

VCEplus.com

Certkiller 300-101 Casey 82q

Number: 300-101
Passing Score: 800
Time Limit: 120 min
File Version: 20.5

Exam Code: 300-101

Exam Name: Implementing Cisco IP Routing (ROUTE v2.0)



Exam A

QUESTION 1

Refer to the exhibit.

```
R2#show ip cef
```

Prefix	Next Hop	Interface
0.0.0.0/0	192.168.201.1	FastEthernet0/0
0.0.0.0/32	receive	
192.168.201.0/27	attached	FastEthernet0/0
192.168.201.0/32	receive	
192.168.201.1/32	192.168.201.1	FastEthernet0/0
192.168.201.2/32	receive	
192.168.201.31/32	receive	
224.0.0.0/4	drop	
224.0.0.0/24	receive	
255.255.255.255/32	receive	

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: (none)

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.

```
Router#show adjacency
```

Protocol	Interface	Address
IP	Serial0	192.168.209.130 (2) (incomplete)
IP	Serial0	192.168.209.131 (7)
IP	Ethernet0	192.168.201.1 (7)

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: (none)

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:

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: (none)

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 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: (none)

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 5

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: (none)

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 6

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: (none)

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 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: (none)

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 ip route {* | {route |  
prefix/length}[next-hop interface]}  
[vrf vrf-name]
```

Example:

```
switch(config)# clear ip route  
10.2.2.2
```

Clears one or more routes from both the unicast RIB and all the module FIBs. The route options are as follows:

*--All routes.

.

route--An individual IP route.

.

prefix/length--Any IP prefix.

.

next-hop--The next-hop address

.

interface--The interface to reach the 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: (none)

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/xrfs cmd5.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 ip tcp window-size 65536.
- B. Set window scaling to be used on the remote host.

- C. Execute the command ip tcp queuemax.
- D. Set TCP options to "enabled" on the remote host.
- E. Execute the command ip tcp adjust-mss.

Correct Answer: AB

Section: (none)

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 ip tcp 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>

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: (none)

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 `ip tcp 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 `ip tcp 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 `ip tcp ecn` 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/x-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: (none)

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#show ip sla 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\_paper09186a008\_02d5efe.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: (none)

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: (none)

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

enable

2.

configure terminal

3.

vpdn enable

4.

vpdn-group name

5.

request-dialin

6.

protocol pppoe

DETAILED STEPS

Command or Action

Step 1

enable

Example:

Router> enable

Enables privileged EXEC mode.

.

Enter your password if prompted.

Step 2

configure terminal

Example:

```
Router# configure terminal
```

Enters global configuration mode.

Step 3

vpdn enable

Example:

```
Router(config)# vpdn enable
```

Enables virtual private dialup networking.

Step 4

vpdn-group name

Example:

```
Router(config)# vpdn-group group1
```

Associates a VPDN group with a customer or VPDN profile.

Step 5

request-dialin

Example:

```
Router(config-vpdn)# request-dialin
```

Creates a request-dialin VPDN subgroup.

Step 6

protocol pppoe

Example:

```
Router(config-vpdn-req-in)# protocol pppoe
```

Enables the VPDN subgroup to establish PPPoE

Reference:

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

QUESTION 14

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: (none)

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/fsecur_c/scfathen.html

QUESTION 15

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: (none)

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 number

Creates a dialer interface and enters interface configuration mode.

Step 2

Router(config-if)# ip address address mask

Specifies the IP address and mask of the dialer interface as a node in the destination network to be called.

Step 3

Router(config-if)# encapsulation type

Specifies the encapsulation type.

Step 4

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

Step 5

Router(config-if)# dialer pool number

Specifies the dialing pool to use for calls to this destination.

Step 6

Router(config-if)# dialer-group group-number

Assigns the dialer interface to a dialer group.

Step 7

Router(config-if)# dialer-list dialer-group protocol protocol-name {permit | deny | list access-list- number}

Specifies an access list by list number or by protocol and list number to define the interesting packets that can trigger a call.

Step 8

Router(config-if)# dialer remote-name user-name

(Optional) Specifies the authentication name of the remote router on the destination subnetwork for a dialer interface.

Step 9

Router(config-if)# dialer persistent [delay [initial] seconds | max-attempts number] Forces a dialer interface to be connected at all times, even in 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 16

Which PPP authentication method sends authentication information in cleartext?

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

Correct Answer: D

Section: (none)

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: (none)

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: (none)

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 ip ospf network command:

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

Network Type

Hello Interval (secs)

Dead Interval (secs)

Point-to-Point

Point-to-Multipoint

Broadcast

Non-Broadcast

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

QUESTION 19

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: (none)

Explanation

Explanation/Reference:

Explanation:

PPPoE is composed of two main phases:

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

QUESTION 20

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: (none)

Explanation

Explanation/Reference:

Explanation:


PPPoE is composed of two main phases:

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 21

Refer to the exhibit.



