

Test-King.300-101.235.Questions

VCEplus.com

Number: 300-101
Passing Score: 800
Time Limit: 120 min
File Version: 21.9



300-101

Implementing Cisco IP Routing

Found this Vce in good condition with detailed explanations, clear exhibits and real exam questions. Best for self preparations to pass the exam.


Sections

1. Network Principles
2. Layer 2 Technologies
3. Layer 3 Technologies
4. VPN Technologies
5. Infrastructure Security
6. Infrastructure Services
7. Mixed Questions

Exam A

QUESTION 1

Refer to the exhibit.



Based on this FIB table, which statement is correct?

- A. There is no default gateway.
- B. The IP address of the router on FastEthernet is 209.168.201.1.
- C. The gateway of last resort is 192.168.201.1.
- D. The router will listen for all multicast traffic.

Correct Answer: C

Section: Network Principles

Explanation


Explanation/Reference:

Explanation:

The 0.0.0.0/0 route is the default route and is listed as the first CEF entry. Here we see the next hop for this default route lists 192.168.201.1 as the default router (gateway of last resort).

QUESTION 2

Refer to the exhibit.



A network administrator checks this adjacency table on a router. What is a possible cause for the incomplete marking?

- A. incomplete ARP information
- B. incorrect ACL
- C. dynamic routing protocol failure
- D. serial link congestion

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

To display information about the Cisco Express Forwarding adjacency table or the hardware Layer 3-switching adjacency table, use the show adjacency command.

Reasons for Incomplete Adjacencies

There are two known reasons for an incomplete adjacency:

The router cannot use ARP successfully for the next-hop interface. After a clear iparp or a clear adjacency command, the router marks the adjacency as in- complete. Then it fails to clear the entry.

In an MPLS environment, IP CEF should be enabled for Label Switching. Interface level command ip route-cache cef

No ARP Entry

When CEF cannot locate a valid adjacency for a destination prefix, it punts the packets to the CPU for ARP resolution and, in turn, for completion of the adjacency.

Reference: <http://www.cisco.com/c/en/us/support/docs/ip/express-forwarding-cef/17812-cef-incomp.html#t4>

QUESTION 3

A network engineer notices that transmission rates of senders of TCP traffic sharply increase and decrease simultaneously during periods of congestion.

Which condition causes this?

- A. global synchronization
- B. tail drop
- C. random early detection
- D. queue management algorithm

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

TCP global synchronization in computer networks can happen to TCP/IP flows during periods of congestion because each sender will reduce their transmission rate at the same time when packet loss occurs.

Routers on the Internet normally have packet queues, to allow them to hold packets when the network is busy, rather than discarding them.

Because routers have limited resources, the size of these queues is also limited. The simplest technique to limit queue size is known as tail drop. The queue is allowed to fill to its maximum size, and then any new packets are simply discarded, until there is space in the queue again. This causes problems when used on TCP/IP routers handling multiple TCP streams, especially when bursty traffic is present. While the network is stable, the queue is constantly full, and there are no problems except that the full queue results in high latency. However, the introduction of a sudden burst of traffic may cause large numbers of established, steady streams to lose packets simultaneously.

Reference: http://en.wikipedia.org/wiki/TCP_global_synchronization

QUESTION 4

Which statement about the use of tunneling to migrate to IPv6 is true?

- A. Tunneling is less secure than dual stack or translation.
- B. Tunneling is more difficult to configure than dual stack or translation.
- C. Tunneling does not enable users of the new protocol to communicate with users of the old protocol without dual-stack hosts.
- D. Tunneling destinations are manually determined by the IPv4 address in the low-order 32 bits of IPv4-compatible IPv6 addresses.

Correct Answer: C

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Using the tunneling option, organizations build an overlay network that tunnels one protocol over the other by encapsulating IPv6 packets within IPv4 packets and IPv4 packets within IPv6 packets. The advantage of this approach is that the new protocol can work without disturbing the old protocol, thus providing connectivity between users of the new protocol. Tunneling has two disadvantages, as discussed in RFC 6144:

- Users of the new architecture cannot use the services of the underlying infrastructure.
- Tunneling does not enable users of the new protocol to

communicate with users of the old protocol without dual-stack hosts, which negates interoperability. Reference: http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/enterprise-ipv6-solution/white_paper_c11-676278.html

QUESTION 5

Which three problems result from application mixing of UDP and TCP streams within a network with no QoS? (Choose three.)

- A. starvation
- B. jitter
- C. latency
- D. windowing
- E. lower throughput

Correct Answer: ACE

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

It is a general best practice not to mix TCP-based traffic with UDP-based traffic (especially streaming video) within a single service provider class due to the behaviors of these protocols during periods of congestion. Specifically, TCP transmitters will throttle-back flows when drops have been detected. Although some UDP applications have application-level windowing, flow control, and retransmission capabilities, most UDP transmitters are completely oblivious to drops and thus never lower transmission rates due to dropping. When TCP flows are combined with UDP flows in a single service provider class and the class experiences congestion, then TCP flows will continually lower their rates, potentially giving up their bandwidth to drop-oblivious UDP flows. This effect is called TCP-starvation/UDP-dominance. This can increase latency and lower the overall throughput.

TCP-starvation/UDP-dominance likely occurs if (TCP-based) mission-critical data is assigned to the same service provider class as (UDP-based) streaming video and the class experiences sustained congestion. Even if WRED is enabled on the service provider class, the same behavior would be observed, as WRED (for the most part) only affects TCP-based flows. Granted, it is not always possible to separate TCP-based flows from UDP-based flows, but it is beneficial to be aware of this behavior when making such application-mixing decisions. Reference: http://www.cisco.com/warp/public/cc/so/neso/vpn/vpnsp/spqsd_wp.htm

QUESTION 6

Which method allows IPv4 and IPv6 to work together without requiring both to be used for a single connection during the migration process?

- A. dual-stack method
- B. 6to4 tunneling
- C. GRE tunneling
- D. NAT-PT

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Dual stack means that devices are able to run IPv4 and IPv6 in parallel. It allows hosts to simultaneously reach IPv4 and IPv6 content, so it offers a very flexible coexistence strategy. For sessions that support IPv6, IPv6 is used on a dual stack endpoint. If both endpoints support IPv4 only, then IPv4 is used.

Benefits:

- Native dual stack does not require any tunneling mechanisms on internal networks
- Both IPv4 and IPv6 run independent of each other
- Dual stack supports gradual migration of endpoints, networks, and applications. Reference: http://www.cisco.com/web/strategy/docs/gov/IPV6at_a_glance_c45-625859.pdf

QUESTION 7

A network administrator executes the command `clear ip route`. Which two tables does this command clear and rebuild? (Choose two.)

- A. IP routing
- B. FIB
- C. ARP cache
- D. MAC address table
- E. Cisco Express Forwarding table
- F. topology table

Correct Answer: AB

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

To clear one or more entries in the IP routing table, use the following commands in any mode:

Command Purpose

`clear iproute {* |`

Clears one or more routes from both the unicast RIB and all the module FIBs. The

`{route |`

route options are as follows:

`prefix/length}[next-`

`hopinterface]]`

· `--All routes.`

`*`

`[vrfvrf-name]`

Example:

· `--An individual IP route.`

`route`

```
switch(config)# clear  
iproute - --Any IP prefix.  
prefix/length  
10.2.2.2  
· --The next-hop address  
next-hop
```

```
· --The interface to reach the  
interface  
next-hop address.
```

The vrf-name can be any case-sensitive, alphanumeric string up to 32 characters.

Reference:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/unicast/5_0_3_N1_1/Cisco_n5k_layer3_ucast_cfg_rel_503_N1_1/l3_manage-routes.html

QUESTION 8

Which switching method is used when entries are present in the output of the command show ip cache?

- A. fast switching
- B. process switching
- C. Cisco Express Forwarding switching
- D. cut-through packet switching

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Fast switching allows higher throughput by switching a packet using a cache created by the initial packet sent to a particular destination. Destination addresses are stored in the high-speed cache to expedite forwarding. Routers offer better packet-transfer performance when fast switching is enabled. Fast switching is enabled by default on all interfaces that support fast switching.

To display the routing table cache used to fast switch IP traffic, use the show ip cache EXEC command.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/12_2/switch/command/reference/fswtch_r/xfscmd5.html#wp1038133

QUESTION 9

Which two actions must you perform to enable and use window scaling on a router? (Choose two.)

- A. Execute the command `iptcp window-size 65536`.
- B. Set window scaling to be used on the remote host.
- C. Execute the command `iptcpqueuemax`.
- D. Set TCP options to "enabled" on the remote host.
- E. Execute the command `iptcp adjust-mss`.

Correct Answer: AB

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

The TCP Window Scaling feature adds support for the Window Scaling option in RFC 1323, TCP Extensions for High Performance . A larger window size is recommended to improve TCP performance in network paths with large bandwidth-delay product characteristics that are called Long Fat Networks (LFNs). The TCP Window Scaling enhancement provides that support. The window scaling extension in Cisco IOS software expands the definition of the TCP window to 32 bits and then uses a scale factor to carry this 32-bit value in the 16-bit window field of the TCP header. The window size can increase to a scale factor of 14. Typical applications use a scale factor of 3 when deployed in LFNs.

The TCP Window Scaling feature complies with RFC 1323. The larger scalable window size will allow TCP to perform better over LFNs. Use the `iptcp window-size` command in global configuration mode to configure the TCP window size. In order for this to work, the remote host must also support this feature and its window size must be increased. Reference: <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/12-4t/iap-12-4t-book/iap-tcp.html#GUID-BD998AC6-F128-47DD-B5F7-B226546D4B08>

Verified

QUESTION 10

Which three TCP enhancements can be used with TCP selective acknowledgments? (Choose three.)

- A. header compression
- B. explicit congestion notification
- C. keepalive
- D. time stamps
- E. TCP path discovery
- F. MTU window

Correct Answer: BCD

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

TCP Selective Acknowledgment

The TCP Selective Acknowledgment feature improves performance if multiple packets are lost from one TCP window of data.

Prior to this feature, because of limited information available from cumulative acknowledgments, a TCP sender could learn about only one lost packet per-round-trip time. An aggressive sender could choose to resend packets early, but such re-sent segments might have already been successfully received.

The TCP selective acknowledgment mechanism helps improve performance. The receiving TCP host returns selective acknowledgment packets to the sender, informing the sender of data that has been received. In other words, the receiver can acknowledge packets received out of order. The sender can then resend only missing data segments (instead of everything since the first missing packet).

Prior to selective acknowledgment, if TCP lost packets 4 and 7 out of an 8-packet window, TCP would receive acknowledgment of only packets 1, 2, and 3. Packets 4 through 8 would need to be re-sent. With selective acknowledgment, TCP receives acknowledgment of packets 1, 2, 3, 5, 6, and 8. Only packets 4 and 7 must be re-sent.

TCP selective acknowledgment is used only when multiple packets are dropped within one TCP window. There is no performance impact when the feature is enabled but not used. Use the `iptcp selective-ack` command in global configuration mode to enable TCP selective acknowledgment. Refer to RFC 2018 for more details about TCP selective acknowledgment.

TCP Time Stamp

The TCP time-stamp option provides improved TCP round-trip time measurements. Because the time stamps are always sent and echoed in both directions and the time-stamp value in the header is always changing, TCP header compression will not compress the outgoing packet. To allow TCP header compression over a serial link, the TCP time-stamp option is disabled. Use the `iptcp timestamp` command to enable the TCP time-stamp option.

TCP Explicit Congestion Notification

The TCP Explicit Congestion Notification (ECN) feature allows an intermediate router to notify end hosts of impending network congestion. It also provides enhanced support for TCP sessions associated with applications, such as Telnet, web browsing, and transfer of audio and video data that are sensitive to delay or packet loss. The benefit of this feature is the reduction of delay and packet loss in data transmissions. Use the `iptcpecn` command in global configuration mode to enable TCP ECN.

TCP Keepalive Timer

The TCP Keepalive Timer feature provides a mechanism to identify dead connections. When a TCP connection on a routing device is idle for too long, the device sends a TCP keepalive packet to the peer with only the Acknowledgment (ACK) flag turned on. If a response packet (a TCP ACK packet) is not received after the device sends a specific number of probes, the connection is considered dead and the device initiating the probes frees resources used by the TCP connection.

Reference: <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/xr-3s/asr1000/iap-xe-3s-asr1000-book/iap-tcp.html#GUID-22A82C5F-631F-4390-9838-F2E48FFEEA01>

QUESTION 11

A network administrator uses IP SLA to measure UDP performance and notices that packets on one router have a higher one-way delay compared to the opposite direction. Which UDP characteristic does this scenario describe?

- A. latency
- B. starvation
- C. connectionless communication
- D. nonsequencing unordered packets
- E. jitter

Correct Answer: A

Section: Network Principles

Explanation**Explanation/Reference:**

Explanation:

Cisco IOS IP SLAs provides a proactive notification feature with an SNMP trap. Each measurement operation can monitor against a pre-set performance threshold. Cisco IOS IP SLAs generates an SNMP trap to alert management applications if this threshold is crossed. Several SNMP traps are available: round trip time, average jitter, one-way latency, jitter, packet loss, MOS, and connectivity tests.

Here is a partial sample output from the IP SLA statistics that can be seen:

```
router#showipsla statistics 1
```

```
Round Trip Time (RTT) for Index 55
```

```
Latest RTT: 1 ms
```

```
Latest operation start time: *23:43:31.845 UTC Thu Feb 3 2005 Latest operation return code: OK
```

```
RTT Values:
```

```
Number Of RTT: 10 RTT Min/Avg/Max: 1/1/1 milliseconds Latency one-way time:
```

```
Number of Latency one-way Samples: 0
```

```
Source to Destination Latency one way Min/Avg/Max: 0/0/0 milliseconds Destination to Source Latency one way Min/Avg/Max: 0/0/0 milliseconds
```

```
Reference:
```

```
http://www.cisco.com/en/US/technologies/tk648/tk362/tk920/technologies\_white\_paper09186a00802d5efe.html
```

QUESTION 12

Under which condition does UDP dominance occur?

- A. when TCP traffic is in the same class as UDP
- B. when UDP flows are assigned a lower priority queue
- C. when WRED is enabled
- D. when ACLs are in place to block TCP traffic

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Mixing TCP with UDP

It is a general best practice to not mix TCP-based traffic with UDP-based traffic (especially Streaming-Video) within a single service-provider class because of the behaviors of these protocols during periods of congestion. Specifically, TCP transmitters throttle back flows when drops are detected. Although some UDP applications have application-level windowing, flow control, and retransmission capabilities, most UDP transmitters are completely oblivious to drops and, thus, never lower transmission rates because of dropping. When TCP flows are combined with UDP flows within a single service-provider class and the class experiences congestion, TCP flows continually lower their transmission rates, potentially giving up their bandwidth to UDP flows that are oblivious to drops. This effect is called TCP starvation/UDP dominance.

TCP starvation/UDP dominance likely occurs if (TCP-based) Mission-Critical Data is assigned to the same service-provider class as (UDP-based) Streaming-Video and the class experiences sustained congestion. Even if WRED is enabled on the service-provider class, the same behavior would be

observed because WRED (for the most part) manages congestion only on TCP-based flows.

Reference:

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book/VPNQoS.html

Topic 2, Layer 2 Technologies

QUESTION 13

Prior to enabling PPPoE in a virtual private dialup network group, which task must be completed?

- A. Disable CDP on the interface.
- B. Execute the vpdn enable command.
- C. Execute the no switchport command.
- D. Enable QoS FIFO for PPPoE support.

Correct Answer: B

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Enabling PPPoE in a VPDN Group

Perform this task to enable PPPoE in a virtual private dial-up network (VPDN) group.

Restrictions

This task applies only to releases prior to Cisco IOS Release 12.2(13)T.

SUMMARY STEPS

1. enable
2. configureterminal
3. vpdn enable
4. vpdn-groupname
5. request-dialin
6. protocolpppoe

DETAILED STEPS

Command or Action Purpose

Step 1 enable Enables privileged EXEC mode.

.

Example: Enter your password if

Router> enable prompted.

Step 2 configureterminal Enters global configuration mode.

Example:

Router# configure terminal

Step 3 vpdn enable
Enables virtual private dialup
Example: networking.
Router(config)# vpdn enable

Step 4 vpdn-groupname Associates a VPDN group with a Example: customer or VPDN profile.
Router(config)# vpdn-group
group1

Step 5 request-dialin Creates a request-dialin VPDN Example: subgroup.
Router(config-vpdn)# request-
dialin

Step 6 protocol pppoe Enables the VPDN subgroup to
Example: establish PPPoE
Router(config-vpdn-req-
in)# protocol pppoe

Reference:

http://www.cisco.com/en/US/docs/ios/12_2t/12_2t2/feature/guide/ftpppoe_support_TSD_Island_of_Content_Chapter.html

QUESTION 14

A corporate policy requires PPPoE to be enabled and to maintain a connection with the ISP, even if no interesting traffic exists. Which feature can be used to accomplish this task?

- A. TCP Adjust
- B. Dialer Persistent
- C. PPPoE Groups
- D. half-bridging
- E. Peer Neighbor Route

Correct Answer: B

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

A new interface configuration command, dialer persistent, allows a dial-on-demand routing (DDR) dialer profile connection to be brought up without being triggered by interesting traffic. When configured, the dialer persistent command starts a timer when the dialer interface starts up and starts the connection when the timer expires. If interesting traffic arrives before the timer expires, the connection is still brought up and set as persistent. The command provides a default timer interval, or you can set a custom timer interval. To configure a dialer interface as persistent, use the following commands beginning in global configuration mode:

Command Purpose

Step 1 Router(config)# interface dialer Creates a dialer interface and number enters interface configuration mode.

Step 2 Router(config-if)# ip Specifies the IP address and mask addressaddress mask of the dialer interface as a node in the destination network to be called.

Step 3 Router(config-if)# encapsulation Specifies the encapsulation type.
type

Step 4 Router(config-if)# dialer string Specifies the remote destination dial-string class class-name to call and the map class that defines characteristics for calls to this destination.

Step 5 Router(config-if)# dialer pool Specifies the dialing pool to use number for calls to this destination.

Step 6 Router(config-if)# dialer- Assigns the dialer interface to a groupgroup-number dialer group.

Step 7 Router(config-if)# dialer-list Specifies an access list by list dialer-group protocol protocol- number or by protocol and list name{permit | deny | list access- number to define the interesting list-number} packets that can trigger a call.

Step 8 Router(config-if)# dialer remote- (Optional) Specifies the name user-name authentication name of the remote router on the destination subnetwork for a dialer interface.

Step 9 Router(config-if)# dialer Forces a dialer interface to be persistent [delay [initial] seconds connected at all times, even in | max-attemptsnumber] the absence of interesting traffic.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/dial/configuration/guide/12_4t/dia_12_4t_book/dia_dialer_persist.html

QUESTION 15

A network engineer has been asked to ensure that the PPPoE connection is established and authenticated using an encrypted password. Which technology, in combination with PPPoE, can be used for authentication in this manner?

- A. PAP
- B. dot1x
- C. IPsec
- D. CHAP
- E. ESP

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

With PPPoE, the two authentication options are PAP and CHAP. When CHAP is enabled on an interface and a remote device attempts to connect to it, the access server sends a CHAP packet to the remote device. The CHAP packet requests or "challenges" the remote device to respond. The challenge packet consists of an ID, a random number, and the host name of the local router. When the remote device receives the challenge packet, it concatenates the ID, the remote device's password, and the random number, and then encrypts all of it using the remote device's password. The remote device sends the results back to the access server, along with the name associated with the password used in the encryption process. When the access server receives the response, it uses the name it received to retrieve a password stored in its user database. The retrieved password should be the same password the remote device used in its encryption process. The access server then encrypts the concatenated information with the newly retrieved password--if the result matches the result sent in the response packet, authentication succeeds.

The benefit of using CHAP authentication is that the remote device's password is never transmitted in clear text (encrypted). This prevents other devices from stealing it and gaining illegal access to the ISP's network.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/12_2/security/configuration/guide/fsec_c/scfathen.html

QUESTION 16

Which PPP authentication method sends authentication information in cleartext?

- A. MS CHAP
- B. CDPCP
- C. CHAP
- D. PAP

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

PAP authentication involves a two-way handshake where the username and password are sent across the link in clear text; hence, PAP authentication does not provide any protection against playback and line sniffing.

CHAP authentication, on the other hand, periodically verifies the identity of the remote node using a three-way handshake. After the PPP link is established, the host sends a "challenge" message to the remote node. The remote node responds with a value calculated using a one-way hash function. The host checks the response against its own calculation of the expected hash value. If the values match, the authentication is acknowledged; otherwise, the connection is terminated.

Reference: <http://www.cisco.com/c/en/us/support/docs/wan/point-to-point-protocol-ppp/10241-ppp-callin-hostname.html>

QUESTION 17

Which protocol uses dynamic address mapping to request the next-hop protocol address for a specific connection?

- A. Frame Relay inverse ARP
- B. static DLCI mapping
- C. Frame Relay broadcast queue
- D. dynamic DLCI mapping

Correct Answer: A

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Dynamic address mapping uses Frame Relay Inverse ARP to request the next-hop protocol address for a specific connection, given its known DLCI. Responses to Inverse ARP requests are entered in an address-to-DLCI mapping table on the router or access server; the table is then used to supply the next-hop protocol address or the DLCI for outgoing traffic.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/12_2/wan/configuration/guide/fwan_c/wcffrely.html

QUESTION 18

What is the default OSPF hello interval on a Frame Relay point-to-point network?

- A. 10
- B. 20
- C. 30
- D. 40

Correct Answer: A

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Before you troubleshoot any OSPF neighbor-related issues on an NBMA network, it is important to remember that an NBMA network can be configured in these modes of operation with the ipospf network command:

Point-to-Point

Point-to-Multipoint

Broadcast

NBMA

The Hello and Dead Intervals of each mode are described in this table:

Hello Interval Dead Interval
Network Type
(secs) (secs)

Point-to-Point 10 40

Point-to-Multipoint 30 120

Broadcast 10 40

Non-Broadcast 30 120

Reference: <http://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13693-22.html>

QUESTION 19

Which statement is true about the PPP Session Phase of PPPoE?

- A. PPP options are negotiated and authentication is not performed. Once the link setup is completed, PPPoE functions as a Layer 3 encapsulation method that allows data to be transferred over the PPP link within PPPoE headers.
- B. PPP options are not negotiated and authentication is performed. Once the link setup is completed, PPPoE functions as a Layer 4 encapsulation method that allows data to be transferred over the PPP link within PPPoE headers.
- C. PPP options are automatically enabled and authorization is performed. Once the link setup is completed, PPPoE functions as a Layer 2 encapsulation method that allows data to be encrypted over the PPP link within PPPoE headers.
- D. PPP options are negotiated and authentication is performed. Once the link setup is completed, PPPoE functions as a Layer 2 encapsulation method that allows data to be transferred over the PPP link within PPPoE headers.

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

PPPoE is composed of two main phases:

Active Discovery Phase--In this phase, the PPPoE client locates a PPPoE server, called an access concentrator. During this phase, a Session ID is assigned and the PPPoE layer is established.


PPP Session Phase--In this phase, PPP options are negotiated and authentication is performed. Once the link setup is completed, PPPoE functions as a Layer 2 encapsulation method, allowing data to be transferred over the PPP link within PPPoE headers.

Reference: <http://www.cisco.com/c/en/us/td/docs/security/asa/asa92/configuration/vpn/asa-vpn- cli/vpn-pppoe.html>

Topic 3, Layer 3 Technologies

QUESTION 20

Refer to the exhibit.



Which one statement is true?

- A. Traffic from the 172.16.0.0/16 network will be blocked by the ACL.
- B. The 10.0.0.0/8 network will not be advertised by Router B because the network statement for the 10.0.0.0/8 network is missing from Router B.
- C. The 10.0.0.0/8 network will not be in the routing table on Router B.
- D. Users on the 10.0.0.0/8 network can successfully ping users on the 192.168.5.0/24 network, but users on the 192.168.5.0/24 cannot successfully ping users on the 10.0.0.0/8 network.
- E. Router B will not advertise the 10.0.0.0/8 network because it is blocked by the ACL.

Correct Answer: E

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

You can filter what individual routes are sent (out) or received (in) to any interface within your EIGRP configuration.

One example is noted above. If you filter outbound, the next neighbor(s) will not know about anything except the 172.16.0.0/16 route and therefore won't send it to anyone else downstream. If you filter inbound, YOU won't know about the route and therefore won't send it to anyone else downstream.

QUESTION 21

A router with an interface that is configured with ipv6 address autoconfig also has a link-local address assigned. Which message is required to obtain a global unicast address when a router is present?

- A. DHCPv6 request
- B. router-advertisement
- C. neighbor-solicitation
- D. redirect

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Autoconfiguration is performed on multicast-enabled links only and begins when a multicast-enabled interface is enabled (during system startup or manually). Nodes (both, hosts and routers) begin the process by generating a link-local address for the interface. It is formed by appending the interface identifier to well-known link-local prefix FE80 :: 0. The interface identifier replaces the right-most zeroes of the link-local prefix. Before the link-local address can be assigned to the interface, the node performs the Duplicate Address Detection mechanism to see if any other node is using the same link-local address on the link. It does this by sending a Neighbor Solicitation message with target address as the "tentative" address and destination address as the solicited-node multicast address corresponding to this tentative address. If a node responds with a Neighbor Advertisement message with tentative address as the target address, the address is a duplicate address and must not be used.

Hence, manual configuration is required.

Once the node verifies that its tentative address is unique on the link, it assigns that link-local address to the interface. At this stage, it has IP-connectivity to other neighbors on this link. The autoconfiguration on the routers stop at this stage, further tasks are performed only by the hosts. The routers will need manual configuration (or stateful configuration) to receive site-local or global addresses.

The next phase involves obtaining Router Advertisements from routers if any routers are present on the link. If no routers are present, a stateful configuration is required. If routers are present, the Router Advertisements notify what sort of configurations the hosts need to do and the hosts receive a global unicast IPv6 address.

Reference: <https://sites.google.com/site/amitsciscozone/home/important-tips/ipv6/ipv6-stateless-autoconfiguration>

QUESTION 22

PPPoE is composed of which two phases?

- A. Active Authentication Phase and PPP Session Phase
- B. Passive Discovery Phase and PPP Session Phase
- C. Active Authorization Phase and PPP Session Phase
- D. Active Discovery Phase and PPP Session Phase

Correct Answer: D

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

PPPoE is composed of two main phases:

Active Discovery Phase--In this phase, the PPPoE client locates a PPPoE server, called an access concentrator. During this phase, a Session ID is assigned and the PPPoE layer is established.

PPP Session Phase--In this phase, PPP options are negotiated and authentication is performed. Once the link setup is completed, PPPoE functions as a Layer 2 encapsulation method, allowing data to be transferred over the PPP link within PPPoE headers. Reference: <http://www.cisco.com/c/en/us/td/docs/security/asa/asa92/configuration/vpn/asa-vpn-cli/vpn-pppoe.html>

QUESTION 23

An engineer has configured a router to use EUI-64, and was asked to document the IPv6 address of the router. The router has the following interface parameters:

mac address C601.420F.0007

subnet 2001:DB8:0:1::/64

Which IPv6 addresses should the engineer add to the documentation?

- A. 2001:DB8:0:1:C601:42FF:FE0F:7
- B. 2001:DB8:0:1:FFFF:C601:420F:7
- C. 2001:DB8:0:1:FE80:C601:420F:7
- D. 2001:DB8:0:1:C601:42FE:800F:7

Correct Answer: A

Section: Layer 3 Technologies

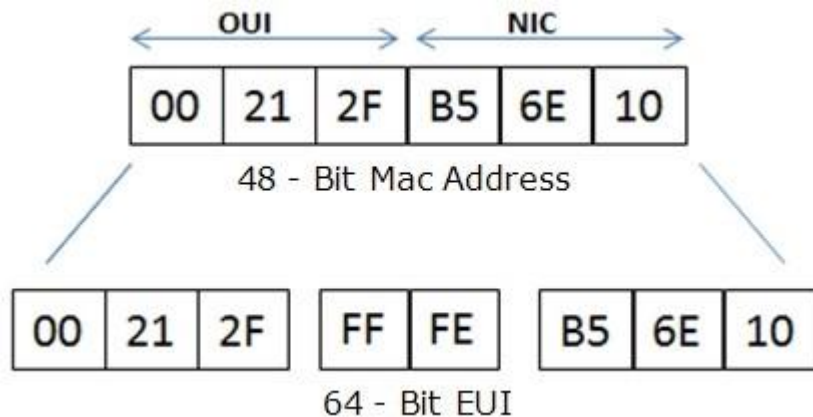
Explanation

Explanation/Reference:

Explanation:

Extended Unique Identifier (EUI), as per RFC2373, allows a host to assign itself a unique 64-Bit IP Version 6 interface identifier (EUI-64). This feature is a key benefit over IPv4 as it eliminates the need of manual configuration or DHCP as in the world of IPv4. The IPv6 EUI-64 format address is obtained through the 48-bit MAC address. The Mac address is first separated into two 24-bits, with one being OUI (Organizationally Unique Identifier) and the other being NIC specific. The 16-bit 0xFFFF is then inserted between these two 24-bits to form the 64-bit EUI address. IEEE has chosen FFFE as a reserved value which can only appear in EUI-64 generated from the EUI-48 MAC address.

Here is an example showing how the Mac Address is used to generate EUI.



Next, the seventh bit from the left, or the universal/local (U/L) bit, needs to be inverted. This bit identifies whether this interface identifier is universally or locally administered. If 0, the address is locally administered and if 1, the address is globally unique. It is worth noticing that in the OUI portion, the globally unique addresses assigned by the IEEE has always been set to 0 whereas the locally created addresses has 1 configured. Therefore, when the bit is inverted, it maintains its original scope (global unique address is still global unique and vice versa). The reason for inverting can be found in RFC4291 section 2.5.1. Reference: <https://supportforums.cisco.com/document/100566/understanding-ipv6-eui-64-bit-address>

QUESTION 24

For security purposes, an IPv6 traffic filter was configured under various interfaces on the local router. However, shortly after implementing the traffic filter, OSPFv3 neighbor adjacencies were lost. What caused this issue?

- A. The traffic filter is blocking all ICMPv6 traffic.
- B. The global anycast address must be added to the traffic filter to allow OSPFv3 to work properly.
- C. The link-local addresses that were used by OSPFv3 were explicitly denied, which caused the neighbor relationships to fail.
- D. IPv6 traffic filtering can be implemented only on SVIs.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

OSPFv3 uses link-local IPv6 addresses for neighbor discovery and other features, so if any IPv6 traffic filters are implemented be sure to include the link local address so that it is permitted in the filter list.

Reference: http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/unicast/configuration/guide/l3_cli_nxos/l3_ospfv3.html

QUESTION 25

What is the purpose of the autonomous-system {autonomous-system-number} command?

- A. It sets the EIGRP autonomous system number in a VRF.
- B. It sets the BGP autonomous system number in a VRF.
- C. It sets the global EIGRP autonomous system number.
- D. It sets the global BGP autonomous system number.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

To configure the autonomous-system number for an Enhanced Interior Gateway Routing Protocol (EIGRP) routing process to run within a VPN routing and forwarding (VRF) instance, use the autonomous-system command in address-family configuration mode. To remove the autonomous-system for an EIGRP routing process from within a VPN VRF instance, use the no form of this command.

autonomous-system autonomous-system-number

no autonomous-system autonomous-system-number

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/iproute_eigrp/command/reference/ire_book/ire_a1.htm#wp1062796

QUESTION 26

Router A and Router B are configured with IPv6 addressing and basic routing capabilities using OSPFv3. The networks that are advertised from Router A do not show up in Router B's routing table. After debugging IPv6 packets, the message "not a router" is found in the output. Why is the routing information not being learned by Router B?

- A. OSPFv3 timers were adjusted for fast convergence.
- B. The networks were not advertised properly under the OSPFv3 process.
- C. An IPv6 traffic filter is blocking the networks from being learned via the Router B interface that is connected to Router A.
- D. IPv6 unicast routing is not enabled on Router A or Router B.

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

show ipv6 traffic Field Descriptions

Field Description

source- Number of source-routed packets.
routed

truncated Number of truncated packets.

format Errors that can result from checks performed on header fields, errors the version number, and packet length.

not a Message sent when IPv6 unicast routing is not enabled.
router

Reference:
http://www.cisco.com/c/en/us/td/docs/ios/ipv6/command/reference/ipv6_book/ipv6_16.html

QUESTION 27

Which type of traffic does DHCP snooping drop?

- A. discover messages
- B. DHCP messages where the source MAC and client MAC do not match
- C. traffic from a trusted DHCP server to client
- D. DHCP messages where the destination MAC and client MAC do not match

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The switch validates DHCP packets received on the untrusted interfaces of VLANs with DHCP snooping enabled. The switch forwards the DHCP packet unless any of the following conditions occur (in which case the packet is dropped):

- The switch receives a packet (such as a DHCP OFFER, DHCP ACK, DHCP NAK, or DHCP REQUEST packet) from a DHCP server outside the network or firewall.

- The switch receives a packet on an untrusted interface, and the source MAC address and the DHCP client hardware address do not match. This check is performed only if the DHCP snooping MAC address verification option is turned on.

- The switch receives a DHCP RELEASE or DHCP DECLINE message from an untrusted host with an entry in the DHCP snooping binding table, and the interface information in the binding table does not match the interface on which the message was received.

- The switch receives a DHCP packet that includes a relay agent IP address that is not 0.0.0.0. To support trusted edge switches that are connected to untrusted aggregation-switch ports, you can enable the DHCP option-82 on untrusted port feature, which enables untrusted aggregation-switch ports to accept DHCP packets that include option-82 information. Configure the port on the edge switch that connects to the aggregation switch as a trusted port.

Reference: <http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/book/snoodhcp.html>

QUESTION 28

Refer to the exhibit.

```
access-list 1 permit 1.0.0.0
0.255.255.255
access-list 2 permit 1.2.3.0
0.0.0.255
!
router rip
```

Which command only announces the 1.2.3.0/24 network out of FastEthernet 0/0?

- A. distribute list 1 out
- B. distribute list 1 out FastEthernet0/0
- C. distribute list 2 out
- D. distribute list 2 out FastEthernet0/0

Correct Answer: D

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Access list 2 is more specific, allowing only 1.2.3.0/24, whereas access list 1 permits all 1.0.0.0/8 networks. This question also asks us to apply this distribute list only to the outbound direction of the fast Ethernet 0/0 interface, so the correct command is distribute list 2 out FastEthernet0/0.

QUESTION 29

Which prefix is matched by the command ip prefix-list name permit 10.8.0.0/16 ge 24 le 24?

- A. 10.9.1.0/24
- B. 10.8.0.0/24
- C. 10.8.0.0/16
- D. 10.8.0.0/23

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

With prefix lists, the ge 24 term means greater than or equal to a /24 and the le 24 means less than or equal to /24, so only a /24 is both greater than or equal to 24 and less than or equal to 24. This translates to any prefix in the 10.8.x.0/24 network, where X is any value in the 0-255 range.

Only the choice of 10.8.0.0.24 matches this.

QUESTION 30

After you review the output of the command show ipv6 interface brief, you see that several IPv6 addresses have the 16-bit hexadecimal value of "FFFE" inserted into the address. Based on this information, what do you conclude about these IPv6 addresses?

- A. IEEE EUI-64 was implemented when assigning IPv6 addresses on the device.
- B. The addresses were misconfigured and will not function as intended.
- C. IPv6 addresses containing "FFFE" indicate that the address is reserved for multicast.
- D. The IPv6 universal/local flag (bit 7) was flipped.
- E. IPv6 unicast forwarding was enabled, but IPv6 Cisco Express Forwarding was disabled.

Correct Answer: A

Section: Layer 3 Technologies

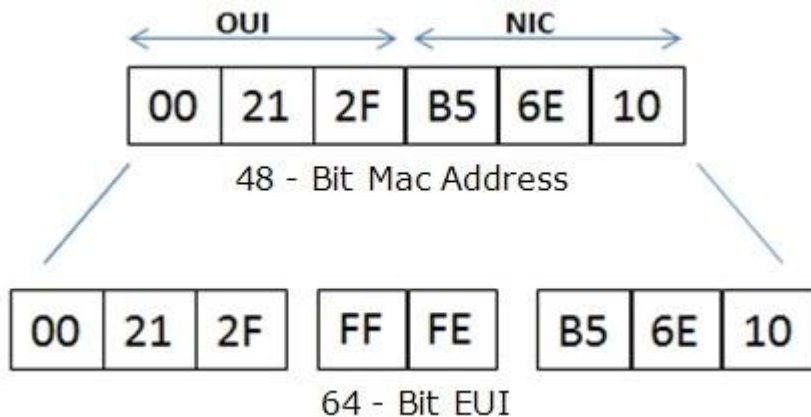
Explanation

Explanation/Reference:


Explanation:

Extended Unique Identifier (EUI), as per RFC2373, allows a host to assign itself a unique 64-Bit IP Version 6 interface identifier (EUI-64). This feature is a key benefit over IPv4 as it eliminates the need of manual configuration or DHCP as in the world of IPv4. The IPv6 EUI-64 format address is obtained through the 48-bit MAC address. The Mac address is first separated into two 24-bits, with one being OUI (Organizationally Unique Identifier) and the other being NIC specific. The 16-bit 0xFFFFE is then inserted between these two 24-bits to form the 64-bit EUI address. IEEE has chosen FFFE as a reserved value which can only appear in EUI-64 generated from an EUI-48 MAC address.

Here is an example showing how a Mac Address is used to generate EUI.



Next, the seventh bit from the left, or the universal/local (U/L) bit, needs to be inverted. This bit identifies whether this interface identifier is universally or locally administered. If 0, the address is locally administered and if 1, the address is globally unique. It is worth noticing that in the OUI portion, the globally unique addresses assigned by the IEEE has always been set to 0 whereas the locally created addresses has 1 configured. Therefore, when the bit is inverted, it maintains its original scope (global unique address is still global unique and vice versa). The reason for inverting can be found in RFC4291 section 2.5.1.



Once the above is done, we have a fully functional EUI-64 format address. Reference: <https://supportforums.cisco.com/document/100566/understanding-ipv6-eui-64-bit-address>

QUESTION 31

A packet capture log indicates that several router solicitation messages were sent from a local host on the IPv6 segment. What is the expected

acknowledgment and its usage?

- A. Router acknowledgment messages will be forwarded upstream, where the DHCP server will allocate addresses to the local host.
- B. Routers on the IPv6 segment will respond with an advertisement that provides an external path from the local subnet, as well as certain data, such as prefix discovery.
- C. Duplicate Address Detection will determine if any other local host is using the same IPv6 address for communication with the IPv6 routers on the segment.
- D. All local host traffic will be redirected to the router with the lowest ICMPv6 signature, which is statically defined by the network administrator.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Router Advertisements (RA) are sent in response to router solicitation messages. Router solicitation messages, which have a value of 133 in the Type field of the ICMP packet header, are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled RA message. Given that router solicitation messages are usually sent by hosts at system startup (the host does not have a configured unicast address), the source address in router solicitation messages is usually the unspecified IPv6 address (0:0:0:0:0:0:0:0). If the host has a configured unicast address, the unicast address of the interface sending the router solicitation message is used as the source address in the message. The destination address in router solicitation messages is the all-routers multicast address with a scope of the link. When an RA is sent in response to a router solicitation, the destination address in the RA message is the unicast address of the source of the router solicitation message. RA messages typically include the following information:

- One or more onlink IPv6 prefixes that nodes on the local link can use to automatically configure their IPv6 addresses

- Lifetime information for each prefix included in the advertisement

- Sets of flags that indicate the type of autoconfiguration (stateless or stateful) that can be completed

- Default router information (whether the router sending the advertisement should be used as a default router and, if so, the amount of time (in seconds) the router should be used as a default router)


- Additional information for hosts, such as the hop limit and MTU a host should use in packets that it originates

Reference:


http://www.cisco.com/c/en/us/td/docs/ios/ipv6/configuration/guide/12_4t/ipv6_12_4t_book/ip6-addrg_bsc_con.html


QUESTION 32


You have been asked to evaluate how EIGRP is functioning in a customer network.

















Traffic from R1 to R6's Loopback address is load shared between R1-R2-R4-R6 and R1-R3-R5-R6 paths. What is the ratio of traffic over each path?

- A. 1:1
- B. 1:5
- C. 6:8
- D. 19:80

Correct Answer: D

Section: Layer 3 Technologies**Explanation****Explanation/Reference:**

Explanation:

First, find the IP address of the loopback0 interface on R6:

R6

```
!  
!  
no ip domain-lookup  
no ipv6 cef  
ipv6 multicast rpf use-bgp  
!  
!  
!  
!  
!  
!  
!  
!  
!  
!  
!  
!  
interface Loopback0  
  ip address 150.1.6.6 255.255.255.255  
!  
interface Loopback1  
  ip address 172.16.6.6 255.255.255.255  
!  
interface Ethernet0/0  
  ip address 192.168.46.6 255.255.255.0
```

We see that it is 150.1.6.6, so we issue the show ip route 150.1.6.6 command from R1 and see this:


```
R1#sh ip route 150.1.6.6
Routing entry for 150.1.6.6/32
  Known via "eigrp 1", distance 90, metric 461056, type internal
  Redistributing via eigrp 1
  Last update from 192.168.13.3 on Ethernet0/1, 00:00:08 ago
  Routing Descriptor Blocks:
  * 192.168.13.3, from 192.168.13.3, 00:00:08 ago, via Ethernet0/1
    Route metric is 1938688, traffic share count is 19
    Total delay is 65730 microseconds, minimum bandwidth is 10000 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 3
  192.168.12.2, from 192.168.12.2, 00:00:08 ago, via Ethernet0/0
    Route metric is 461056, traffic share count is 80
    Total delay is 8010 microseconds, minimum bandwidth is 10000 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 3
```


R1#


Notice the traffic share count shows 19 for the first path, and 80 for the second path.


QUESTION 33


You have been asked to evaluate how EIGRP is functioning in a customer network.
















What is the advertised distance for the 192.168.46.0 network on R1?

