networks are originated by this box. This is a different concept from what you are used to configuring with IGRP and RIP. With this command we are not trying to run BGP on a certain interface, rather we are trying to indicate to BGP what networks it should originate from this box.
The network command is one way to advertise your networks via BGP. Another way is to redistribute your IGP (IGRP, OSPF, RIP, EIGRP, etc.) into BGP. Careful filtering should be applied to make sure you are sending to the internet only routes that you want to advertise and not everything you have.

QUESTION NO: 28
What feature monitors the level of each traffic type in 1-second intervals?

A. Port Fast
B. Uplink Fast
C. Storm Control
D. Port Aggregation Protocol
E. Link Aggregation Configuration Options

Answer: C

Explanation:
Traffic storm control (also called traffic suppression) monitors incoming traffic levels over a 1-second traffic storm control interval and, during the interval, compares the traffic level with the traffic storm control level that you configure. The traffic storm control level is a percentage of the total available bandwidth of the port. Each port has a single traffic storm control level that is used for all types of traffic (broadcast, multicast, and unicast).

QUESTION NO: 29
IP multicast addresses in which range are used for Scope Relative multicast?

A. The lowest (numerically) 256 multicast addresses of each administratively scoped address range are automatically reserved for Scope Relative multicast.
B. Scope Relative multicast addresses must be chosen from the administratively scoped address range by the network administrator and configured on every router.
C. The highest (numerically) 256 addresses of each administratively scoped address range are automatically reserved for Scope Relative multicast.
D. The highest (numerically) 32 addresses of each administratively scoped address range are automatically reserved for Scope Relative multicast.

Answer: C

Explanation:
Multicast addresses may be allocated in any of three ways:
Static:
Statically allocated addresses are allocated by IANA for specific protocols that require well-known addresses to work. Examples of static addresses are 224.0.1.1 which is used for the Network Time Protocol and 224.2.127.255 which is used for global scope multicast session announcements.
Scope-relative:
RFC 2365 reserves the highest 256 addresses in every administrative scope range for relative assignments. Relative assignments are made by IANA and consist of an offset which is valid in every scope.
Dynamic:

For most purposes, the correct way to use multicast is to obtain a dynamic multicast address. These addresses are provided on demand and have a specific lifetime.
Reference: http://www.ietf.org/rfc/rfc2908.txt

QUESTION NO: 30
Refer to the exhibit. Which protocol will load-balance traffic across all gateways in a group by dynamically assigning responsibility for a Virtual IP address and multiple virtual MAC addresses to each member of the group?

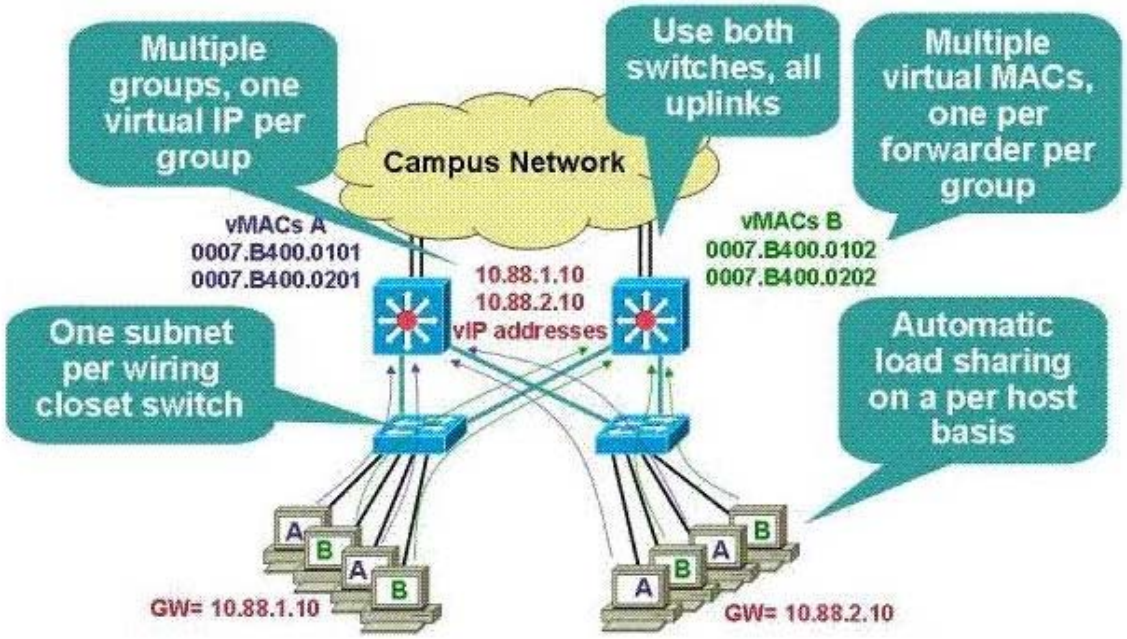

A. Hot Standby Router Protocol
B. Gateway Load Balancing Protocol
C. Virtual Router Redundancy Protocol
D. Simple Network Management Protocol
E. Spanning Tree Protocol

Answer: B
Explanation:
The Gateway Load Balancing Protocol feature provides automatic router backup for IP hosts configured with a single default gateway on an IEEE 802.3 LAN. Multiple first hop routers on the LAN combine to offer a single virtual first hop IP router while sharing the IP packet forwarding load. Other routers on the LAN may act as redundant GLBP routers that will become active if any of the existing forwarding routers fail.
GLBP performs a similar, but not identical, function for the user as the HSRP and the VRRP. HSRP and VRRP protocols allow multiple routers to participate in a virtual router group configured with a virtual IP address. One member is elected to be the active router to forward packets sent to the virtual IP address for the group. The other routers in the group are redundant until the active router fails. These standby routers have unused bandwidth that the protocol is not using. Although multiple virtual router groups can be configured for the same set of routers, the

hosts must be configured for different default gateways, which results in an extra administrative burden. GLBP provides load balancing over multiple routers (gateways) using a single virtual IP address and multiple virtual MAC addresses. Each host is configured with the same virtual IP address, and all routers in the virtual router group participate in forwarding packets.
Reference: GLBP - Gateway Load Balancing Protocolhttp://www.cisco.com/en/US/docs/ios/12_2t/12_2t15/feature/guide/ft_glbp.html