```
router eigrp 1
  network 172.16.0.0
  network 192.168.5.0
  distribute-list 7 out s0
!
access-list 7 permit 172.16.0.0 0.0.255.255
```

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: (none)

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 22

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: (none)

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 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?

0

- 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: (none)


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 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

1

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: (none)

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: (none)

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.html#wp1062796

QUESTION 26

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: (none)

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

snooping MAC address verification option is turned on.

.
The switch receives a DHCPRELEASE or DHCPDECLINE 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 27

Refer to the exhibit.

3

```
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: (none)

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 28

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: (none)

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 translate 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.

4

QUESTION 29

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: (none)

Explanation

Explanation/Reference:

Explanation:

show ipv6 traffic Field Descriptions

Field

Description

source-routed

Number of source-routed packets.

truncated

Number of truncated packets.

format errors

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

not a router

Message sent when IPv6 unicast routing is not enabled.

Reference:

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

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.
5
- 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: (none)


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 0xFFFE 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 an EUI-48 MAC address.

Here is an example showing how a 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.



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: (none)

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

7

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

CORRECT TEXT

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 8 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 Secre

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



Seattle con0 is now available

ActualTests

Press RETURN to get started.

Seattle>

CiscoTerminal

Chicago con0 is now available

Press RETURN to get started.

Chicago>

ActualTests

CiscoTerminal

```
NewYork con0 is now available
```

```
Press RETURN to get started.
```

```
NewYork#
```

ActualTests

- A.
- B.
- C.
- D.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: Here is the solution below:

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

1

2

3

4

5

0

4

8

12

16

...

3

7

11

15

19

...

es

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 33

CORRECT TEXT

JS 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.



ActualTests

```
% 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

ActualTests

```
% 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>
```


ActualTests

```
% 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>
```

- A.
- B.
- C.
- D.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: Here are the solution as below:

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
```

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  
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  
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 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 34

CORRECT TEXT

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>

% 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>


```
% 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>
```



```
% 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>
```

- A.
- B.
- C.
- D.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

Answer: Here is the solution below:

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
```

```
!
```

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
```

```
|
```

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 35


CORRECT TEXT

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.

Topology



R1



ActualTests

Press RETURN to get started!

R1>

- A.
- B.
- C.
- D.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

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.

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 36


CORRECT TEXT

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.



- A.
- B.

- C.
- D.

Correct Answer:

Section: (none)

Explanation

Explanation/Reference:

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:

```
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 37

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



R1



ActualTests

R1#



R2



ActualTests

R2#



R3



ActualTests

R3#



R4



ActualTests

R4#



R5



ActualTests

R5#



R6



ActualTests

R6#



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: (none)

Explanation

Explanation/Reference:

Explanation:

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

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 38

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



R1



ActualTests

R1#



R2



ActualTests

R2#



R3



ActualTests

R3#



R4



ActualTests

R4#



R5



ActualTests

R5#



R6



ActualTests

R6#



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: (none)

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
```

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

QUESTION 39

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



R1



ActualTests

R1#



R2



ActualTests

R2#



R3



ActualTests

R3#



R4



ActualTests

R4#



R5



ActualTests

R5#



R6



ActualTests

R6#



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: (none)

Explanation


Explanation/Reference:

Explanation: R4 and R2 configs are as shown below:

Clearly we see the actual key chain is named CISCO.