- A. 333056
- B. 1938688
- C. 1810944
- D. 307456

Correct Answer: C

Section: Layer 3 Technologies
Explanation

Explanation/Reference:

Explanation:

R1's routing table is as follows

R1

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
+ - replicated route, % - next hop override


Gateway of last resort is not set


```
150.1.0.0/32 is subnetted, 2 subnets
C      150.1.1.1 is directly connected, Loopback0
D      150.1.6.6 [90/1938688] via 192.168.13.3, 00:13:02, Ethernet0/1
192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.12.0/24 is directly connected, Ethernet0/0
L      192.168.12.1/32 is directly connected, Ethernet0/0
192.168.13.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.13.0/24 is directly connected, Ethernet0/1
L      192.168.13.1/32 is directly connected, Ethernet0/1
D      192.168.24.0/24 [90/1862144] via 192.168.13.3, 00:13:02, Ethernet0/1
D      192.168.35.0/24 [90/1785088] via 192.168.13.3, 00:13:08, Ethernet0/1
D      192.168.46.0/24 [90/1810944] via 192.168.13.3, 00:13:02, Ethernet0/1
D      192.168.56.0/24 [90/1810688] via 192.168.13.3, 00:13:03, Ethernet0/1
```


The numbers after the route specify the administrative distance of the route (90 for EIGRP) and the distance metric of that particular route, which is shown as 1810944 for the 192.168.46.0 route.


QUESTION 34

You have been asked to evaluate how EIGRP is functioning in a customer network.

















What type of route filtering is occurring on R6

- A. Distribute-list using an ACL
- B. Distribute-list using a prefix-list
- C. Distribute-list using a route-map
- D. An ACL using a distance of 255

Correct Answer: A

Section: Layer 3 Technologies**Explanation****Explanation/Reference:**

Explanation:


The configuration on R6 is as follows:


```
router eigrp 1
  distribute-list 1 out
  network 150.1.6.6 0.0.0.0
  network 172.16.6.6 0.0.0.0
  network 192.168.46.0
  network 192.168.56.0
!
!
!
no ip http server
!
access-list 1 permit 192.168.46.0
access-list 1 permit 192.168.56.0
access-list 1 permit 150.1.6.6
access-list 1 deny 172.16.6.6
access-list 2 permit 192.168.47.1
access-list 2 permit 192.168.13.1
access-list 2 permit 192.168.12.1
access-list 2 deny 150.1.1.1
!
```


This is a standard distribute list using access list number 1.


QUESTION 35

You have been asked to evaluate how EIGRP is functioning in a customer network.

















What percent of R1's interfaces bandwidth is EIGRP allowed to use?

- A. 10
- B. 20
- C. 30
- D. 40

Correct Answer: B

Section: Layer 3 Technologies**Explanation****Explanation/Reference:**

Explanation:

The relevant configuration of R1 is shown below:

R1

```
!  
interface Ethernet0/0  
  description Link to R2  
  ip address 192.168.12.1 255.255.255.0  
  ip bandwidth-percent eigrp 1 20  
!  
interface Ethernet0/1  
  description Link to R3  
  ip address 192.168.13.1 255.255.255.0  
  ip bandwidth-percent eigrp 1 20  
  delay 5773  
!  
interface Ethernet0/2  
  description Not Currently Used  
  no ip address  
  shutdown  
!  
interface Ethernet0/3  
  description Not Currently Used  
  no ip address  
  shutdown  
!  
!  
router eigrp 1
```


ip bandwidth-percent eigrp 1 20


1 = the EIGRP AS


20 = 20% of the bandwidth


QUESTION 36


Scenario:


You have been asked to evaluate an OSPF network setup in a test lab and to answer questions a customer has about its operation. The customer has disabled your access to the show running- config command.
















How old is the Type 4 LSA from Router 3 for area 1 on the router R5 based on the output you have examined?

- A. 1858
- B. 1601
- C. 600
- D. 1569

Correct Answer: A

Section: Layer 3 Technologies
Explanation

Explanation/Reference:

Explanation:

Part of the show ipospf topology command on R5 shows this:


Link ID	ADV Router	Age	Seq#	Checksum
1.1.1.1	4.4.4.4	600	0x80000002	0x007ED6
2.2.2.2	4.4.4.4	1858	0x80000009	0x004208
3.3.3.3	4.4.4.4	1858	0x80000009	0x00E8FB
4.4.4.4	4.4.4.4	1858	0x80000009	0x00F716
6.6.6.6	4.4.4.4	1601	0x80000009	0x008766
6.6.66.6	4.4.4.4	1601	0x80000009	0x00C7D4
192.168.13.0	4.4.4.4	600	0x80000002	0x006182
192.168.23.0	4.4.4.4	1858	0x80000009	0x00E4ED
192.168.34.0	4.4.4.4	1858	0x80000009	0x004026
192.168.46.0	4.4.4.4	1858	0x80000009	0x00BB9E


R5#


The Link ID of R3 (3.3.3.3) shows the age is 1858.


QUESTION 37


You have been asked to evaluate how EIGRP is functioning in a customer network.

















Which key chain is being used for authentication of EIGRP adjacency between R4 and R2?

- A. CISCO
- B. EIGRP
- C. key
- D. MD5

Correct Answer: A

Section: Layer 3 Technologies**Explanation****Explanation/Reference:**

R4 and R2 configs are as shown below:

R4

```
!  
no ip domain-lookup  
no ipv6 cef  
ipv6 multicast rpf use-bgp  
!  
key chain CISCO  
  key 1  
    key-string firstkey  
!  
!  
!  
!  
!  
!  
!  
!  
!  
!  
!  
interface Loopback0  
  ip address 150.1.4.4 255.255.255.255  
!  
interface Ethernet0/0  
  description Link to R6  
  
  ip address 192.168.46.4 255.255.255.0
```

R2


```
!  
no aaa new-model  
clock timezone PST -8 0  
!  
ip cef  
!  
!  
no ipv6 cef  
ipv6 multicast rpf use-bgp  
!  
key chain CISCO  
  key 1  
    key-string firstkey  
key chain FIRSTKEY  
  key 1  
    key-string CISCO  
key chain R3  
  key 1  
    key-string R3  
  key 2  
    key-string R1  
!  
!  
!  
  
--- More (164) ---
```


Clearly we see the actual key chain is named CISCO.


QUESTION 38


Scenario:

You have been asked to evaluate an OSPF network setup in a test lab and to answer questions a customer has about its operation. The customer has disabled your access to the show running- config command.
















How many times was SPF algorithm executed on R4 for Area 1?

- A. 1
- B. 5
- C. 9
- D. 20
- E. 54
- F. 224

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:


This can be found using the show ipospf command on R4. Look for the Area 1 stats which shows this:


```
Flood list length 0
Area 1
  Number of interfaces in this area is 2 (1 loopback)
  This area has transit capability: Virtual Link Endpoint
  Area has no authentication
  SPF algorithm last executed 04:32:05.765 ago
  SPF algorithm executed 9 times
  Area ranges are
  Number of LSA 15. Checksum Sum 0x05538F
  Number of opaque link LSA 0. Checksum Sum 0x000000
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0
Area 2
  Number of interfaces in this area is 1
  It is a NSSA area
  Perform type-7/type-5 LSA translation
  Area has no authentication
```


QUESTION 39


Scenario:


You have been asked to evaluate an OSPF network setup in a test lab and to answer questions a customer has about its operation. The customer has disabled your access to the show running- config command.
















Which of the following statements is true about the serial links that terminate in R3

- A. The R1-R3 link needs the neighbor command for the adjacency to stay up
- B. The R2-R3 link OSPF timer values are 30, 120, 120
- C. The R1-R3 link OSPF timer values should be 10,40,40
- D. R3 is responsible for flooding LSUs to all the routers on the network.

Correct Answer: B

Section: Layer 3 Technologies**Explanation****Explanation/Reference:**

Explanation:

We can see the configured timers using the following command:

```
R3#show ip ospf interface serial 1/0
Serial1/0 is up, line protocol is up
Internet Address 192.168.13.3/24, Area 0, Attached via Network Statement
Process ID 100, Router ID 3.3.3.3, Network Type NON_BROADCAST, Cost: 1943
Topology-MTID      Cost      Disabled      Shutdown      Topology Name
      0            1943         no            no            Base
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 3.3.3.3, Interface address 192.168.13.3
Backup Designated router (ID) 1.1.1.1, Interface address 192.168.13.1
Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
  oob-resync timeout 120
  Hello due in 00:00:06
Supports Link-local Signaling (LLS)
Cisco NSF helper support enabled
IETF NSF helper support enabled
Index 2/3, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 2, maximum is 11
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 1.1.1.1  (Backup Designated Router)
Suppress hello for 0 neighbor(s)
```

R3#

QUESTION 40

Refer to the following output:

Router#showipnhrp detail

10.1.1.2/8 via 10.2.1.2, Tunnel1 created 00:00:12, expire 01:59:47 TypE.dynamic, Flags: authoritative unique nat registered used NBMA address: 10.12.1.2

What does the authoritative flag mean in regards to the NHRP information?

- A. It was obtained directly from the next-hop server.
- B. Data packets are process switches for this mapping entry.
- C. NHRP mapping is for networks that are local to this router.
- D. The mapping entry was created in response to an NHRP registration request.
- E. The NHRP mapping entry cannot be overwritten.

Correct Answer: A

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Show NHRP: Examples

The following is sample output from the show ipnhrp command:

```
Router# show ipnhrp
```

```
10.0.0.2 255.255.255.255, tunnel 100 created 0:00:43 expire 1:59:16 Type: dynamic Flags: authoritative
```

```

NBMA address: 10.1111.1111.1111.1111.1111.1111.1111.1111.11 10.0.0.1 255.255.255.255, Tunnel0 created 0:10:03 expire 1:49:56 Type: static

```

Flags: authoritative

NBMA address: 10.1.1.2

The fields in the sample display are as follows:

•

The IP address and its network mask in the IP-to-NBMA address cache. The mask is always 255.255.255.255 because Cisco does not support aggregation of NBMA information through NHRP.

•

The interface type and number and how long ago it was created (hours:minutes:seconds).

■

The time in which the positive and negative authoritative NBMA address will expire (hours:minutes:seconds). This value is based on the ipnhrpholdtime command.

•

Type of interface:

dynamic--NBMA address was obtained from the NHRP Request packet.

static--NBMA address was statically configured.

.

Flags:

authoritative--Indicates that the NHRP information was obtained from the Next Hop Server or router that maintains the NBMA-to-IP address mapping for a particular destination.


Reference:


http://www.cisco.com/c/en/us/td/docs/ios/12_4/ip_addr/configuration/guide/hadnhrp.html


QUESTION 41


Scenario:

You have been asked to evaluate an OSPF network setup in a test lab and to answer questions a customer has about its operation. The customer has disabled your access to the show running- config command.















Areas of Router 5 and 6 are not normal areas, inspect their routing tables and determine which statement is true?

- A. R5's Loopback and R6's Loopback are both present in R5's Routing table
- B. R5's Loopback and R6's Loopback are both present in R6's Routing table
- C. Only R5's loopback is present in R5's Routing table
- D. Only R6's loopback is present in R5's Routing table
- E. Only R5's loopback is present in R6's Routing table

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Here are the routing tables of R5 and R6:

R5

```
1.0.0.0/32 is subnetted, 1 subnets
O IA      1.1.1.1 [110/2544] via 192.168.45.4, 00:46:34, Ethernet0/0
2.0.0.0/32 is subnetted, 1 subnets
O IA      2.2.2.2 [110/2544] via 192.168.45.4, 04:57:48, Ethernet0/0
3.0.0.0/32 is subnetted, 1 subnets
O IA      3.3.3.3 [110/601] via 192.168.45.4, 04:57:48, Ethernet0/0
4.0.0.0/32 is subnetted, 1 subnets
O IA      4.4.4.4 [110/301] via 192.168.45.4, 04:57:48, Ethernet0/0
5.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
C          5.5.1.0/24 is directly connected, Loopback1
L          5.5.1.1/32 is directly connected, Loopback1
C          5.5.2.0/24 is directly connected, Loopback2
L          5.5.2.1/32 is directly connected, Loopback2
C          5.5.3.0/24 is directly connected, Loopback3
L          5.5.3.1/32 is directly connected, Loopback3
C          5.5.4.0/24 is directly connected, Loopback4
L          5.5.4.1/32 is directly connected, Loopback4
C          5.5.5.5/32 is directly connected, Loopback0
6.0.0.0/32 is subnetted, 2 subnets
O IA      6.6.6.6 [110/1600] via 192.168.45.4, 04:56:43, Ethernet0/0
O IA      6.6.66.6 [110/601] via 192.168.45.4, 04:56:43, Ethernet0/0
O IA     192.168.13.0/24 [110/2543] via 192.168.45.4, 00:46:44, Ethernet0/0
O IA     192.168.23.0/24 [110/2543] via 192.168.45.4, 04:57:48, Ethernet0/0
O IA     192.168.34.0/24 [110/600] via 192.168.45.4, 04:57:48, Ethernet0/0

192.168.45.0/24 is variably subnetted, 2 subnets, 2 masks
```



```
R6#show ip route
```

```
R6#show ip route
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
       E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```
       ia - IS-IS inter area, * - candidate default, U - per-user static route
```

```
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
```

```
       + - replicated route, % - next hop override
```

```
Gateway of last resort is 192.168.46.4 to network 0.0.0.0
```

```
O*IA  0.0.0.0/0 [110/301] via 192.168.46.4, 05:09:56, Ethernet0/0  
      6.0.0.0/32 is subnetted, 2 subnets
```

```
C      6.6.6.6 is directly connected, Loopback0
```

```
C      6.6.66.6 is directly connected, Loopback1
```

```
      192.168.46.0/24 is variably subnetted, 2 subnets, 2 masks
```

```
C      192.168.46.0/24 is directly connected, Ethernet0/0
```

```
L      192.168.46.6/32 is directly connected, Ethernet0/0
```

```
R6#
```

Topic 4, VPN Technologies

QUESTION 42

A company has just opened two remote branch offices that need to be connected to the corporate network. Which interface configuration output can be applied to the corporate router to allow communication to the remote sites?

- A. interface Tunnel0
bandwidth 1536
ip address 209.165.200.230 255.255.255.224
tunnel source Serial0/0
tunnel mode gre multipoint
- B. interface fa0/0
bandwidth 1536
ip address 209.165.200.230 255.255.255.224
tunnel mode gre multipoint
- C. interface Tunnel0
bandwidth 1536
ip address 209.165.200.231 255.255.255.224
tunnel source 209.165.201.1
tunnel-mode dynamic
- D. interface fa 0/0
bandwidth 1536
ip address 209.165.200.231 255.255.255.224
tunnel source 192.168.161.2
tunnel destination 209.165.201.1
tunnel-mode dynamic

Correct Answer: A

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The configuration of mGRE allows a tunnel to have multiple destinations. The configuration of mGRE on one side of a tunnel does not have any relation to the tunnel properties that might exist at the exit points. This means that an mGRE tunnel on the hub may connect to a p2p tunnel on the branch. Conversely, a p2p GRE tunnel may connect to an mGRE tunnel. The distinguishing feature between an mGRE interface and a p2p GRE interface is the tunnel destination. An mGRE interface does not have a configured destination. Instead the GRE tunnel is configured with the command tunnel mode gre multipoint. This command is used instead of the tunnel destination x.x.x.x found with p2p GRE tunnels. Besides allowing for multiple destinations, an mGRE tunnel requires NHRP to resolve the tunnel endpoints. Note, tunnel interfaces by default are point-to-point (p-p) using GRE encapsulation, effectively they have the tunnel mode gre command, which is not seen in the configuration because it is the default.

The mGRE configuration is as follows:

!

```
interface Tunnel0
bandwidth 1536
ip address 10.62.1.10 255.255.255.0
tunnel source Serial0/0
tunnel mode gre multipoint
```

Reference:

http://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/WAN_and_MAN/DMVPDG/DMVP_N_2_Phase2.html

QUESTION 43

A network engineer executes the show crypto ipseca command. Which three pieces of information are displayed in the output? (Choose three.)

- A. inbound crypto map
- B. remaining key lifetime
- C. path MTU
- D. tagged packets
- E. untagged packets
- F. invalid identity packets

Correct Answer: ABC

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

show crypto ipseca

This command shows IPsec SAs built between peers. The encrypted tunnel is built between 12.1.1.1 and 12.1.1.2 for traffic that goes between networks 20.1.1.0 and 10.1.1.0. You can see the two Encapsulating Security Payload (ESP) SAs built inbound and outbound. Authentication Header (AH) is not used since there are no AH SAs.

This output shows an example of the show crypto ipseca command (bolded ones found in answers for this question).

interface: FastEthernet0

Crypto map tag: test, local addr. 12.1.1.1

local ident (addr/mask/prot/port): (20.1.1.0/255.255.255.0/0/0) remoteident (addr/mask/prot/port): (10.1.1.0/255.255.255.0/0/0) current_peer: 12.1.1.2
PERMIT, flags={origin_is_acl,}

#pktsencaps: 7767918, #pkts encrypt: 7767918, #pkts digest 7767918 #pktsdecaps: 7760382, #pkts decrypt: 7760382, #pkts verify 7760382 #pkts
compressed: 0, #pkts decompressed: 0

#pkts not compressed: 0, #pktscompr. failed: 0,

#pkts decompress failed: 0, #send errors 1, #recv errors 0 local crypto endpt.: 12.1.1.1, remote crypto endpt.: 12.1.1.2 pathmtu 1500, media mtu 1500
current outbound spi: 3D3

inboundespsas:

spi: 0x136A010F(325714191)

transform: esp-3des esp-md5-hmac ,

in use settings ={Tunnel, }

slot: 0, conn id: 3442, flow_id: 1443, crypto map: test sa timing: remaining key lifetime (k/sec): (4608000/52) IV size: 8 bytes
replay detection support: Y
inboundahsas:
inboundpcpsas:
inboundpcpsas:
outboundespsas:
spi: 0x3D3(979)
transform: esp-3des esp-md5-hmac ,
in use settings ={Tunnel, }
slot: 0, conn id: 3443, flow_id: 1444, crypto map: test sa timing: remaining key lifetime (k/sec): (4608000/52) IV size: 8 bytes
replay detection support: Y
outboundahsas:
outboundpcpsas:
Reference: <http://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/5409-ipsec-debug-00.html>

QUESTION 44

Which common issue causes intermittent DMVPN tunnel flaps?

- A. a routing neighbor reachability issue
- B. a suboptimal routing table
- C. interface bandwidth congestion
- D. that the GRE tunnel to hub router is not encrypted

Correct Answer: A

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

DMVPN Tunnel Flaps Intermittently

Problem

DMVPN tunnel flaps intermittently.

Solution

When DMVPN tunnels flap, check the neighborship between the routers as issues with neighborship formation between routers may cause the DMVPN tunnel to flap. In order to resolve this problem, make sure the neighborship between the routers is always up. Reference: <http://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/29240-dcmvpn.html#Prblm1>

QUESTION 45

Which encapsulation supports an interface that is configured for an EVN trunk?

- A. 802.1Q
- B. ISL

- C. PPP
- D. Frame Relay
- E. MPLS
- F. HDLC

Correct Answer: A

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Restrictions for EVN

- An EVN trunk is allowed on any interface that supports 802.1q encapsulation, such as Fast Ethernet, Gigabit Ethernet, and port channels. · A single IP infrastructure can be virtualized to provide up to 32 virtual networks end-to-end. · If an EVN trunk is configured on an interface, you cannot configure VRF-Lite on the same interface.

- OSPFv3 is not supported; OSPFv2 is supported.

Reference: <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/evn/configuration/xr-3s/evn-xr-3s-book/evn-overview.pdf>

QUESTION 46

A user is having issues accessing file shares on a network. The network engineer advises the user to open a web browser, input a prescribed IP address, and follow the instructions. After doing this, the user is able to access company shares. Which type of remote access did the engineer enable?

- A. EZVPN
- B. IPsec VPN client access
- C. VPDN client access
- D. SSL VPN client access

Correct Answer: D

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The Cisco AnyConnect VPN Client provides secure SSL connections to the security appliance for remote users. Without a previously installed client, remote users enter the IP address in their browser of an interface configured to accept SSL VPN connections. Unless the security appliance is configured to redirect http:// requests to https://, users must enter the URL in the form https://<address>.

After entering the URL, the browser connects to that interface and displays the login screen. If the user satisfies the login and authentication, and the security appliance identifies the user as requiring the client, it downloads the client that matches the operating system of the remote computer. After downloading, the client installs and configures itself, establishes a secure SSL connection and either remains or uninstalls itself (depending on the security appliance configuration) when the connection terminates.

Reference: <http://www.cisco.com/c/en/us/support/docs/security/asa-5500-x-series-next-generation-firewalls/100936-asa8x-split-tunnel-anyconnect->

config.html

QUESTION 47

Which three characteristics are shared by subinterfaces and associated EVNs? (Choose three.)

- A. IP address
- B. routing table
- C. forwarding table
- D. access control lists
- E. NetFlow configuration

Correct Answer: ABC

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

A trunk interface can carry traffic for multiple EVNs. To simplify the configuration process, all the subinterfaces and associated EVNs have the same IP address assigned. In other words, the trunk interface is identified by the same IP address in different EVN contexts. This is accomplished as a result of each EVN having a unique routing and forwarding table, thereby enabling support for overlapping IP addresses across multiple EVNs. Reference: <http://www.cisco.com/en/US/docs/ios-xml/ios/evn/configuration/xe-3sg/evn-overview.pdf>

QUESTION 48

Which Cisco IOS VPN technology leverages IPsec, mGRE, dynamic routing protocol, NHRP, and Cisco Express Forwarding?

- A. FlexVPN
- B. DMVPN
- C. GETVPN
- D. Cisco Easy VPN

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

Dynamic Multipoint Virtual Private Network (DMVPN) is a dynamic tunneling form of a virtual private network (VPN) supported on Cisco IOS-based routers and Unix-like Operating Systems based on the standard protocols, GRE, NHRP and IPsec. This DMVPN provides the capability for creating a dynamic-mesh VPN network without having to pre-configure (static) all possible tunnel end-point peers, including IPsec (Internet Protocol Security) and ISAKMP (Internet Security Association and Key Management Protocol) peers. DMVPN is initially configured to build out a hub-and-spoke network by statically configuring the hubs (VPN headends) on the spokes, no change in the configuration on the hub is required to accept new spokes. Using this

initial hub-and-spoke network, tunnels between spokes can be dynamically built on demand (dynamic-mesh) without additional configuration on the hubs or spokes. This dynamic-mesh capability alleviates the need for any load on the hub to route data between the spoke networks.

DMVPN is combination of the following technologies:

Multipoint GRE (mGRE)

Next-Hop Resolution Protocol (NHRP)

Dynamic Routing Protocol (EIGRP, RIP, OSPF, BGP)

Dynamic IPsec encryption

Cisco Express Forwarding (CEF)

Reference: http://en.wikipedia.org/wiki/Dynamic_Multipoint_Virtual_Private_Network Topic 5, Infrastructure Security

Updated reference

QUESTION 49

Which traffic does the following configuration allow?

```
ipv6 access-list cisco
permit ipv6 host 2001:DB8:0:4::32 any eqssh
linevty 0 4
ipv6 access-class cisco in
```

- A. all traffic to vty 0 4 from source 2001:DB8:0:4::32
- B. only ssh traffic to vty 0 4 from source all
- C. only ssh traffic to vty 0 4 from source 2001:DB8:0:4::32
- D. all traffic to vty 0 4 from source all

Correct Answer: C

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Here we see that the IPv6 access list called cisco is being applied to incoming VTY connections to the router. IPv6 access list has just one entry, which allows only the single IPv6 IP address of 2001:DB8:0:4::32 to connect using SSH only.

QUESTION 50

Refer to the following access list.

```
access-list 100 permit ip any any log
```

After applying the access list on a Cisco router, the network engineer notices that the router CPU utilization has risen to 99 percent. What is the reason for this?

- A. A packet that matches access-list with the "log" keyword is Cisco Express Forwarding switched.
- B. A packet that matches access-list with the "log" keyword is fast switched.
- C. A packet that matches access-list with the "log" keyword is process switched.
- D. A large amount of IP traffic is being permitted on the router.

Correct Answer: C

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Logging-enabled access control lists (ACLs) provide insight into traffic as it traverses the network or is dropped by network devices. Unfortunately, ACL logging can be CPU intensive and can negatively affect other functions of the network device. There are two primary factors that contribute to the CPU load increase from ACL logging: process switching of packets that match log-enabled access control entries (ACEs) and the generation and transmission of log messages.

Reference: <http://www.cisco.com/web/about/security/intelligence/acl-logging.html#4>

QUESTION 51

For troubleshooting purposes, which method can you use in combination with the debug ip packet command to limit the amount of output data?

- A. You can disable the IP route cache globally.
- B. You can use the KRON scheduler.
- C. You can use an extended access list.
- D. You can use an IOS parser.
- E. You can use the RITE traffic exporter.

Correct Answer: C

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

The debug ip packet command generates a substantial amount of output and uses a substantial amount of system resources. This command should be used with caution in production networks. Always use with the access-list command to apply an extended ACL to the debug output. Reference: <http://www.cisco.com/c/en/us/support/docs/security/dynamic-multipoint-vpn-dmvpn/111976-dmvpn-troubleshoot-00.html>

QUESTION 52

Which address is used by the Unicast Reverse Path Forwarding protocol to validate a packet against the routing table?

- A. source address

- B. destination address
- C. router interface
- D. default gateway

Correct Answer: A

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

The Unicast RPF feature helps to mitigate problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network by discarding IP packets that lack a verifiable IP source address. For example, a number of common types of denial-of-service (DoS) attacks, including Smurf and Tribal Flood Network (TFN), can take advantage of forged or rapidly changing source IP addresses to allow attackers to thwart efforts to locate or filter the attacks. For Internet service providers (ISPs) that provide public access, Unicast RPF deflects such attacks by forwarding only packets that have source addresses that are valid and consistent with the IP routing table. This action protects the network of the ISP, its customer, and the rest of the Internet.

Reference: http://www.cisco.com/en/US/docs/ios/12_2/security/configuration/guide/scfrpf.html

QUESTION 53

What are the three modes of Unicast Reverse Path Forwarding?

- A. strict mode, loose mode, and VRF mode
- B. strict mode, loose mode, and broadcast mode
- C. strict mode, broadcast mode, and VRF mode
- D. broadcast mode, loose mode, and VRF mode

Correct Answer: A

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

Network administrators can use Unicast Reverse Path Forwarding (Unicast RPF) to help limit the malicious traffic on an enterprise network. This security feature works by enabling a router to verify the reachability of the source address in packets being forwarded. This capability can limit the appearance of spoofed addresses on a network. If the source IP address is not valid, the packet is discarded. Unicast RPF works in one of three different modes: strict mode, loose mode, or VRF mode. Note that not all network devices support all three modes of operation. Unicast RPF in VRF mode will not be covered in this document. When administrators use Unicast RPF in strict mode, the packet must be received on the interface that the router would use to forward the return packet. Unicast RPF configured in strict mode may drop legitimate traffic that is received on an interface that was not the router's choice for sending return traffic. Dropping this legitimate traffic could occur when asymmetric routing paths are present in the network. When administrators use Unicast RPF in loose mode, the source address must appear in the routing table. Administrators can change this behavior using the allow-default option, which allows the use of the default route in the source verification process. Additionally, a packet that contains a source address for which the return route points to the Null 0 interface will be dropped. An access list may also be specified that permits or denies certain

source addresses in Unicast RPF loose mode.

Care must be taken to ensure that the appropriate Unicast RPF mode (loose or strict) is configured during the deployment of this feature because it can drop legitimate traffic. Although asymmetric traffic flows may be of concern when deploying this feature, Unicast RPF loose mode is a scalable option for networks that contain asymmetric routing paths. Reference: <http://www.cisco.com/web/about/security/intelligence/unicast-rpf.html>

QUESTION 54

What does the following access list, which is applied on the external interface FastEthernet 1/0 of the perimeter router, accomplish?

```
router(config)#access-list 101 deny ip 10.0.0.0 0.255.255.255 any log
router (config)#access-list 101 deny ip 192.168.0.0 0.0.255.255 any log
router (config)#interface fastEthernet 1/0
router (config-if)#ip access-group 101 in
```

- A. It prevents incoming traffic from IP address ranges 10.0.0.0-10.0.0.255, 172.16.0.0- 172.31.255.255, 192.168.0.0-192.168.255.255 and logs any intrusion attempts.
- B. It prevents the internal network from being used in spoofed denial of service attacks and logs any exit to the Internet.
- C. It filters incoming traffic from private addresses in order to prevent spoofing and logs any intrusion attempts.
- D. It prevents private internal addresses to be accessed directly from outside.

Correct Answer: C

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

The private IP address ranges defined in RFC 1918 are as follows:

10.0.0.0 - 10.255.255.255

172.16.0.0 - 172.31.255.255

192.168.0.0 - 192.168.255.255

These IP addresses should never be allowed from external networks into a corporate network as they would only be able to reach the network from the outside via routing problems or if the IP addresses were spoofed. This ACL is used to prevent all packets with a spoofed reserved private source IP address to enter the network. The log keyword also enables logging of this intrusion attempt.

QUESTION 55

Refer to the following command:

```
router(config)# ip http secure-port 4433
```

Which statement is true?

- A. The router will listen on port 4433 for HTTPS traffic.
- B. The router will listen on port 4433 for HTTP traffic.
- C. The router will never accept any HTTP and HTTPS traffic.

D. The router will listen to HTTP and HTTP traffic on port 4433.

Correct Answer: A

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

To set the secure HTTP (HTTPS) server port number for listening, use the ip http secure-port command in global configuration mode. To return the HTTPS server port number to the default, use the no form of this command.

iphttpsecure-portport-number

noiphttpsecure-port

Syntax Description

port- Integer in the range of 0 to 65535 is accepted, but the port number must be number higher than 1024 unless the default is used. The default is 443.

Reference: <http://www.cisco.com/en/US/docs/ios-xml/ios/https/command/nm-https-cr-cl-sh.html#wp3612805529>

QUESTION 56

A network engineer is configuring a routed interface to forward broadcasts of UDP 69, 53, and 49 to 172.20.14.225. Which command should be applied to the configuration to allow this?

- A. router(config-if)#ip helper-address 172.20.14.225
- B. router(config-if)#udp helper-address 172.20.14.225
- C. router(config-if)#ipudp helper-address 172.20.14.225
- D. router(config-if)#ip helper-address 172.20.14.225 69 53 49

Correct Answer: A

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

To let a router forward broadcast packet the command ip helper-address can be used. The broadcasts will be forwarded to the unicast address which is specified with the ip helper command.

ip helper-address {ip address}

When configuring the ip helper-address command, the following broadcast packets will be forwarded by the router by default:

- TFTP - UDP port 69
- Domain Name System (DNS) UDP port 53
- Time service - port 37
- NetBIOS Name Server - port 137

- NetBIOS Datagram Server - port 138
- Bootstrap Protocol (BOOTP) - port 67
- TACACS UDP port 49

Reference: http://www.cisco-faq.com/163/forward_udp_broadcasts.html Topic 6, Infrastructure Services

QUESTION 57

When using SNMPv3 with NoAuthNoPriv, which string is matched for authentication?

- A. username
- B. password
- C. community-string
- D. encryption-key

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The following security models exist: SNMPv1, SNMPv2, SNMPv3. The following security levels exist: noAuthNoPriv (no authentication and no encryption - noauth keyword in CLI), AuthNoPriv (messages are authenticated but not encrypted auth keyword in CLI), AuthPriv (messages are authenticated and encrypted priv keyword in CLI). SNMPv1 and SNMPv2 models only support the noAuthNoPriv model since they use plain community string to match the incoming packets. The SNMPv3 implementations could be configured to use either of the models on per-group basis (in case if noAuthNoPriv is configured, username serves as a replacement for community string).

Reference: <http://blog.ine.com/2008/07/19/snmpv3-tutorial/>

QUESTION 58

After a recent DoS attack on a network, senior management asks you to implement better logging functionality on all IOS-based devices. Which two actions can you take to provide enhanced logging results? (Choose two.)

- A. Use the msec option to enable service time stamps.
- B. Increase the logging history
- C. Set the logging severity level to 1.
- D. Specify a logging rate limit.
- E. Disable event logging on all noncritical items.

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The optional msec keyword specifies the date/time format should include milliseconds. This can aid in pinpointing the exact time of events, or to correlate the order that the events happened. To limit syslog messages sent to the router's history table and to an SNMP network management station based on severity, use the logging history command in global configuration mode. By default, Cisco devices Log error messages of severity levels 0 through 4 (emergency, alert, critical, error, and warning levels); in other words, "saving level warnings or higher." By increasing the severity level, more granular monitoring can occur, and SNMP messages will be sent by the less severe (5-7) messages.

QUESTION 59

A network engineer is trying to implement broadcast-based NTP in a network and executes the ntp broadcast client command. Assuming that an NTP server is already set up, what is the result of the command?

- A. It enables receiving NTP broadcasts on the interface where the command was executed.
- B. It enables receiving NTP broadcasts on all interfaces globally.
- C. It enables a device to be an NTP peer to another device.
- D. It enables a device to receive NTP broadcast and unicast packets.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The NTP service can be activated by entering any ntp command. When you use the ntp broadcast client command, the NTP service is activated (if it has not already been activated) and the device is configured to receive NTP broadcast packets on a specified interface simultaneously.

Command Description

ntp broadcast Allows the system to receive NTP broadcast packets on an client interface.

Reference: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/bsm/command/bsm-xe-3se-3850-cr-book/bsm-xe-3se-3850-cr-book_chapter_00.html

QUESTION 60

A network engineer finds that a core router has crashed without warning. In this situation, which feature can the engineer use to create a crash collection?

- A. secure copy protocol
- B. core dumps
- C. warm reloads
- D. SNMP
- E. NetFlow

Correct Answer: B

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

When a router crashes, it is sometimes useful to obtain a full copy of the memory image (called a core dump) to identify the cause of the crash. Core dumps are generally very useful to your technical support representative.

Four basic ways exist for setting up the router to generate a core dump:

- Using Trivial File Transfer Protocol (TFTP)

- Using File Transfer Protocol (FTP)

- Using remote copy protocol (rcp)

- Using a Flash disk

Reference: <http://www.cisco.com/en/US/docs/internetworking/troubleshooting/guide/tr19aa.html>

QUESTION 61

A network engineer is configuring SNMP on network devices to utilize one-way SNMP notifications. However, the engineer is not concerned with authentication or encryption. Which command satisfies the requirements of this scenario?

- A. `router(config)#snmp-server host 172.16.201.28 traps version 2c CISCORO`
- B. `router(config)#snmp-server host 172.16.201.28 informs version 2c CISCORO`
- C. `router(config)#snmp-server host 172.16.201.28 traps version 3 auth CISCORO`
- D. `router(config)#snmp-server host 172.16.201.28 informs version 3 auth CISCORO`

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Most network admins and engineers are familiar with SNMPv2c which has become the dominant SNMP version of the past decade. It's simple to configure on both the router/switch-side and just as easy on the network monitoring server. The problem of course is that the SNMP statistical payload is not encrypted and authentication is passed in cleartext. Most companies have decided that the information being transmitted isn't valuable enough to be worth the extra effort in upgrading to SNMPv3, but I would suggest otherwise. Like IPv4 to IPv6, there are some major changes under the hood. SNMP version 2 uses community strings (think cleartext passwords, no encryption) to authenticate polling and trap delivery. SNMP version 3 moves away from the community string approach in favor of user-based authentication and view-based access control. The users are not actual local user accounts, rather they are simply a means to determine who can authenticate to the device. The view is used to define what the user account may access on the

IOS device. Finally, each user is added to a group, which determines the access policy for its users. Users, groups, views. Reference: <http://www.ccnpguide.com/snmp-version-3/>

QUESTION 62

IPv6 has just been deployed to all of the hosts within a network, but not to the servers. Which feature allows IPv6 devices to communicate with IPv4 servers?

- A. NAT
- B. NATng
- C. NAT64
- D. dual-stack NAT
- E. DNS64

Correct Answer: C

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

NAT64 is a mechanism to allow IPv6 hosts to communicate with IPv4 servers. The NAT64 server is the endpoint for at least one IPv4 address and an IPv6 network segment of 32-bits (for instance 64:ff9b::/96, see RFC 6052, RFC 6146). The IPv6 client embeds the IPv4 address it wishes to communicate with using these bits, and sends its packets to the resulting address. The NAT64 server then creates a NAT-mapping between the IPv6 and the IPv4 address, allowing them to communicate.

Reference: <http://en.wikipedia.org/wiki/NAT64>

QUESTION 63

What is a function of NPTv6?

- A. It interferes with encryption of the full IP payload.
- B. It maintains a per-node state.
- C. It is checksum-neutral.
- D. It rewrites transport layer headers.

Correct Answer: C

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

RFC 6296 describes a stateless IPv6-to-IPv6 Network Prefix Translation (NPTv6) function, designed to provide address independence to the edge network. It is transport-agnostic with respect to transports that do not checksum the IP header, such as SCTP, and to transports that use the TCP/UDP/

DCCP (Datagram Congestion Control Protocol) pseudo-header and checksum NPTv6 provides a simple and compelling solution to meet the address-independence requirement in IPv6. The address-independence benefit stems directly from the translation function of the network prefix translator. To avoid as many of the issues associated with NAT44 as possible, NPTv6 is defined to include a two-way, checksum-neutral, algorithmic translation function, and nothing else.

Reference: <http://tools.ietf.org/html/rfc6296>

QUESTION 64

A network engineer initiates the ipsla responder tcp-connect command in order to gather statistics for performance gauging. Which type of statistics does the engineer see?

- A. connectionless-oriented
- B. service-oriented
- C. connection-oriented
- D. application-oriented

Correct Answer: C

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Configuration Examples for IP SLAs TCP Connect Operations The following example shows how to configure a TCP Connection-oriented operation from Device B to the Telnet port (TCP port 23) of IP Host 1 (IP address 10.0.0.1), as shown in the "TCP Connect Operation" figure in the "Information About the IP SLAs TCP Connect Operation" section. The operation is scheduled to start immediately. In this example, the control protocol is disabled on the source (Device B). IP SLAs uses the control protocol to notify the IP SLAs responder to enable the target port temporarily. This action allows the responder to reply to the TCP Connect operation. In this example, because the target is not a Cisco device and a well-known TCP port is used, there is no need to send the control message.

Device A (target device) Configuration
configure terminal

ipsla responder tcp-connect ipaddress 10.0.0.1 port 23 Reference: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipsla/configuration/15-mt/sla-15-mt-book/sla_tcp_conn.html

QUESTION 65

A network engineer executes the ipv6 flowset command. What is the result?

- A. Flow-label marking in 1280-byte or larger packets is enabled.
- B. Flow-set marking in 1280-byte or larger packets is enabled.
- C. IPv6 PMTU is enabled on the router.
- D. IPv6 flow control is enabled on the router.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Enabling Flow-Label Marking in Packets that Originate from the Device This feature allows the device to track destinations to which the device has sent packets that are 1280 bytes or larger.

SUMMARY STEPS

1. enable
2. configure terminal
3. ipv6flowset
4. exit
5. clear ipv6 mtu

DETAILED STEPS

Command or Action Purpose

Step 1 enable Enables privileged EXEC mode.
Enter your password if prompted.

Example:
Device> enable

Step 2 configure terminal Enters global configuration mode.

Example:
Device# configure
terminal

Step 3 ipv6 flowset Configures flow-label marking in 1280-byte or larger packets sent by the device.

Example:
Device(config)# ipv6
flowset

Reference: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_basic/configuration/15-mt/ipv6b-15-mt-book/ipv6-mtu-path-disc.html

QUESTION 66

A network engineer executes the show ip flow export command. Which line in the output indicates that the send queue is full and export packets are not being sent?

A. output drops

- B. enqueueing for the RP
- C. fragmentation failures
- D. adjacency issues

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Table 5 show ip flow export Field Descriptions

Field Description

Exporting flows to 10.1.1.1 Specifies the export destinations and ports.
(1000) and 10.2.1.1 The ports are in parentheses.

Exporting using source Specifies the source address or interface.
IP address 10.3.1.1

Version 5 flow records Specifies the version of the flow.

11 flows exported in 8 udp The total number of export packets sent, and datagrams the total number of flows contained within them.

0 flows failed due to lack of No memory was available to create an export export packet packet.

0 export packets were sent The packet could not be processed by CEF or up to process level by fast switching, possibly because another feature requires running on the packet.

0 export packets were Indicates that CEF was unable to switch the dropped due to no fib packet or forward it up to the process level.

0 export packets were
dropped due to adjacency
issues

0 export packets were Indicates that the packet was dropped because dropped due to of problems constructing the IP packet.
fragmentation failures

0 export packets were
dropped due to
encapsulation fixup failures

0 export packets were Indicates that there was a problem transferring dropped enqueueing for the the export packet between the RP and the line RP
card.

0 export packets were

dropped due to IPC rate
limiting

0 export packets were Indicates that the send queue was full while dropped due to output the packet was being transmitted.
drops

Reference: http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/oaggnf.html

QUESTION 67

A network engineer is asked to configure a "site-to-site" IPsec VPN tunnel. One of the last things that the engineer does is to configure an access list (access-list 1 permit any) along with the command ipnat inside source list 1 int s0/0 overload. Which functions do the two commands serve in this scenario?

- A. The command access-list 1 defines interesting traffic that is allowed through the tunnel.
- B. The command ipnat inside source list 1 int s0/0 overload disables "many-to-one" access for all devices on a defined segment to share a single IP address upon exiting the external interface.
- C. The command access-list 1 permit any defines only one machine that is allowed through the tunnel.
- D. The command ipnat inside source list 1 int s0/0 overload provides "many-to-one" access for all devices on a defined segment to share a single IP address upon exiting the external interface.

Correct Answer: D

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Configuring NAT to Allow Internal Users to Access the Internet Using Overloading

NAT Router

```
interface ethernet 0
ip address 10.10.10.1 255.255.255.0
ipnat inside
```

!--- Defines Ethernet 0 with an IP address and as a NAT inside interface.

```
interface ethernet 1
ip address 10.10.20.1 255.255.255.0
ipnat inside
```

!--- Defines Ethernet 1 with an IP address and as a NAT inside interface.

```
interface serial 0
ip address 172.16.10.64 255.255.255.0
```

ipnat outside

!--- Defines serial 0 with an IP address and as a NAT outside interface.

ipnat pool ovrlld 172.16.10.1 172.16.10.1 prefix 24

!

!--- Defines a NAT pool named ovrlld with a range of a single IP

!--- address, 172.16.10.1.

ipnat inside source list 7 pool ovrlld overload

!

!

!

!

!--- Indicates that any packets received on the inside interface that

!--- are permitted by access-list 7 has the source address !--- translated to an address out of the NAT pool named ovrlld.

!--- Translations are overloaded, which allows multiple inside

!--- devices to be translated to the same valid IP

address.

access-list 7 permit 10.10.10.0 0.0.0.31

access-list 7 permit 10.10.20.0 0.0.0.31

!--- Access-list 7 permits packets with source addresses ranging from

!--- 10.10.10.0 through 10.10.10.31 and 10.10.20.0

through 10.10.20.31.

Note in the previous second configuration, the NAT pool "ovrlld" only has a range of one address. The keyword overload used in the ipnat inside source list 7 pool ovrlld overload command allows NAT to translate multiple inside devices to the single address in the pool.

Reference:

http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080094e77.shtml

QUESTION 68

A network engineer is configuring a solution to allow failover of HSRP nodes during maintenance windows, as an alternative to powering down the active router and letting the network respond accordingly. Which action will allow for manual switching of HSRP nodes?

- A. Track the up/down state of a loopback interface and shut down this interface during maintenance.
- B. Adjust the HSRP priority without the use of preemption.
- C. Disable and enable all active interfaces on the active HSRP node.
- D. Enable HSRPv2 under global configuration, which allows for maintenance mode.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The standby track command allows you to specify another interface on the router for the HSRP process to monitor in order to alter the HSRP priority for a given group. If the line protocol of the specified interface goes down, the HSRP priority is reduced. This means that another HSRP router with higher priority can become the active router if that router has standby preempt enabled. Loopback interfaces can be tracked, so when this interface is shut down the HSRP priority for that router will be lowered and the other HSRP router will then become the active one.

Reference: <http://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol- hsrp/13780-6.html>

QUESTION 69

A network engineer is notified that several employees are experiencing network performance related issues, and bandwidth-intensive applications are identified as the root cause. In order to identify which specific type of traffic is causing this slowness, information such as the source/destination IP and Layer 4 port numbers is required. Which feature should the engineer use to gather the required information?

- A. SNMP
- B. Cisco IOS EEM
- C. NetFlow
- D. Syslog
- E. WCCP

Correct Answer: C

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

NetFlow Flows Key Fields

A network flow is identified as a unidirectional stream of packets between a given source and destination--both are defined by a network-layer IP address and transport-layer source and destination port numbers. Specifically, a flow is identified as the combination of the following key fields:

Source IP address

Destination IP address

Source Layer 4 port number

Destination Layer 4 port number

Layer 3 protocol type

Type of service (ToS)

Input logical interface

Reference: <http://www.cisco.com/en/US/docs/ios-xml/ios/netflow/configuration/12-4t/cfg-nflow- data-expt.html>

QUESTION 70

An organization decides to implement NetFlow on its network to monitor the fluctuation of traffic that is disrupting core services. After reviewing the

output of NetFlow, the network engineer is unable to see OUT traffic on the interfaces. What can you determine based on this information?

- A. Cisco Express Forwarding has not been configured globally.
- B. NetFlow output has been filtered by default.
- C. Flow Export version 9 is in use.
- D. The command ip flow-capture fragment-offset has been enabled.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

We came across a recent issue where a user setup a router for NetFlow export but was unable to see the OUT traffic for the interfaces in NetFlow Analyzer. Every NetFlow configuration aspect was checked and nothing incorrect was found. That is when we noticed the no ipcef' command on the router. CEF was enabled at the global level and within seconds, NetFlow Analyzer started showing OUT traffic for the interfaces. This is why this topic is about Cisco Express Forwarding.

What is switching?

A Router must make decisions about where to forward the packets passing through. This decision-making process is called switching. Switching is what a router does when it makes the following decisions:

1. Whether to forward or not forward the packets after checking that the destination for the packet is reachable.
2. If the destination is reachable, what is the next hop of the router and which interface will the router use to get to that destination.

What is CEF?

CEF is one of the available switching options for Cisco routers. Based on the routing table, CEF creates its own table, called the Forwarding Information Base (FIB). The FIB is organized differently than the routing table and CEF uses the FIB to decide which interface to send traffic from. CEF offers the following benefits:

1. Better performance than fast-switching (the default) and takes less CPU to perform the same task.
2. When enabled, allows for advanced features like NBAR
3. Overall, CEF can switch traffic faster than route-caching using fast-switching

How to enable CEF?

CEF is disabled by default on all routers except the 7xxx series routers. Enabling and Disabling CEF is easy. To enable CEF, go into global configuration mode and enter the CEF command.

```
Router#config t
```

```
Router(config)#ipcef
```

```
Router(config)#
```

To disable CEF, simply use the no' form of the command, ie. noipcef.

Why CEF Needed when enabling NetFlow ?

CEF is a prerequisite to enable NetFlow on the router interfaces. CEF decides through which interface traffic is exiting the router. Any NetFlow analyzer product will calculate the OUT traffic for an interface based on the Destination Interface value present in the NetFlow packets exported from the router. If the CEF is disabled on the router, the NetFlow packets exported from the router will have Destination interface as null and this leads NetFlow Analyzer to show no OUT traffic for the interfaces. Without enabling the CEF on the router, the NetFlow packets did not mark the destination interfaces and so NetFlow Analyzer was not able to show the OUT traffic for the interfaces.

Reference: <https://blogs.manageengine.com/network-2/netflowanalyzer/2010/05/19/need-for-cef- in-netflow-data-export.html>

QUESTION 71

A company's corporate policy has been updated to require that stateless, 1-to-1, and IPv6 to IPv6 translations at the Internet edge are performed. What is the best solution to ensure compliance with this new policy?

- A. NAT64
- B. NAT44
- C. NATv6
- D. NPTv4
- E. NPTv6

Correct Answer: E

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

NPTv6 provides a mechanism to translate the private internal organization prefixes to public globally reachable addresses. The translation mechanism is stateless and provides a 1:1 relationship between the internal addresses and external addresses. The use cases for NPTv6 outlined in the RFC include peering with partner networks, multi homing, and redundancy and load sharing.

Reference:

http://www.cisco.com/c/dam/en/us/td/docs/solutions/SBA/August2012/Cisco_SBA_BN_IPv6AddressingGuide-Aug2012.pdf

QUESTION 72

A network engineer has left a NetFlow capture enabled over the weekend to gather information regarding excessive bandwidth utilization. The following command is entered:

```
switch#show flow exporter Flow_Exporter-1
```

What is the expected output?

- A. configuration of the specified flow exporter
- B. current status of the specified flow exporter
- C. status and statistics of the specified flow monitor
- D. configuration of the specified flow monitor

Correct Answer: B

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

show flow exporter exporter-name (Optional) Displays the current status of the specified flow exporter.

Example:

```
Device# show flow exporter  
FLOW_EXPORTER-1
```

Reference: <http://www.cisco.com/en/US/docs/ios-xml/ios/fnetflow/configuration/15-mt/cfg-de-fnflow-exprts.html>

QUESTION 73

Which two functions are completely independent when implementing NAT64 over NAT-PT? (Choose two.)

- A. DNS
- B. NAT
- C. port redirection
- D. stateless translation
- E. session handling

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Network Address Translation IPv6 to IPv4, or NAT64, technology facilitates communication between IPv6-only and IPv4-only hosts and networks (whether in a transit, an access, or an edge network). This solution allows both enterprises and ISPs to accelerate IPv6 adoption while simultaneously handling IPv4 address depletion. The DNS64 and NAT64 functions are completely separated, which is essential to the superiority of NAT64 over NAT-PT. Reference: http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/enterprise-ipv6-solution/white_paper_c11-676278.html

QUESTION 74

Which two methods of deployment can you use when implementing NAT64? (Choose two.)

- A. stateless
- B. stateful
- C. manual
- D. automatic
- E. static
- F. functional
- G. dynamic

Correct Answer: AB

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

While stateful and stateless NAT64 perform the task of translating IPv4 packets into IPv6 packets and vice versa, there are important differences. The following table provides a high-level overview of the most relevant differences.

Table 2. Differences Between Stateless NAT64 and Stateful NAT64

Stateless NAT64	Stateful NAT64
-----------------	----------------

1:1 translation	1:N translation
-----------------	-----------------

No conservation of IPv4 address	Conserves IPv4 address
---------------------------------	------------------------

Assures end-to-end address transparency	Uses address overloading, hence lacks transparency and scalability in end-to-end address transparency
---	---

No state or bindings created on the State	or bindings are created on every translation
---	--

Requires IPv4-translatable IPv6 requirement	No requirement on the nature of IPv6 addresses assignment (mandatory address assignment requirement)
---	--

Requires either manual or DHCPv6	Free to choose any mode of IPv6 based address assignment for IPv6 address assignment viz. Manual, hosts DHCPv6, SLAAC
----------------------------------	---

Reference: http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/enterprise-ipv6-solution/white_paper_c11-676277.html

QUESTION 75

Which NetFlow component is applied to an interface and collects information about flows?

- A. flow monitor
- B. flow exporter
- C. flow sampler
- D. flow collector

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

Flow monitors are the NetFlow component that is applied to interfaces to perform network traffic monitoring. Flow monitors consist of a record and a cache. You add the record to the flow monitor after you create the flow monitor. The flow monitor cache is automatically created at the time the flow monitor is applied to the first interface. Flow data is collected from the network traffic during the monitoring process based on the key and nonkey fields in the record, which is configured for the flow monitor and stored in the flow monitor cache.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/fnetflow/command/reference/fnf_book/fnf_01.html#w p1314030

QUESTION 76

What is the result of the command `ip flow-export destination 10.10.10.1 5858`?

- A. It configures the router to export cache flow information to IP 10.10.10.1 on port UDP/5858.
- B. It configures the router to export cache flow information about flows with destination IP 10.10.10.1 and port UDP/5858.
- C. It configures the router to receive cache flow information from IP 10.10.10.1 on port UDP/5858.
- D. It configures the router to receive cache flow information about flows with destination IP 10.10.10.1 and port UDP/5858.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

To enable the exporting of information in NetFlow cache entries, use the `flow-export destination` command in global configuration mode.

Syntax Description

`ip`- IP address of the workstation to which you want to send the address NetFlow information.

`udp-port` UDP protocol-specific port number.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/12s_mdnf.html#wp1023091

QUESTION 77

Refer to the exhibit.

Sampler : mysampler, id : 1, packets matched : 10, mode : random sampling mode

Which statement about the output of the `show flow-sampler` command is true?

- A. The sampler matched 10 packets, each packet randomly chosen from every group of 100 packets.
- B. The sampler matched 10 packets, one packet every 100 packets.
- C. The sampler matched 10 packets, each one randomly chosen from every 100-second interval.
- D. The sampler matched 10 packets, one packet every 100 seconds.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

The sampling mode determines the algorithm that selects a subset of traffic for NetFlow processing. In the random sampling mode that Random Sampled NetFlow uses, incoming packets are randomly selected so that one out of each n sequential packets is selected on average for NetFlow processing. For example, if you set the sampling rate to 1 out of 100 packets, then NetFlow might sample the 5th, 120th, 199th, 302nd, and so on packets. This sample configuration provides NetFlow data on 1 percent of total traffic. The n value is a parameter from 1 to 65535 packets that you can configure.

Table 2 show flow-sampler Field Descriptions

Field Description

Sampler Name of the flow sampler

id Unique ID of the flow sampler

packets matched Number of packets matched for the flow sampler

mode Flow sampling mode

sampling interval is Flow sampling interval (in packets)

Reference:

http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/nfstatsa.html#wp1084291

QUESTION 78

Route.com is a small IT corporation that is attempting to implement the network shown in the exhibit. Currently the implementation is partially completed. OSPF has been configured on routers Chicago and NewYork. The SO/O interface on Chicago and the SO/1 interface on NewYork are in Area 0. The loopbackO interface on NewYork is in Area 1. However, they cannot ping from the serial interface of the Seattle router to the loopback interface of the NewYork router. You have been asked to complete the implementation to allow this ping. ROUTE.com's corporate implementation guidelines require:

- The OSPF process ID for all routers must be 10.
- The routing protocol for each interface must be enabled under the routing process.

- The routing protocol must be enabled for each interface using the most specific wildcard mask possible.
- The serial link between Seattle and Chicago must be in OSPF area 21.
- OSPF area 21 must not receive any inter-area or external routes.

Network Information

Seattle

S0/0 192.168.16.5/30 - Link between Seattle and Chicago

Secret Password: cisco

Chicago

S0/0 192.168.54.9/30 - Link between Chicago and NewYork

S0/1 192.168.16.6/30 - Link between Seattle and Chicago

Secret Password: cisco

NewYork

S0/1 192.168.54.10/30 - Link between Chicago and NewYork


Loopback0 172.16.189.189

Secret Password: cisco

Name : Seattle
S0/0 : 192.168.16.5/30
Secret Password : cisco

Name : Chicago
S0/0 : 192.168.54.9/30
S0/1 : 192.168.16.6/30
Secret Password : cisco

Name : NewYork
S0/1 : 192.168.54.10/30
Loopback0 : 172.16.189.189/32



CiscoTerminal

Seattle con0 is now available

Press RETURN to get started.


Seattle>

CiscoTerminal

Chicago con0 is now available

Press RETURN to get started.

Chicago>



Correct Answer: Answer: Here is the solution below:

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

Note: In actual exam, the IP addressing, OSPF areas and process ID, and router hostnames may change, but the overall solution is the same.

Seattle's S0/0 IP Address is 192.168.16.5/30. So, we need to find the network address and wildcard mask of 192.168.16.5/30 in order to configure the OSPF.

IP Address: 192.168.16.5 /30
Subnet Mask: 255.255.255.252

Here subtract 252 from 256, $256 - 252 = 4$, hence the subnets will increment by 4.

First, find the 4th octet of the Network Address:

Subnet	Network	Broadcast
0	0	3
1	4	7
2	8	11
3	12	15
4	16	19
5

The 4th octet of IP address (192.168.16.5) belongs to subnet 1 (4 to 7).

Network Address: 192.168.16.4
Broadcast Address: 192.168.16.7

Lets find the wildcard mask of /30.

Subnet Mask: (Network Bits – 1's, Host Bits – 0's)

Lets find the wildcard mask of /30.

Subnet Mask: (Network Bits – 1's, Host Bits – 0's)

/30	11111111	11111111	11111111	11111100
	255	255	255	252

Wildcard Mask : (Network Bits – 0's, Host Bits – 1's)

/30	00000000	00000000	00000000	00000011
	0	0	0	3

Now we configure OSPF using process ID 10 (note the process ID may change to something else in real exam).

```
Seattle>enable
Password:
Seattle#conf t
Seattle(config)#router ospf 10
```

```
Seattle(config-router)#network 192.168.16.4 0.0.0.3 area 21
```

One of the tasks states that area 21 should not receive any external or inter-area routes (except the default route).

```
Seattle(config-router)#area 21 stub
Seattle(config-router)#end
Seattle#copy run start
```

Chicago Configuration:

```
Chicago>enable
Password: cisco
Chicago#conf t
Chicago(config)#router ospf 10
```

We need to add Chicago's S0/1 interface to Area 21

```
Chicago(config-router)#network 192.168.16.4 0.0.0.3 area 21
```

Again, area 21 should not receive any external or inter-area routes (except the default route).

In order to accomplish this, we must stop LSA Type 5 if we don't want to send external routes. And if we don't want to send inter-area routes, we have to stop LSA Type 3 and Type 4. Therefore we want to configure area 21 as a totally stubby area.

```
Chicago(config-router)#area 21 stub no-summary
```

```
Chicago(config-router)#end
```

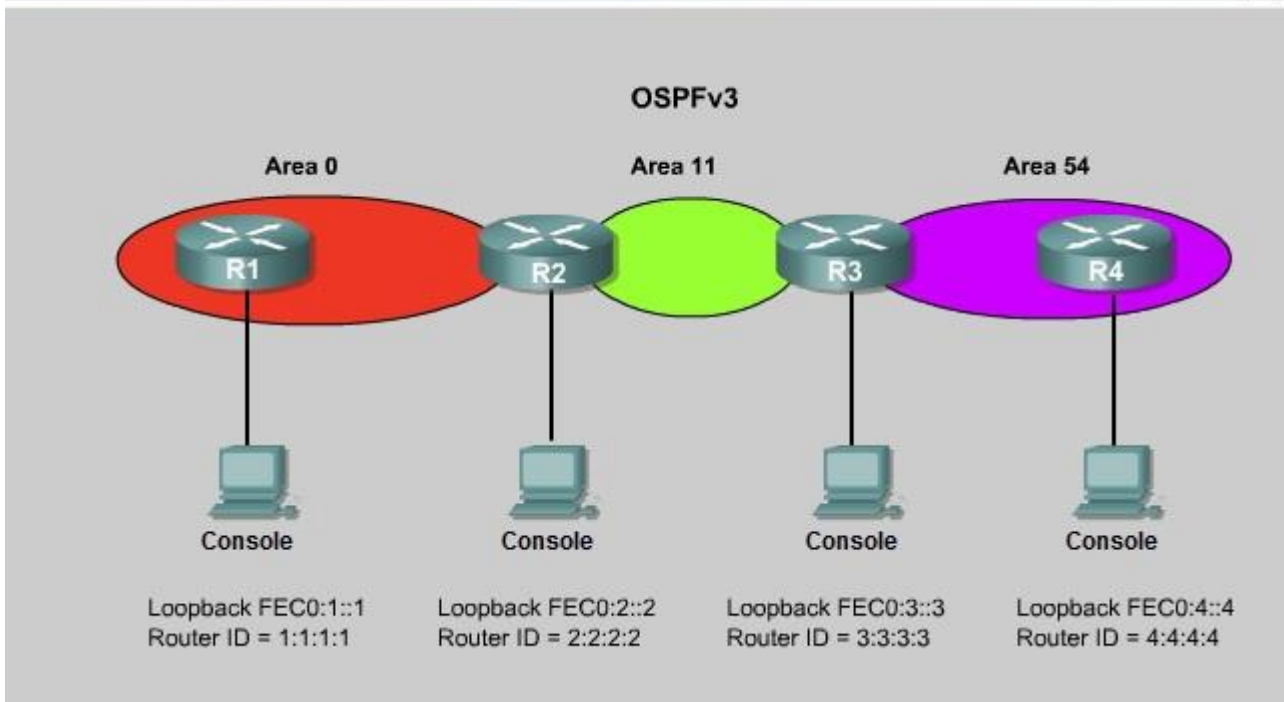
```
Chicago#copy run start
```

The other interface on the Chicago router is already configured correctly in this scenario, as well as the New York router so there is nothing that needs to be done on that router.

QUESTION 79

ROUTE.com is a small IT corporation that has an existing enterprise network that is running IPv6 OSPFv3. Currently OSPF is configured on all routers. However, R4's loopback address (FEC0:4:4) cannot be seen in R1's IPv6 routing table. You are tasked with identifying the cause of this fault and implementing the needed corrective actions that uses OPSF features and does not change the current area assignments. You will know that you have corrected the fault when R4's loopback address (FEC0:4:4) can be seen in RTs IPv6 routing table.

Special Note: To gain the maximum number of points you must remove all incorrect or unneeded configuration statements related to this issue.



```
R1
% Some configuration options may have changed
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to administratively down
*Wed Oct 15 15:22:47.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/1 from FULL to DOWN, Neighbor Down: Interface down or detached
%LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
*Wed Oct 15 15:22:57.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/0 from LOADING to FULL, Loading Done
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to up
*Wed Oct 15 15:22:57.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/0 from LOADING to FULL, Loading Done
Press RETURN to get started!
R1>
```

```
R2
% Some configuration options may have changed
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to administratively down
*Wed Oct 15 15:22:47.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/1 from FULL to DOWN, Neighbor Down: Interface down or detached
%LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
*Wed Oct 15 15:22:57.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/0 from LOADING to FULL, Loading Done
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to up
*Wed Oct 15 15:22:57.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/0 from LOADING to FULL, Loading Done
Press RETURN to get started!
R2>
```

R3

```
% Some configuration options may have changed
*Wed Oct 15 15:22:47.367: %OSPFv3-5-ADJCHG: Process 1, Nbr 4.4.4.4 on OSPFv3_VL0
  from LOADING to FULL, Loading Done
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to administratively down
*Wed Oct 15 15:22:47.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/1
  from FULL to DOWN, Neighbor Down: Interface down or detached
%LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
*Wed Oct 15 15:22:57.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/0
  from LOADING to FULL, Loading Done
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to up
*Wed Oct 15 15:22:57.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/0
  from LOADING to FULL, Loading Done
Press RETURN to get started!
R3>
```



```
R4

% Some configuration options may have changed
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to administratively down
*Wed Oct 15 15:22:47.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/1 from FULL to DOWN, Neighbor Down: Interface down or detached
*Wed Oct 15 15:22:47.367: %OSPFv3-5-ADJCHG: Process 1, Nbr 3.3.3.3 on OSPFv3_VL0 from LOADING to FULL, Loading Done
%LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
*Wed Oct 15 15:22:57.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/0 from LOADING to FULL, Loading Done
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to up
*Wed Oct 15 15:22:57.273: %OSPFv3-5-ADJCHG: Process 1, Nbr 2.2.2.2 on FastEthernet0/0 from LOADING to FULL, Loading Done
Press RETURN to get started!
R4>
```

The output of the "show running-config" command of R3:

```
<output omitted>
!
ipv6 router ospf 1
router-id 3.3.3.3
log-adjacency-changes
area 54 virtual-link 4.4.4.4
!
<output omitted>
```

We knew that all areas in an Open Shortest Path First (OSPF) autonomous system must be physically connected to the backbone area (Area 0). In

some cases, where this is not possible, we can use a virtual link to connect to the backbone through a non-backbone area. The area through which you configure the virtual link is known as a transit area. In this case, the area 11 will become the transit area. Therefore, routers R2 and R3 must be configured with the area <area id> virtual-link <neighbor router-id>command. + Configure virtual link on R2 (from the first output above, we learned that the OSPF process ID of R2 is 1):

```
R2>enable
```

```
R2#configure terminal
```

```
R2(config)#ipv6 router ospf 1
```

```
R2(config-rtr)#area 11 virtual-link 3.3.3.3
```

Save the configuration:

```
R2(config-rtr)#end
```

```
R2#copy running-config startup-config
```

(Notice that we have to use neighbor router-id 3.3.3.3, not R2's router-id 2.2.2.2) + Configure virtual link on R3 (from the second output above, we learned that the OSPF process ID of R3 is 1 and we have to disable the wrong configuration of "area 54 virtual-link 4.4.4.4"):

```
R3>enable
```

```
R3#configure terminal
```

```
R3(config)#ipv6 router ospf 1
```

```
R3(config-rtr)#no area 54 virtual-link 4.4.4.4
```

```
R3(config-rtr)#area 11 virtual-link 2.2.2.2
```

Save the configuration:

```
R3(config-rtr)#end
```

```
R3#copy running-config startup-config
```

You should check the configuration of R4, too. Make sure to remove the incorrect configuration statements to get the full points.

```
R4(config)#ipv6 router ospf 1
```

```
R4(config-router)#no area 54 virtual-link 3.3.3.3
```

```
R4(config-router)#end
```

After finishing the configuration doesn't forget to ping between R1 and R4 to make sure they work.

Note. If you want to check the routing information, use the show ipv6 route command, not "show ip route".

Correct Answer: Answer: Here is the solution below:

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

To troubleshoot the problem, first issue the show running-config on all of 4 routers. Pay more attention to the outputs of routers R2 and R3 The output of the "show running-config" command of R2:

```
<output omitted>
!
ipv6 router ospf 1
router-id 2.2.2.2
log-adjacency-changes
!
<output omitted>
```

The output of the "show running-config" command of R3:

```
<output omitted>
!
ipv6 router ospf 1
router-id 3.3.3.3
log-adjacency-changes
area 54 virtual-link 4.4.4.4
!
<output omitted>
```

We knew that all areas in an Open Shortest Path First (OSPF) autonomous system must be physically connected to the backbone area (Area 0). In some cases, where this is not possible, we can use a virtual link to connect to the backbone through a non-backbone area. The area through which you configure the virtual link is known as a transit area. In this case, the area 11 will become the transit area. Therefore, routers R2 and R3 must be configured with the area <area id> virtual-link <neighbor router-id> command. + Configure virtual link on R2 (from the first output above, we learned that the OSPF process ID of R2 is 1):

R2>enable

R2#configure terminal

R2(config)#ipv6 router ospf 1

R2(config-rtr)#area 11 virtual-link 3.3.3.3

Save the configuration:

R2(config-rtr)#end

R2#copy running-config startup-config

(Notice that we have to use neighbor router-id 3.3.3.3, not R2's router-id 2.2.2.2) + Configure virtual link on R3 (from the second output above, we learned that the OSPF process ID of R3 is 1 and we have to disable the wrong configuration of "area 54 virtual-link 4.4.4.4"):

R3>enable

R3#configure terminal

R3(config)#ipv6 router ospf 1

R3(config-rtr)#no area 54 virtual-link 4.4.4.4

R3(config-rtr)#area 11 virtual-link 2.2.2.2

Save the configuration:

R3(config-rtr)#end

R3#copy running-config startup-config

You should check the configuration of R4, too. Make sure to remove the incorrect configuration statements to get the full points.

```
R4(config)#ipv6 router ospf 1
```

```
R4(config-router)#no area 54 virtual-link 3.3.3.3
```

```
R4(config-router)#end
```

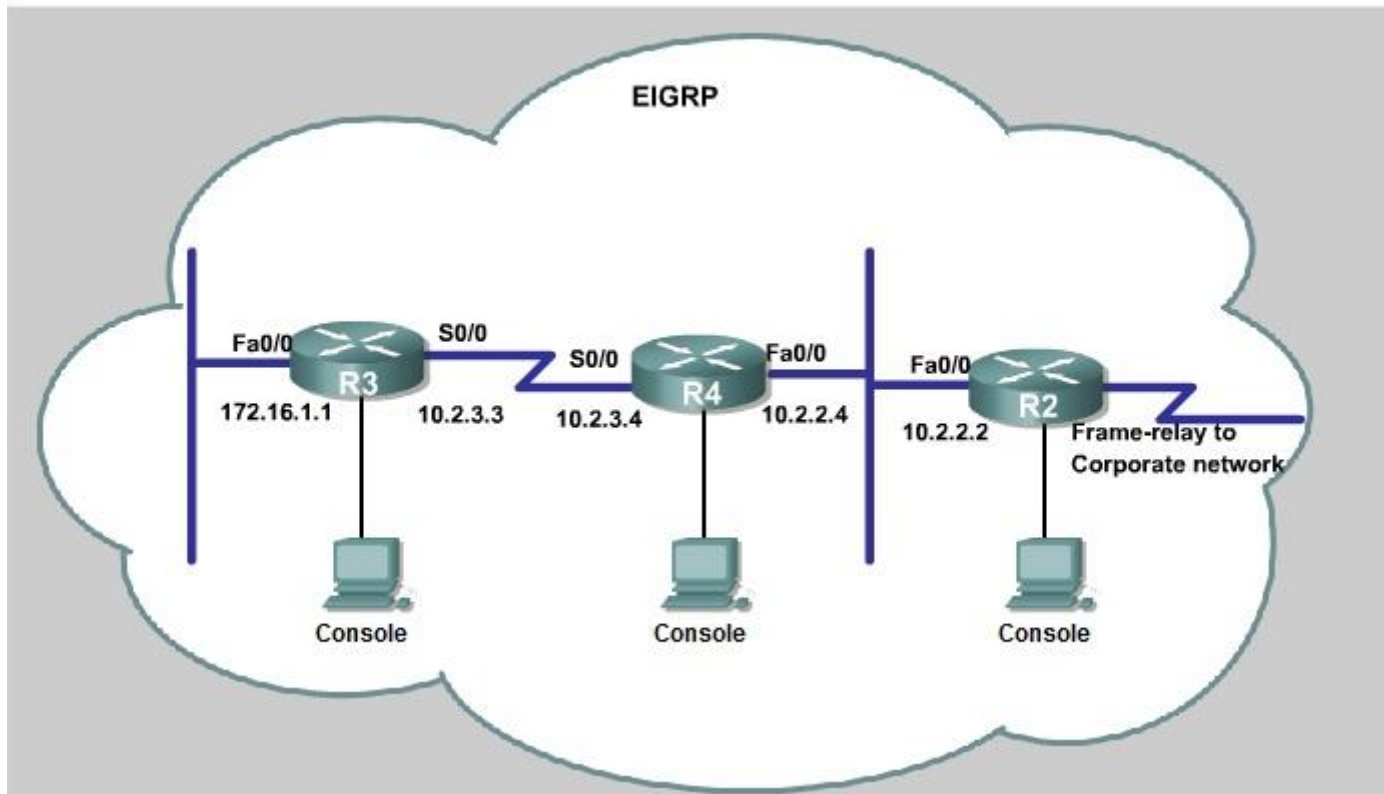
After finishing the configuration doesn't forget to ping between R1 and R4 to make sure they work.

Note. If you want to check the routing information, use the show ipv6 route command, not "show ip route".

QUESTION 80

S Industries has expanded their business with the addition of their first remote office. The remote office router (R3) was previously configured and all corporate subnets were reachable from R3. JS Industries is interested in using route summarization along with the EIGRP Stub Routing feature to increase network stability while reducing the memory usage and bandwidth utilization to R3. Another network professional was tasked with implementing this solution. However, in the process of configuring EIGRP stub routing connectivity with the remote network devices off of R3 has been lost. Currently EIGRP is configured on all routers R2, R3, and R4 in the network. Your task is to identify and resolve the cause of connectivity failure with the remote office router R3. Once the issue has been resolved you should complete the task by configuring route summarization only to the remote office router R3.

You have corrected the fault when pings from R2 to the R3 LAN interface are successful, and the R3 IP routing table only contains 2 10.0.0.0 subnets.



R3

```
% Some configuration options may have changed
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to administratively down
%LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to up
%LINK-3-UPDOWN: Interface Serial0/0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0, changed state to up
Press RETURN to get started!
R3>
```

R4

```
% Some configuration options may have changed
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to administratively down
%LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to up
%LINK-3-UPDOWN: Interface Serial0/0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0, changed state to up
Press RETURN to get started!
R4>
```

```
R2

% Some configuration options may have changed
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to administratively down
%LINK-3-UPDOWN: Interface Serial0/0, changed state to up
%LINK-3-UPDOWN: Interface Serial0/0.1, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0.1, changed state to up
%LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state to up
Press RETURN to get started!
R2>
```

Correct Answer: Answer: Here are the solution as below:

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

First we have to figure out why R3 and R4 can not communicate with each other. Use the show running-config command on router R3.

```
R3#show run
```

```
<output omitted>
```

```
!
```

```
!
```

```
router eigrp 123
```

```
network 10.0.0.0
```

```
network 172.16.0.0
```

```
no auto-summary
```

```
eigrp stub receive-only
```

```
!
```

```
<output omitted>
```

Notice that R3 is configured as a stub receive-only router. The receive-only keyword will restrict the router from sharing any of its routes with any other router in that EIGRP autonomous system. This keyword will also prevent any type of route from being sent. Therefore we will remove this command and replace it with the eigrp stub command:

```
R3# configure terminal R3(config)# router eigrp 123 R3(config-router)# no eigrp stub receive-only R3(config-router)# eigrp stub
```

```
R3(config-router)# end
```

Now R3 will send updates containing its connected and summary routes to other routers. Notice that the eigrp stub command equals to the eigrp stub connected summary because the connected and summary options are enabled by default.

Next we will configure router R3 so that it has only 2 subnets of 10.0.0.0 network. Use the show ip route command on R3 to view its routing table:


```
R3#show ip route
```

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
```

```
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
```

```
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
```

```
* - candidate default, U - per-user static route, o - ODR
```

```
P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
```

```
D 10.2.2.0/24 [90/30720] via 10.2.3.4, 00:00:06, Serial0/0
```

```
C 10.2.3.0/24 is directly connected, Serial0/1
```

```
D 10.2.4.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
```

```
D 10.2.5.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
```

```
D 10.2.6.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
```

```
D 10.2.7.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
```

```
D 10.2.8.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
```

```
D 10.2.9.0/24 [90/161280] via 10.2.3.4, 00:00:03, Serial0/0
```

```
172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
```

```
D 172.16.0.0/16 is a summary, 02:04:06, Null0
```

```
C 172.16.1.0/24 is directly connected, FastEthernet0/0
```

Because we want the routing table of R3 only have 2 subnets so we have to summary sub-networks at the interface which is connected with R3, the s0/0 interface of R4.

There is one interesting thing about the output of the show ip route shown above: the 10.2.3.0/24, which is a directly connected network of R3. We can't get rid of it in the routing table no matter what technique we use to summary the networks. Therefore, to make the routing table of R3 has only 2 subnets we have to summary other subnets into one subnet.

In the output if we don't see the summary line (like 10.0.0.0/8 is a summary...) then we should use the command ip summary-address eigrp 123 10.2.0.0 255.255.0.0 so that all the ping can work well.

In conclusion, we will use the ip summary-address eigrp 123 10.2.0.0 255.255.0.0 at the interface s0/0 of R4 to summary.

```
R4> enable R4# conf t
```

```
R4(config)# interface s0/0 R4(config-if)# ip summary-address eigrp 123 10.2.0.0 255.255.0.0
```

Now we jump back to R3 and use the show ip route command to verify the effect, the output is shown below:

```
R3#show ip route
```

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area

* - candidate default, U - per-user static route, o - ODR

P - periodic downloaded static route

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
D    10.0.0.0/8 is a summary, 00:18:43, Null0
D    10.2.0.0/16 [90/161280] via 10.2.3.4, 00:00:11, Serial0/0
C    10.2.3.0/24 is directly connected, Serial0/1
D    172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D    172.16.0.0/16 is a summary, 02:04:06, Null0
C    172.16.1.0/24 is directly connected, FastEthernet0/0
```

Note: Please notice that the IP addresses and the subnet masks in your real exam might be different so you might use different ones to solve this question.

Just for your information, notice that if you use another network than 10.0.0.0/8 to summary, for example, if you use the command ip summary-address eigrp 123 10.2.0.0 255.255.0.0 you will leave a /16 network in the output of the show ip route command.

```
R3#show ip route
```

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area

* - candidate default, U - per-user static route, o - ODR

P - periodic downloaded static route

Gateway of last resort is not set

```
10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
D    10.0.0.0/8 is a summary, 00:18:43, Null0
D    10.2.0.0/16 [90/161280] via 10.2.3.4, 00:00:11, Serial0/0
C    10.2.3.0/24 is directly connected, Serial0/1
D    172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D    172.16.0.0/16 is a summary, 02:04:06, Null0
C    172.16.1.0/24 is directly connected, FastEthernet0/0
```

But in your real exam, if you don't see the line "10.0.0.0/8 is a summary, Null0" then you can summarize using the network 10.2.0.0/16. This summarization is better because all the pings can work well.

Finally don't forget to use the copy run start command on routers R3 and R4 to save the configurations.

```
R3(config-if)# end
```

```
R3# copy run start
```

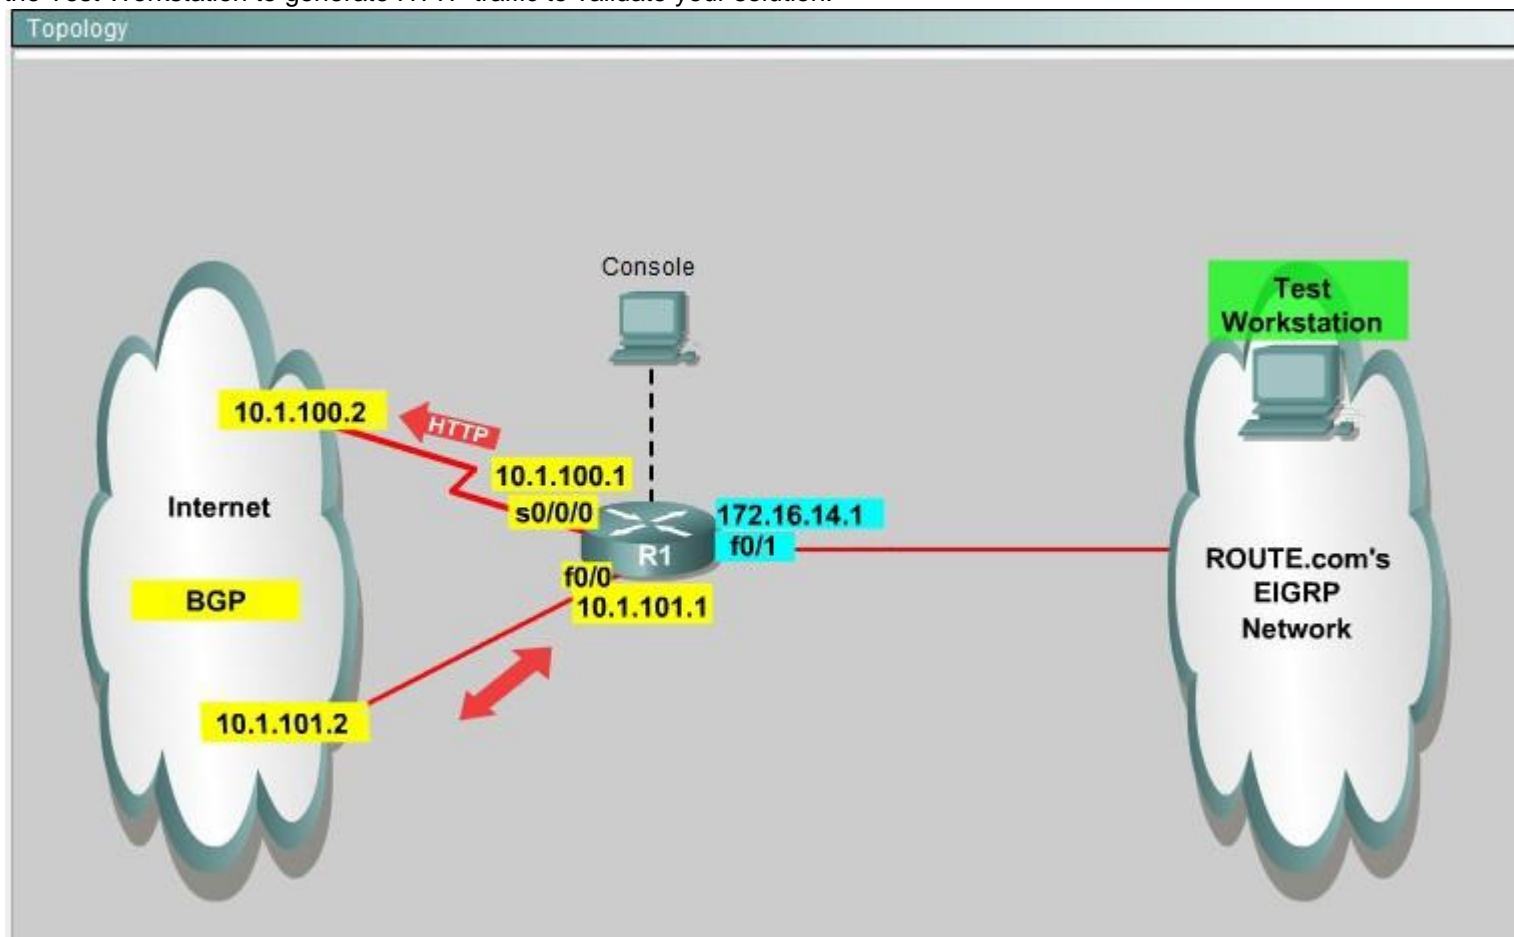
```
R4(config-if)# end
R4# copy run start
```


If the “copy run start” command doesn’t work then use “write memory.”

QUESTION 81

You are a network engineer with ROUTE.com, a small IT company. ROUTE.com has two connections to the Internet; one via a frame relay link and one via an EoMPLS link. IT policy requires that all outbound HTTP traffic use the frame relay link when it is available. All other traffic may use either link. No static or default routing is allowed.

Choose and configure the appropriate path selection feature to accomplish this task. You may use the Test Workstation to generate HTTP traffic to validate your solution.





Correct Answer: Answer: We need to configure policy based routing to send specific traffic along a path that is different from the best path in the routing table.

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Here are the step by Step Solution for this:

1) First create the access list that catches the HTTP traffic:

R1(config)#access-list 101 permit tcp any any eq www

2) Configure the route map that sets the next hop address to be ISP1 and permits the rest of the traffic:

R1(config)#route-map pbr permit 10

R1(config-route-map)#match ip address 101

R1(config-route-map)#set ip next-hop 10.1.100.2

```
R1(config-route-map)#exit
R1(config)#route-map pbr permit 20
3) Apply the route-map on the interface to the server in the EIGRP Network:
R1(config-route-map)#exit
R1(config)#int fa0/1
R1(config-if)#ip policy route-map pbr
R1(config-if)#exit
R1(config)#exit
```

Explanation:

First you need to configure access list to HTTP traffic and then configure that access list. After that configure the route map and then apply it on the interface to the server in EIGRP network.

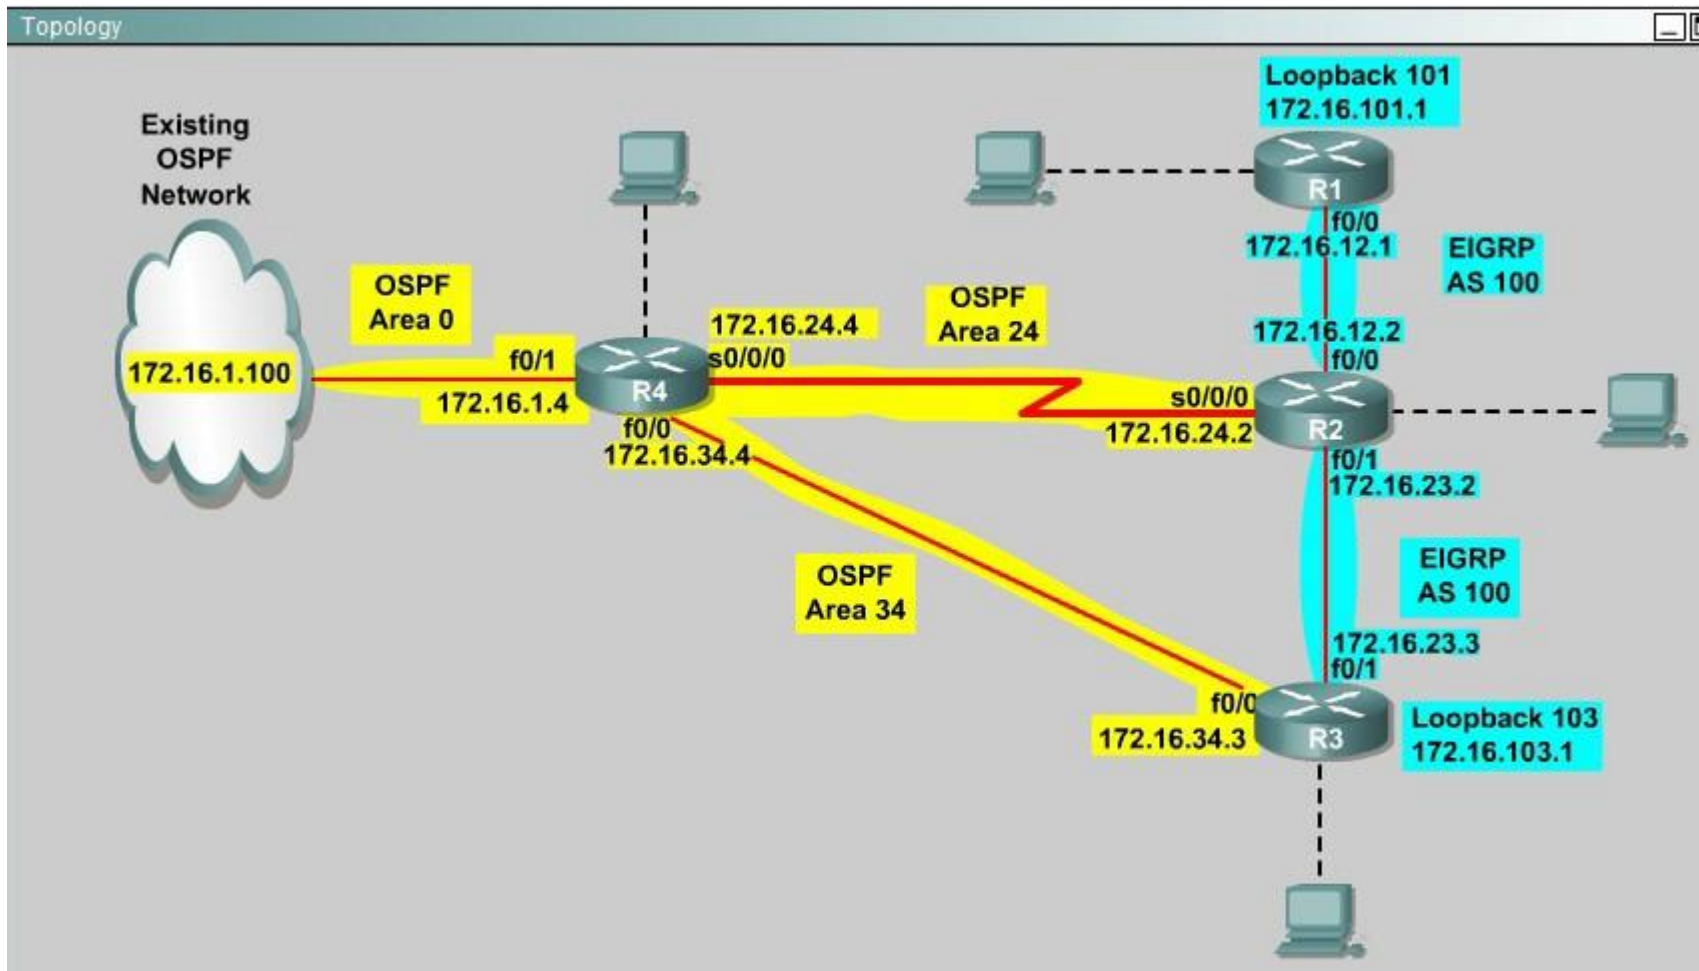
QUESTION 82

You are a network engineer with ROUTE.com, a small IT company. They have recently merged two organizations and now need to merge their networks as shown in the topology exhibit. One network is using OSPF as its IGP and the other is using EIGRP as its IGP. R4 has been added to the existing OSPF network to provide the interconnect between the OSPF and EIGRP networks.

Two links have been added that will provide redundancy.

The network requirements state that you must be able to ping and telnet from loopback 101 on R1 to the OPSF domain test address of 172.16.1.100. All traffic must use the shortest path that provides the greatest bandwidth. The redundant paths from the OSPF network to the EIGRP network must be available in case of a link failure. No static or default routing is allowed in either network.

A previous network engineer has started the merger implementation and has successfully assigned and verified all IP addressing and basic IGP routing. You have been tasked with completing the implementation and ensuring that the network requirements are met. You may not remove or change any of the configuration commands currently on any of the routers. You may add new commands or change default values.



Correct Answer: Answer: First we need to find out 5 parameters (Bandwidth, Delay, Reliability, Load, MTU) of the s0/0/0 interface (the interface of R2 connected to R4) for redistribution:

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

R2#show interface s0/0/0

Write down these 5 parameters, notice that we have to divide the Delay by 10 because the metric unit is in tens of microsecond. For example, we get Bandwidth=1544 Kbit, Delay=20000 us,

Reliability=255, Load=1, MTU=1500 bytes then we would redistribute as follows:

R2#config terminal

R2(config)# router ospf 1

R2(config-router)# redistribute eigrp 100 metric-type 1 subnets

R2(config-router)#exit

R2(config-router)#router eigrp 100

R2(config-router)#redistribute ospf 1 metric 1544 2000 255 1 1500

Note: In fact, these parameters are just used for reference and we can use other parameters with no problem.

If the delay is 20000us then we need to divide it by 10, that is $20000 / 10 = 2000$

For R3 we use the show interface fa0/0 to get 5 parameters too

R3#show interface fa0/0

For example we get Bandwidth=10000 Kbit, Delay=1000 us, Reliability=255, Load=1, MTU=1500 bytes

R3#config terminal

R3(config)#router ospf 1

R3(config-router)#redistribute eigrp 100 metric-type 1 subnets

R3(config)#exit

R3(config-router)#router eigrp 100

R3(config-router)#redistribute ospf 1 metric 10000 100 255 1 1500

Finally you should try to "show ip route" to see the 172.16.100.1 network (the network behind R4) in the routing table of R1 and make a ping from R1 to this network.

Note: If the link between R2 and R3 is FastEthernet link, we must put the command below under EIGRP process to make traffic from R1 to go through R3 (R1 -> R2 -> R3 -> R4), which is better than R1 -> R2 -> R4.

R2(config-router)# distance eigrp 90 105

This command sets the Administrative Distance of all EIGRP internal routes to 90 and all EIGRP external routes to 105, which is smaller than the Administrative Distance of OSPF (110) -> the link between R2 & R3 will be preferred to the serial link between R2 & R4.

Note: The actual OPSF and EIGRP process numbers may change in the actual exam so be sure to use the actual correct values, but the overall solution is the same.

QUESTION 83

After DUAL calculations, a router has identified a successor route, but no routes have qualified as a feasible successor. In the event that the current successor goes down, what process will EIGRP use in the selection of a new successor?

- A. EIGRP will find the interface with the lowest MAC address
- B. The route will transition to the active state
- C. The route will transition to the passive state

D. EIGRP will automatically use the route with the lowest feasible distance(FD)

Correct Answer: B

Section: Mixed Questions

Explanation

Explanation/Reference:

When a route (current successor) goes down, the router first checks its topology table for a feasible successor but it can't find one. So it goes active on the that route to find a new successor by sending queries out to its neighbors requesting a path to the lost route.

QUESTION 84

Which three statements about the EIGRP routing protocol are true? (Choose three)

- A. EIGRP sends periodic hello packets to the multicast IP address 224.0.0.10.
- B. EIGRP will not form a neighbor relationship with another peer when their AS number and K values, either or both are mismatched.
- C. EIGRP will form a neighbor relationship with another peer even when their K values are mismatched.
- D. EIGRP supports five generic packet types, including Hello, Update, Query, Reply, and ACK packets.

Correct Answer: ABD

Section: Mixed Questions

Explanation

Explanation/Reference:

Reference: <http://www.ietf.org/id/draft-savage-eigrp-00.txt> (see eigrp packets)
<http://www.ciscopress.com/articles/article.asp?p=27839>

QUESTION 85

Which two routing protocols require a metric to be configured when redistributing routes from other protocols? (Choose two.)

- A. RIP
- B. IS-IS
- C. OSPF
- D. EIGRP

Correct Answer: AD

Section: Mixed Questions

Explanation

Explanation/Reference:

Metrics must be set manually via configuration when redistributing into RIP and EIGRP, whereas OSPF uses a default value of 20.

QUESTION 86

Which condition must be satisfied before an EIGRP neighbor can be considered a feasible successor?

- A. The neighbor's advertised distance must be less than or equal to the feasible distance of the current successor.
- B. The neighbor's advertised distance must be less than the feasible distance of the current successor.
- C. The neighbor's advertised distance must be greater than the feasible distance of the current successor.
- D. The neighbor's advertised distance must be equal to the feasible distance of the current successor.

Correct Answer: B

Section: Mixed Questions

Explanation

Explanation/Reference:

The feasible successor route is a route which has a higher metric than the successor route to reach a subnet but meets the feasibility condition and can be used in the event that the successor route goes down. This route does NOT get installed in the routing table but is kept in the topology table. The feasibility condition states that the AD from a neighbor must be less than the metric of the successor route (the feasible distance [FD]) because routing through a feasible successor when the AD > FD may cause a routing loop.

Updated Question

QUESTION 87

Your network consists of a large hub-and-spoke Frame Relay network with a CIR of 56 kb/s for each spoke. Which statement about the selection of a dynamic protocol is true?

- A. EIGRP would be appropriate if LMI type ANSI is NOT used.
- B. EIGRP would be appropriate, because the Frame Relay spokes could be segmented into their own areas.
- C. EIGRP would be appropriate, because by default, queries are not propagated across the slow speed Frame Relay links.
- D. EIGRP would be appropriate, because you can manage how much bandwidth is consumed over the Frame Relay interface.

Correct Answer: D

Section: Mixed Questions

Explanation

Explanation/Reference:

By default, EIGRP will limit itself to using no more than 50% of the interface bandwidth. The primary benefit of controlling EIGRP's bandwidth usage is to avoid losing EIGRP packets, which could occur when EIGRP generates data faster than the interface line can absorb it. This is of particular benefit on Frame Relay networks, where the access interface bandwidth and the PVC capacity may be very different.

QUESTION 88

You have implemented mutual route redistribution between OSPF and EIGRP on a border router. When checking the routing table on one of the OSPF routers within the OSPF routing domain, you are seeing some, but not all of the expected routes.

Which two things should you verify to troubleshoot this problem? (Choose two.)

- A. The border router is using a proper seed metric for EIGRP.
- B. The administrative distance is set for OSPF and EIGRP.
- C. The missing EIGRP routes are present in the routing table of the border router.
- D. The subnet keyword on the border router in the redistribute EIGRP command.

Correct Answer: CD

Section: Mixed Questions

Explanation

Explanation/Reference:

Reviewed and updated

QUESTION 89

Which command will display EIGRP packets sent and received, as well as statistics on hello packets, updates, queries, replies, and acknowledgments?

- A. debug eigrp packets
- B. show ip eigrp traffic
- C. debug ip eigrp
- D. show ip eigrp interfaces

Correct Answer: B

Section: Infrastructure Services

Explanation

Explanation/Reference:

The show ip eigrp traffic command displays the number of Enhanced IGRP (EIGRP) packets sent and received.

Example:

The following is sample output from the show ip eigrp traffic command.

Router# show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 77

Hellos sent/received. 218/205

Updates sent/received. 7/23

Queries sent/received. 2/0

Replies sent/received. 0/2

Acks sent/received. 21/14

Reference:http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_command_reference_chapter09186a00800ca5a9.html#wp1018815

Answer Corrected and explained.

QUESTION 90

Which of the below mentioned conditions form a neighbor relationship in EIGRP? (Choose three)

- A. Hello or ACK received
- B. AS number match
- C. Hello timer match
- D. Identical metric (k values)

Correct Answer: ABD

Section: Mixed Questions


Explanation

Explanation/Reference:

Fixed the answers.

QUESTION 91

Refer to the exhibit.



Which statement is correct regarding the operation of NAT-PT between the IPv4 and IPv6 networks shown?

- A. The router will determine the IPv4 destination address.
- B. The source IPv6 host can use DNS to determine the IPv6-to-IPv4 address mapping.
- C. The host is statically configured with the IPv6-to-IPv4 address mapping.
- D. ICMP can be used to determine the IPv6-to-IPv4 address mapping.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 92

Router R1, a branch router, connects to the Internet using DSL. Some traffic flows through a GRE and IPsec tunnel, over the DSL connection, and into the core of an Enterprise network. The branch also allows local hosts to communicate directly with public sites in the Internet over this same DSL connection. Which of the following answers defines how the branch NAT config avoids performing NAT for the Enterprise directed traffic but does perform NAT for the Internet-directed traffic?

- A. By not enabling NAT on the IPsec tunnel interface
- B. By not enabling NAT on the GRE tunnel interface
- C. By configuring the NAT-referenced ACL to not permit the Enterprise traffic
- D. By asking the ISP to perform NAT in the cloud

Correct Answer: C

Section: Network Principles


Explanation

Explanation/Reference:

The NAT configuration acts only on packets permitted by a referenced ACL. As a result, the ACL can permit packets destined for the Internet, performing NAT on those packets. The ACL also denies packets going to the Enterprise, meaning that the router does not apply NAT to those packets.

QUESTION 93

Refer to the exhibit.



Router B1 Configuration

```
crypto ipsec transform-set 10 esp-sha-hmac esp-3des
crypto map tunnel 1 ipsec-isakmp
 set transform-set 10
 set peer 192.168.2.1
 match address 102
crypto isakmp policy 1
 authentication pre-share
 group 2
crypto isakmp key ***** address 192.168.2.1
access-list 102 permit esp host 192.168.1.1 host
192.168.2.1
access-list 102 permit udp host 192.168.1.1 eq isakmp
 host 192.168.2.1
deny ip any any log
```

Router B1 Configuration (con't)

```
Interface f0/0
 Ip address 192.168.1.1 255.255.255.0
interface Tunnel0
 ip address 172.16.1.1 255.255.255.0
 crypto map tunnel
 tunnel source F0/0
 tunnel destination 192.168.2.1
 tunnel path-mtu-discovery
 ip ospf mtu-ignore
router ospf 200
 network 10.10.1.1 0.0.0.224 area 1
 network 172.16.10.1 0.0.0.240 area 1
 network 192.168.1.0 0.0.0.255 area 0
```

A new TAC engineer came to you for advice. A GRE over IPsec tunnel was configured, but the tunnel is not coming up. What did the TAC engineer configure incorrectly?

- A. The crypto map is not configured correctly.
- B. The crypto ACL is not configured correctly.

- C. The crypto map is not applied to the correct interface.
- D. The OSPF network is not configured correctly.

Correct Answer: B

Section: Mixed Questions

Explanation

Explanation/Reference:

QUESTION 94

Drag drop

Select and Place:

Click and drag the associated EIGRP functionality on the left to the corresponding topology characteristic on the right.	
redistribution	low-speed WAN links
bandwidth management	WAN link to an external supplier
authentication	integrating two merging companies
stubs	256 kb/s CIR FR hub and spokes

Correct Answer:

Click and drag the associated EIGRP functionality on the left to the corresponding topology characteristic on the right.

bandwidth management

authentication

redistribution

stubs

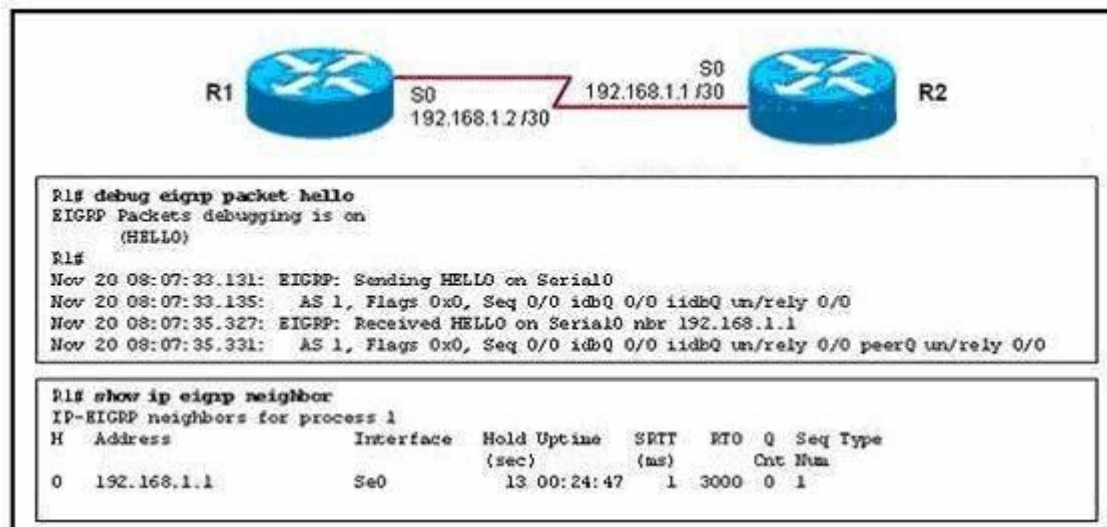
Section: Network Principles

Explanation

Explanation/Reference:

QUESTION 95

Refer to the exhibit.



Routers R1 and R2 have established a neighbor relationship and are exchanging routing information. The network design requires that R1 receive routing updates from R2, but not advertise any routes to R2. Which configuration command sequence will successfully accomplish this task?

- A. R1(config)# router eigrp 1
R1(config-router)# passive-interface serial 0
- B. R2(config)# router eigrp 1
R2(config-router)# passive-interface serial 0
- C. R1(config)# access-list 20 deny any
R1(config)# router eigrp 1
R1(config-router)# distribute-list 20 out serial 0
- D. R2(config)# access-list 20 deny any
R2(config)# router eigrp 1
R2(config-router)# distribute-list 20 out serial 0

Correct Answer: C

Section: Network Principles

Explanation

Explanation/Reference:

We can not use passive-interfaces to accomplish this task because the “passive-interface...” command (in EIGRP or OSPF) will shut down the neighbor relationship of these two routers (no hello packets are exchanged). And to filter routing updates we should configure a distribute list on R1 with an access list that deny all and apply it to the outbound direction so that R1 can receive but cannot send routing updates.

QUESTION 96

EIGRP has been configured to operate over Frame Relay multipoint connections. What should the bandwidth command be set to?

- A. the CIR rate of the lowest speed connection multiplied by the number of circuits
- B. the CIR rate of the lowest speed connection
- C. the CIR rate of the highest speed connection
- D. the sum of all the CIRs divided by the number of connections

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

If the multipoint network has different speeds allocated to the VCs, take the lowest CIR and simply multiply it by the number of circuits. This is because in Frame-relay all neighbors share the bandwidth equally, regardless of the actual CIR of each individual PVC, so we have to get the lowest speed CIR rate and multiply it by the number of circuits. This result will be applied on the main interface (or multipoint connection interface).

QUESTION 97

Refer to the exhibit.

```
R1# show ip eigrp topology
<output omitted>
P 10.1.2.0/24, 1 successors, FD is 281600
   via Connected, FastEthernet0/0
A 10.6.1.0/24, 0 successors, FD is 3385160704, Q
   1 replies, active 00:00:41, query-origin: Local origin
   Remaining replies:
     via 10.1.2.1, r. FastEthernet0/0
```

EIGRP is configured on all routes in the network. On a basis of the show ip eigrp topology output provided, what conclusion can be derived?

- A. Router R1 can send traffic destined for network 10.6.1.0/24 out of interface FastEthernet0/0
- B. Router R1 is waiting for a reply from the neighbor 10.1.2.1 to the hello message sent out before it declares the neighbor unreachable
- C. Router R1 is waiting for a reply from the neighbor 10.1.2.1 to the hello message sent out inquiring for a second successor to network 10.6.1.0/24
- D. Router R1 is waiting for a reply from the neighbor 10.1.2.1 in response to the query sent about network 10.6.1.0/24

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

The “show ip eigrp topology” command lists all routes that EIGRP is aware of and shows whether EIGRP is actively processing information on that route. Under most normal conditions, the routes should all be in a passive state and no EIGRP process are running for that route. If the routes are active, this could indicate the dreaded stuck in active, or SIA, state.

The fields to note in this output are as follows:

P— Passive; no EIGRP computation is being performed. This is the ideal state.

A— Active; EIGRP computations are "actively" being performed for this destination. Routes constantly appearing in an active state indicate a neighbor or query problem. Both are symptoms of the SIA problem.

U— Update; an update packet was sent to this destination.


Q— Query; a query packet was sent to this destination.

R— Reply; a reply packet was sent to this destination.

Route information— IP address of the route or network, its subnet mask, and the successor, or next hop to that network, or the feasible successor.

QUESTION 98

Refer to the exhibit.



EIGRP has been configured on all routers in the network. What additional configuration statement should be included on router R4 to advertise a default route to its neighbors?

- A. R4(config)# ip default-network 10.0.0.0
- B. R4(config)# ip route 0.0.0.0 0.0.0.0 10.1.1.1
- C. R4(config)# ip route 10.0.0.0 255.0.0.0 10.1.1.1
- D. R4(config-router)# default-information originate

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

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.

Gateways of last resort selected using the ip default-network command are propagated differently depending on which routing protocol is propagating the default route. For IGRP and EIGRP to propagate the route, the network specified by the ip default-network command must be known to IGRP or EIGRP. This means the network must be an IGRP- or EIGRP-derived network in the routing table, or the static route used to generate the route to the network must be redistributed into IGRP or EIGRP, or advertised into these protocols using the network command. In this case, the 10.0.0.0 network is indeed being advertised via EIGRP.

Reference: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094374.shtml#ipnetwork

QUESTION 99

Which two methods use IPsec to provide secure connectivity from the branch office to the headquarters office? (Choose two.)

- A. DMVPN
- B. MPLS VPN
- C. Virtual Tunnel Interface (VTI)
- D. SSL VPN

Correct Answer: AC

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Reference: <http://www.ccnpguide.com/ccnp-route-642-902-vpns-and-ipsec/>

QUESTION 100

What is the purpose of configuring the router as a PPPoE client?

- A. to provide VPN access over L2TP
- B. to enable PPP session from the router to the termination device at the headend for metro Ethernet connectivity
- C. for DSL connectivity and removing the need for the end-user PC to run the PPPoE client software
- D. for connecting the router to a cable modem, which bridges the Ethernet frames from the router to the cable modem termination system

Correct Answer: C

Section: Network Principles

Explanation

Explanation/Reference:

DSL Technology used PPPoE protocol (service provide end) and user end required to be used same Protcol running as client to communicate with it

QUESTION 101

What is the international standard for transmitting data over a cable system?

- A. PPPoE
- B. DOCSIS
- C. CMTS
- D. AAL5

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

http://www.cablelabs.com/news/pr/1998/1998_03_19.html (see first para)

QUESTION 102

What is a key benefit of using a GRE tunnel to provide connectivity between branch offices and headquarters?

- A. authentication, integrity checking, and confidentiality
- B. less overhead
- C. dynamic routing over the tunnel
- D. granular QoS support

Correct Answer: C

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Generic routing encapsulation. Tunneling protocol developed by Cisco that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to Cisco routers at remote points over an IP internetwork.

QUESTION 103

Which DSL encapsulation method requires client software running on the end-user PC that is directly connected to a DSL modem?

- A. PPPoA
- B. PPPoE
- C. PPP
- D. L2TP

Correct Answer: B

Section: Network Principles

Explanation

Explanation/Reference:

<http://www.vicomsoft.com/learning-center/pppoe/>

QUESTION 104

The Cisco SA 500 Series Security Appliances are built specifically for businesses with less than 100 employees. What are three important benefits of this device? (Choose three)

- A. business-grade firewall

- B. premium support via SMART net
- C. site-to-site VPN for remote offices
- D. email security

Correct Answer: ACD

Section: Mixed Questions

Explanation

Explanation/Reference:

http://www.cisco.com/en/US/prod/collateral/vpndevc/ps6032/ps6094/ps9932/at_a_glance_c45-562587.pdf (Page 1, see key features and benefits)

QUESTION 105

For a GRE tunnel to be up between two routers, which of the following must be configured?

- A. Loopback Interface
- B. IP reachability between the loopback interfaces
- C. Dynamic Routing between routers.
- D. Tunnel interfaces must be in the same subnet.

Correct Answer: D

Section: Mixed Questions


Explanation

Explanation/Reference:

GRE tunnels don't require loopbacks. They work quite well using physical interfaces as the source and destination. They also don't require dynamic routing; static routes work just fine. Place the IP addresses assigned to the tunnels in different subnets and there won't be any connectivity over the tunnels... that is unless you place static routes at both endpoints pointing to the remote tunnel IP address via the tunnel. Host routes work just fine.

QUESTION 106

Refer to the exhibit.



Router RTA is the hub router for routers RTB and RTC. The Frame Relay network is configured with EIGRP, and the entire network is in autonomous system 1. However, router RTB and RTC are not receiving each other's routes. What is the solution?

- A. Configure the auto summary command under router eigrp 1 on router RTA.
- B. Issue the no ip split horizon command on router RTA.
- C. Configure subinterfaces on the spoke routers and assign different IP address subnets for each subinterface.
- D. Check and change the access lists on router RTA.
- E. Issue the no ip split horizon eigrp 1 command on router RTA.
- F. Configure a distribute list on router RTA that allows it to advertise all routes to the spoke routers.

Correct Answer: E

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Split horizon controls the sending of EIGRP update and query packets. When split horizon is enabled on an interface, these packets are not sent for destinations for which this interface is the next hop. This reduces the possibility of routing loops.

By default, split horizon is enabled on all interfaces. Split horizon blocks route information from being advertised by a router out of any interface from which that information originated. This behavior usually optimizes communications among multiple routing devices, particularly when links are broken. However, with nonbroadcast networks (such as Frame Relay and SMDs), situations can arise for which this behavior is less than ideal. For these situations, you may want to disable split horizon. In this example, routes received by RTB and RTC are not being sent back out the same serial interface on RTA, so they are not receiving each other's routes. Disabling Split horizons on interface S0/0 on RTA will fix this issue.

QUESTION 107

When troubleshooting an EIGRP connectivity problem, you notice that two connected EIGRP routers are not becoming EIGRP neighbors. A ping between the two routers was successful.

What is the next thing that should be checked?

- A. Verify that the EIGRP hello and hold timers match exactly.
- B. Verify that EIGRP broadcast packets are not being dropped between the two routers with the show ip EIGRP peer command.
- C. Verify that EIGRP broadcast packets are not being dropped between the two routers with the show ip EIGRP traffic command.
- D. Verify that EIGRP is enabled for the appropriate networks on the local and neighboring router.

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

The point of this question is about the condition of establish EIGRP neighbor. You can use these ways to troubleshoot the EIGRP connectivity problem.


1. Whether EIGRP is enabled for the proper networks.
2. Whether the K values of EIGRP neighbors is the same.
3. Whether EIGRP autonomous number is the same.

Incorrect answers:

- *. EIGRP use multicast, not broadcast.
- *. EIGRP use multicast, not broadcast.
- *. Hello and hold timers match is the condition of establish OSPF neighbor, not EIGRP.

QUESTION 108

Refer to the exhibit.



You are the network administrator of the Route.com company. You have been tasked to implement a hub and spoke EIGRP topology over Frame Relay to provide connectivity between the networks at headquarters and all 300 spokes. Before you begin the actual implementation, which three pieces of information are more important to know than the others? (Choose three.)

- A. the Committed Information Rate of all the Frame Relay PVCs
- B. the Cisco IOS version running on all the routers
- C. the router model number of all the spoke routers
- D. the number of HQ networks connected behind the headquarter routers
- E. the routing policy, such as whether or not the spokes can be used as backup transient point between the two headquarter routers

Correct Answer: ABE

Section: Mixed Questions

Explanation


Explanation/Reference:

Explanation: You have to know the committed information rate because it is a bandwidth associated with logical connection in a PVC. You also need to know the IOS version on all routers so that there is no conflict in versions. As per the topology, you need to know the routing policy because it will be

used as a backup transient point between headquarter routers

QUESTION 109

Refer to the exhibit.



The Route.com company is running EIGRP between all the routers. Currently, if one of the LAN links (LAN1 or LAN2) at the headquarters flaps (goes up and down), the HQ-RTR1 and HQ-RTR2 routers will experience high CPU usage and have a long EIGRP convergence time. As the new network administrator, you are asked to investigate this situation and determine if there is a quick way to resolve this issue. Which is the most important thing that you can quickly verify first to resolve this issue?

- A. Verify that the bandwidth setting on all WAN links is correct.
- B. Verify that the HQ-RTR1 and HQ-RTR2 routers are configured to send only a default route to all the spoke routers.
- C. Verify that the HQ-RTR1 and HQ-RTR2 routers are configured for EIGRP Nonstop Forwarding.
- D. Verify that all the spoke routers are configured for auto summarization.
- E. Verify that all the spoke routers are configured as EIGRP stub.

Correct Answer: E

Section: Network Principles

Explanation

Explanation/Reference:

Explanation: Stub routing is commonly used in a hub and spoke network topology. In a hub and spoke network, one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is adjacent only to one or more distribution routers. The only route for IP traffic to follow into the remote router is through a distribution router. This type of configuration is commonly used in WAN topologies where the distribution router is directly connected to a WAN. The distribution router can be connected to many more remote routers. Often, the distribution router will be connected to 100 or more remote routers. In a hub and spoke topology, the remote router must forward all nonlocal traffic to a distribution router, so it becomes unnecessary for the remote router to hold a complete routing table. Generally, the distribution router need not send anything more than a default route to the remote router.


When using the EIGRP Stub Routing feature, you need to configure the distribution and remote routers to use EIGRP, and to configure only the remote router as a stub. Only specified routes are propagated from the remote (stub) router. The router responds to queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message "inaccessible." A router that is configured as a stub will send a special peer information packet to all neighboring routers to report its status as a stub router.

Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router will depend on the distribution router to send the proper updates to all peers.

Reference: http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/eigrpstb.html

QUESTION 110

Refer to the exhibit.



When you examine the routing tables of R1 and R4, you are not able to see the R1 Ethernet subnet on the R4 routing table. You are also not able to see the R4 Ethernet subnet on the R1 routing table.

Which two configuration changes should be made to resolve this issue? Select the routers where the configuration change will be required, and select the required EIGRP configuration command(s). Choose two answers. (Choose two.)

- A. R1 and R4
- B. R2 and R3
- C. ip summary-address eigrp 1 10.1.1.0 255.255.255.0 and ip summary-address eigrp 1 10.2.2.0 255.255.255.0
- D. variance 2
- E. eigrp stub connected
- F. no auto-summary

Correct Answer: BF

Section: Network Principles


Explanation

Explanation/Reference:

Explanation: Of course, the routing is going through R2 and R3 to reach R4. So the two routers that need configuration change are R2 and R3. Also you need to set auto-summary to No. The no auto-summary command configures classless routing protocols such as RIPv2 and EIGRP to really act as classless because by default they're classfull.

QUESTION 111

Refer to the exhibit.



The actual speed of the serial links between R2 and R3 are 256 kb/s and 512 kb/s. When configuring EIGRP on routers R2 and R3, the network administrator configured the bandwidth of both serial interfaces to 512 kb/s.

What will be the effect?

- A. EIGRP will over utilize the 512 kb/s link.
- B. The interface "delay" value used in the EIGRP metric calculation will be inaccurate on the 256 kb/s serial interface.
- C. The amount of bandwidth used for EIGRP routing protocol traffic on the 256 kb/s link can become excessive.
- D. EIGRP can load balance between the two serial links only if the variance is set to 2 or higher.
- E. Unequal cost load balancing will be disabled.

Correct Answer: C

Section: Network Principles


Explanation

Explanation/Reference:

Explanation: If you assign more bandwidth than what is available between R2 and R3, the EIGRP traffic will become excessive because it uses only the actual bandwidth.

QUESTION 112

Refer to the exhibit.



ROUTE.com has just implemented this EIGRP network. A network administrator came to you for advice while trying to implement load balancing across part of their EIGRP network. If the variance value is configured as 2 on all routers and all other metric and K values are configured to their default values, traffic from the Internet to the data center will be load balanced across how many paths?

- A. 1
- B. 2
- C. 3
- D. 4
- E. 5

Correct Answer: C

Section: Network Principles

Explanation

Explanation/Reference:

Explanation

First we should list all the paths from the Internet to the data center:

- + A-B-C-H with a metric of 70 ($40 + 15 + 15$)
- + A-B-E-H with a metric of 60 ($40+10+10$)
- + A-D-E-H with a metric of 30 ($10+10+10$)
- + A-D-E-B-C-H with a metric of 60 ($10+10+10+15+15$)
- + A-D-E-F-G-H with a metric of 70 ($10+10+10+20+20$)
- + A-F-G-H with a metric of 60 ($20+20+20$)
- + A-F-E-H with a metric of 40 ($20+10+10$)

So the path A-D-E-H will be chosen because it has the best metric. But EIGRP can support unequal cost path load balancing. By configuring the variance value of 2, the minimum metric is increased to 60 ($30 * 2$) and all the routes that have a metric of less than or equal to 60 and satisfy the feasibility condition will be used to send traffic.

Besides the main path A-D-E-H we have 4 more paths that have the metric of less than or equal to 60 (we also include the Advertised Distances of these routes for later comparison):

- + A-B-E-H with an AD of 20
- + A-D-E-B-C-H with an AD of 50
- + A-F-G-H with an AD of 40
- + A-F-E-H with an AD of 20

Now the last thing we need to consider is the feasible condition. The feasible condition states:

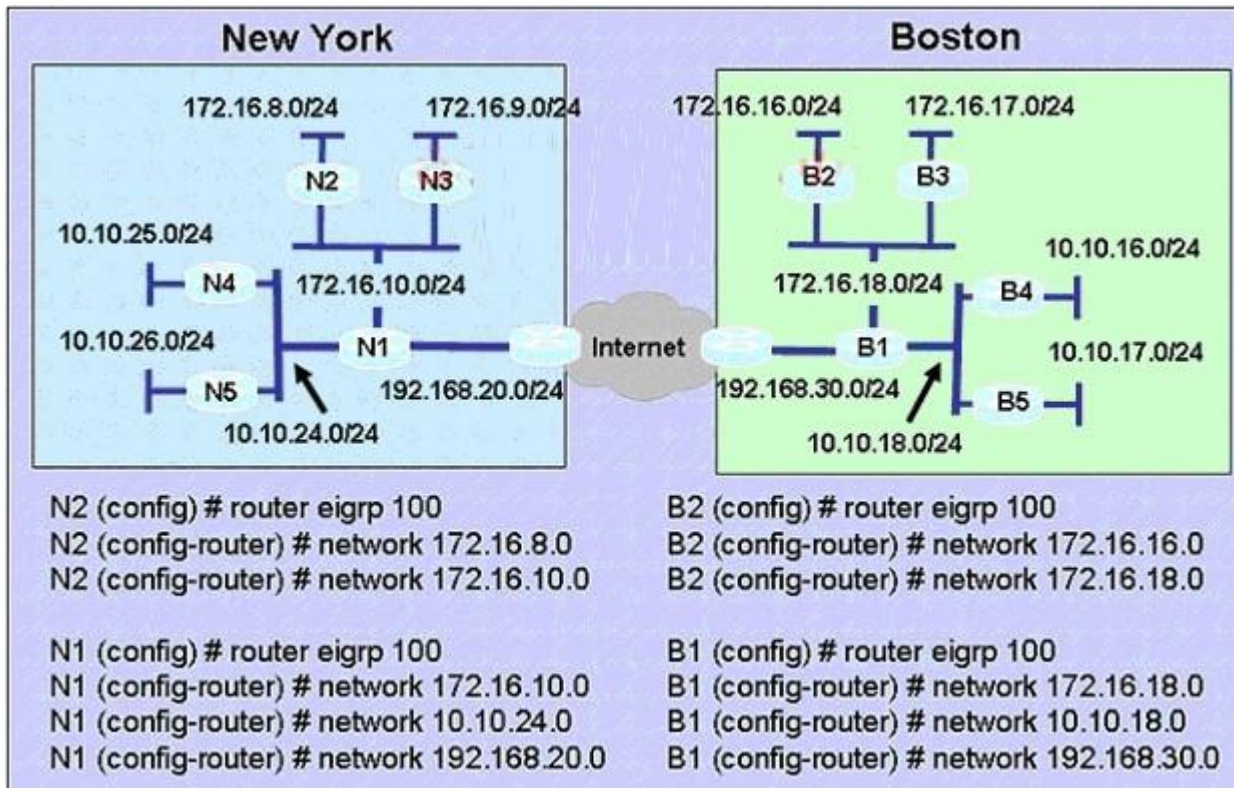
"To qualify as a feasible successor, a router must have an AD less than the FD of the current successor route"

The FD of the current successor route here is 30 (notice that the variance number is not calculated here). Therefore there are only 2 paths that can satisfy this conditions: the path A- B-E-H & A-F-E-H.

In conclusion, traffic from the Internet to the data center will be load balanced across 3 paths, including the main path (successor path)

QUESTION 113

Refer to the exhibit.



A Boston company bought the assets of a New York company and is trying to route traffic between the two data networks using EIGRP. The show command output shows that traffic will not flow between the networks. As a network consultant, you were asked to modify the configuration and certify the interoperability of the two networks. For traffic to flow from subnet 172.16.8.0/24 to the 172.16.16.0/24 subnet.

Which configuration change do you recommend?

- A. Turn off autosummarization on routers N1 and B1.
- B. Add IP summary addresses to the Internet-pointing interfaces of routers N1 and B1.
- C. Turn off auto summarization on routers N2 and B2.
- D. Add wildcard masks to the network commands on routers N2 and B2.

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation: Basically auto route summarization happens at the classful network boundary...so that would make N1 and B1 the locations that summarization would occur for the 172.16.0.0/16 classful networks.

So if you left auto-summarization enabled on those 2 routers, you would have an issue with discontinuous networks being advertised by both routers N1 and B1 with their classful mask (172.16.0.0/16 and 10.0.0.0/8), which will cause you issues.

Turning off auto-summarization on N2 and B2 wouldn't make any difference, as their networks wouldn't be summarized due to the fact that they are not meeting a classful boundary on their perspective routers. N1 will receive the 172.16.8.0/24 network from N2 with auto-summarization enabled.


QUESTION 114

Refer to the exhibit.

Gateway of last resort is not set

172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks

- D 172.16.8.0/24 [90/30720] via 172.16.10.2, 00:20:43, FastEthernet0/0
- C 172.16.10.0/24 is directly connected, FastEthernet0/0
- D 172.16.0.0/16 is a summary, 00:19:05, Null0
- C 192.168.20.0/24 is directly connected, FastEthernet0/1
- D 10.0.0.0/8 [90/30720] via 192.168.20.2, 00:14:51, FastEthernet0/1



Gateway of last resort is not set

172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks

- D 172.16.16.0/24 [90/30720] via 172.16.18.2, 00:06:04, FastEthernet0/0.172
- C 172.16.18.0/24 is directly connected, FastEthernet0/0.172
- D 172.16.0.0/16 is a summary, 00:20:05, Null0
- C 192.168.20.0/24 is directly connected, FastEthernet0/1
- 10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
- D 10.0.0.0/8 is a summary, 00:15:51, Null0
- D 10.10.16.0/24 [90/30720] via 10.10.18.3, 00:04:28, FastEthernet0/0.10
- C 10.10.18.0/24 is directly connected, FastEthernet0/0.10

A Boston company bought the assets of a New York company and is trying to route traffic between the two data networks using EIGRP over EoMPLS. As a network consultant, you were asked to verify the interoperability of the two networks.

From the show ip route command output, what can you tell the customer about the traffic flow between the subnet in New York (172.16.8.0/24) and the subnets in Boston (172.16.16.0/24 and 10.10.16.0/24)?

- A. Traffic is flowing between the 172.16.8.0 subnet and subnets 172.16.16.0 and 10.10.16.0 and no configuration changes are needed.
- B. Auto-summary must be disabled on N1 and B1 before traffic can flow between the 172.16.8.0 subnet and subnets 172.16.16.0 and 10.10.16.0.
- C. Traffic will flow between the 172.16.8.0 subnet and 172.16.16.0 without any further configuration changes. However, auto-summary must be disabled on N1 and B1 before traffic can flow between the 172.16.8.0 subnet and the 10.10.16.0 subnet.
- D. Auto-summary must be disabled on N1 and B1 before traffic can flow between the 172.16.8.0 subnet and the 172.16.16.0 subnet. However, traffic will flow between the 172.16.8.0 subnet and 10.10.16.0 without any further configuration changes.

Correct Answer: B

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation: Basically auto route summarization happens at the classful network boundary...so that would make N1 and B1 the locations that summarization would occur for the 172.16.0.0/16 classful networks.

So if you left auto-summarization enabled on those 2 routers, you would have an issue with discontinuous networks being advertised by both routers N1 and B1 with their classful mask (172.16.0.0/16 and 10.0.0.0/8), which will cause you issues.

Turning off auto-summarization on N2 and B2 wouldn't make any difference, as their networks wouldn't be summarized due to the fact that they are not meeting a classful boundary on their perspective routers.

QUESTION 115

Which statement about a non-zero value for the load metric (k2) for EIGRP is true?

- A. A change in the load on an interface will cause EIGRP to recalculate the routing metrics and send a corresponding update out to each of its neighbors.
- B. EIGRP calculates interface load as a 5-minute exponentially weighted average that is updated every 5 minutes.
- C. EIGRP considers the load of an interface only when sending an update for some other reason.
- D. A change in the load on an interface will cause EIGRP to recalculate and update the administrative distance for all routes learned on that interface.

Correct Answer: C

Section: Mixed Questions

Explanation

Explanation/Reference:**Explanation**

The load metric (k2) represents the worst load on a link between source and destination.

EIGRP routing updates are triggered only by a change in network topology (like links, interfaces go up/down, router added/removed), and not by change in interface load or reliability.

The load is a five minute exponentially weighted average that is updated every five seconds (not five minutes) .

EIGRP considers the load of an interface only when sending an update for some other reason (like a link failure)

QUESTION 116

Refer to the exhibit. Why are the EIGRP neighbors for this router not learning the routes redistributed from OSPF?

```
router eigrp 123
 redistribute ospf 123
 network 116.16.35.0 0.0.0.255
 network 130.130.0.0
 auto-summary
!
router ospf 123
 log-adjacency-changes
 network 116.16.34.0 0.0.0.255 area 0
 neighbor 116.16.34.4
```

- A. Redistribution must be enabled mutually (in both directions) to work correctly.
- B. Auto-summary causes the OSPF routes redistributed into EIGRP to be summarized; thus the OSPF network 116.16.34 is summarized to 116.34.0.0, which is already covered by the EIGRP protocol.
- C. Default metrics are not configured under EIGRP.
- D. Both routing protocols must have unique autonomous system numbers for redistribution to function correctly.

Correct Answer: C

Section: VPN Technologies


Explanation**Explanation/Reference:****Explanation**

Same as RIP, when redistribute into EIGRP from OSPF, the default metric is infinite -> We must set a seed metric when redistributing into EIGRP. Below lists the default seed metrics when redistributing from a routing protocol into another:

Redistributed Protocol	Default Seed Metric
RIP	Infinity
IGRP/EIGRP	Infinity
OSPF	20 for all (except for BGP, which is 1)
BGP	is set to IGP metric value

QUESTION 117

Refer to the exhibit.



Which three statements are true? (Choose three.)

- A. On the routing table of R4, the 10.1.1.0/24 route appears as an O E2 route.
- B. On R4, the 172.16.1.0/24 route has a metric of 20.
- C. The R3 S0/0 interface should not need the no ip split-horizon eigrp 1 configuration command for the 172.16.1.0/24 route to appear in the routing table of R2 as an D EX route.
- D. The administrative distance of the 172.16.1.0/24 route in the routing table of R3 is 170.
- E. On R5, the 4.0.0.0/8 route will have an administrative distance of 120 and a hop count of 6.

Correct Answer: ABD

Section: Network Principles

Explanation

Explanation/Reference:

Explanation

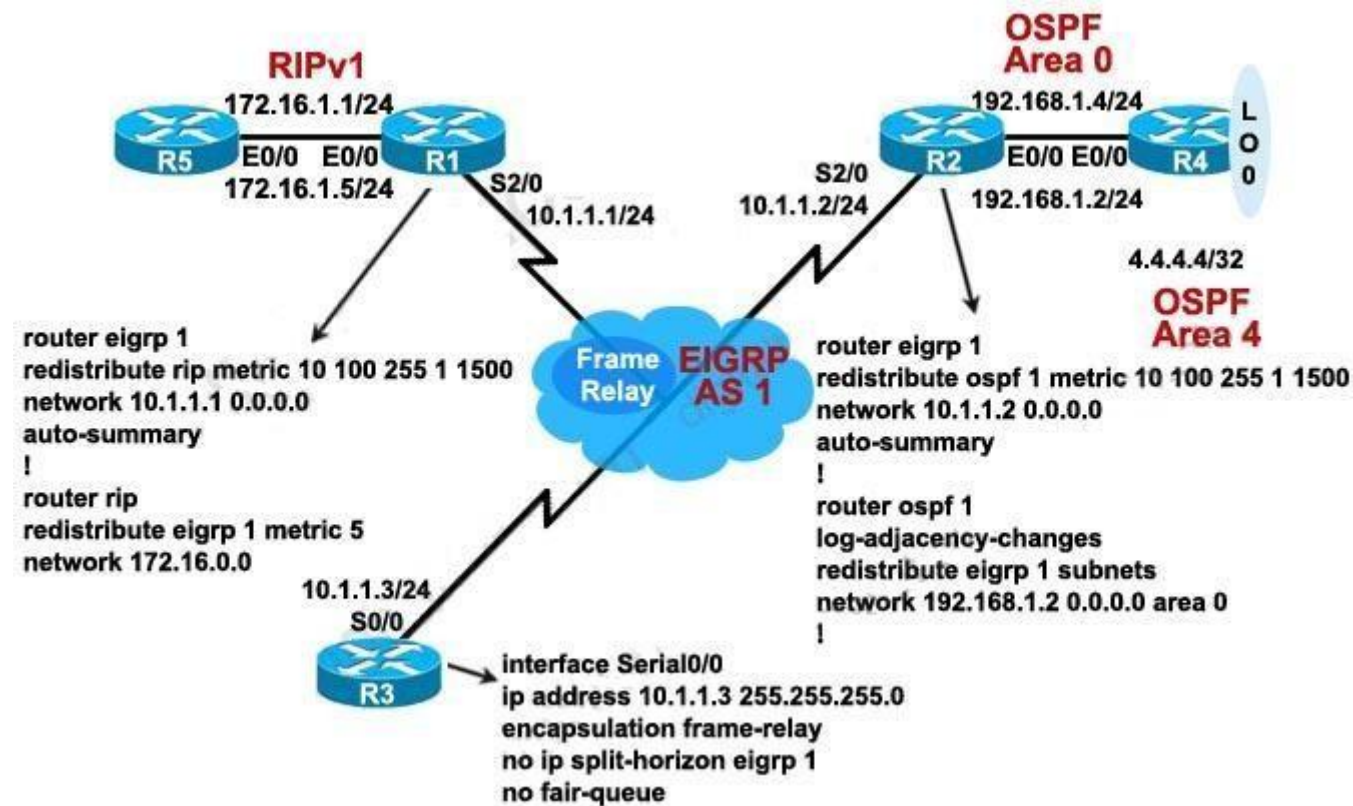
When redistributing into OSPF, the default route type is E2. Notice that the cost of E2 type is always the cost of external route only.

Also, the default seed metric when redistributing into OSPF is always 20 (except for BGP, which is 1)

When redistributing into EIGRP, the external EIGRP routes have an administrative distance of 170 by default

QUESTION 118

Refer to the exhibit.



Looking at the topology diagram and the partial router configurations shown, which statement is true?

A. A routing loop will occur due to mutual route redistribution occurring on R1 and R2.

- B. Suboptimal routing will occur due to mutual route redistribution occurring on R1 and R2.
- C. Additional route filtering configurations using route maps and ACLs are required on the R1 and R2 routers to prevent routing loops.
- D. R2 will not be able to redistribute the EIGRP subnets into OSPF, because R2 is missing the default seed metric for OSPF.
- E. The 10.1.1.0/24 subnet will appear as 10.0.0.0/8 in the R5 routing table.

Correct Answer: E

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

RIPv1 is a classful routing protocol so the subnet 10.1.1.0/24 will be summarized to 10.0.0.0/8 in the R5 routing table. If we use RIPv2 on R1, R5 and use the no auto-summary command on R1 then the 10.1.1.0 subnet will appear in the routing table of R5. Notice that even if the auto-summary command is configured under router eigrp 1 of R1 but when redistributing into another routing protocol EIGRP still advertises the detailed network.

QUESTION 119

You have implemented mutual route redistribution between OSPF and EIGRP on a border router. When checking the routing table on one of the EIGRP routers within the EIGRP routing domain, you are seeing some, but not all of the expected routes. What should you verify to troubleshoot this problem?

- A. The border router is using a proper seed metric for OSPF.
- B. The border router is using a proper seed metric for EIGRP.
- C. The administrative distance is set for OSPF and EIGRP.
- D. The missing OSPF routes are present in the routing table of the border router.
- E. The subnet keyword on the border router in the redistribute OSPF command.

Correct Answer: D

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

We are checking the routing table on EIGRP routers not OSPF so we don't need to check the seed metric for OSPF. Besides OSPF doesn't need to specify seed metric as all external routes get a default metric of 20 (except for BGP, which is 1).

We must specify seed metrics when redistributing into EIGRP (and RIP). If not all the redistributed routes will not be seen but the question says only some routes are missing.

The default administrative distance for external routes redistributed into EIGRP is 170 so we don't need to set it .

The `summary` keyword is only used when redistributing into OSPF, not to other routing protocols.

We should check the routing table of the border router to see the missing OSPF routes are there or not. An incorrect `distribute-list` can block some routes and we can't see it in other EIGRP routers.

QUESTION 120

Refer to the exhibit.

```
R1#show ip eigrp topology | section 0.0.0.0
P 0.0.0.0/0, 2 successors, FD is 2174976
  via 212.50.185.125 (2174976/2169856), Ethernet0/0
  via 212.50.185.126 (2174976/2169856), Ethernet0/0
  via 212.50.185.65 (2180096/2172416), Ethernet1/0
  via 212.50.185.66 (2180096/2172416), Ethernet1/0
  via 212.50.185.33 (2180096/2172416), Ethernet2/0
  via 212.50.185.34 (2180096/2172416), Ethernet2/0
R1#show ip route 0.0.0.0
Routing entry for 0.0.0.0/0, supernet
  known via "eigrp 212", distance 170, metric 2174976, candidate default path, type external
  Redistributing via eigrp 212
  Last update from 212.50.185.126 on Ethernet0/0, 00:00:32 ago
  Routing Descriptor Blocks:
    * 212.50.185.126, from 212.50.185.126, 00:00:32 ago, via Ethernet0/0
      Route metric is 2174976, traffic share count is 1
      Total delay is 20200 microseconds, minimum bandwidth is 1544 kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 3/255, Hops 1
    212.50.185.125, from 212.50.185.125, 00:00:32 ago, via Ethernet0/0
      Route metric is 2174976, traffic share count is 1
      Total delay is 20200 microseconds, minimum bandwidth is 1544 kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 3/255, Hops 1
```

All EIGRP load balancing parameters are set to their defaults. You want to use all the routes in the EIGRP topology for IP load balancing. Which two EIGRP subcommands would you use to accomplish this goal? (Choose two.)

- A. `traffic-share balanced`
- B. `distance`
- C. `maximum-paths`
- D. `default-network`
- E. `variance`

Correct Answer: CE

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

Notice that the `maximum-paths` command is used to share traffic to equal cost path while the `variance` command can share traffic to unequal cost path.

In the output above we learn that EIGRP is using 2 successors to send traffic. By using the variance 2 command we can share traffic to other feasible successor routes. But by default, EIGRP only shares traffic to 4 paths. So we need to use the maximum-paths 6 to make sure all of these routes are used.

QUESTION 121

Refer to the exhibit.

```
R1#show ip route
 1.0.0.0/24 is subnetted, 1 subnets
C    1.1.1.0 is directly connected, Loopback0
D EX 212.50.167.0/24 [170/2172416] via 190.0.0.1 00:45:34, Serial1/0
    [170/2172416] via 191.0.0.1, 00:45:34, Serial2/0
 191.0.0.0/30 is subnetted, 1 subnets
C    191.0.0.0 is directly connected, Serial2/0
D EX 212.50.166.0/24 [170/2172416] via 192.0.0.1, 00:45:34, Serial1/0
    [170/2172416] via 191.0.0.1, 00:45:34, Serial2/0
 20.0.0.0/24 is subnetted, 1 subnets
C    20.20.20.0 is directly connected, Ethernet0/0
 212.50.185.0/27 is subnetted, 3 subnets
D EX 212.50.185.64 [170/2172416] via 192.0.0.1, 00:45:34, Serial1/0
    [170/2172416] via 191.0.0.1, 00:45:34, Serial2/0
D EX 212.50.185.96 [170/2172416] via 192.0.0.1, 00:45:34, Serial1/0
    [170/2172416] via 191.0.0.1, 00:45:34, Serial2/0
D EX 212.50.185.32 [170/2172416] via 192.0.0.1, 00:45:34, Serial1/0
    [170/2172416] via 191.0.0.1, 00:45:34, Serial2/0
 192.0.0.0/30 is subnetted, 1 subnets
C    192.0.0.0 is directly connected, Serial1/0
```

R1 accesses the Internet using E0/0. You have been asked to configure R1 so that a default route is generated to its downstream neighbors (191.0.0.1 and 192.0.0.1). Which commands would create this configuration?

- A. router eigrp 190
 redistribute static
 !
 ip route 0.0.0.0 0.0.0.0 Null0
- B. ip default-network 20.0.0.0
- C. router eigrp 190
 redistribute static
 !
 ip route 0.0.0.0 255.255.255.255 Null0

D. ip default-network 20.20.20.0

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Since you are running EIGRP and you have other routers that need a default route, you can use EIGRP to distribute that without having to program static routes in each. Since 2 are default routes they are only used on the router that they are configured on. The first option has you configure the static route as well as a way to redistribute that route to other routers connecting to you via EIGRP. This will essentially publish this route the same as if it were programmed in with the network x.x.x.x sub-command in the router eigrp 100 routing table.

QUESTION 122

Refer to the exhibit.

```
router eigrp 190
 redistribute eigrp 212
 network 192.0.0.0 0.0.0.3
!
router eigrp 212
 redistribute eigrp 190 route-map default_route
 network 212.50.185.96 0.0.0.31
!
route-map default_route permit 10
 match ip address 100
```

A partial routing configuration is shown. Complete the configuration so that only the default- network is redistributed from EIGRP 190 into EIGRP 212. Which ACL statement completes the configuration correctly?

- A. access-list 100 permit ip 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
- B. access-list 100 permit ip host 0.0.0.0 any
- C. access-list 100 permit ip any host 0.0.0.0
- D. A default-network cannot be redistributed between routing processes.

Correct Answer: C

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

The command `access-list 100 permit ip any host 0.0.0.0` means permit any source address with the destination of 0.0.0.0/0, which is the default route


Note:

any equals 0.0.0.0 255.255.255.255

host 0.0.0.0 equals 0.0.0.0 0.0.0.0

QUESTION 123

Refer to the exhibit.



Which three commands should be used on router B1 to redistribute the EIGRP AS 10 routes into RIP? (Choose three.)

- A. `router rip`
- B. `router eigrp 10`
- C. `redistribute eigrp 10`
- D. `redistribute rip`
- E. `default-metric 10000 100 255 1 1500`
- F. `default-metric 5`

Correct Answer: ACF

Section: VPN Technologies

Explanation


Explanation/Reference:

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009487e.shtml#ri p

QUESTION 124

You want the redistributed EIGRP AS 10 routes to have an administrative distance of 121 when they appear as RIP routes in the routing table of A1. Which command should you use on a router to accomplish this goal?



- A. redistribute eigrp 10 metric 121
- B. redistribute rip metric 121
- C. default-metric 121
- D. distance 121 10.1.1.6 0.0.0.0

Correct Answer: D

Section: Network Principles

Explanation


Explanation/Reference:

Explanation:

If you add that route back with an administrative distance of 121, the ASA will still prefer the route learned via RIP because it prefers the route with a lower administrative distance.

QUESTION 125

Refer to the exhibit.



In a redundant hub-and-spoke deployment using EIGRP, what feature can be used to ensure that routers C through F are not used as transit routers for

data traveling from router B to network 10.1.1.0? Select the best response

- A. Use address summarization at routers C, D, E, and F.
- B. Use the EIGRP Stub feature on routers C, D, E, and F.
- C. Use passive-interface on the spoke links in routers A and B.
- D. Change the administrative distance in routers A and B for routes learned from routers C, D, E, and F.

Correct Answer: B

Section: Network Principles

Explanation

Explanation/Reference:

Explanation

By configuring stub feature on routers C D E and F, routers A and B will not try to transit traffic through these routers. For example, if the network connecting from routers A and B is down, router B will not send to network 10.1.1.0/24 from router B -> router C/D/E/F -> router A -> network 10.1.1.0/24.

QUESTION 126

ACME Rocket Sleds is growing, and so is their network. They have determined that they can no longer continue using static routes and must implement a dynamic routing protocol. They want to have data use multiple paths to the destinations, even if the paths are not equal cost.

Which routing protocol has the ability to do this?

- A. EIGRP
- B. OSPF
- C. RIPv1
- D. RIPv2
- E. BGP
- F. IS-IS

Correct Answer: A

Section: Mixed Questions

Explanation

Explanation/Reference:


Explanation:

Unlike most internal routing protocols, EIGRP has a really cool feature that allows you to share the load of your traffic across multiple links and not just links that have the same cost values. EIGRP allows you to make full use of your redundant links that could be in place just to have for back up but you are paying out a lot of money just to sit there and do nothing. EIGRP makes it easy for us the network engineers to make this happen. Before jumping in the in's and out lets run through a few things first when it comes to EIGRP Load Sharing, also refereed to as Load Balancing sometimes.

Reference: <http://ericleahy.com/index.php/eigrp-equal-and-unequal-cost-load-sharing/>

QUESTION 127

Refer to the exhibit.



ROUTE Enterprises has many stub networks in their enterprise network, such as router B and its associated network. EIGRP is to be implemented on router A so that neither the prefix for the S0/0/0 interface nor the prefixes from router B appear in the routing tables for the router in the enterprise network. Which action will accomplish this goal?

- A. Declare router B a stub router using the `eigrp stub` command.
- B. Use the `passive-interface` command for interface Serial0/0/0.
- C. Use a mask with the `network` command to exclude interface Serial0/0/0.
- D. Implement a distribute list to exclude the link prefix from the routing updates.

Correct Answer: C

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

If we declare router B a stub router then the routers in Enterprise Network still learn about the network for S0/0/0 interface and the network behind router B.

If we use the passive-interface command on s0/0/0 interface then router A & B can not become neighbor because they don't exchange hello messages
-> A can not send traffic to the network behind B .


Theoretically, we can use a distribute list to exclude both the link prefix and the prefix from router B but it is not efficient because:

+ We have many stub networks so we will need a long distribute list. + We declare networks in stub routers (like router B) while filter them out at router A
-> it is a waste.

Not totally sure about answer C because if we use a mask with the network command to exclude interface Serial0/0/0 then router A and B can not become neighbors and the situation is same as answer B. But from many discussions about this question, maybe C is the best answer.

QUESTION 128

Refer to the exhibit.



Which statement about dynamic routing protocols for this network is true?

- A. No dynamic interior routing protocol can summarize as shown.
- B. Unless configured otherwise, EIGRP would automatically summarize the prefixes as shown in the exhibit.
- C. With this IP addressing scheme, EIGRP can be manually configured to summarize prefixes at the specified summarization points.
- D. The IP address design lends itself to OSPF. Each summarizing router would be an ABR, summarizing to the next area in the address hierarchy.

Correct Answer: C

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

Summarization may be manually applied at any point in the network. You can configure manual summarization on any router interface. Consider summarization for both upstream and downstream neighbors. Upstream neighbors should receive a consolidated route, and downstream neighbors can receive a default route.

QUESTION 129

After implementing EIGRP on your network, you issue the show ip eigrp traffic command on router C. The following output is shown:

```
RouterC#show ip eigrp traffic
IF-EIGRP Traffic Statistics for process 1
Hellos sent/received: 481/444
Updates sent/received: 41/32
Queries sent/received: 5/1
Replies sent/received: 1/4
Acks sent/received: 21/25
Input queue high water mark 2, 0 drops
SIA-Queries sent/received: 0/0
SIA-Replies sent/received: 0/0
```

Approximately 25 minutes later, you issue the same command again. The following output is shown:

```
RouterC#show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 1
Hellos sent/received: 1057/1020
Updates sent/received: 41/32
Queries sent/received: 5/1
Replies sent/received: 1/4
Acks sent/received: 21/25
Input queue high water mark 2, 0 drops
SIA-Queries sent/received: 0/0
SIA-Replies sent/received: 0/0
```

Approximately 25 minutes later, you issue the same command a third time. The following output is shown:

```
RouterC#show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 1
Hellos sent/received: 1754/1717
Updates sent/received: 41/32
Queries sent/received: 5/1
Replies sent/received: 1/4
Acks sent/received: 21/25
```

Input queue high water mark 2, 0 drops
SIA-Queries sent/received: 0/0
SIA-Replies sent/received: 0/0

What can you conclude about this network?

- A. The network has been stable for at least the last 45 minutes.
- B. There is a flapping link or interface, and router C knows an alternate path to the network.
- C. There is a flapping link or interface, and router A does not know an alternate path to the network.
- D. EIGRP is not working correctly on router C.
- E. There is not enough information to make a determination.

Correct Answer: A

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

In three times using the command, the Queries sent/received & Replies sent/received are still the same -> the network is stable.

QUESTION 130

After implementing EIGRP on your network, you issue the show ip eigrp traffic command on router C. The following output is shown:

```
RouterC#show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 1
Hellos sent/received: 2112/2076
Updates sent/received: 47/38
Queries sent/received: 5/3
Replies sent/received: 3/4
Acks sent/received: 29/33
Input queue high water mark 2, 0 drops
SIA-Queries sent/received: 0/0
SIA-Replies sent/received: 0/0
```

Moments later, you issue the same command a second time and the following output is shown:

```
RouterC#show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 1
Hellos sent/received: 2139/2104
Updates sent/received: 50/39
Queries sent/received: 5/4
```

Replies sent/received: 4/4
Acks sent/received: 31/37
Input queue high water mark 2, 0 drops
SIA-Queries sent/received: 0/0
SIA-Replies sent/received: 0/0

Moments later, you issue the same command a third time and the following output is shown:

```
RouterC#show ip eigrp traffic
IP-EIGRP Traffic Statistics for process 1
Hellos sent/received: 2162/2126
Updates sent/received: 53/42
Queries sent/received: 5/5
Replies sent/received: 5/4
Acks sent/received: 35/41
Input queue high water mark 2, 0 drops
SIA-Queries sent/received: 0/0
SIA-Replies sent/received: 0/0
```

What information can you determine about this network?

- A. The network is stable.
- B. There is a flapping link or interface, and router C knows an alternate path to the network.
- C. There is a flapping link or interface, and router C does not know an alternate path to the network.
- D. EIGRP is not working correctly on router C.
- E. There is not enough information to make a determination.

Correct Answer: B

Section: Mixed Questions

Explanation


Explanation/Reference:

Explanation

We notice that the Queries received number is increased so router C has been asked for a route. The Replies sent number is also increased -> router C knows an alternate path to the network.

QUESTION 131

Refer to the exhibit.



EIGRP has been configured on router D. Router C is performing mutual redistribution between EIGRP and OSPF. While verifying that the redistribution is functioning properly, you discover that while router C has all of the EIGRP routes in its routing table, router A does not have any routes from the EIGRP domain. What on router C may be the cause of the problem?

- A. The no auto-summary command needs to be added under router eigrp 1.
- B. The subnets keyword was not included in the redistribute command under router ospf 1.
- C. The metric specified for the redistributed EIGRP routes is too large; making the EIGRP routes unreachable by router A.
- D. The default-information originate command needs to be added under router ospf 1.
- E. The administrative distance of either OSPF or EIGRP must be changed so that EIGRP has a higher administrative distance than OSPF.

Correct Answer: B

Section: Mixed Questions

Explanation


Explanation/Reference:

Explanation

If we don't use the subnets keyword when redistributing routes learned from another routing process into OSPF, only classful routes will be redistributed. This is an important thing to remember when redistributing into OSPF.

QUESTION 132

During the redistribution process configured on RTA, some of the EIGRP routes, such as 10.1.1.0/24 and 10.2.2.0/24, are not being redistributed into the OSPF routing domain.



Which two items could be a solution to this problem? (Choose two.)

- A. Change the metric-type to 2 in the redistribute command.
- B. Configure the redistribute command under router eigrp 1 instead.
- C. Change the EIGRP AS number from 100 to 1 in the redistribute command.
- D. Add the subnets option to the redistribute command.
- E. Change the metric to an EIGRP compatible metric value (bandwidth, delay, reliability, load, MTUs) in the redistribute command.

Correct Answer: CD

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

In this example, the router is configured for EIGRP AS 1, but EIGRP AS 100 is being redistributed into OSPF so the EIGRP AS needs to be changed from 100 to 1. The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword, only networks that are not subnetted are

redistributed by OSPF.

Example:

```
RTA(config)#router ospf 109
```

```
RTA(config-router)#redistribute rip subnets
```

```
RTA(config-router)#network 130.10.62.0 0.0.0.255 area 0 RTA(config-router)#network 130.10.63.0 0.0.0.255 area 0
```

The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword, only networks that are not subnetted are redistributed by OSPF.

QUESTION 133

Given the accompanying output, which additional command is needed to redistribute IGRP into EIGRP?

```
Router eigrp 123
```

```
Network 10.10.10.0
```

```
No auto-summary
```

```
!
```

```
Router igrp 123
```

```
Network 172.16.0.0
```

```
Network 172.17.0.0
```

- A. Under the router igrp mode add redistribute eigrp 123
- B. Under the router eigrp mode add redistribute igrp 123
- C. Under the router eigrp mode add redistribute igrp 123 subnets
- D. None, EIGRP and IGRP are automatically redistributed in this instance.

Correct Answer: D

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

The point of this question is redistribute IGRP into EIGRP. When redistributing IGRP into EIGRP, there is a feature that they are automatically redistributed if they have same autonomous system number; in opposite, they need to manually redistributed if they have different autonomous system number..

QUESTION 134

Which three statements are true about EIGRP operation? (Choose three.)

- A. When summarization is configured, the router will also create a route to null 0.
- B. The summary route remains in the route table, even if there are no more specific routes to the network.
- C. Summarization is configured on a per-interface level.
- D. The maximum metric for the specific routes is used as the metric for the summary route.

E. Automatic summarization across major network boundaries is enabled by default.

Correct Answer: ACE

Section: Mixed Questions

Explanation

Explanation/Reference:

Reference: <http://astorinonetworks.com/2011/07/20/summary-routes-to-null0-the-protocols-that-love-them/>

QUESTION 135

Which two statements about the EIGRP DUAL process are correct? (Choose two.)

- A. An EIGRP route will go active if there are no successors or feasible successors in the EIGRP topology table.
- B. An EIGRP route will go passive if there are no successors in the EIGRP topology table.
- C. DUAL will trigger an EIGRP query process while placing the flapping routes in the holddown state.
- D. A feasible successor in the EIGRP topology table can become the successor only after all the query requests have been replied to.
- E. The stuck in active state is caused when the wait for the query replies have timed out.
- F. EIGRP queries are sent during the loading state in the EIGRP neighbor establishment process.

Correct Answer: AE

Section: Mixed Questions

Explanation

Explanation/Reference:

Reference: <http://routemyworld.com/category/routing-protocols/eigrp/>

QUESTION 136

What are three key concepts that apply when configuring the EIGRP stub routing feature in a hub and spoke network? (Choose three.)

- A. A hub router prevents routes from being advertised to the remote router.
- B. Only remote routers are configured as stubs.
- C. Stub routers are not queried for routes.
- D. Spoke routers connected to hub routers answer the route queries for the stub router.
- E. A stub router should have only EIGRP hub routers as neighbors.
- F. EIGRP stub routing should be used on hub routers only.

Correct Answer: BCE

Section: Mixed Questions

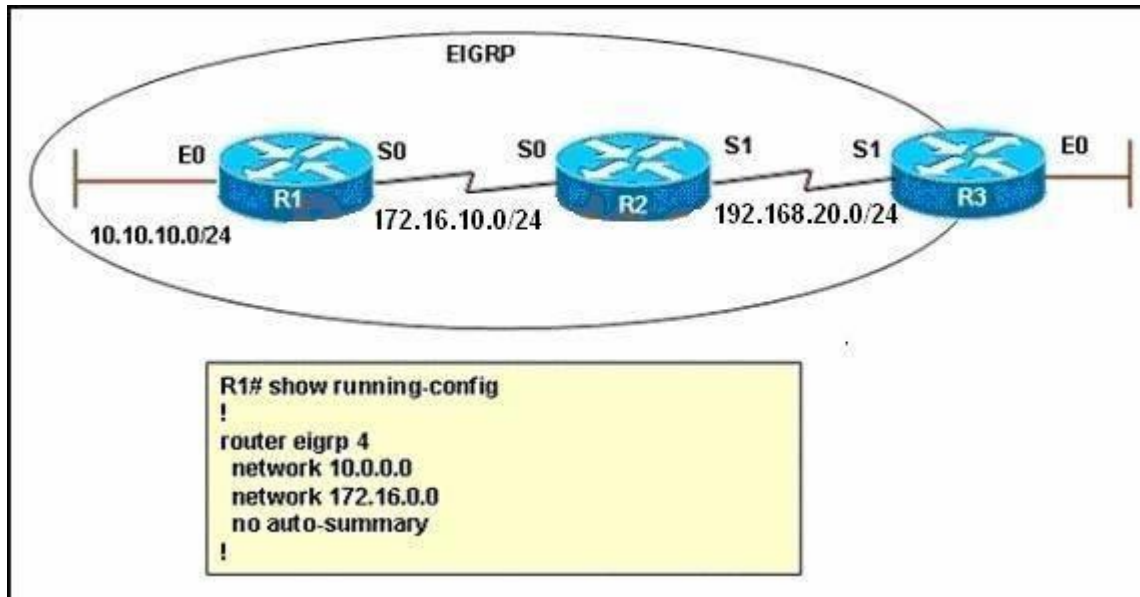
Explanation

Explanation/Reference:

Reference: <http://astorinonetworks.com/2011/06/14/eigrp-stub-routing/> http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/eigrpstb.html

QUESTION 137

Refer to the exhibit. EIGRP is configured with the default configuration on all routers. Autosummarization is enabled on routers R2 and R3, but it is disabled on router R1. Which two EIGRP routes will be seen in the routing table of router R3? (Choose two.)



- A. 10.0.0.0/8
- B. 10.10.0.0/16
- C. 10.10.10.0/24
- D. 172.16.0.0/16
- E. 172.16.0.0/24
- F. 172.16.10.0/24

Correct Answer: CD

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

EIGRP performs an auto-summarization each time it crosses a border between two different major networks. In this case all different networks are in different major networks so EIGRP will perform auto-summarization when it exits an interface. But R1 has been configured with no auto-summary command so EIGRP will not summarize on S0 interface of R1. So the routing table of R2 will have the network 10.10.10.0/24 (not be summarized).

When exiting S1 interface of R2, EIGRP summarizes network 172.16.10.0/24 into the major 172.16.0.0/16 network but it does not summarize network 10.10.10.0/24 because it is not directly connected with this network. Therefore in the routing table of R3 there will have:

- + Network 10.10.10.0/24 (not summarized)
- + Network 172.16.0.0/16 (summarized)

-> C and D are correct.

Note: I simulated this question on GNS3, you can see the final outputs of the show ip route commands on these routers (I connected these routers via FastEthernet, not Serial interfaces so the outputs are slightly different but the main points are not changed).

```
R1#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

    172.16.0.0/24 is subnetted, 1 subnets
C       172.16.10.0 is directly connected, FastEthernet0/0
D       192.168.20.0/24 [90/30720] via 172.16.10.2, 00:02:15, FastEthernet0/0
    10.0.0.0/24 is subnetted, 1 subnets
C       10.10.10.0 is directly connected, Loopback0
```

```

R2#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

    172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C       172.16.10.0/24 is directly connected, FastEthernet0/0
D       172.16.0.0/16 is a summary, 00:02:33, Null0
C       192.168.20.0/24 is directly connected, FastEthernet0/1
    10.0.0.0/24 is subnetted, 1 subnets
D       10.10.10.0 [90/156160] via 172.16.10.1, 00:02:36, FastEthernet0/0

```

```

R3#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route


Gateway of last resort is not set

D       172.16.0.0/16 [90/30720] via 192.168.20.2, 00:03:38, FastEthernet0/0
C       192.168.20.0/24 is directly connected, FastEthernet0/0
    10.0.0.0/24 is subnetted, 1 subnets
D       10.10.10.0 [90/158720] via 192.168.20.2, 00:03:38, FastEthernet0/0

```

QUESTION 138

Refer to the exhibit and the partial configuration on router R2. On router R4 all RIP routes are redistributed into the OSPF domain. A second redistribution is configured on router R2 using a route map. Based on the configuration on router R2, which EIGRP external routes will be present in the routing table of R1?



- A. the routes originating from the RIP routing domain
- B. the routes originating from the OSPF stub area
- C. all OSPF inter and intra-area routes
- D. all routes originating from RIP and OSPF routing domains

Correct Answer: A

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

R2 sees the routes from RIP domain as external routes while it sees the routes from OSPF Stub Area as internal routes. From the output we learn that the route-type external is redistributed from OSPF to EIGRP (via route-map ABC) so we will see the routes from the RIP domain (external) in the routing table of R1.

In the case we want to redistribute routes from OSPF Stub Area (Area 1) to EIGRP we need to use the match route-type internal command instead.

QUESTION 139

Refer to the exhibit. EIGRP has been configured on routers R1 and R2. However, R1 does not show R2 as a neighbor and does not accept routing updates from R2. What could be the cause of the problem?



```
hostname R1
!
interface Ethernet0
!
 ip address 10.1.1.1 255.255.255.0
!
router eigrp 4
 network 10.0.0.0
!
end
```

```
hostname R2
!
interface Ethernet0
!
 ip address 10.1.2.2 255.255.255.0
 ip address 10.1.1.2 255.255.255.0 secondary
!
router eigrp 4
 network 10.0.0.0
!
end
```

```
R1#show ip eigrp neighbor
IP-EIGRP neighbors for process
01:20:54: IP-EIGRP: Neighbor 10.1.2.2 not on common subnet for Ethernet0 (10.1.)
01:21:08: IP-EIGRP: Neighbor 10.1.2.2 not on common subnet for Ethernet0 (10.1.)
```

```
R2# show ip eigrp neighbor
IP-EIGRP neighbors for process 4
```

H	Address	Interface	Hold	Uptime (sec)	SRTT (ms)	RTO	Q Cnt	Seq Num	Type
0	10.1.1.1	Et0	12	00:00:35	1	5000	1	0	

- A. The no auto-summary command has not been issued under the EIGRP process on both routers.
- B. Interface E0 on router R1 has not been configured with a secondary IP address of 10.1.2.1/24.
- C. EIGRP cannot exchange routing updates with a neighbor's router interface that is configured with two IP addresses.
- D. EIGRP cannot form neighbor relationship and exchange routing updates with a secondary address.

Correct Answer: D

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

EIGRP updates always use the primary IP address of the outgoing interface as the source address. In this case R2 will use the 10.1.2.2/24 address, which is not in the same subnet of R1, to send EIGRP update to R1. Therefore R1 does not accept this update and generates the not on common subnet error message.

Answer D is a bit unclear. It should state that EIGRP cannot form neighbor relationship and exchange routing updates if the two primary addresses on two routers are not in the same subnet.

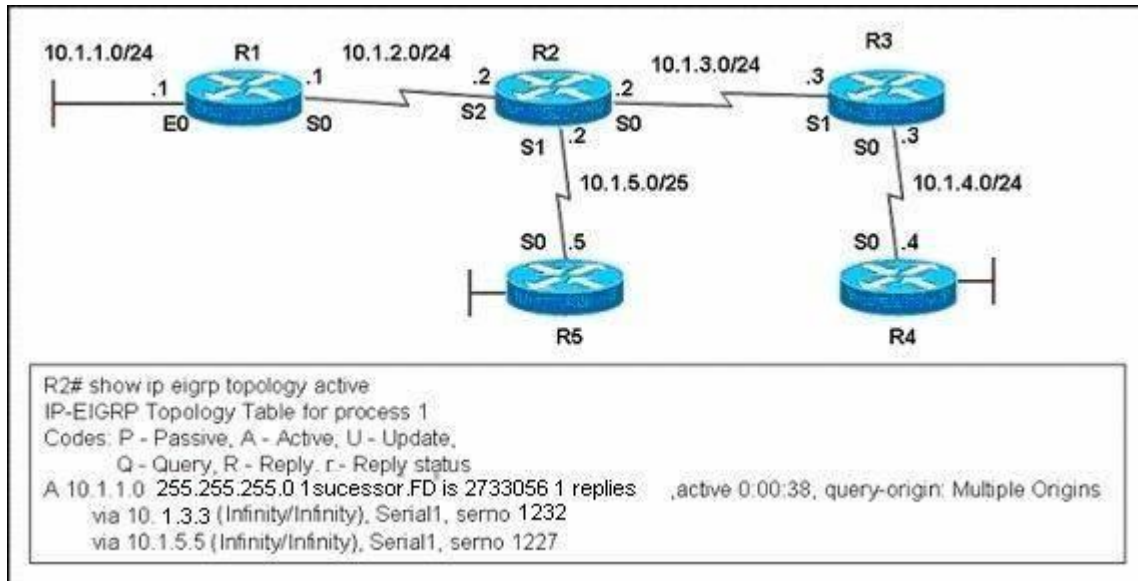
Notice that although R1 does not accept R2 as its EIGRP neighbors but R2 accepts R1 as its EIGRP neighbor accepts R1 hello packets..

For more information about this problem, please read

http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080093f09.shtml.

QUESTION 140

Refer to the exhibit.



EIGRP had converged in AS 1 when the link between router R1 and R2 went down. The console on router R2 generated the following messages:

*Mar 20 12:12:06: %DUAL-5-NBRCHANGE. IP-EIGRP 1: Neighbor 10.1.4.3 (Serial0) is down: stuck in active

*Mar 20 12:15:23: %DUAL-3-SIA. Route 10.1.1.0/24 stuck-in-active state in IP-EIGRP 1. Cleaning up The network administrator issued the show ip eigrp topology active command on R2 to check the status of the EIGRP network. Which statement best describes the reason for the error messages?

- A. Incorrect bandwidth configuration on router R3 prevents R2 from establishing neighbor adjacency.
- B. Incorrect bandwidth configuration on router R5 prevents R2 from establishing neighbor adjacency.
- C. Router R3 did not reply to the query about network 10.1.1.0/24 sent by router R2 .
- D. Router R5 did not reply to the query about network 10.1.1.0/24 sent by router R2 .

Correct Answer: C

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

When the link between R1 and R2 is down, R2 loses its successor for the network 10.1.1.0/24. R2 checks its topology table for a feasible successor but it can't find one. So R2 goes active on the that route to find a new successor by sending queries out to its neighbors (R3 and R5) requesting a path to the lost route. Both R3 and R5 also go active for the that route. But R5 doesn't have any neighbor to ask besides R2 so it will send an unreachable message to indicate it has no alternative path for that route and has no other neighbor to query. R3 also checks its EIRGP topology table for a feasible successor but it has none, too. Unlike R5, R3 has a neighbor (R4) so it continues to query this router.

Now suppose there is a problem on the link between R3 and R4 so R4 never receives the query from R3 and of course, R3 also never receives a reply back from R4. Therefore, R3 can't reply back to R2. After about 3 minutes, the Stuck in active (SIA) timer on R2 expires and R2 marks the route 10.1.1.0/24 as stuck in active route.

The output line via 10.1.3.3 (Infinity/Infinity), r, Serial0, serial 1232 indicates R2 has sent a query to 10.1.3.3 and is waiting for a reply (the lowercase r).

(Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008010f016.shtml)

QUESTION 141

Which EIGRP packet statement is true?

- A. On high-speed links, hello packets are broadcast every 5 seconds for neighbor discovery.
- B. On low-speed links, hello packets are broadcast every 15 seconds for neighbor discovery.
- C. Reply packets are multicast to IP address 224.0.0.10 using RTP.
- D. Update packets route reliable change information only to the affected routers.
- E. Reply packets are used to send routing updates.

Correct Answer: D

Section: Network Principles

Explanation


Explanation/Reference:

Reference:

http://docwiki.cisco.com/wiki/Enhanced_Interior_Gateway_Routing_Protocol#EIGRP_Packet_Types

QUESTION 142

Refer to the exhibit. EIGRP has been configured on all routers in the network. The command metric weights 0 0 1 0 0 has been added to the EIGRP process so that only the delay metric is used in the path calculations. Which router will R1 select as the successor and feasible successor for Network A?



- A. R4 becomes the successor for Network A and will be placed in the routing table. R2 becomes the feasible successor for Network A.
- B. R4 becomes the successor for Network A and will be included in the routing table. No feasible successor will be selected as the advertised distance from R2 is higher than the feasible distance.
- C. R2 becomes the successor and will be placed in the routing table. R4 becomes the feasible successor for Network A.
- D. R2 becomes the successor and will be placed in the routing table. No feasible successor will be selected as the reported distance from R4 is lower than the feasible distance.

Correct Answer: B

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

The point of the question is DUAL of EIGRP.

$FD = 5 + 10 + 5 = 20$

R4 is successor

No FS, because no other router's AD is lower the FD

QUESTION 143

During a redistribution of routes from OSPF into EIGRP, an administrator notices that none of the OSPF routes are showing in EIGRP. What are two possible causes? (Choose two.)

- A. incorrect distribute lists have been configured
- B. missing ip classless command
- C. CEF not enabled
- D. no default metric configured for EIGRP

Correct Answer: AD

Section: Infrastructure Services

Explanation

Explanation/Reference:


Explanation

An incorrect distribute list can filter out updates therefore none of the OSPF routes are showing in EIGRP.

The default metric when redistributing into EIGRP is infinite so we must specify a seed metric for EIGRP to work with.

QUESTION 144

You are a network technician, study the exhibit carefully.



What must be done on router A in order to make EIGRP work effectively in a Frame Relay multipoint environment?

- A. Issue the command bandwidth 56 on the physical interface.
- B. Issue the command bandwidth 56 on each subinterface.
- C. Issue the command bandwidth 224 on each subinterface.
- D. Issue the command bandwidth 224 on the physical interface.

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

Explanation

In Frame Relay, all neighbors share the same bandwidth, regardless of the actual CIR of each individual PVC. In this case the CIR of each PVC is the same so we can find the bandwidth of the main interface (multipoint connection interface) by $56 \times 4 = 224$.

Notice that if the bandwidth on each PVC is not equal then we get the lowest bandwidth to multiply.

QUESTION 145

Based on the exhibited output, which three statements are true? (Choose three.)

```
R1#show ip route

Gateway of last resort is 10.1.1.2 to network 0.0.0.0

C    1.0.0.0/8 is directly connected, Loopback0
    172.17.0.0/24 is subnetted, 1 subnets
D      172.17.1.0 [90/25632000] via 10.1.1.2, 00:05:20, Serial0/0
    172.16.0.0/24 is subnetted, 1 subnets
D      172.16.1.0 [90/23072000] via 10.1.1.2, 00:05:20, Serial0/0
        [90/20640000] via 10.1.1.3, 00:00:13, Serial0/0
D    172.19.0.0/16 [90/391248640] via 10.1.1.3, 00:05:20, Serial0/0
D    172.22.0.0/16 [90/20640000] via 10.1.1.3, 00:05:21, Serial0/0
D EX 172.25.0.0/16 [170/32032000] via 10.1.1.2, 00:00:10, Serial0/0
    10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D      10.2.0.0/16 is a summary, 00:06:18, Null0
C      10.2.1.0/24 is directly connected, FastEthernet0/0
C      10.1.1.0/24 is directly connected, Serial0/0
D*EX 0.0.0.0/0 [170/20514560] via 10.1.1.2, 00:00:11, Serial0/0
R1#
```

- A. R1 is configured with the variance command.
- B. The route to 10.2.0.0/16 was redistributed into EIGRP.
- C. A default route has been redistributed into the EIGRP autonomous system.
- D. R1 is configured with the ip summary-address command.
- E. The router at 10.1.1.2 is configured with the ip default-network 0.0.0.0 command.
- F. R1 is sourcing an external EIGRP route from Null0.

Correct Answer: ACD
Section: Network Principles
Explanation

Explanation/Reference:
Explanation

From the routing table above, we see that network 172.16.1. can be reached via 2 unequal paths (with FD of 23072000 & 20640000) so surely R1 has been configured with the variance command

By configuring a default route and redistribute it into EIGRP you will get the line D *EX 0.0.0.0/0 ... line in the routing table of that router

From the line 10.2.0.0/16 is a summary, 00:16:18, Null0 we know that this network has been summarized with the ip summarary-address command (notice that 10.2.0.0 is not the major network of net

QUESTION 146

Examine the exhibit carefully.

```
R1# show ip eigrp topology
<output omitted>
P 10.1.2.0/24, 1 successors, FD is 281600
   via Connected, FastEthernet0/0
A 10.6.1.0/24, 0 successors, FD is 3385160704, Q
   1 replies, active 00:00:41, query-origin: Local origin
   Remaining replies:
     via 10.1.2.1, r. FastEtherent0/0
```

EIGRP is configured on all routers in the network. What conclusion can be derived from the show ip eigrp topology output provided?

- A. Router R1 is waiting for a reply from the neighbor 10.1.2.1 to the hello message sent out inquiring for a second successor to network 10.6.1.0/24.
- B. Router R1 can send traffic destined for network 10.6.1.0/24 out of interface FastEthernet0/0.
- C. Router R1 is waiting for a reply from the neighbor 10.1.2.1 to the hello message sent out before it declares the neighbor unreachable.
- D. Router R1 is waiting for a reply from the neighbor 10.1.2.1 in response to the query sent out about network 10.6.1.0/24.

Correct Answer: D
Section: Network Principles
Explanation

Explanation/Reference:

Explanation:

The show ip eigrp topology command lists all routes that EIGRP is aware of and shows whether EIGRP is actively processing information on that route. Under most normal conditions, the routes should all be in a passive state and no EIGRP process are running for that route. If the routes are active, this could indicate the dreaded stuck in active, or SIA, state.

The fields to note in this output are as follows:

P-- Passive; no EIGRP computation is being performed. This is the ideal state. A-- Active; EIGRP computations are "actively" being performed for this destination. Routes constantly appearing in an active state indicate a neighbor or query problem.

Both are symptoms of the SIA problem.

U-- Update; an update packet was sent to this destination. Q-- Query; a query packet was sent to this destination. R-- Reply; a reply packet was sent to this destination. Route information-- IP address of the route or network, its subnet mask, and the successor, or next hop to that network, or the feasible successor.

QUESTION 147

Which three statements are true about EIGRP route summarization? (Choose three.)

- A. Manual route summarization is configured in router configuration mode when the router is configured for EIGRP routing.
- B. Manual route summarization is configured on the interface.
- C. When manual summarization is configured, the summary route will use the metric of the largest specific metric of the summary routes.
- D. The ip summary-address eigrp command generates a default route with an administrative distance of 90.
- E. The ip summary-address eigrp command generates a default route with an administrative distance of 5.
- F. When manual summarization is configured, the router immediately creates a route that points to null0 interface.

Correct Answer: BEF

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

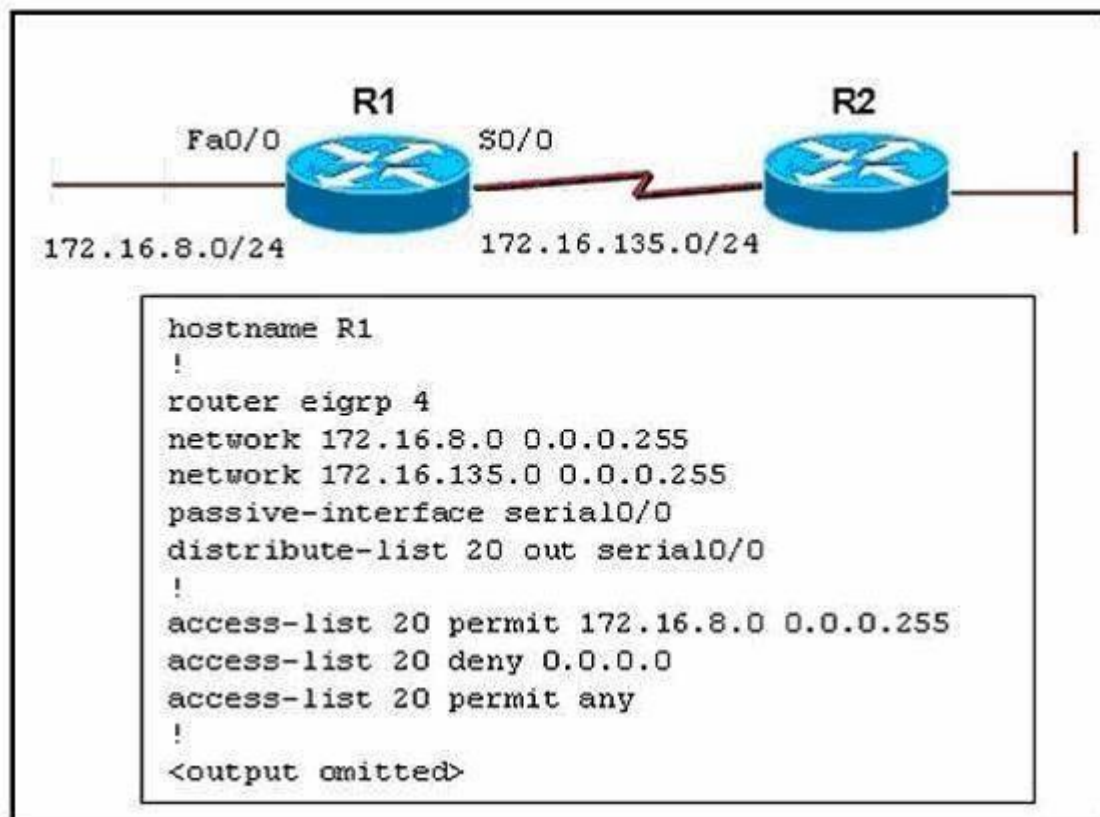
The purpose of route summarization is small routing tables, smaller updates. On major network boundaries, subnetworks are summarized to a single classful network and automatic route summarization is enabled by default. Manual route summarization can be configured on per interface basis. When summarization is configured on an interface, the router immediately creates a route pointing to null0.

Route summarization works in conjunction with the ip summary-address eigrp interface configuration command, in which additional summarization can be performed. If automatic summarization is in effect, there usually is no need to configure network level summaries using the ip summary-address eigrp command. You can configure a summary aggregate address for a specified interface. If there are any more specific routes in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes

Reference: http://www.cisco.com/en/US/docs/ios/12_0/np1/configuration/guide/1ceigrp.html

QUESTION 148

Refer to the exhibit.



Routers R1 and R2 are running EIGRP and have converged. On the basis of the information that is presented, which statement is true?

- A. All outgoing routing updates from router R1 to router R2 will be suppressed, but the inbound updates will continue to be received.
- B. All incoming routing updates from R2 will be suppressed, but the outgoing updates will continue to be sent.
- C. Both outgoing and incoming routing updates on R1 will be stopped because of the passive-interface Serial0/0 configuration statement.
- D. Both outgoing and incoming routing updates on R1 will be permitted because the distribute-list 20 out Serial0/0 command cannot be used with association with the outgoing interface.

Correct Answer: C

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

You can use the passive-interface command to control the advertisement of routing information. The command enables the suppression of routing updates over some interfaces while it allows updates to be exchanged normally over other interfaces. With most routing protocols, the passive-interface command restricts outgoing advertisements only. However, when used with Enhanced Interior Gateway Routing Protocol (EIGRP), the effect is slightly different. With EIGRP running on a network, the passive-interface command stops both outgoing and incoming routing updates, since the effect of the command causes the router to stop sending and receiving hello packets over an interface.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0a.shtml

QUESTION 149

What does the default value of the EIGRP variance command of 1 mean?

- A. Load balancing is disabled on this router.
- B. The router performs equal-cost load balancing.
- C. Only the path that is the feasible successor should be used.
- D. The router only performs equal-cost load balancing on all paths that have a metric greater than 1.

Correct Answer: B

Section: Mixed Questions

Explanation

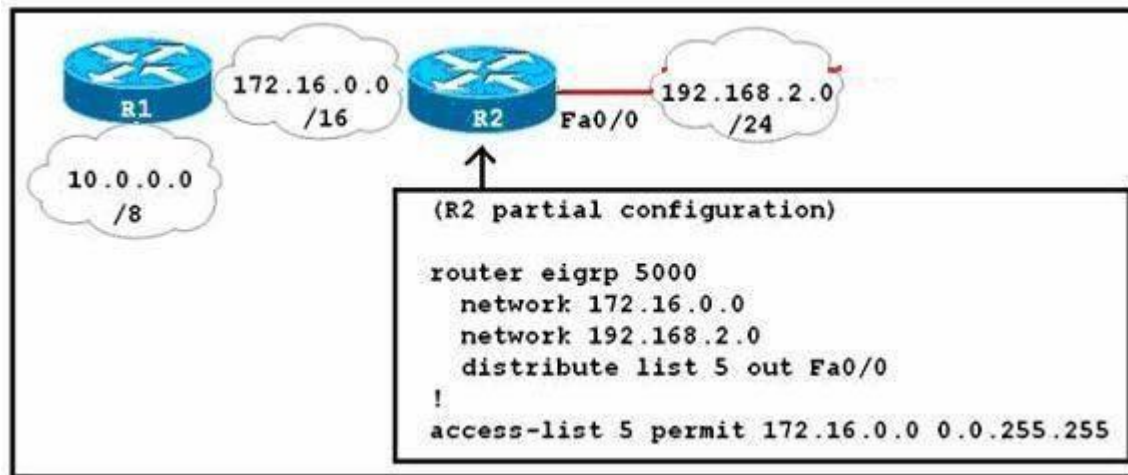
Explanation/Reference:

Explanations

The point of the question is about the balance configuration of EIGRP. If variance is 1, it means that it support Equal cost path.

QUESTION 150

Refer to the exhibit.



R1 and R2 have been configured to share routing information via EIGRP. What will be the result of the configuration section shown for R2?

- A. Any routes learned by R2 from the interface tied to the 172.16.0.0 network will not be advertised to neighbors on the 192.168.2.0 network.
- B. Only routes learned by R2 from the interface tied to the 172.16.0.0 network will be advertised to neighbors on the 192.168.2.0 network.
- C. Only the 172.16.0.0 network will be advertised to neighbors on the 192.168.2.0 network.
- D. All networks, except the 172.16.0.0 network will be advertised to neighbors on the 192.168.2.0 network.

Correct Answer: C


Section: Mixed Questions

Explanation

Explanation/Reference:

QUESTION 151

Refer to the exhibit.



Which two statements are true? (Choose two.)

- A. The eigrp stub command prevents queries from being sent from R2 to R1.
- B. The eigrp stub command will automatically enable summarization of routes on R2.
- C. The eigrp stub command prevents all routes except a default route from being advertised to R1.
- D. Router R1 will advertise connected and summary routes only.
- E. Router R1 will advertise connected and static routes. The sending of summary routes will not be permitted.
- F. Router R1 is configured as a receive-only neighbor and will not send any connected, static or summary routes.

Correct Answer: AD

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

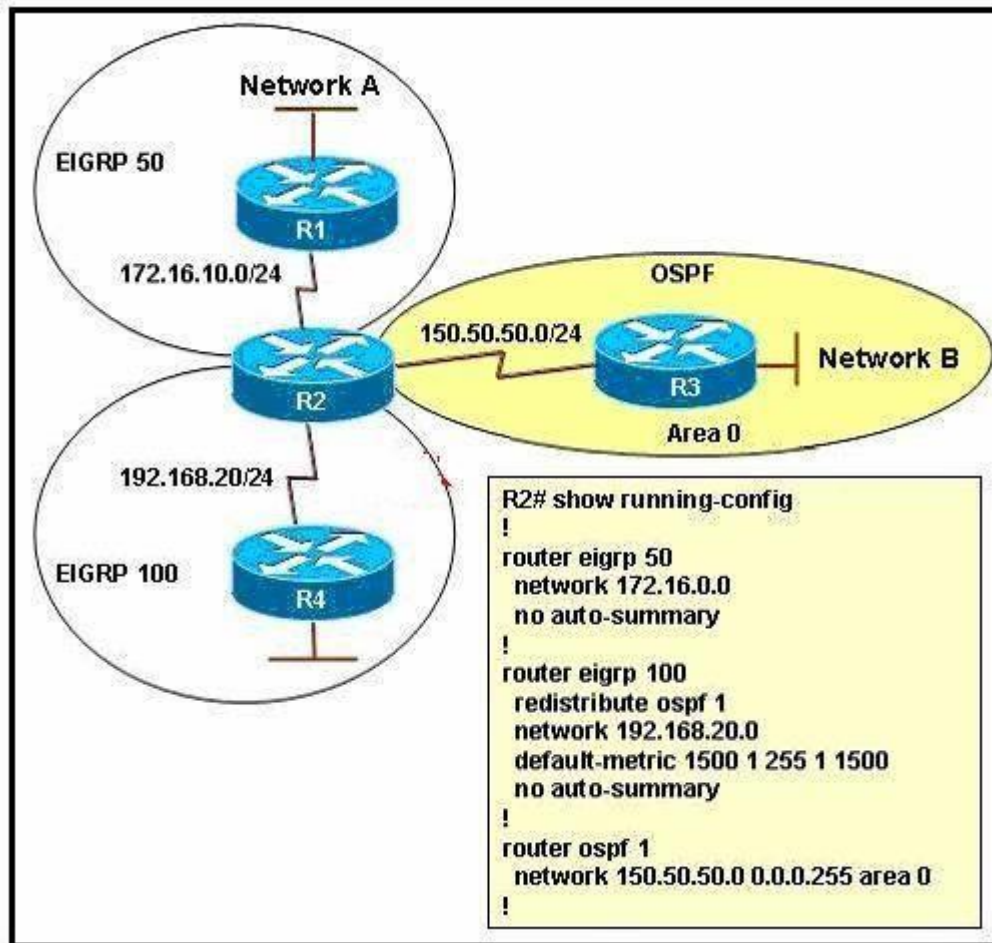
The command `eigrp stub` turns R1 into a stub router so R2 will never send any query to R1 because R2 knows that a stub router will only route packets for networks it has explicitly advertised.

The command `eigrp stub` is same as `eigrp stub connected summary` command because connected and summarized routes are advertised by default.

Note: Because the network 192.168.50.0 is not advertised by network statement, it is necessary to redistribute connected route with the `redistribute connected` command.

QUESTION 152

Refer to the exhibit.



The routing protocols EIGRP and OSPF have been configured as indicated in the exhibit. Given the partial configuration of router R2, which network will be present in the routing table of R4?

- A. Network A
- B. Network B
- C. Network A and Network B
- D. neither Network A nor Network B

Correct Answer: B

Section: Mixed Questions**Explanation****Explanation/Reference:**

Explanation:

In this exhibit the OSPF domain is redistributed into the EIGRP 100 domain so Network B will present into Router R-4. However, the Network A network will not be seen on router R-4 (The bottom router which is improperly labeled Network B) because EIGRP 50 was not redistributed into EIGRP 100.

QUESTION 153

When an EIGRP topology change is detected, what is the correct order of events when there is a FS?

- A. The neighbor adjacency is deleted.
The feasible route is used.
DUAL is notified.
Remove all topology entries learned from that neighbor.
- B. DUAL is notified.
Remove all topology entries learned from that neighbor.
The neighbor adjacency is deleted.
Routes enter the Active state and the feasible route is used.
- C. The neighbor adjacency is deleted.
Routes enter the Active state and the feasible route is used.
DUAL is notified.
Remove all topology entries learned from that neighbor.
- D. DUAL is notified.
The neighbor adjacency is deleted.
Remove all topology entries learned from that neighbor.
The feasible route is used.

Correct Answer: D

Section: Mixed Questions**Explanation****Explanation/Reference:**

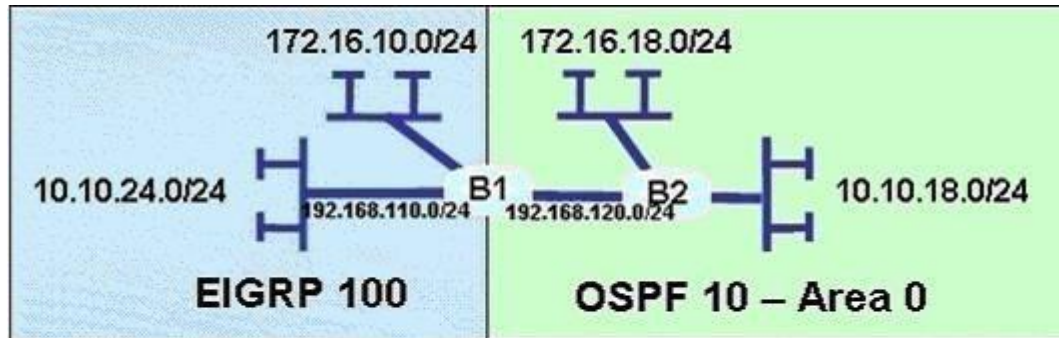
Explanation:

If a packet is not received before the expiration of the hold time, the neighbor adjacency is deleted, and all topology table entries learned from that neighbor are removed, as if the neighbor had sent an update stating that all the routes are unreachable. If the neighbor is a successor for any destination networks, those networks are removed from the routing table, and alternative paths, if available, are computed. This lets the routes quickly reconverge if an alternative feasible route is available.

QUESTION 154

Refer to the exhibit. The network administrator is trying to configure mutual redistribution between EIGRP and OSPF. Autosummarization in EIGRP 100 AS is disabled. After adding OSPF configuration to router B1, the network administrator checked the routing table of router B2, but none of the EIGRP

routes appeared there. To redistribute the EIGRP AS 100 routes into OSPF, which command should be added, or edited, on router B1 under router ospf 10?



```
B2# show ip route
 10.0.0.0/24 is subnetted, 1 subnets
C    10.10.18.0 is directly connected, Ethernet1/0
 172.16.0.0/24 is subnetted, 1 subnets
C    172.16.18.0 is directly connected, FastEthernet0/1
C    192.168.20.0/24 is directly connected, FastEthernet0/0
```

- A. redistribute eigrp 100 metric-type 1
- B. redistribute eigrp 100 subnets
- C. no auto-summary 10.0.0.0 255.0.0.0
- D. area 0 range 10.10.0.0 255.255.0.0

Correct Answer: B

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation

When redistributing into OSPF without keyword subnets, only classful networks will be redistributed. Classful networks here mean networks with the default major subnet masks (for example 10.0.0.0/8; 180.1.0.0/16; 200.200.200.0/24...).

In fact, the routing table on the exhibit above is not totally correct. The network 192.168.110.0/24 will be redistributed and shown in the routing table of

B2 even if the keyword subnets is not used because it belongs to class C with the default subnet mask of class C.

To make all the networks, including subnets appear in the routing table of B2 we must use keyword subnets when redistributing into OSPF. This is also an important thing to remember when redistributing into OSPF.

QUESTION 155

Which of the following are methods EIGRP uses to initially populate (seed) its EIGRP topology table, before learning topology data from neighbors? (Choose two.)

- A. By adding all subnets listed by the show ip route connected command
- B. By adding the subnets of working interfaces over which static neighbors have been defined
- C. By adding subnets redistributed on the local router from another routing source
- D. By adding all subnets listed by the show ip route static command

Correct Answer: BC

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation:

Other than the two listed correct answers, the local router also adds connected routes for which the network command matches the corresponding interfaces, so it may not add all connected routes. Also, EIGRP does not add static routes to the EIGRP topology table, unless those routes are redistributed.

QUESTION 156

An engineer has added the following configuration snippet to an implementation planning document. The configuration will be added to Router R1, whose Fa0/0 interface connects to a LAN to which Routers R2 and R3 also connect. R2 and R3 are already EIGRP neighbors with each other. Assuming the snippet shows all commands on R1 related to EIGRP authentication, which answer lists an appropriate comment to be made during the implementation plan peer review?

```
key chain fred
key 3
key-string whehew
interface fa0/0
ip authentication key-chain eigrp 9 fred
```

- A. The configuration is missing one authentication-related configuration command.
- B. The configuration is missing two authentication-related configuration commands.
- C. Authentication type 9 is not supported; type 5 should be used instead.
- D. The key numbers must begin with key 1, so change the key 3 command to key 1.

Correct Answer: A

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation:

The configuration requires the ip authentication mode eigrp asn md5 command, which is currently missing. This command enables MD5-style authentication, rather than the default of no authentication. Adding this one command completes the configuration. Any valid key numbers can be used. Also, the 9 in the ip authentication key-chain eigrp 9 fred command refers to the EIGRP ASN, not an authentication type.

QUESTION 157

Which of the following settings could prevent two potential EIGRP neighbors from becoming neighbors? (Choose two answers.)

- A. The interface used by one router to connect to the other router is passive in the EIGRP process.
- B. Duplicate EIGRP router IDs
- C. Mismatched Hold Timers.
- D. IP addresses of 10.1.1.1/24 and 10.2.2.2/24, respectively.

Correct Answer: AD

Section: Infrastructure Security

Explanation

Explanation/Reference:

Reference: <http://smitley.net/?p=167> (see configuration settings that could prevent neighbor relationships')

QUESTION 158

Based on the need to limit processing and bandwidth utilization due to dynamic routing protocol operation, the following routing requirements have been specified for your network.

- Partial and incremental routing updates
- Only the devices affected by a topology change perform route recomputation
- Route recomputation only occurs for routes that were affected

Which dynamic routing protocol should be deployed in your network to best meet these requirements?

- A. BGP
- B. OSPF
- C. IS-IS
- D. EIGRP
- E. RIPv2

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:


The bandwidth utilization issue has been addressed by implementing partial and incremental updates. Therefore, only when a topology change occurs does routing information get sent. Regarding processor utilization, the feasible successor technology greatly reduces the total processor utilization of an AS by requiring only the routers that were affected by a topology change to perform the route recomputation. Furthermore, the route recomputation only occurs for routes that were affected. Only those data structures are accessed and used. This greatly reduces search time in complex data structures.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f07.shtml (See frequently asked questions)

QUESTION 159

Refer to the exhibits. Router B should advertise the network connected to the E0/0/0 interface to router A and block all other network advertisements. The IP routing table on router A indicates that it is not receiving this prefix from router B. What is the probable cause of the problem?



```
RouterB#show ip eigrp neighbors
IP-EIGRP neighbors for process 1
H   Address                Interface    Hold Uptime   SRTT   RTO   Q   Seq Type
  (sec)          (ms)
0   10.255.1.1              Se0/0       138 00:15:26   16    582   0   3
RouterB#

RouterB#debug ip eigrp
IP-EIGRP Route Events debugging is on
02:17:54: IP-EIGRP: 192.168.3.16/28 - denied by distribute list
02:17:54: IP-EIGRP: 192.168.3.32/28 - denied by distribute list
02:17:54: IP-EIGRP: 192.168.3.0/30 - denied by distribute list
02:17:54: IP-EIGRP: 192.168.3.0/24 - denied by distribute list
02:17:54: IP-EIGRP: 10.0.0.0/8 - denied by distribute list
```

- A. An access list on router B is causing the 192.168.3.16/28 network to be denied.
- B. An access list on router B is causing the 192.168.3.32/28 network to be denied.
- C. The distribute list on router B is referencing a numbered access list that does not exist on router B.
- D. The distribute list on router B is referencing the wrong interface.

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:**Explanation**

This is an unclear question. The question says Router B should advertise the network connected to the E0/0/0 interface to router A and block all other network advertisements. The IP routing table on router A indicates that it is not receiving this prefix from router B. That means the network 192.168.3.16/28 (including the IP 192.168.3.21/28) is not received on router A -> A is the most suitable answer.

Note: Distribute list are used to filter routing updates and they are based on access lists.

QUESTION 160

Refer to the exhibit. Router B and router C are performing mutual redistribution between OSPF and EIGRP, and their default metrics are configured the same. Router D has equal cost paths to networks where both paths are not really equal cost. For example, network 172.16.54.0 shows equal cost through both router B and router C, though in reality the cost is greater using router C. Other routers, though not shown, are connected to the 172.16.54.0 and 172.16.55.0 networks, and the same issues exist to those routers and the networks connected to them. What can be done so that data will be routed along the most optimal path in the network?

- A. Redistribute connected interfaces on router B and router C.
- B. Set the maximum number of equal cost paths to 1 in all routers.
- C. When redistributing EIGRP into OSPF, set the external metric type to type E1.
- D. Adjust the default metrics in router B and router C so that the values are different in each router.
- E. None of these solutions will fix the problem. Migrate to a single dynamic routing protocol.

Correct Answer: E

Section: Network Principles

Explanation**Explanation/Reference:****Explanation**

From the output, we learn that all the External OSPF routes have metrics of 100 (the second parameters in [110/100]). This is not the default metric of OSPF Type 2 External route (the default value is 20) so the metrics of redistributed routes have been modified. Maybe when redistributing into OSPF, the metric in the redistribute command or the default-metric command was used on router B & C to assign the metric of these routes. Something like this:

```
router ospf 1
redistribute eigrp 1 metric 100 subnets
```

or

```
router ospf 1
.....
default-metric 100
```

Therefore even if we use the metric type E1 the problem still exists because the link B-D & C-D seems to have the same metric -> the total metrics remains the same .

We can use route-map and set different metrics for each networks but some unshown networks will have the same issues -

QUESTION 161

Based on the exhibited output,

```
R1# show ip eigrp topology

IP-EIGRP Topology Table for process 200

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status

P 192.168.1.64/28 1 successors, FD is 281600
    via Connected, Ethernet
  P 192.168.1.48/28, 1 successors, FD is 40512000
    via Connected, Serial1
P 192.168.1.48/28, 1 successors, FD is 40537600
    via 192.168.1.66 (40537600/40512000), Ethernet0
    via 192.168.1.17 (41024000/40512000), Serial0
    via 192.168.1.33 (41024000/40512000), Serial1
P 192.168.1.16/28 1 successors, FD is 40512000
    via Connected, Serial0
```

Which three statements are true? (Choose three.)

- A. R1 is in AS 200.
- B. R1 will load balance between three paths to reach the 192.168.1.48/28 prefix because all three paths have the same advertised distance (AD) of 40512000.
- C. The best path for R1 to reach the 192.168.1.48/28 prefix is via 192.168.1.66.
- D. 40512000 is the advertised distance (AD) via 192.168.1.66 to reach the 192.168.1.48/28 prefix.
- E. All the routes are in the passive mode because these routes are in the hold-down state.
- F. All the routes are in the passive mode because R1 is in the query process for those routes.

Correct Answer: ACD

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

It can be determined that AS 200 is used, from the fact that the IS-IS process ID is labeled as 200. The best path to reach the network 192.168.1.48/28 is the first one displayed in the routing table. This can be further demonstrated by the fact that the metric is less than the alternative route, via serial 0. Finally, the AD can be found by viewing the second number within the parentheses, which in this case is 40512000.

QUESTION 162

Study the exhibit carefully.

```
R1#show running-config
<Output omitted>
!
router eigrp 100
 network 172.16.0.0
 distribute-list prefix TEST out
 auto-summary
 no eigrp log-neighbor-changes
!
ip prefix-list TEST seq 5 permit 172.16.1.0/26
!
<Output omitted>
```

Router R1 is connected to networks 172.16.1.0 /26 and 172.16.1.64 /27. Based on the partial output in the exhibit, which description is correct?

- A. Router R1 should be reconfigured with an ACL instead of an ip prefix-list command.
- B. Router R1 will advertise both routes.
- C. Router R1 will deny the 172.16.1.0/27 route while permitting the 172.16.1.0/26 route to be advertised.
- D. Router R1 will deny the 172.16.1.0/26 route while permitting the 172.16.1.64/27 route to be advertised.

Correct Answer: C

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation

Prefix lists are configured with permit or deny keywords to either permit or deny the prefix based on the matching condition. A prefix list consists of an IP address and a bit mask. The IP address can be a classful network, a subnet, or a single host route. The bit mask is entered as a number from 1 to 32.

Prefix lists are configured to match an exact prefix length or a prefix range. The ge and le keywords are used to specify a range of the prefix lengths to match, providing more flexible configuration than can be configured with just the network/length argument. The prefix list is processed using an exact match when neither ge nor le keyword is entered.

Therefore in this case the exact 172.16.1.0/26 network is permitted while other networks are denied.

(Reference:

http://www.cisco.com/en/US/docs/ios/12_3t/ip_route/command/reference/ip2_i2gt.html)

QUESTION 163

Refer to the exhibit.

Core 1#show ip eigrp topology all-links

ip EIGRP Topology table for AS(65001) / ID (172.17.10.0)

Codes: P - P Passive, A - Active, U - Update, Q - Query,
R - Reply, r - reply, s - sales

P 172.17.3.128/25, 2 successors, FD is 30720, semo 9
via 172.17.10.2 (30720/28160), FastEthernet0/1

via 172.17.10.2 (30720/28160), FastEthernet0/3

P 10.140.0.0/24, 1 successors, FD is 28160, semo 16

via 172.17.3.2 (156160/128256), FastEthernet0/3

via 172.17.10.2 (157720/155160), FastEthernet0/1

P 172.17.10.0/24, 1 successors, FD is 28160, semo1

via Connected, FastEthernet0/1

P 172.17.0.0/30, 1 successors, FD is 20514560, semo 15

via 172.17.1.1 (20514560/205122000), FastEthernet0/2

via 172.17.10.2 (20516120/20513560), FastEthernet0/1

P 172.17.1.0/24, 1 successors, FD is 28160, semo2

via Connected, FastEthernet0/2

P 172.17.2.0/24, 1 successors, FD is 30720, semo 8

via 172.17.10.2 (30720/28160), FastEthernet0/1

via 172.17.3.2 (33280/30720), FastEthernet0/3

P 172.17.3.0/25, 1 successors, FD is 28160, semo 3

via Connected, FastEthernet0/3

Core 1#

BigBids Incorporated is a worldwide auction provider. The network uses EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's question.

Which two networks does the Core1 device have feasible successors for? (Choose two)

A. 172.17.0.0/30

B. 172.17.1.0/24

- C. 172.17.2.0/24
- D. 172.17.3.0/25
- E. 172.17.3.128/25
- F. 10.140.0.0/24

Correct Answer: AF

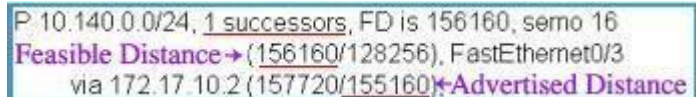
Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation

To understand the output of the show ip eigrp topology all-links command, let's analyze an entry (we choose the second entry because it is better for demonstration than the first one)



```
P: 10.140.0.0/24, 1 successors, FD is 156160, serno 16
Feasible Distance -> (156160/128256), FastEthernet0/3
via 172.17.10.2 (157720/155160) *Advertised Distance
```

The first line tells us there is only 1 successor for the path to 10.140.0.0/24 network but there are 2 lines below. So we can deduce that one line is used for successor and the other is used for another route to that network. Each of these two lines has 2 parameters: the first one (156160 or 157720) is the Feasible Distance (FD) and the second (128256 or 155160) is the Advertised Distance (AD) of that route.

The next thing we want to know is: if the route via 172.17.10.2 (the last line) would become the feasible successor for the 10.140.0.0/24 network. To figure out, we have to compare the Advertised Distance of that route with the Feasible Distance of the successor's route, if $AD < FD$ then it will become the feasible successor. In this case, because $AD (155160) < FD (156160)$ so it will become the feasible successor. Therefore we can conclude the network 10.140.0.0/24 has 1 feasible successor.

Because the question asks about feasible successor so we just need to focus on entries which have more paths than the number of successor. In this case, we find 3 entries that are in blue boxes because they have only 1 successor but has 2 paths, so the last path can be the feasible successor.

By comparing the value of AD (of that route) with the FD (of successor's route) we figure out there are 2 entries will have the feasible successor: the first and the second entry. The third entry has $AD = FD (30720)$ so we eliminate it.

QUESTION 164

Refer to the exhibit.

Core 1#show ip eigrp topology all-links

ip EIGRP Topology table for AS(65001) / ID (172.17.10.0)

Codes: P - P Passive, A - Active, U - Update, Q - Query,
R - Reply, r - reply, s - sales

P 172.17.3.128/25, 2 successors, FD is 30720, semo 9
via 172.17.10.2 (30720/28160), FastEthernet0/1

via 172.17.10.2 (30720/28160), FastEthernet0/3

P 10.140.0.0/24, 1 successors, FD is 28160, semo 16

via 172.17.3.2 (156160/128256), FastEthernet0/3

via 172.17.10.2 (157720/155160), FastEthernet0/1

P 172.17.10.0/24, 1 successors, FD is 28160, semo1

via Connected, FastEthernet0/1

P 172.17.0.0/30, 1 successors, FD is 20514560, semo 15

via 172.17.1.1 (20514560/205122000), FastEthernet0/2

via 172.17.10.2 (20516120/20513560), FastEthernet0/1

P 172.17.1.0/24, 1 successors, FD is 28160, semo2

via Connected, FastEthernet0/2

P 172.17.2.0/24, 1 successors, FD is 30720, semo 8

via 172.17.10.2 (30720/28160), FastEthernet0/1

via 172.17.3.2 (33280/30720), FastEthernet0/3

P 172.17.3.0/25, 1 successors, FD is 28160, semo 3

via Connected, FastEthernet0/3

Core 1#

BigBids Incorporated is a worldwide auction provider. The network uses EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's question.

Which three EIGRP routes will be installed for the 172.17.3.128/25 and 172.17.2.0/24 networks? (Choose three)

- A. 172.17.3.128.25 [90/28160] via 172.17.1 2, 01:26:35, FastEthernet0/2
- B. 172.17.3.128/25 [90/30720] via 172.17.3.2, 01:26:35, FastEthernet0/3

- C. 172.17.3.128/25 [90/30720] via 172.17.10.2, 01:26:35. FastEthernet0/1
- D. 172.17.2.0/24 [90/30720] via 172.17.10.2, 02:10:11, FastEthernet0/1
- E. 172.17.2.0/24 [90/28160] via 172.17.10.2, 02:10:11. FastEthernet0/1
- F. 172.17.2.0/24 [90/33280] via 172.17.3.2, 02:10:11. FastEthernet0/3

Correct Answer: BCD

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation

First indicate the positions of these networks:

Network 172.17.3.128/25 has 2 successors, therefore the two paths below are both successors.

Network 172.17.2.0/24 has only 1 successor, therefore the path lies right under it is the successor.

QUESTION 165

Refer to the exhibit.

Core 1#show ip eigrp topology all-links

ip EIGRP Topology table for AS(65001) / ID (172.17.10.0)

Codes: P - P Passive, A - Active, U - Update, Q - Query,
R - Reply, r - reply, s - sales

P 172.17.3.128/25, 2 successors, FD is 30720, semo 9
via 172.17.10.2 (30720/28160), FastEthernet0/1

via 172.17.10.2 (30720/28160), FastEthernet0/3

P 10.140.0.0/24, 1 successors, FD is 28160, semo 16

via 172.17.3.2 (156160/128256), FastEthernet0/3

via 172.17.10.2 (157720/155160), FastEthernet0/1

P 172.17.10.0/24, 1 successors, FD is 28160, semo1

via Connected, FastEthernet0/1

P 172.17.0.0/30, 1 successors, FD is 20514560, semo 15

via 172.17.1.1 (20514560/205122000), FastEthernet0/2

via 172.17.10.2 (20516120/20513560), FastEthernet0/1

P 172.17.1.0/24, 1 successors, FD is 28160, semo2

via Connected, FastEthernet0/2

P 172.17.2.0/24, 1 successors, FD is 30720, semo 8

via 172.17.10.2 (30720/28160), FastEthernet0/1

via 172.17.3.2 (33280/30720), FastEthernet0/3

P 172.17.3.0/25, 1 successors, FD is 28160, semo 3

via Connected, FastEthernet0/3

Core 1#

BigBids Incorporated is a worldwide auction provider. The network uses EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's question.

Which three networks is the router at 172.17.10.2 directly connected to? (Choose three)

A. 172.17.0.0/30

B. 172.17.1.0/24

- C. 172.17.2.0/24
- D. 172.17.3.0/25
- E. 172.17.3.128/25
- F. 172.17.10.0/24

Correct Answer: CEF

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

```
Core1#show ip eigrp topology all-links
IP EIGRP Topology table for AS(65001) / ID (172.17.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.17.3.128/25, 2 successors, FD is 30720, serno 9
   via 172.17.10.2 (30720/28160), FastEthernet0/1
   via 172.17.3.2 (30720/28160), FastEthernet0/3
P 10.140.0.0/24, 1 successors, FD is 156160, serno 16
   via 172.17.3.2 (156160/128256), FastEthernet0/3
   via 172.17.10.2 (157720/155160), FastEthernet0/1
P 172.17.10.0/24, 1 successors, FD is 28160, serno 1
   via Connected, FastEthernet0/1
P 172.17.0.0/30, 1 successors, FD is 20514560, serno 15
   via 172.17.1.1 (20514560/205122000), FastEthernet0/2
   via 172.17.10.2 (20516120/20513560), FastEthernet0/1
P 172.17.1.0/24, 1 successors, FD is 28160, serno 2
   via Connected, FastEthernet0/2
P 172.17.2.0/24, 1 successors, FD is 30720, serno 8
   via 172.17.10.2 (30720/28160), FastEthernet0/1
   via 172.17.3.2 (33280/30720), FastEthernet0/3
P 172.17.3.0/25, 1 successors, FD is 28160, serno 3
   via Connected, FastEthernet0/3
Core1#
```

First, we should notice about the entry in the orange box, it shows that the network 172.17.10.0/24 is directly connected with this router and has a FD of 28160. So we can guess the networks that directly connected with router at 172.17.10.2 will be shown with an AD of 28160. From that, we find out 3 networks which are directly connected to the router at 172.17.10.2 (they are green underlined). The network 172.17.10.0/24 is surely directly connected to the router at 172.17.10.2 (in fact it is the network that links the router at 172.17.10.2 with Core1 router).

QUESTION 166

Which two statements are true about EIGRP manual summarization? (Choose two.)

- A. Manual summarization is configured on a per interface basis.
- B. Manual summaries can be configured with the classful mask only.
- C. When manual summarization is configured, autosummarization is automatically disabled by default.
- D. The summary address is assigned an administrative distance of 10 by default.
- E. The summary address is entered into the routing table and is shown to be sourced from the Null0 interface.

Correct Answer: AE

Section: Infrastructure Services

Explanation


Explanation/Reference:

Explanation:

You can configure a summary addresses on a per-interface basis. You need to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary or if you want to use summary addresses on a ASA with automatic route summarization disabled. If any more specific routes are in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

QUESTION 167

Refer to the exhibit.



Implementation Plan

1. Establish a traffic throughput baseline.
2. Configure variance on R1 and R4.
3. Use traceroute to validate load balancing has been activated.
4. Establish a new traffic throughput baseline.
5. Compare the new and old baselines and verify that load balancing is implemented as desired.

ROUTE.com is planning to implement EIGRP load balancing for traffic between hosts on the 172.16.10.0/24 and 172.16.20.0/24 networks. You have been asked to review the implementation plan for this project. Which statement about the plan is true?

- A. It is complete as written.
- B. It should include a task to configure multipath to equal a value of 2 on R1 and R4.
- C. It should use a ping instead of a traceroute to validate that load balancing has been activated.
- D. It should contain a task that documents the changes made to the configurations.

Correct Answer: A

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

This implementation plan is complete because it has all the requirements for an EIGRP load balancing process.

QUESTION 168

A network administrator is managing a hub-and-spoke network with EIGRP routing that has been enabled. The hub router is trying to query a remote router. However, delays are occurring that are caused by certain paths being stuck in active (SIA). How should the administrator configure EIGRP in order to limit the scope of the query range and prevent SIA from occurring?

- A. Configure the hub router with a scope limit of 1.
- B. Configure the remote router with a scope limit of 1.
- C. Configure the hub to indicate that the remote router is a stub router.
- D. Configure the hub and remote router as stub routers.
- E. Configure the remote router as a stub router.
- F. Disable the SIA feature of EIGRP on the remote router.

Correct Answer: E

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Configuring a router as a stub also helps the rest of the network. Queries are responded to much quicker and convergence happens much faster. Sometimes a query can cause delays that result in the path being SIA. If the stub configuration is applied, the router responds to queries as inaccessible, thus limiting the scope of the query range and preventing SIA from occurring.

QUESTION 169

What administrative distance is given to EIGRP summary routes?

- A. 0
- B. 1
- C. 5
- D. 90
- E. 95
- F. 170

Correct Answer: C

Section: Network Principles

Explanation

Explanation/Reference:

Reference:

http://www.cisco.com/en/US/docs/ios/iproute_eigrp/command/reference/ire_i1.html (See usage guidelines)

QUESTION 170

What are two possible causes for EIGRP Stuck-In-Active routers? (Choose Two)

- A. Some query or reply packets are lost between the routers.
- B. The neighboring router starts receiving route updates from this router.
- C. A failure causes traffic on a link between two neighboring routers to flow in only one direction (unidirectional link).
- D. The neighboring router stops receiving ACK packets from this router.

Correct Answer: AC

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

Generally, a route shown as Active is going to be there for a very short period of time by the time you repeat the command, hopefully that Active route has gone Passive. Sometimes that doesn't happen, though, and the route becomes SIA - Stuck In Active.

A route becomes SIA when a query goes unanswered for so long that the neighbor relationship is reset. From experience, I can tell you that troubleshooting SIA routes is more of an art form than a science, but there are four main reasons a route becomes SIA:

The link is unidirectional, so the query can't possibly be answered. The queried router's resources are unavailable, generally due to high CPU utilization. The queried router's memory is corrupt or otherwise unable to allow the router to answer the query.

The link between the two routers is of low quality, allowing just enough packets through to keep the neighbor relationship intact, but not good enough to allow the replies through.

To sum it up, routes generally become SIA when a neighbor either doesn't answer a query, or either the query or reply took a wrong turn somewhere. I told you it wasn't the easiest thing to troubleshoot!

QUESTION 171

When configuring EIGRP to run across a 56 Kbps serial PPP link, what command do you need to put under the serial interface ensure proper convergence of EIGRP routes?

- A. bandwidth 56
- B. bandwidth 56000
- C. ip bandwidth-percent eigrp 1 56
- D. ip bandwidth-percent eigrp 1 56000

Correct Answer: A
Section: VPN Technologies
Explanation

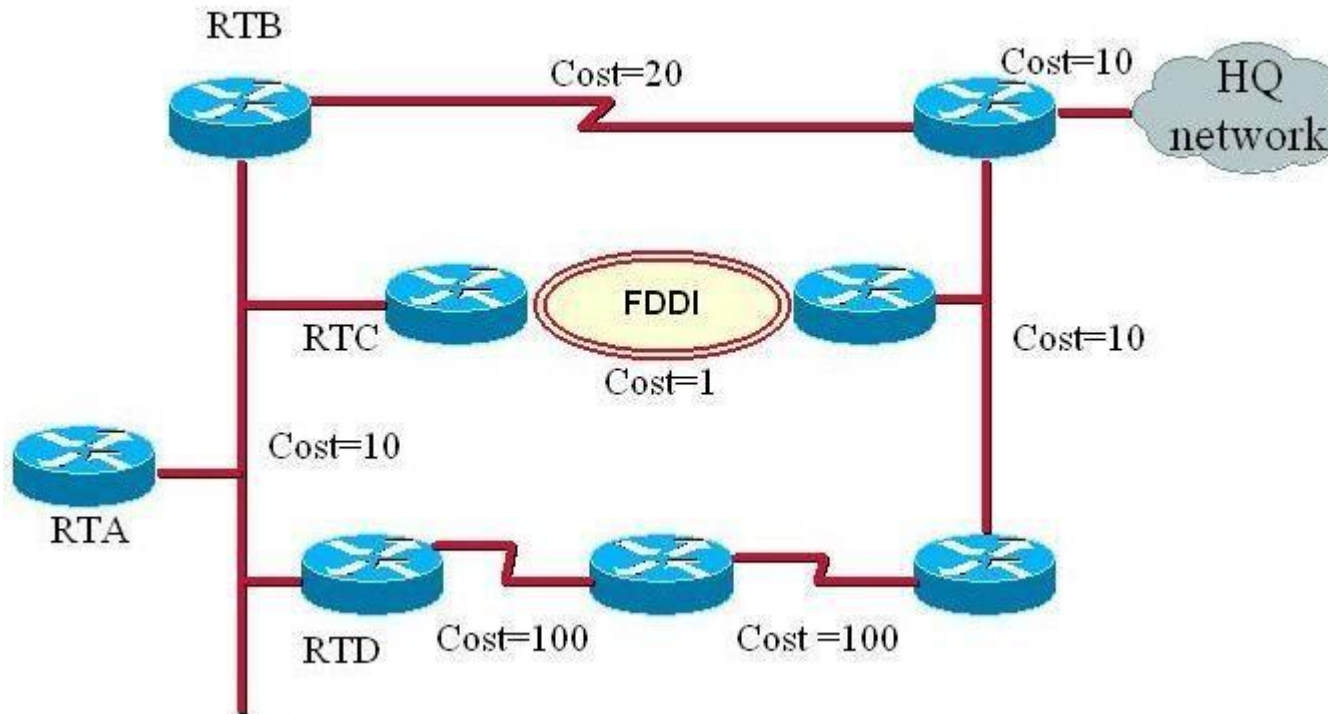
Explanation/Reference:

Explanation:

When configuring serial links using EIGRP it is important to configure the bandwidth setting on the interface. If the bandwidth setting is not changed for these interfaces EIGRP assumes the default bandwidth on the link instead of the true bandwidth. If the link is slower, the router may not be able to converge, routing updates might become lost, or suboptimal path selection may result. Router(config-if)#bandwidth kilobits The value, kilobits, indicates the intended bandwidth in kilobits per second. For generic serial interfaces, such as PPP or HDLC, set the bandwidth to the line speed.

QUESTION 172

Refer to the Exhibit.



Routers in the Diagram are configured with EIGRP. If RTB and RTC fail, which action will RTA take with respect to the HQ network?

- A. RTA will automatically route packets via RTD to the HQ network.
- B. RTA will place the route via RTD into the hold down state.
- C. RTA will go into the active state for all routers.
- D. RTA will go into the active state for the route to HQ network.

Correct Answer: D

Section: Infrastructure Services

Explanation


Explanation/Reference:

Explanation:

When RTB and RTC fails, RTA will go into active state for the HQ network route.

QUESTION 173

Exhibit:



Refer to the topology diagram R3 is redistributing the EIGRP routers into OSPF. What will the EIGRP routes appear in the routing table of R1?

- A. O
- B. O IA
- C. E2
- D. D
- E. D EX

Correct Answer: C

Section: Infrastructure Services**Explanation****Explanation/Reference:**

Reference: http://kowon.dongseo.ac.kr/~nok60/lecture/CCNP1-V50/Labs/CCNP1v50_L5-2.pdf (page 10 and 11)

QUESTION 174

In EIGRP, when the IP default-network command is configured on a router, what is generated in the router's configuration?

- A. A static route
- B. A directly connected route
- C. An EIGRP route
- D. A default route

Correct Answer: A

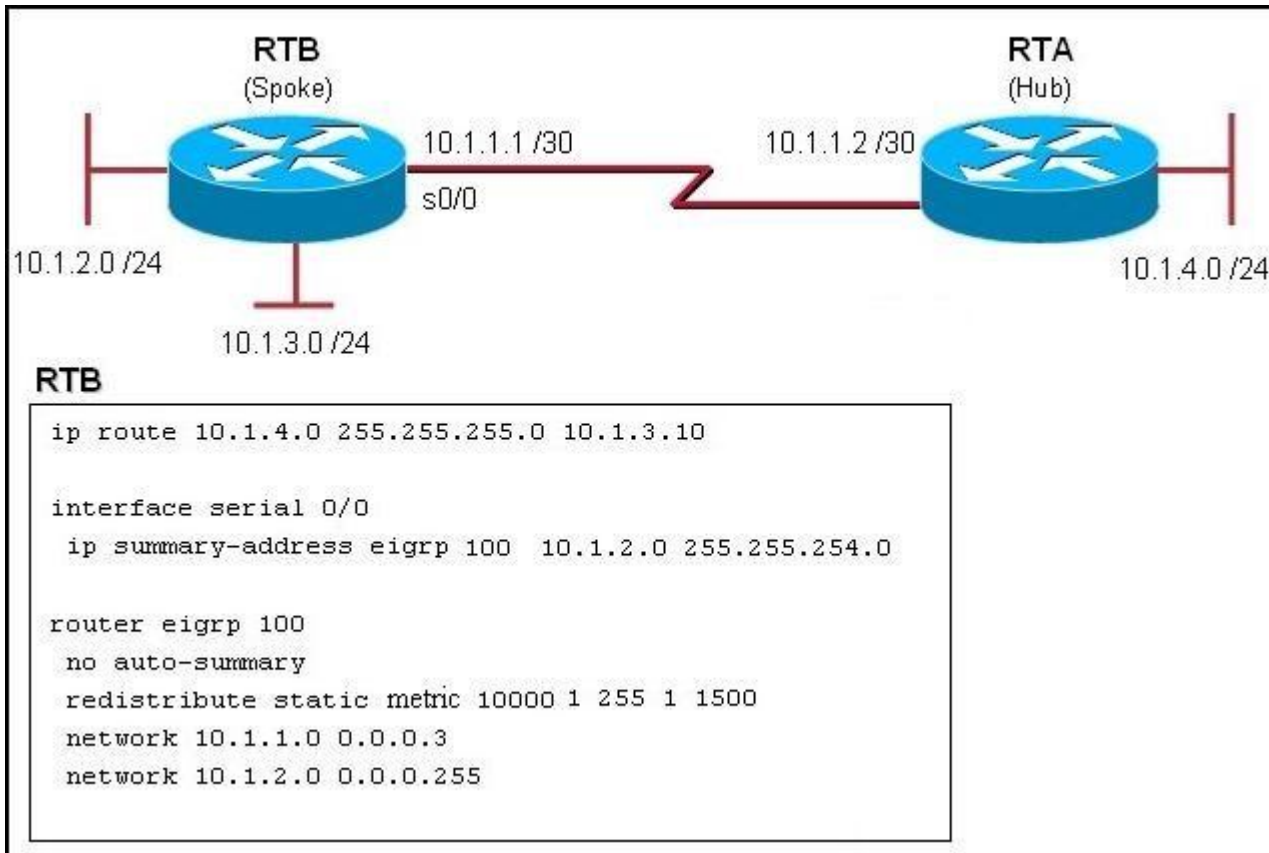
Section: Infrastructure Services**Explanation****Explanation/Reference:**

Explanation:

When you configure the ip default-network command and specify a subnet, a static route (the ip route command) is generated in the router's configuration; however, the IOS does not display a message to indicate that this has been done. The entry appears as a static route in the routing table of the router where the command is configured. This can be confusing when you want to remove the default network; the configuration must be removed with the no ip route command, not with the no ip default-network command.

QUESTION 175

Refer to the exhibit.



Which router configuration command can be given that will restrict router RTB from sharing its routing information with router RTA?

- A. the eigrp stub command on router RTA
- B. the eigrp stub command on router RTB
- C. the eigrp stub connected command on router RTA
- D. the eigrp stub connected command on router RTB
- E. the eigrp stub receive-only command on router RTA
- F. the eigrp stub receive-only command on router RTB

Correct Answer: F

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

This is a hub and spoke network, so EIGRP stub receive-only command on RTB will restrict the router from sharing its routing information with RTA.

QUESTION 176

The following command was issued on R2

```
R2#sho ip rou
<output omitted>
C    10.1.1.0 is directly connected, FastEthernet0/0
D    172.16.0.0/16 [90/156160] via 10.1.1.1, 00:07:48, FastEthernet0/0
D EX 192.168.1.0/24 [170/1308160] via 10.1.1.1, 00:00:11, FastEthernet0/0
```

Given the above output, which statement is true?

- A. 192.168.1.0 is a static route.
- B. 192.168.1.0 is a summarized route.
- C. 192.168.1.0 is a redistributed route into EIGRP.
- D. 192.168.1.0 is equal path load balancing with 172.16.1.0.

Correct Answer: C

Section: Mixed Questions

Explanation

Explanation/Reference:

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009487e.shtml (administrative distance, second para)

QUESTION 177

Which two types of routes will be advertised with the EIGRP configuration as shown? (Choose two.)

```
router eigrp 100
network 10.0.0.0
eigrp stub
```

- A. static
- B. receive-only

- C. summary
- D. stub
- E. connected
- F. dynamic

Correct Answer: CE

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

The Enhanced Interior Gateway Routing Protocol (EIGRP) Stub Routing feature improves network stability, reduces resource utilization, and simplifies stub router configuration. Stub routing is commonly used in a hub and spoke network topology. In a hub and spoke network, one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is adjacent only to one or more distribution routers. The only route for IP traffic to follow into the remote router is through a distribution router. This type of configuration is commonly used in WAN topologies where the distribution router is directly connected to a WAN. The distribution router can be connected to many more remote routers. Often, the distribution router will be connected to 100 or more remote routers. In a hub and spoke topology, the remote router must forward all nonlocal traffic to a distribution router, so it becomes unnecessary for the remote router to hold a complete routing table. Generally, the distribution router need not send anything more than a default route to the remote router.

When using the EIGRP Stub Routing feature, you need to configure the distribution and remote routers to use EIGRP, and to configure only the remote router as a stub. Only specified routes are propagated from the remote (stub) router. The router responds to queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message "inaccessible." A router that is configured as a stub will send a special peer information packet to all neighboring routers to report its status as a stub router.

Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router will depend on the distribution router to send the proper updates to all peers. Router(config- router)#eigrp stub [receive-only | connected | static | summary] :Configures a remote router as an EIGRP stub router.

QUESTION 178

Refer to the exhibit.

```
Router# show ip route
```

```
...
```

```
C    10.1.3.0/24 is directly connected, Serial2
D    10.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1
D    10.0.0.0/8 is a summary, 00:23:20, Null0
C    10.1.1.0/24 is directly connected, Serial1
S    192.168.20.0/24 is directly connected, Ethernet0
```

What happens when the router stops receiving advertisements for the 10.1.2.0/24 network?

- A. The summary route will be removed from the table.
- B. The summary route will remain in the table.
- C. The more specific routes will be advertised from the table.
- D. 10.1.2.0/24 will still be advertised but packets destined for it will be dropped when they reach this router.

Correct Answer: B

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:


If you look very closely at the routing table output, we can conclude that R1 has "auto- summary" enabled under the EIGRP routing process.

D 10.0.0.0/8 is a summary, 00:23:20, Null 0

Anytime the "auto-summary" is enabled under the routing process the router will install a summary route to "null 0" as long as the router has one or more subnets within the "classful" network. In the case above, you have two directly connected interfaces (Serial1 & Serial2) that are within the "classful" network. Therefore, regardless of whether you learn a route via EIGRP that is in the "classful" network, R1 will still install this summary route to "null 0".

QUESTION 179

Refer to the exhibit.



On all routers in the network, EIGRP has been configured for load balancing across the three links. However, traffic destined for Network B from R1 is only load balanced over paths R1- R2-R5 and R1-R3-R5. What is the cause of the problem?

- A. EIGRP will not select more than two links for unequal cost path load balancing.
- B. Because the path has a different link type, EIGRP will not select path R1-R4-R5 for load balancing.
- C. Because Router R4 is not a feasible successor, EIGRP will not select path R1-R4-R5 for load balancing.
- D. EIGRP will not select path R1-R4-R5 for load balancing unless the value of the variance parameter is increased.

Correct Answer: C

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

Since R4 is not configured as a feasible successor, EIGRP will not select that path for load balancing. IN EIGRP, you need to configure feasible successor to enable load balancing on the path.

QUESTION 180

Identify three characteristics of EIGRP feasible successors? (Choose three.)

- A. A feasible successor is selected by comparing the advertised distance of a non-successor route to the feasible distance of the best route.
- B. If the advertised distance of the non-successor route is less than the feasible distance of best route, then that route is identified as a feasible successor.
- C. If the successor becomes unavailable, then the feasible successor can be used immediately without recalculating for a lost route.
- D. The feasible successor can be found in the routing table.
- E. Traffic will be load balanced between feasible successors with the same advertised distance.

Correct Answer: ABC

Section: Mixed Questions

Explanation

Explanation/Reference:

Reference: <http://packetlife.net/blog/2010/aug/9/eigrp-feasible-successor-routes/>

QUESTION 181

Which are three features of EIGRP? (Choose three)

- A. Support VLSM and discontinuous subnets
- B. Link-state protocol
- C. Partial routing updates
- D. External Administrative distance is 100
- E. Fast convergence.
- F. Used by other vendors than Cisco

Correct Answer: ACE

Section: Mixed Questions


Explanation

Explanation/Reference:

Reference: http://cisco.jjc.edu/cnt205/ch2/2_1_1/index.html (See first three bullets)

QUESTION 182

Refer to exhibit.



RouterA (DR) failed, and after 10 minutes it came back. Which two statements are true? (Choose two)

- A. RouterA is a DR
- B. RouterA is a BDR
- C. RouterA is a DROTHER
- D. RouterB is a DR
- E. RouterB is a BDR
- F. RouterC is a DROTHER

Correct Answer: CD

Section: Mixed Questions

Explanation

Explanation/Reference:

Reference: <http://packetlife.net/blog/2011/jun/2/ospf-designated-router-election/>

QUESTION 183

Which of the below mentioned conditions form a neighbor relation in EIGRP? (Choose Three)

- A. Hello or ACK received
- B. AS number match
- C. Hello timer match
- D. Identical metric(k values)
- E. Dead Timer Match
- F. Network Time Match

Correct Answer: ABD

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

To form neighbor relationship in EIGRP, these conditions must be met:

- * Pass the authentication process
- * Have the same configured AS number
- * Must believe that the source IP address of a received Hello is in that router's primary connected subnet on that interface
- * Match K values

The third item means that the primary ip address of the neighbor must be in the same subnet with the primary ip address of the received interface. But in this case the primary ip address of router A is 10.10.10.1/30 and it is not in the same subnet with the primary ip address of router B 10.10.10.6/30 -> no EIGRP neighbor relationship is formed.

QUESTION 184

Refer to exhibit.

```
R1#show ip eigrp topology all-links
ip-eigrp topology table for AS(1)/ID(192.168.1.0)
codes: P-Passive , A-Active , U-Update , Q-Query , R-Reply , r-reply status , s-sia status
P 192.168.1.0/24, 1 successors, FD is 21152000, sermo 4
    via summary (21152000/0), Null 0
    via 172.16.3.2(41024000/30118400) seria 0/0/0
P 192.168.1.4/30, 1 successors, FD is 21152000, sermo 2 via connected, serial 0/0/1
P 192.168.1.0/24, 1 successors, FD is 2297856, sermo 6
    via 198.18.10.6 (2297856/39260), s0/0/1
    via 172.16.3.2 (41026560/3128695) , serial 0/0/0
P 192.168.1.8/30, 1 successors, FD is 3523840, sermo 2
    via 192.168.1.6 (3523840/3011840), serial 0/0/1
```

The exhibit shows R1 topology table to reach 192.168.1.0/24 network. Which route(s) will be installed in routing table of R1 to reach network 192.168.1.0/24 after configuring R1 with the following command?

Router(config-router)# variance 2

- A. R2 only
- B. R2 and R3
- C. R2 and R4

D. R2, R3 and R4

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation:

EIGRP will only use equal-cost load-balancing feature even when the variance command is used. However, if you use both the traffic-share min command and variance command, even though traffic is sent over the minimum-cost path only, all feasible routes get installed into the routing table, which decreases the convergence times

QUESTION 185

Refer to the following.

Router # sh ip route eigrp

13.0.0.0/8 is variably subnetted, 2 subnets, 2 masks

D 13.0.0.0/8 is a summary, 00:00:32, Null0

What happens to packets that are forwarded from the 13.0.0.0/8 network to the Null0 interface?

- A. Flagged
- B. Accepted
- C. Summarized
- D. Dropped

Correct Answer: D

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

When an EIGRP router summarizes, it automatically builds a route to null0 for the summarized route. The route to null0 prevents packets that do not match a specific entry in the routing table from following a default route. (The route to null0 causes the packet to be dropped).

QUESTION 186

In which state do DR and BDR establish adjacency with each OSPF router in the network?

- A. Init State
- B. Exstart State
- C. Exchange State

D. Loading State

Correct Answer: A

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation: DR and BDR will only establish adjacency with each OSPF router on broadcast multiaccess networks. So Init state is the correct answer.

QUESTION 187

A stub area is typically created using what kind of topology?

- A. Broadcast
- B. Point-to-point
- C. Hub and spoke
- D. Full Mesh

Correct Answer: C

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

A stub area is typically created when you have a hub-and-spoke topology, with the spoke being the stub area, such as a branch office. In this case, the branch office does not need to know about every network at the headquarters site; instead, it can use a default route to get there.

QUESTION 188

A network administrator is troubleshooting an EIGRP connection between RouterA, IP address 10.1.2.1, and RouterB, IP address 10.1.2.2.

```
RouterA# debug eigrp packets
...
01:39:13: EIGRP: Received HELLO on Serial0/0 nbr 10.1.2.2
01:39:13: AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iadbQ un/rely 0/0 peerQ un/rely 0/0
01:39:13:          K-value mismatch
```

Given the debug output on RouterA, which two statements are true?

- A. RouterA received a hello packet with mismatched autonomous system numbers.

- B. RouterA received a hello packet with mismatched hello timers.
- C. RouterA received a hello packet with mismatched authentication parameters.
- D. RouterA received a hello packet with mismatched metric-calculation mechanisms.
- E. RouterA will form an adjacency with RouterB.
- F. RouterA will not form an adjacency with RouterB.

Correct Answer: DF

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

If the k-value mismatch occurs, Router A will never form an adjacency with Router B since it is one of the basic requirements of adjacency. If you see the exhibit, Router A received HELLO packet with a mismatched metric.

QUESTION 189

Refer to the exhibit.

Correct Answer: B

Section: Network Principles

Explanation

Explanation/Reference:

Explanation:

A router that is configured as a stub with the `eigrp stub` command shares connected and summary routing information with all neighbor routers by default. Four optional keywords can be used with the `eigrp stub` command to modify this behavior:

- receive-only
- connected
- static
- summary
-

This section provides configuration examples for all forms of the `eigrp stub` command. The `eigrp stub` command can be modified with several options, and these options can be used in any combination except for the `receive-only` keyword. The `receive-only` keyword will restrict the router from sharing any of its routes with any other router in that EIGRP autonomous system, and the `receive-only` keyword will not permit any other option to be specified because it prevents any type of route from being sent. The three other optional keywords (`connected`, `static`, and `summary`) can be used in any combination but cannot be used with the `receive-only` keyword. If any of these three keywords is used individually with the `eigrp stub` command, `connected` and `summary` routes will not be sent automatically.

The `connected` keyword will permit the EIGRP Stub Routing feature to send connected routes. If the connected routes are not covered by a network statement, it may be necessary to redistribute connected routes with the `redistribute connected` command under the EIGRP process. This option is enabled by default.

The `static` keyword will permit the EIGRP Stub Routing feature to send static routes. Without the configuration of this option, EIGRP will not send any static routes, including internal static routes that normally would be automatically redistributed. It will still be necessary to redistribute static routes with the `redistribute static` command.

The `summary` keyword will permit the EIGRP Stub Routing feature to send summary routes. Summary routes can be created manually with the `summary address` command or automatically at a major network border router with the `auto-summary` command enabled. This option is enabled by default.

In the following example, the `eigrp stub` command is used to configure the router as a stub that advertises connected and summary routes:

```
router eigrp 1
network 10.0.0.0
eigrp stub
```

In the following example, the `eigrp stub` command is issued with the `connected` and `static` keywords to configure the router as a stub that advertises connected and static routes (sending summary routes will not be permitted):

```
router eigrp 1
network 10.0.0.0
eigrp stub connected static
```

In the following example, the `eigrp stub` command is issued with the `receive-only` keyword to configure the router as a receive-only neighbor (Connected, summary, and static routes will not be sent):

```
router eigrp 1
network 10.0.0.0 eigrp
eigrp stub receive-only
```

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00_80087026.html

QUESTION 190

Which configuration command is used to enable EIGRP unequal-cost path load balancing?

- A. `maximum-paths`
- B. `distance`
- C. `metric`
- D. `variance`
- E. `default-metric`

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009437d.shtml (See traffic sharing
<http://networkninja.co.za/page/2/?s=botha&cat=plus-5-results>

QUESTION 191

If the primary path goes down, what will EIGRP use to reach a destination?

- A. administrative distance

- B. advertised successor
- C. successor
- D. feasible successor

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

The key to this question is the four terminology about DUAL. Enhanced Interior Gateway Routing Protocol (EIGRP) is an enhanced distance-vector protocol based on the diffusing update algorithm (DUAL). It is capable of (conservatively) finding all loop-free paths to any given destination based on route advertisements from neighbors. The neighbor (or neighbors) with the best path to a destination is called the successor. The remaining neighbors with loop-free paths to the destination are called feasible successors. To reduce traffic load on the network, EIGRP maintains neighbor relationships and exchanges routing information only as needed, using a query process to find alternate paths when all loop-free paths to a destination have failed.

QUESTION 192

Refer to the output.

```
Routing Process "ospfv3 1" with ID 172.16.3.3
It is an autonomous system boundary router
Redistributing External Routes from,
static
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
Numbers of external LSA 1. Checksum sum 0x218D
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Area 1
Number of interfaces in this area is 2
SPF algorithm executed 9 times
Number of LSA 15. Checksum Sum 0x67581
Number of DCbitless LSA 0
Number of indication LSA 0
Number of DoNotAge LSA 0
Flood list length 0
```

What IOS command produces this output?

Select the best response

- A. show ip ospf
- B. show ip ospf interface
- C. show ipv6 ospf interface
- D. show ipv6 ospf

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

Reference:

http://www.cisco.com/en/US/docs/ios/ipv6/command/reference/ipv6_15.html#wp2439467

QUESTION 193

How is authentication handled with OSPFv3?

- A. OSPFv3 for IPv6 authentication is supported by IPv6 IPsec.
- B. OSPFv3 for IPv6 authentication is supported by MD5 authentication.
- C. OSPFv3 for IPv6 authentication is supported by IPv4 IPsec.
- D. OSPFv3 for IPv6 authentication is supported by SHA-1 authentication.

Correct Answer: A

Section: Layer 2 Technologies


Explanation

Explanation/Reference:

Reference: http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_ospf/configuration/xen-3s/ip6-route-ospfv3-auth-ipsec-xe.html

QUESTION 194

Refer to the exhibit.



Routers R1 and R2 have been configured to operate with OSPF. Routers R1 and R3 have been configured to operate with RIP. After configuring the redistribution between OSPF and RIP on R1, no OSPF routes are distributed into RIP. What should be done to correct this problem?

- A. The redistribution command should be reentered with the match route-type parameter included.
- B. The redistribution command should be reentered with the route-map map-tag parameter included.
- C. The redistribution command should be reentered with the metric metric-value parameter included.
- D. Routes will first need to be distributed into another protocol, and then into RIP.

Correct Answer: C

Section: Layer 2 Technologies

Explanation

Explanation/Reference:

Explanation

Notice that RIP metric is based on hop count only, and the maximum valid metric is 15. Anything above 15 is considered infinite. By default, when no metric is assigned when redistributing from EIGRP, OSPF, IS-IS, BGP into RIP, the default metric will be infinite. Therefore we must define a metric that is understandable to the receiving protocol. Usually, we should use a small value (like 1, 2, 3) so that after redistributing, that route can be advertised through many routers (because the limit is 15).

QUESTION 195

By default, which statement is correct regarding the redistribution of routes from other routing protocols into OSPF?

- A. They will appear in the OSPF routing table as type E1 routes.
- B. They will appear in the OSPF routing table as type E2 routes.
- C. Summarized routes are not accepted.
- D. All imported routes will be automatically summarized when possible.
- E. Only routes with lower administrative distances will be imported.

Correct Answer: B

Section: Layer 3 Technologies

Explanation

Explanation/Reference:

Explanation

Type E1 external routes calculate the cost by adding the external cost to the internal cost of each link that the packet crosses while the external cost of E2 packet routes is always the external cost only. E2 is useful if you do not want internal routing to determine the path. E1 is useful when internal routing should be included in path selection. E2 is the default external metric when redistributing routes from other routing protocols into OSPF.

QUESTION 196

When implementing OSPFv3, which statement describes the configuration of OSPF areas?

- A. In interface configuration mode, the OSPFv3 area ID combination assigns interfaces to OSPFv3 areas.
- B. In router configuration mode, the network wildcard area ID combination assigns networks to OSPFv3 areas.
- C. In interface configuration mode, the IPv6 OSPF process area ID combination assigns interfaces to OSPFv3 areas.
- D. In router configuration mode, the IPv6 OSPF interface area ID combination assigns interfaces to OSPFv3 areas.

Correct Answer: C

Section: Layer 3 Technologies

Explanation

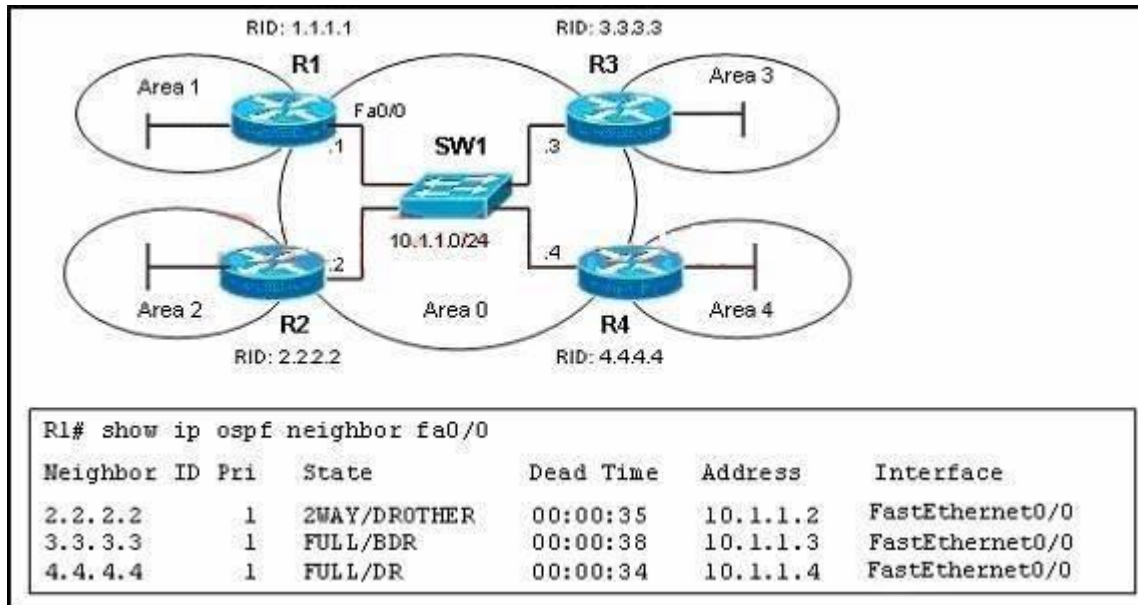
Explanation/Reference:

Reference:

[http://www.hh.se/download/18.4cf286ee134f03ddb7b800015/1326882212358/Chapter3_VT 2012.pdf](http://www.hh.se/download/18.4cf286ee134f03ddb7b800015/1326882212358/Chapter3_VT+2012.pdf) (slide 42)

QUESTION 197

Refer to the exhibit. OSPF is configured on all routers in the network. On the basis of the show ip ospf neighbor output, what prevents R1 from establishing a full adjacency with R2?



- A. Router R1 will only establish full adjacency with the DR and BDR on broadcast multiaccess networks.
- B. Router R2 has been elected as a DR for the broadcast multiaccess network in OSPF area
- C. Routers R1 and R2 are configured as stub routers for OSPF area 1 and OSPF area 2.
- D. Router R1 and R2 are configured for a virtual link between OSPF area 1 and OSPF area 2.
- E. The Hello parameters on routers R1 and R2 do not match.

Correct Answer: A

Section: VPN Technologies

Explanation

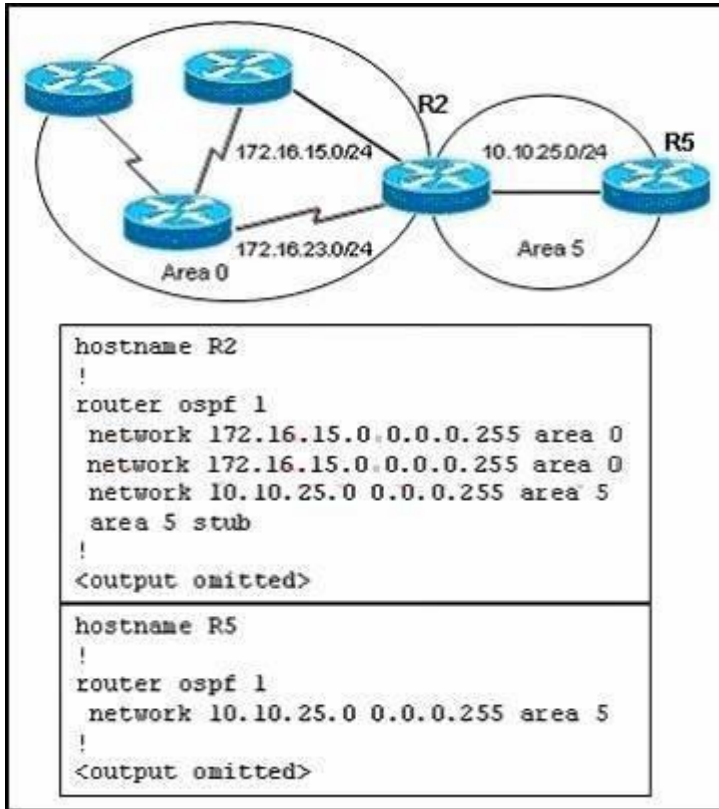
Explanation/Reference:

Explanation

From the output, we learn that R4 is the DR and R3 is the BDR so other routers will only establish full adjacency with these routers. All other routers have the two-way adjacency established.

QUESTION 198

Refer to the exhibit.



On the basis of the configuration provided, how are the Hello packets sent by R2 handled by R5 in OSPF area 5?

- A. The Hello packets will be exchanged and adjacency will be established between routers R2 and R5.
- B. The Hello packets will be exchanged but the routers R2 and R5 will become neighbors only.
- C. The Hello packets will be dropped and no adjacency will be established between routers R2 and R5.
- D. The Hello packets will be dropped but the routers R2 and R5 will become neighbors.

Correct Answer: C

Section: VPN Technologies

Explanation

Explanation/Reference:

Explanation:

The point of this question is the conditions of OSPF establish adjacency relationship. For ospf, the optional capabilities must set the same between

neighbors, but from the exhibit, R5 was configured as a stub area while R2 in area 0 is a normal area. So there will be no adjacent relationship established between routers R2 and R5.

QUESTION 199

Which statement is true about OSPF Network LSAs?

- A. They are originated by every router in the OPSF network. They include all routers on the link, interfaces, the cost of the link, and any known neighbor on the link.
- B. They are originated by the DR on every multi-access network. They include all attached routers including the DR itself.
- C. They are originated by Area Border Routers and are sent into a single area to advertise destinations outside that area.
- D. They are originated by Area Border Router and are sent into a single area to advertise an Autonomous System Border Router.

Correct Answer: B

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation:

The point of this question is OSPF Network LSAs

The feature of OSPF Network LSAs is that they are generated by DR, and DR only exist on multi-access network, the use of OSPF Network LSAs is that it list all neighbors around and send it to every router which run OSPF.

Incorrect Answer: OSPF Network LSAs are not originated by Area Border Routers.

QUESTION 200

Refer to the exhibit.

```
NProuter#debug ip ospf events
```

```
OSPF events debugging is on
```

```
NProuter #
```

```
00:02:03: OSPF: Rcv hello from 172.16.1.1 area 0 from Serial0/0 10.1.1.1
```

```
00:02:03: OSPF: Mismatched hello parameters from 10.1.1.1
```

```
00:02:03: OSPF: Dead R 120 C 10, Hello R 30 C 30
```

```
00:02:26: OSPF: Rcv hellofrom 192.168.1.2 area 0 from Serial0/0 10.1.1.2
```

```
00:02:26: OSPF: Mismatched hello parameters from 10.1.1.2
```

```
00:02:26: OSPF: Dead R 120 C 10, Hello R 30 C 30
```

You are the network administrator responsible for the NRouter, the 10.1.1.1 router, and the 10.1.1.2 router. What can you determine about the OSPF operations from the debug output?

- A. The NRouter has two OSPF neighbors in the "Full" adjacency state.
- B. The NRouter serial0/0 interface has the OSPF dead timer set to 10 seconds.
- C. The NRouter serial0/0 interface has been configured with an OSPF network type of "point-to-point".
- D. The 10.1.1.1 and 10.1.1.2 routers are not using the default OSPF dead and hello timers setting.
- E. The "Mismatched" error is caused by the expiration of the OSPF timers.

Correct Answer: B

Section: Infrastructure Security

Explanation

Explanation/Reference:

Explanation

First we should understand clearly about the line

Dead R 120 C 10, Hello R 30 C 30

The R here means Received and C means Configured. In other words, Dead R is the Dead Timer Received from the neighbor and the Dead C is the Dead Timer of the local router.

Therefore in this case Dead R 120 C 10 means the Death Timer of the neighbor is 120 seconds while the local Dead Timer is 10 seconds, which causes a mismatch. Also we can learn that the local OSPF dead timer is set to 10 seconds.

For your information, by default, OSPF uses a 10-second hello timer and 40-second hold timer on broadcast and point-to-point links, and a 30-second hello timer and 120-second hold timer for all other network types.

QUESTION 201

You have just completed an OSPF implementation. While executing your verification plan, you determine that R1 is not able to establish full OSPF adjacency with R2. The show ip ospf neighbor command output on R1 shows that R2 is stuck in the INIT state.

What could be the cause of this problem?

- A. DR and BDR election errors between R1 and R2.
- B. The R2 router has not received the OSPF hello packets from the R1 router.
- C. Mismatched interface maximum transmission unit (MTU) configuration between the R1 and R2.
- D. Mismatched OSPF hello interval configuration between the R1 and R2.
- E. Corrupted LSAs exchanges between the R1 and R2.

Correct Answer: B

Section: Infrastructure Services

Explanation

Explanation/Reference:**Explanation**


When a router receives an OSPF Hello from a neighbor, it sends the Hello packet by including that neighbor's router ID in the Hello packet. If the neighbor does not receive this packet (means that it doesn't see itself in this packet), it will be stuck in INIT state. INIT state can be understood as a one-way Hello. An example of a router stuck in INIT state is shown below:

```
R1# show ip ospf neighbor fa0/0
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.2	1	INIT/-	00:00:35	10.1.1.2	Fast Ethernet0/0

QUESTION 202

Refer to the exhibit. You have completed an OSPF implementation, and you are verifying OSPF operation. You notice that router A and router B are stuck in the two-way state. From the show ip ospf interface command output, what is the cause of this issue?



```

RouterA# show ip ospf int s1/0
Serial 1/0 is up, line protocol is up
Internet Address 10.100.3.253/30, Area 1
Process ID 1, Router ID 1.1.1.1, Network Type BROADCAST, Cost:64
Transmit Delay is 1 sec, State DROTHER, Priority 0
No designated router on this network
No backup designated router on this network
Old designated Router (ID) 2.2.2.2, Interface address 10.100.3.254
Flush timer for old DR LSA due in 00:01:12
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  oob-resync timeout 40
  Hello due in 00:00:01
Supports Link-Local Signaling (LLS)
Index 1/2, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 3
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 1, Adjacent neighbor count is 0
Suppress hello for 0 neighbor(s)
  
```

```

RouterB# show ip ospf int s1/0
Serial 1/0 is up, line protocol is up
Internet Address 10.100.3.254/30, Area 1
Process ID 1, Router ID 2.2.2.2, Network Type BROADCAST, Cost:64
Transmit Delay is 1 sec, State DROTHER, Priority 0
No designated router on this network
No backup designated router on this network
Old designated Router (ID) 2.2.2.2, Interface address 10.100.3.254
Flush timer for old DR LSA due in 00:01:58
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  oob-resync timeout 40
  Hello due in 00:00:03
Supports Link-Local Signaling (LLS)
Index 1/1, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 1
Last flood scan time is 0 msec, maximum is 4 msec
Neighbor Count is 1, Adjacent neighbor count is 0
Suppress hello for 0 neighbor(s)
  
```

- A. All OSPF implementations must have at least one interface in area 0.
- B. You are attempting to run in the broadcast mode over an NBMA interface.
- C. Both routers are configured to function as a BDR; therefore, there is no DR router.
- D. Someone has changed the OSPF router ID; therefore you must clear the OSPF process.
- E. The OSPF priority is set to 0 on both routers; therefore neither can become the DR.

Correct Answer: E

Section: Infrastructure Services

Explanation

Explanation/Reference:

Explanation

When OSPF adjacency is formed, a router goes through several state changes before it becomes fully adjacent with its neighbor. The states are Down, Attempt, Init, 2-Way, Exstart, Exchange, Loading, and Full.

An OSPF neighbor reaches the 2-way state when bidirectional communication is established (each router has seen the other's hello packet). This is the beginning of an OSPF adjacency. On broadcast media and non-broadcast multiaccess networks, the DR and BDR are elected in this state. But the priority on both routers are 0 so no DR and BDR are elected -> These routers stay in the 2-way state.

(Reference and a good resource of OSPF Neighbor states:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0e.shtml)

QUESTION 203

You have completed an OSPF implementation, and you are verifying OSPF operation. During this verification, you notice that the OSPF route of 172.16.10.0 is repeatedly appearing and disappearing from the routing table. Further investigation finds that the OSPF CPU utilization is very high and the routers are constantly performing SPF calculations. You determine that 172.16.20.2 is the source of the 172.16.10.0 route. Using the show ip ospf database router 172.16.20.1 command, you notice that when this show command is performed repeatedly, the contents of the LSA change every few seconds.

What could be the cause of this problem?

- A. OSPF authentication errors between some of the routers.
- B. Two routers have the same OSPF router ID.
- C. Issues with mistuned OSPF timers.
- D. OSPF LSA pacing issues between some of the routers.
- E. OSPF neighbor adjacency problems between some of the routers.

Correct Answer: B

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

When two routers use the same router ID in an OSPF domain, routing possibly does not work correctly. Cisco bug IDs CSCdr61598 and CSCdu08678 enhance the detection and reporting mechanisms of duplicate router IDs. Access the Bug Toolkit (registered customers only) in order to view additional information about these Cisco bug ID.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080117102.shtml

QUESTION 204

When an OSPF design is planned, which implementation can help a router not have memory resource issues?

- A. Have a backbone area (area 0) with 40 routers and use default routes to reach external destinations.
- B. Have a backbone area (area 0) with 4 routers and 30,000 external routes injected into OSPF.
- C. Have less OSPF areas to reduce the need for interarea route summarizations.
- D. Have multiple OSPF processes on each OSPF router. Example, router ospf 1, router ospf 2.

Correct Answer: A

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation:

Memory issues usually come up when too many external routes are injected in the OSPF domain. A backbone area with 40 routers and a default route to the outside world would have less memory issues compared with a backbone area with 4 routers and 33,000 external routes being injected into OSPF. Router memory could also be conserved by using a good OSPF design. Summarization at the area border routers and use of stub areas could further minimize the number of routes exchanged.

The total memory used by OSPF is the sum of the memory used in the routing table (show ip route summary) and the memory used in the LSDB. The following numbers are a "rule of thumb" estimate. Each entry in the routing table will consume between approximately 200 and 280 bytes plus 44 bytes per extra path. Each LSA will consume a 100 byte overhead plus the size of the actual LSA, possibly another 60 to 100 bytes (For router links, this depends on the number of interfaces on the router). These amounts should be added to memory already used by other processes and by the IOS itself.

If you really want to know the exact number, you can do a show memory with and without OSPF being turned on. The difference in the processor memory used would be the answer.

QUESTION 205

The maximum number of routers per OSPF area typically depends on which three factors? (Choose three.)

- A. the kind of OSPF areas being implemented
- B. the number of external LSAs in the network

- C. the number of DRs and BDRs in the areas
- D. the number of virtual links in the areas
- E. how well the areas can be summarized
- F. the use of LSA filters

Correct Answer: ABE

Section: Mixed Questions

Explanation

Explanation/Reference:

Reference:

https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=6&cad=rja&ved=0C FwQFjAF&url=http%3A%2F%2Ffaculty.valenciacollege.edu%2Fwyousif%2FCCNP%2FSe mester5%2FPresentations%2FMAOSPF_P2.ppt&ei=VUurUbmOA9OThgeDhYDoCg&usg=AFQjCNE5mLCAUIWCzou_vUX_DGhOOwcYxw&sig2=_7fgBDpXZCFi0Tay60wYmw&bvm=bv.47244034,d.ZG4 (Slide 85)

QUESTION 206

You are troubleshooting an OSPF problem where external routes are not showing up in the OSPF database. Which two options are valid checks that should be performed first to verify proper OSPF operation? (Choose two.)

- A. Are the ASBRs trying to redistribute the external routes into a totally stubby area?
- B. Are the ABRs configured with stubby areas?
- C. Is the subnets keyword being used with the redistribution command?
- D. Is backbone area (area 0) contiguous?
- E. Is the CPU utilization of the routers high?

Correct Answer: AC

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

A totally stubby area cannot have an ASBR so it will discard this type of LSA (LSA Type 5) -> A is a valid check.

Each stubby area needs an ABR to communicate with other areas so it is normal -> B is not a valid check.

When pulling routes into OSPF, we need to use the keyword subnets so that subnets will be redistributed too. For example, if we redistribute these EIGRP routes into OSPF:

+ 10.0.0.0/8
+ 10.10.0.0/16

+ 10.10.1.0/24

without the keyword subnets

```
router ospf 1  
redistribute eigrp 1
```

Then only 10.0.0.0/8 network will be redistributed because other routes are not classful routes, they are subnets. To redistribute subnets we must use the keyword subnets

```
router ospf 1  
redistribute eigrp 1 subnets
```

-> C is a valid check.

We don't need to care if area 0 is contiguous or not -> D is not a valid check. CPU utilization cannot be the cause for this problem -> E is not a valid check.

QUESTION 207

When verifying the OSPF link state database, which type of LSAs should you expect to see within the different OSPF area types? (Choose three.)

- A. All OSPF routers in stubby areas can have type 3 LSAs in their database.
- B. All OSPF routers in stubby areas can have type 7 LSAs in their database.
- C. All OSPF routers in totally stubby areas can have type 3 LSAs in their database.
- D. All OSPF routers in totally stubby areas can have type 7 LSAs in their database.
- E. All OSPF routers in NSSA areas can have type 3 LSAs in their database.
- F. All OSPF routers in NSSA areas can have type 7 LSAs in their database.

Correct Answer: AEF

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

Below summarizes the LSA Types allowed and not allowed in area types:

Area Type	Type 1 & 2 (within area)	Type 3 (from other areas)	Type 4	Type 5	Type 7
Standard & backbone	Yes	Yes	Yes	Yes	No
Stub	Yes	Yes	No	No	No
Totally stubby	Yes	No	No	No	No
NSSA	Yes	Yes	No	No	Yes
Totally stubby NSSA	Yes	No	No	No	Yes

Popular LSA Types are listed below:

LSA Type	Description	Details
1	Router LSA	Generated by all routers in an area to describe their directly attached links
2	Network LSA	Advertised by the DR of the broadcast network (does not cross ABR)
3	Summary LSA	Advertised by the ABR of originating area
4	Summary LSA	Generated by the ABR of the originating area to advertise an ASBR to all other areas in the autonomous system
5	AS external LSA	Used by the ASBR to advertise networks from other autonomous systems
7	Defined for NSSAs	Generated by an ASBR inside a Not-so-stubby area (NSSA) to describe routes redistributed into the NSSA

QUESTION 208

When verifying OSPF virtual link problems, which is an important item to check on the two transit OSPF routers?

- A. OSPF process ID
- B. OSPF router ID
- C. OSPF network type
- D. OSPF memory usage
- E. OSPF CPU utilization
- F. OSPF stub area configurations

Correct Answer: B

Section: Mixed Questions

Explanation


Explanation/Reference:

Explanation

The OSPF router IDs of the two transit OSPF routers are used to form the virtual link (with the area area-id virtual-link neighbor-router-id command) so it is an important item to check - > B is correct.

QUESTION 209

Refer to the exhibit.




A network administrator wants to reduce the number of OSPF routes advertised from Area 1 into Area 2. As the router configuration specialist, what two things would you do to accomplish this goal? (Choose two.)

- A. Enter the configuration on router B1.
- B. Enter the configuration on router B4.
- C. On the same router, enter the Summary-address 10.10.1.0 255.255.255.128 subcommand.
- D. On the same router, enter the Area 1 range 10.10.1.0 255.255.255.128 subcommand.
- E. On the same router, enter the Area 2 range 10.10.1.0 255.255.255.128 subcommand.

Correct Answer: AD
Section: Mixed Questions
Explanation

Explanation/Reference:

QUESTION 210
Refer to the exhibit.



A company would prefer all Internet-bound OSPF routed traffic to use ISP ABC with ISP DEF as a backup. As the network consultant, what three configuration changes should you make? (Choose three.)

- A. The default-information originate command should be configured on router B1 and B4.
- B. The default-information originate command should be configured on router B2 and B3.

- C. If the metric value for ISP ABC is set at the default, the ISP DEF metric value should be set to 1.
- D. If the metric value for ISP ABC is set at the default, the ISP DEF metric value should be set to 25.
- E. The metric type value should be set to type 1.
- F. The metric type value should be set to type 2.

Correct Answer: BDF

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

Routers B2 & B3 need to advertise a default route to the Internet for inside OSPF routers so we should use the default-information originate command with a default route (something like ip route 0.0.0.0 0.0.0.0) pointing to the ISP router.

If no metric is specified, OSPF puts a default value of 20 when redistributing routes from all protocols except BGP routes (BGP routes get a metric of 1). We use ISP DEF as a backup so its metric value should be set to a higher value than 20.

There are two types of external routes: external type 1 and external type 2. The difference between the two is in the way the cost (metric) of the route is being calculated:

- + The cost of a type 2 route is always the external cost, irrespective of the interior cost to reach that route.
- + Type 1 cost is the addition of the external cost and the internal cost used to reach that route.

-> We should configure the type 2 external route to make sure the ISP ABC is always referred over ISP DEF because internal routing does not determine the path. Note: E2 is the default external metric, but E1 is preferred over E2 if two equal-cost paths exist.

QUESTION 211

The administrator wants to verify the current state of the OSPF database loading process.
Which show command should the administrator use?

- A. show ip ospf [process-id] interface
- B. show ip ospf neighbor
- C. show ip ospf [process-id]
- D. show ip ospf [process-id area-id] database

Correct Answer: B

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation


The show ip ospf neighbor command can be used to view the current state of the OSPF database loading process. In the output below we can see router 2.2.2.2 is in 2way state, router 3.3.3.3 is elected as the BDR & router 4.4.4.4 is the BR.

```
R1# show ip ospf neighbor fa0/0
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.2	1	2WAY/DR0THER	00:00:35	10.1.1.2	Fast Ethernet0/0
3.3.3.3	1	FULL/BDR	00:00:38	10.1.1.3	Fast Ethernet0/0
4.4.4.4	1	FULL/BR	00:00:34	10.1.1.4	Fast Ethernet0/0

QUESTION 212

Refer to the exhibit. OSPF is running throughout the network. You want to minimize the propagation of LSAs into and out of Area 1. Which OSPF feature would best achieve this goal?



- A. stub
- B. totally stubby
- C. NSSA
- D. totally NSSA

Correct Answer: D

Section: Mixed Questions

Explanation

Explanation/Reference:

Explanation

We need to redistribute RIP from R1 to Area 1 so Area 3 cannot be a stub or totally stubby area. To minimize the propagation of LSAs into and out of

Area 1 we should configure it as a totally NSSA. Notice that a NSSA allows LSA Type 3 & 7 while a Totally NSSA only allows LSA Type 7

Note:

Both Totally Stubby Area & Totally Stubby NSSA do not accept external AS routes or inter- area routes (LSA Types 3, 4 and 5). They recognize only intra-area routes and the default route 0.0.0.0. The main difference between them is Totally Stubby NSSA accepts routes from other AS while Totally Stubby Area does not.

Below summarizes the LSA Types allowed and not allowed in area types:

Area Type	Type 1 & 2 (within area)	Type 3 (from other areas)	Type 4	Type 5	Type 7
Standard & backbone	Yes	Yes	Yes	Yes	No
Stub	Yes	Yes	No	No	No
Totally stubby	Yes	No	No	No	No
NSSA	Yes	Yes	No	No	Yes
Totally (stubby) NSSA	Yes	No	No	No	Yes

QUESTION 213

Refer to the exhibit.

```
ASBR
router ospf 123
 redistribute eigrp 1 route-map eigrp-to-ospf
 !
 route-map eigrp-to-ospf permit 10
  match ip address prefix-list private
  set tag 255
 route-map eigrp-to-ospf permit 20
 !
 ip prefix-list private permit 10.0.0.0/8 ge 8 le 30
 ip prefix-list private permit 172.16.0.0/11 ge 11 le 30
 ip prefix-list private permit 192.168.0.0/16 ge 16 le 30
```

```
Area 0 ABR
router ospf 123
 distribute-list route-map private-filter in
 !
 route-map private-filter deny 10
  match tag 255
 route-map private-filter permit 20
```

The partial configuration for an OSPF ASBR and an Area 0 ABR is shown. Assume the OSPF configurations throughout the network are operable. Which statement about these configurations is true?

- A. The ASBR route-maps are basically useless, because there are no deny prefix-lists.
- B. LSA Type 5s will not be received by the ABR from the ASBR.
- C. The OSPF backbone will not learn any RFC 1918 addresses.
- D. The matched prefix-list addresses will be given a metric of 255, which is essentially unreachable.

Correct Answer: C

Section: Network Principles

Explanation

Explanation/Reference:

Explanation

The ASBR accepts RFC 1918 addresses and set these networks to tag 255 but when advertising into Area 0, the ABR Area 0 filters out these networks because they match tag 255 so the OSPF backbone will not learn any RFC 1918 addresses.

Note that if you use an ACL in a route-map deny clause, routes that are permitted by the ACL are not redistributed.

All the networks with tag 255 are blocked by the clause 10 while all other networks are permitted by the clause 20 of the route-map (if a match command is not present, all routes match the clause).

Note:

RFC 1918 addresses include:

+ Class A: 10.0.0.0 10.255.255.255 (10/8 prefix)

+ Class B: 172.16.0.0 172.31.255.255 (172.16/12 prefix) + Class C: 192.168.0.0 192.168.255.255 (192.168/16 prefix)

QUESTION 214

Refer to the exhibit.

```
Router eigrp 10
  Network 10.0.0.0
Router ospf 100
  redistribute eigrp 10 subnets
  network 172.16.0.0 0.0.255.255 area 1
```

.....

A partial routing table

- C 10.10.10.16/28 is directly connected, Ethernet1/0
- D 10.10.10.64/26 [90/284160] via 10.10.24.2, 01:54:39, Ethernet1/0
- C 172.16.10.0/24 is directly connected, FastEthernet0/1

Examine the partial configuration and the routing table excerpt. Which routes would be redistributed into OSPF area 1?

- A. 10.10.10.16/28 only
- B. 10.10.10.16/28 and 10.10.10.64/26
- C. 10.10.10.16/28, 10.10.10.64/26, and 172.16.10.0/24
- D. 10.10.10.64/26 only

Correct Answer: D

Section: Network Principles

Explanation

Explanation/Reference:

QUESTION 215

Which two statements about route redistribution when implementing OSPF are true? (Choose two.)

- A. Routes learned using any IP routing protocol can only be redistributed into non IP routing protocols.
- B. OSPF can import routes learned using EIGRP, RIP, and IS-IS.
- C. OSPF routes cannot be exported into EIGRP, RIP, and IS-IS.
- D. At the interdomain level, OSPF cannot import routes learned using BGP.
- E. OSPF routes can be exported into BGP.

Correct Answer: BE

Section: Layer 2 Technologies**Explanation****Explanation/Reference:**

Reference:

http://www.cisco.com/en/US/docs/switches/lan/catalyst3750x_3560x/software/release/12.2_5_5_se/configuration/guide/swiprout.html (See configuring OSPF)

QUESTION 216

Which two routing interface parameters are supported in OSPF implementations? (Choose two.)

- A. retransmit-interval
- B. dead-interval
- C. stub area
- D. virtual link
- E. NSSA area

Correct Answer: AB

Section: Layer 2 Technologies**Explanation****Explanation/Reference:**

Explanation

When OSPF sends an advertisement to an adjacent router, it expects to receive an acknowledgment from that neighbor. If no acknowledgment is received, the router will retransmit the advertisement to its neighbor. The retransmit-interval timer controls the number of seconds between retransmissions. To edit the retransmit-interval, use the `ip ospf retransmit-interval seconds` in interface configuration mode.

Dead-interval is the number of seconds without hello packets before an adjacency is declared down. To edit the dead-interval, use the `ip ospf dead-interval seconds` in interface configuration mode.

Other answers are not correct because they are not interface parameters.

QUESTION 217

RouterA#

~~~~~

!

router ospf 1

log-adjacency-changes

network 10.0.0.0 0.255.255.255 area 1

network 172.16.1.0 0.0.0.255 area 1

!




~~~~~  
RouterB#
~~~~~

```
router ospf 1
log-adjacency-changes
network 10.0.0.0 0.255.255.255 area 2
network 172.16.2.0 0.0.0.255 area 2
!
```

~~~~~  
RouterC#

```
~~~~~
!
router ospf 1
log-adjacency-changes
network 10.0.0.0 0.255.255.255 area 0
!
```

~~~~~  
Refer to the exhibits. You are verifying your OSPF implementation, and it does not seem to be functioning properly. What can you conclude from the exhibit and the show running- configuration command output?



- A. The OSPF areas are not configured correctly.
- B. The wildcard masks for the 10.x.x.x networks are incorrect.
- C. The 172.16.x.x networks need to be connected to area 0 using virtual links.
- D. The 172.16.x.x networks are discontinuous. OSPF is automatically summarizing them to 172.16.0.0/16 and data is being "black holed."
- E. There is not enough information to make a determination.

**Correct Answer:** A

**Section:** Layer 3 Technologies

**Explanation**

**Explanation/Reference:**

Explanation

The E0/0 & E0/1 interfaces of router C belong to area 0 while E0/0 of router A belongs to area 1; E0/0 of router B belongs to area 2 -> it is not correct. Both E0/0 interfaces of router A & B should be in area 0 -> A is not correct.

#### **QUESTION 218**

OSPF is enabled on router A. You execute the following command on router A and receive the accompanying output:

```
RouterA#ping 224.0.0.5 repeat 1
```

```
Type escape sequence to abort.
```

```
Sending 1, 100-byte ICMP Echos to 224.0.0.5, timeout is 2 seconds:
```

```
Reply to request 0 from 10.100.100.1, 4 ms
```

10.100.100.1 is the IP address of a loopback interface on router A.

What can you conclude about router A?

- A. Only the router A loopback interface is participating in the OSPF routing process.
- B. None of the router A interfaces are participating in the OSPF routing process.
- C. Router A is using the loopback interface IP address as its OSPF router ID.
- D. Router A does not have any reachable OSPF neighbors.

**Correct Answer:** D

**Section:** Layer 3 Technologies

**Explanation**


**Explanation/Reference:**

Explanation

The 224.0.0.5 address is the multicast for OSPF routers. Therefore when you ping to this address all OSPF routers will reply but from the output we learn that only the local router responded -> Router A does not have any reachable OSPF neighbors.

**QUESTION 219**

Refer to the exhibit.



What is the effect of the OSPF configuration on router B?

- A. All interfaces will be in area 0.
- B. The router will be an ABR with s1/0 in area 0 and f0/0 and 0/1 in area 1.
- C. The router will be an ABR with s1/0 in area 0 and f0/0 and 0/1 in area 2.
- D. The router will be an ABR with s1/0 in area 0, f0/0 in area 1, and f0/1 in area 2.

**Correct Answer:** D

**Section:** Layer 3 Technologies

**Explanation**


**Explanation/Reference:**

Reference:

[https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=4&cad=rja&ved=0C EIQFjAD&url=https%3A%2F%2Flearningnetwork.cisco.com%2F servlet%2FJiveServlet%2Fdownload%2F4907-2332%2FCase%2520..&ei=sU-rUeGvFlanhAfR-IHwDQ&usg=AFQjCNH4zqER\\_WNM rHEY2VhuMzDTy7D1xQ&sig2=v1xIJWspSrBE6sz oQAP8cg&bvm=bv.47244034,d.ZG4](https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=4&cad=rja&ved=0C EIQFjAD&url=https%3A%2F%2Flearningnetwork.cisco.com%2F servlet%2FJiveServlet%2Fdownload%2F4907-2332%2FCase%2520..&ei=sU-rUeGvFlanhAfR-IHwDQ&usg=AFQjCNH4zqER_WNM rHEY2VhuMzDTy7D1xQ&sig2=v1xIJWspSrBE6sz oQAP8cg&bvm=bv.47244034,d.ZG4)

**QUESTION 220**

Refer to the exhibit. Two routers are connected by Frame Relay and are running OSPF between them. Each router has been configured with the appropriate network statements under router ospf 1, but the routers are not forming an adjacency. Which of the following three commands could be configured on each router to correct this problem? (Choose three.)



- A. RouterC(config-if)#ip ospf network broadcast  
RouterD(config-if)#ip ospf network broadcast
- B. RouterC(config-if)#ip ospf network point-to-point

- RouterD(config-if)#ip ospf network point-to-point
- C. RouterC(config-router)#neighbor 10.100.100.4  
RouterD(config-router)#neighbor 10.100.100.3
  - D. RouterC(config-router)#neighbor 10.255.255.254  
RouterD(config-router)#neighbor 10.255.255.253

**Correct Answer:** ABD

**Section:** VPN Technologies

**Explanation**

**Explanation/Reference:**

Explanation

By default, Frame Relay is classified as a non-broadcast network, meaning it doesn't send any broadcasts/multicasts like RIP, OSPF or EIGRP updates across the network (hello packets of OSPF are multicast to 224.0.0.5). Hence, in NBMA network, the neighbors are not discovered automatically; they must be configured manually. There are two ways to simulate a broadcast model on an NBMA network:

- + Define the network type as broadcast with the ip ospf network broadcast interface sub- command
- + Configure the neighbor statements under router ospf mode (though configuring the neighbor statement on one end is sufficient to form adjacency, it is a good practice to have it configured on both the ends)

Besides these two ways, another way for making OSPF work with Frame Relay is configuring the network as a Point-to-Point network (or Point-to-Multipoint, notice that OSPF treats Point-to-Multipoint network as a collective of point-to-point links).

Some information about Point-to-Multipoint (or Point-to-Point) network:

Note: Point-to-Multipoint networks do not maintain a DR/BDR relationship.

#### **QUESTION 221**

Your network has a mixture of Fast Ethernet and Gigabit Ethernet links. What needs to be done to ensure optimal data routing when using OSPF?

- A. Nothing. OSPF will determine the most optimal path for routing data by default.
- B. Adjust the hello and dead timers for more rapid detection of link failures.
- C. Increase the reference-bandwidth used to calculate the interface default metrics, on all routers in your network.
- D. Set the priority values on every broadcast interface to ensure that the designated and backup designated routers are the routers with the most processor and memory resources.

**Correct Answer:** C

**Section:** VPN Technologies

**Explanation**

**Explanation/Reference:**

Explanation

The default formula to calculate the cost for the OSPF metric is (10 /BW). Therefore when using default reference bandwidth (100M) to calculate ospf,

Gi and Fa interfaces could have the same cost (1). This problem can be fixed by configuring reference bandwidth to 1000M.

| Default Bandwidth (100M)                            | Configured Bandwidth (1000M) |
|-----------------------------------------------------|------------------------------|
| Fa = $100M / 100M = 1$                              | Fa = $1000M / 100M = 10$     |
| Gi = $100M / 1000M = 0.1 \rightarrow$ round up to 1 | Gi = $1000M / 1000M = 1$     |

This is how to configure reference bandwidth to 1000Mbps:

```
Router(config)#router ospf 1
Router(config-router)#auto-cost reference-bandwidth 1000
```

#### QUESTION 222

To make OSPF area 4 a totally stubby area, which two things need to be done? (Choose two.)

- A. Apply the area 4 stub command to all routers in the area.
- B. On the ABR, use the area 4 stub command with the no-summary keyword.
- C. On the ABR, specify a default cost for the area with the area default-cost command.
- D. On the ABR, use the default-information originates command to inject a default route into area 4.
- E. Use the auto-cost command on each router in the area to automatically determine the cost to other OSPF areas.

**Correct Answer:** AB

**Section:** VPN Technologies

**Explanation**

**Explanation/Reference:**

Explanation

To make an area totally stubby we must apply the area area-id stub no-summary command on the ABR and area area-id stub commands to all other routers in that area.

Note: The ABR in a totally stubby area does not create Type 3 summary LSA. It only creates a default route to outside destinations.

#### QUESTION 223

Refer to the exhibit. Router C was configured so that it could form an adjacency with three OSPF neighbors, one connected to each of its three physical interfaces.

```
RouterC#show ip interface brief
```

| Interface   | IP-Address     | OK? | Method | Status | Protocol |
|-------------|----------------|-----|--------|--------|----------|
| Ethernet0/0 | 10.1.0.1       | YES | NVRAM  | up     | up       |
| Serial0/0   | 10.255.255.253 | YES | NVRAM  | up     | up       |
| Ethernet0/1 | 10.2.0.1       | YES | NVRAM  | up     | up       |
| Loopbac k0  | 10.100.100.3   | YES | NVRAM  | up     | up       |

```
RouterC#show ip ospf
Routing Process "ospf 1" with ID 10.100.100.3
Supports only single TOS(TOS0) routes
Supports opaque LSA
It is an area border router
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
Number of external LSA 0. Checksum Sum 0x000000
Number of opaque AS LSA 0. Checksum Sum 0x000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 2. 1 normal 1 stub 0 nssa
External flood list length 0
Area BACKBONE(0)
  Number of interfaces in this area is 3
  Area has no authentication
  SPF algorithm executed 10 times
  Area ranges are
  Number of LSA 14. Checksum Sum 0x053D21
  Number of opaque link LSA 0. Checksum Sum 0x000000
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0
Area 4
  Number of interfaces in this area is 1
  It is a stub area, no summary LSA in this area
    generates stub default route with cost 1
  Area has no authentication
  SPF algorithm executed 9 times
  Area ranges are
  Number of LSA 4. Checksum Sum 0x01F200
  Number of opaque link LSA 0. Checksum Sum 0x000000
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0
```



Which statement is correct about router C?

- A. It is configured and functioning correctly as an OSPF internal router.
- B. It is configured and functioning correctly as an ABR attached to stub area 4.
- C. It is configured and functioning correctly as an ASBR attached to external area 4.
- D. It is configured and functioning correctly as an ABR attached to totally stubby area 4.
- E. It is not configured correctly to function as specified.

**Correct Answer: D**

**Section: Infrastructure Security**

**Explanation**

**Explanation/Reference:**

Explanation

From the output of the show ip ospf command we notice 2 lines:

- + It is an area border router
- + It is a stub area, no summary LSA in this area (in Area 4)

Therefore we can conclude RouterC is an ABR and it is attached to a totally stubby area (with no summary LSA advertised).

#### **QUESTION 224**

Refer to the exhibits. What can be done to fix the problem?

```
RouterA#debug ip ospf events
```

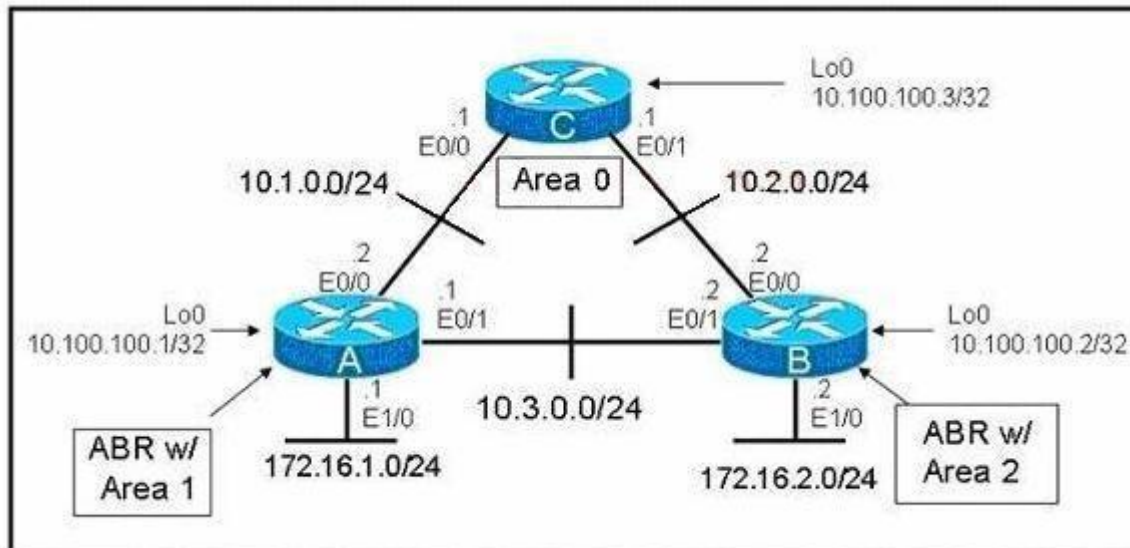
```
OSPF events debugging is on
```

```
RouterA#
```

```
04:43:16: OSPF. Rcv pkt from 10.3.0.2, Ethernet0/1, area 0.0.0.1 mismatch area 0.0.0.2 in the header
```

```
04:43:19: OSPF. Rcv hello from 10.100.100.3 area 0 from Ethernet0/0 10.1.0.1
```

```
04:43:19: OSPF. End of hello processing
```



- Change router B E0/1 interface to area 0.
- Change router A interface E0/1 to area 0.0.0.2.
- Configure the E0/1 interfaces of router A and router B to be in area 0.
- Shut down the E0/1 interfaces in router A and router B as OSPF does not allow "back doors" between areas.
- Remove the E0/1 interfaces in router A and router B from the OSPF process. Use static routes to route data directly from router A to router B to avoid passing data through router C in area 0.

**Correct Answer:** C

**Section:** Infrastructure Security

**Explanation**

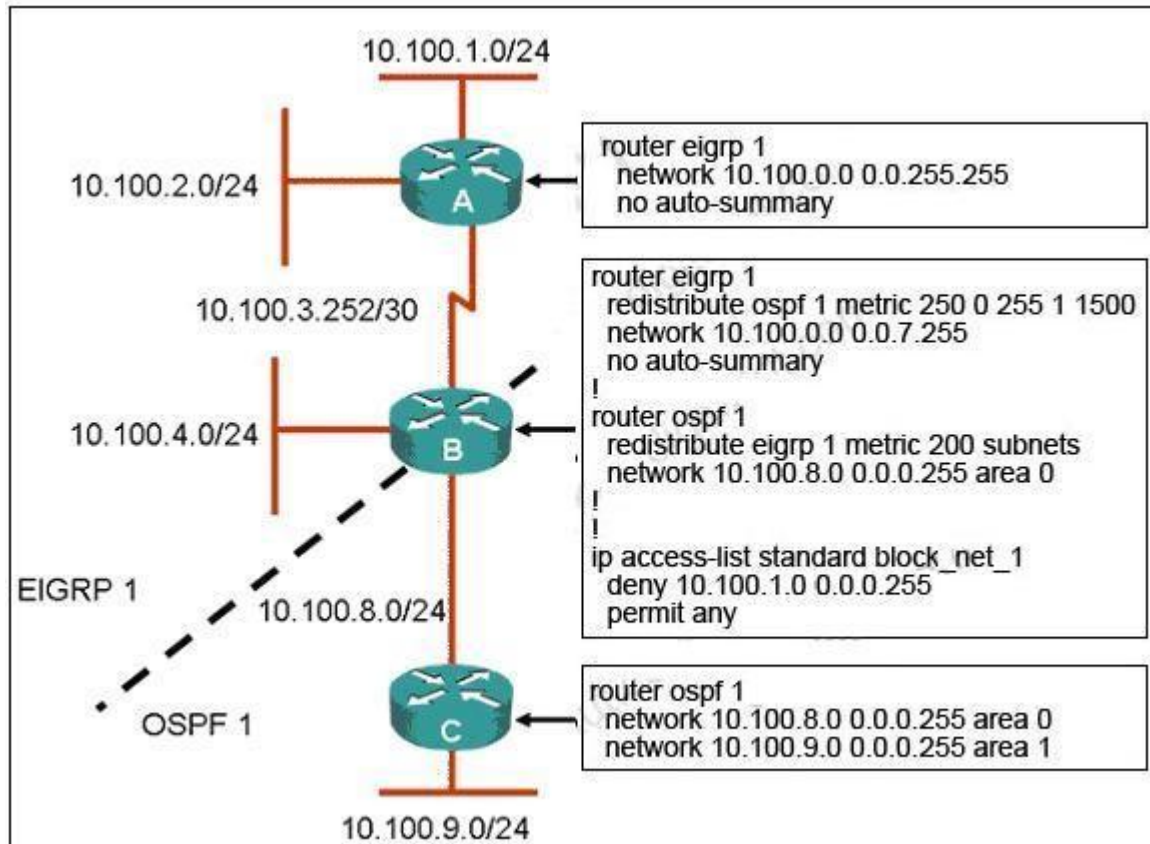
**Explanation/Reference:**

Reference:

[http://www.cisco.com/en/US/tech/tk365/technologies\\_white\\_paper09186a0080094e9e.shtml](http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml) (See topic: the backbone and area 0)

#### QUESTION 225

Refer to the exhibit.



Router B is performing bidirectional redistribution between EIGRP and OSPF. The network 10.100.1.0/24 should not be reachable from the 10.100.9.0/24 network. However, it needs to be reachable from any network within the EIGRP domain. All other networks should be seen in both domains. Which change to router B would accomplish these goals?

- A. Under the EIGRP process, insert the distribute-list block\_net\_1 out ospf 1 command.
- B. Under the OSPF process, insert the distribute-list block\_net\_1 in serial1/0 command.
- C. Under the EIGRP process, insert the distribute-list block\_net\_1 in serial1/0 command.
- D. Under the OSPF process, insert the distribute-list block\_net\_1 out eigrp 1 command.

**Correct Answer:** D

**Section:** Infrastructure Security

**Explanation**


**Explanation/Reference:**

Explanation

The network 10.100.1.0/24 should not be reachable from network 10.100.9.0/24 -> we need to filter updates from EIGRP to OSPF so that the routing table of router C doesn't have network 10.100.1.0/24 -> we need to filter it under OSPF process and the direction should be out of router B ->.

**QUESTION 226**

Refer to the exhibit. ROUTE.com is planning to implement a new secure OSPF network to support traffic between clients on the 172.16.10.0/24 network and the file server on the 172.16.20.0/24 network. You have been asked to review the implementation plan for the OSPF project. Which statement about the plan is true?

**Implementation Plan**

1. Enable OSPF process 1 on all routers.
2. Enable area 0 on R2, R3, and R4.
3. Enable area 10 on R1, R2 and R3.
4. Enable area 20 on R4 and R5.
5. Verify that all routers contain a complete routing table.
6. Verify that the clients can successfully access the file server.
7. Document configuration changes.

- A. It is complete as written.
- B. It should include a task that shuts down all unused interfaces.
- C. It should include tasks that enable and verify OSPF authentication.
- D. It should include a task that establishes a file transfer baseline before and after the configuration is changed.

**Correct Answer: C**

**Section: Infrastructure Security**

**Explanation**

**Explanation/Reference:**


Explanation

The complete implementation plan should be

1. Enable OSPF process 1 on all routers
2. Enable area 0 on R2, R3 and R4
3. Enable area 10 on R1, R2 and R3
4. Enable area 20 on R4 and R5
5. Enable and verify OSPF authentication
6. Verify that all routers contain a complete routing table
7. Verify that the clients can successfully access the server
8. Document configuration changes

**QUESTION 227**

Refer to the exhibit. ROUTE.com is planning to implement a secure OSPF network to support traffic between clients on the 172.16.10.0/24 network and the file server on the 172.16.20.0/24 network. You have been asked to review the implementation and verification plans for this OSPF project. Which statement about the plan is true?



#### Verification Plan

1. On R1 and R5, verify end-to-end connectivity between prefixes 172.16.10.0 and 172.16.20.0 using the extended ping command.
2. Verify that all routers have established a full neighbor relationship with the appropriate neighbors.
3. Verify the proper authentication method is active for each neighbor relationship.
4. Verify that each prefix is assigned to the appropriate OSPF area.
5. Verify the appropriate routes and summaries are in the routing table of each router.
6. Verify end-to-end connectivity between the clients and the file server.

- A. It is complete as written.
- B. It should include a task that verifies that the interarea routes are using the proper MED.
- C. It should include a task that verifies that load sharing is active on R1 and R4.
- D. It should include a task that verifies that all redundant links will become active when the primary links are shut down.

**Correct Answer: D**

**Section: Infrastructure Security**

**Explanation**


**Explanation/Reference:**

Reference:

[http://www.cisco.com/en/US/docs/routers/access/as5350/software/feature/guide/pull\\_rlm.htm](http://www.cisco.com/en/US/docs/routers/access/as5350/software/feature/guide/pull_rlm.htm)

#### QUESTION 228

Refer to the exhibit. Which two Cisco IOS commands on R2 would verify its OSPF neighbor relationships? (Choose two.)



- A. `show ip ospf`
- B. `show ip ospf interface`
- C. `show ip ospf neighbor`
- D. `show ip ospf database`
- E. `show ip ospf statistics`
- F. `show running-config | begin router ospf`



**Correct Answer:** BC

**Section:** Mixed Questions

**Explanation**

**Explanation/Reference:**

Explanation

The show ip ospf interface command shows us information about the neighbor count and adjacent neighbor count:

```
RouterA# show ip ospf int s1/0
Serial 1/0 is up, line protocol is up
 Internet Address 10.100.3.253/30, Area 1
 Process ID 1, Router ID 1.1.1.1, Network Type BROADCAST, Cost:64
 Transmit Delay is 1 sec, State DROTHER, Priority 0
 No designated router on this network
 No backup designated router on this network
 Old designated Router (ID) 2.2.2.2, Interface address 10.100.3.254
 Flush timer for old DR LSA due in 00:01:12
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
   oob-resync timeout 40
   Hello due in 00:00:01
 Supports Link-Local Signaling (LLS)
 Index 1/2, flood queue length 0
 Next 0x0(0)/0x0(0)
 Last flood scan length is 1, maximum is 3
 Last flood scan time is 0 msec, maximum is 0 msec
 Neighbor Count is 1, Adjacent neighbor count is 1
 Suppress hello for 0 neighbor(s)
```

The show ip ospf neighbor command shows us the role of each neighbor (DR, BDR, DROTHER).



Dev-1#sh ip ospf neighbor

| Neighbor ID   | Pri | State    | Dead Time | Address  | Interface    |
|---------------|-----|----------|-----------|----------|--------------|
| 10.200.200.13 | 1   | FULL/BDR | 00:00:33  | 10.1.1.3 | Ethernet 0/0 |

Dev-3#sh ip ospf neighbor

| Neighbor ID | Pri | State   | Dead Time | Address  | Interface    |
|-------------|-----|---------|-----------|----------|--------------|
| 172.31.1.1  | 2   | FULL/DR | 00:00:31  | 10.1.1.1 | Ethernet 0/0 |

#### QUESTION 229

Which three statements about OSPF areas are true? (Choose three.)

- A. Areas introduce a boundary on the link-state updates.
- B. Areas are logical definitions specific to any given router.
- C. All routers within an area have the exact link-state database.
- D. The calculation of the Dijkstra algorithm on a router is limited to changes within an area.
- E. The area designated router will always have a priority of 0.

**Correct Answer:** ACD

**Section:** Mixed Questions

**Explanation**

**Explanation/Reference:**

Reference:

[http://www.cisco.com/en/US/tech/tk365/technologies\\_white\\_paper09186a0080094e9e.shtml](http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml) (areas and border routers)

#### QUESTION 230

Given the following partial configuration for Router A:

```
interface serial 0
ip address 10.1.1.1 255.255.255.0
encapsulation frame-relay
ip ospf network point-to-multipoint
router ospf 7
network 10.1.1.0 0.0.0.255 area 0
```

Which two statements are correct? (Choose two.)

- A. DR/BDR elections do not take place.
- B. The area 0 NBMA cloud is configured as more than one subnet.
- C. The router is restricted to a hub and spoke topology.
- D. OSPF neighbor statements are not necessary.

**Correct Answer:** AD

**Section:** Mixed Questions

**Explanation**

**Explanation/Reference:**

Explanation:

In an OSPF Point-to-Multipoint environment, DR/BDR elections do not take place. The neighbor command became somewhat obsolete with the introduction of the capability to configure other network modes for the interface, regardless of the underlying physical topology.

Reference: Building Scalable Cisco Networks (Cisco Press) page 130 and 181 Point-to-Multipoint Network:

Point-to-multipoint is a single interface that connects to multiple destinations. The underlying network treats the network as a series of point-to-point circuits. It replicates LSA packets for each circuit. OSPF traffic is sent as multicast. There is no DR or BDR election. This technology uses one IP subnet for all endpoints on the network.

By default, the network is considered to be a series of point-to-point interfaces. There is no need to specify neighbors, because the neighbors will see each other and simply become adjacent, with no need for the election of a DR or a BDR. Point-to-multipoint does not try to reduce adjacencies using a DR. Instead, it accepts the extra overhead of having a full set of adjacencies for the sake of stability. Point-to-multipoint forms an adjacency automatically along any PVC, which causes more overhead but is more resilient than NBMA.

#### **QUESTION 231**

Which command displays the number of times that the OSPF Shortest Path First (SPF) algorithm has been executed?

- A. show ip protocol
- B. show ip ospf
- C. show ip ospf database
- D. show ip ospf interface

**Correct Answer:** B

**Section:** Mixed Questions

**Explanation**

**Explanation/Reference:**

Explanation:

The following table describes the output of the show ip ospf command and their meanings:

Table52 show ip ospf Field Descriptions

| Field                                          | Description                                                             |
|------------------------------------------------|-------------------------------------------------------------------------|
| Routing process "ospf 201" with ID 10.0.0.1    | Process ID and OSPF router ID.                                          |
| Supports...                                    | Number of types of service supported (Type 0 only).                     |
| SPF schedule delay                             | Delay time of SPF calculations.                                         |
| Minimum LSA interval                           | Minimum interval between link-state advertisements.                     |
| LSA group pacing timer                         | Configured LSA group pacing timer (in seconds).                         |
| Interface flood pacing timer                   | Configured LSA flood pacing timer (in milliseconds).                    |
| Retransmission pacing timer                    | Configured LSA retransmission pacing timer (in milliseconds).           |
| Number of...                                   | Number and type of link-state advertisements that have been received.   |
| Number of external LSA                         | Number of external link-state advertisements.                           |
| Number of opaque AS LSA                        | Number of opaque link-state advertisements.                             |
| Number of DCbitless external and opaque AS LSA | Number of demand circuit external and opaque link-state advertisements. |
| Number of DoNotAge external and opaque AS LSA  | Number of do not age external and opaque link-state advertisements.     |
| Number of areas in this router is              | Number of areas configured for the router.                              |
| External flood list length                     | External flood list length.                                             |

Reference:[http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/iprrp\\_r/ip2\\_s3g.htm#wp1036469](http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/iprrp_r/ip2_s3g.htm#wp1036469)

**QUESTION 232**

The Dev-1 and Dev-3 routers are OSPF neighbors over the Ethernet 0/0 connection. Based on the show ip ospf neighbor output from the Dev-1 and Dev-3 routers, which statement is true?

|                           |     |          |           |          |             |
|---------------------------|-----|----------|-----------|----------|-------------|
| Dev-1#sh ip ospf neighbor |     |          |           |          |             |
| Neighbor ID               | Pri | State    | Dead Time | Address  | Interface   |
| 10.200.200.13             | 1   | FULL/BDR | 00:00:33  | 10.1.1.3 | Ethernet0/0 |
| Dev-3#sh ip ospf neighbor |     |          |           |          |             |
| Neighbor ID               | Pri | State    | Dead Time | Address  | Interface   |
| 172.31.1.1                | 2   | FULL/DR  | 00:00:31  | 10.1.1.1 | Ethernet0/0 |

- A. Dev-1 is the DR because it has a higher OSPF router priority.
- B. Dev-1 is the DR because it has a lower OSPF router ID.
- C. Dev-3 is the DR because it has a higher OSPF router priority.
- D. Dev-3 is the DR because it has a lower OSPF router ID.
- E. Both Dev-1 and Dev-3 are using the default OSPF router priority.

**Correct Answer:** A

**Section:** Mixed Questions

**Explanation**

**Explanation/Reference:**

Explanation:

The output shown above displays information about the neighbors, so from Dev-3 we see that the priority of Dev-1 is 2, and that Dev-1 is the DR. From Dev-1 we can see that its neighbor (Dev-3) is the BDR and has an OSPF priority of 1.

**QUESTION 233**

When other routing protocol routes are being redistributed into OSPF, what is one of the most common problems?

- A. missing the tag option in the redistribute command.
- B. missing the subnet option in the redistribute command.
- C. missing the metric option in the redistribute command.
- D. misconfiguring the metric-type option in the redistribute command to type-1.

E. misconfiguring the metric-type option in the redistribute command to type-2.

**Correct Answer:** B

**Section:** Mixed Questions

**Explanation**


**Explanation/Reference:**

Explanation:

The point of this question is about the redistribution of OSPF. To redistribute routes from one routing domain into another routing domain, use the redistribute command in router configuration mode. redistribute protocol [process-id] {level- 1 | level-1-2 | level-2} [as-number] [metric {metric-value | transparent}] [metric-type type- value] [match {internal | external 1 | external 2}] [tag tag-value] [route-map map-tag] [subnets] The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword, only networks that are not subnetted are redistributed by OSPF.

**QUESTION 234**

Refer to the diagram. Which OSPF configuration command is required to configure Area 3 as a totally stubby area?



- A. On the ABR  
router ospf 10  
area 3 nssa
- B. On the ABR  
router ospf 10  
area 3 stub
- C. On all Area 3 routers  
router ospf 10

area 3 stub no-summary

- D. On all Area 3 routers  
router ospf 10  
area 3 stub
- E. On the ABR  
router ospf 10  
area 3 stub no-summary

**Correct Answer:** E

**Section:** Mixed Questions

**Explanation**

**Explanation/Reference:**

Explanation

To make an area totally stubby we must apply the area area-id stub no-summary command on the ABR (the middle router in this case) and area area-id stub commands to all other routers in that area.

#### **QUESTION 235**

Which show command will display only the Type 5 LSAs in the OSPF topology database?

- A. show ip ospf database external
- B. show ip ospf database nssa-external
- C. show ip route ospf
- D. show ip ospf database summary
- E. show ip route

**Correct Answer:** A

**Section:** Mixed Questions

**Explanation**

**Explanation/Reference:**

Explanation

The show ip ospf database external command displays information only about external LSAs (Type 5 LSAs).

Below is an example of the show ip ospf database external command

Notice the line LS Type: AS External Link, which means LSA Type 5. For your information, the Link State ID: 143.105.0.0 indicates the network being advertised; the Advertising Router: 10.187.70.6 indicates the router that originated this LSA.

Note:

- + The show ip ospf database summary command displays information only about the summary LSAs.
- + The show ip ospf database nssa-external command displays information only about the not so stubby area (NSSA) external LSAs.

```
Router# show ip ospf database external
```

```
OSPF Router with id(192.168.239.66) (Autonomous system 300)
```

```
Displaying AS External Link States
```

```
LS age: 280
```

```
Options: (No TOS-capability)
```

```
LS Type: AS External Link
```

```
Link State ID: 143.105.0.0 (External Network Number)
```

```
Advertising Router: 10.187.70.6
```

```
LS Seq Number: 80000AFD
```

```
Checksum: 0xC3A
```

```
Length: 36
```

```
Network Mask: 255.255.0.0
```

```
    Metric Type: 2 (Larger than any link state path)
```

```
    TOS: 0
```

```
    Metric: 1
```

```
    Forward Address: 0.0.0.0
```

```
    External Route Tag: 0
```