QUESTION 40

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



R1



ActualTests

R1#



R2



ActualTests

R2#



R3



ActualTests

R3#



R4



ActualTests

R4#



R5



ActualTests

R5#



R6



ActualTests

R6#



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: (none)

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

R1#

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 41

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



R1



ActualTests

R1#



R2



ActualTests

R2#



R3



ActualTests

R3#



R4



ActualTests

R4#



R5



ActualTests

R5#



R6



ActualTests

R6#



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

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

Correct Answer: B

Section: (none)

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
```

```
ip bandwidth-percent eigrp 1 20
```

1 = the EIGRP AS


20 = 20% of the bandwidth

QUESTION 42

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.



R1



ActualTests

R1#



R2



ActualTests

R2#



R3



ActualTests

R3#



R4



ActualTests

R4#



R5



ActualTests

R5#



R6



ActualTests

R6#



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: (none)

Explanation

Explanation/Reference:

Explanation:

Part of the "show ip ospf 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 43

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.



R1



ActualTests

R1#



R2



ActualTests

R2#



R3



ActualTests

R3#



R4



ActualTests

R4#



R5



ActualTests

R5#



R6



ActualTests

R6#



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: (none)

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 44

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.



R1



ActualTests

R1#



R2



ActualTests

R2#



R3



ActualTests

R3#



R4



ActualTests

R4#



R5



ActualTests

R5#



R6



ActualTests

R6#



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: (none)

Explanation

Explanation/Reference:

Explanation:

This can be found using the "show ip ospf" 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 45

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.



R1



ActualTests

R1#



R2



ActualTests

R2#



R3



ActualTests

R3#



R4



ActualTests

R4#



R5



ActualTests

R5#



R6



ActualTests

R6#



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: (none)

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
```



```
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 46

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: (none)

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/DMVPN_2_Phase2.html

QUESTION 47

A network engineer executes the show crypto ipsec sa 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: (none)

Explanation

Explanation/Reference:

Explanation:

```
show crypto ipsec sa
```

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 ipsec sa 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) remote ident (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,}
```

```
#pkts encaps: 7767918, #pkts encrypt: 7767918, #pkts digest 7767918 #pkts decaps: 7760382, #pkts decrypt: 7760382, #pkts verify 7760382 #pkts compressed: 0, #pkts decompressed: 0
```

```
#pkts not compressed: 0, #pkts compr. 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
```

path mtu 1500, media mtu 1500
current outbound spi: 3D3
inbound esp sas:
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
inbound ah sas:
inbound pcp sas:
inbound pcp sas:
outbound esp sas:
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
outbound ah sas:
outbound pcp sas:
Reference: <http://www.cisco.com/c/en/us/support/docs/security-vpn/ipsec-negotiation-ike-protocols/5409-ipsec-debug-00.html>

QUESTION 48

Refer to the following output:

```
Router#show ip nhrp 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: (none)

Explanation

Explanation/Reference:

Explanation:

Show NHRP: Examples

The following is sample output from the show ip nhrp command:

Router# show ip nhrp

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

NBMA address: 10.1111.1111.1111.1111.1111.1111.1111.1111.1111.11 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 ip nhrp holdtime 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 "Pass Any 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 49

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: (none)

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 50

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: (none)

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 51

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: (none)

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/xr-3sg/evn-overview.pdf>

QUESTION 52

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: (none)

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 53

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: (none)

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:

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

Topic 5, Infrastructure Security

QUESTION 54

Which traffic does the following configuration allow?

```
ipv6 access-list cisco
```

```
permit ipv6 host 2001:DB8:0:4::32 any eq ssh
```

```
line vty 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: (none)

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 55

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: (none)

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 56

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: (none)

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 57

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: (none)

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 58

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: (none)

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 59

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)#access-list 101 deny ip 172.16.0.0 0.15.255.255 any log
```

```
router (config)#access-list 101 permit ip any any
```

```
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: (none)

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 60

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: (none)

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.

```
ip http secure-port port-number
```

```
no ip http secure-port
```

Syntax Description

port-number

Integer in the range of 0 to 65535 is accepted, but the port number must be 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 61

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

Correct Answer: A

Section: (none)

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_broadcas.html

Topic 6, Infrastructure Services

QUESTION 62

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: (none)

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 63

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: (none)

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 64

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: (none)

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 65

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: (none)

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 66

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: (none)

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 client`

Allows the system to receive NTP broadcast packets on an 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 67

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: (none)

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

NAPT44 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 68

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: (none)

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 69

A network engineer initiates the ip sla 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: (none)

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

ip sla 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 70

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: (none)

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. ipv6 flowset
 4. exit
 5. clear ipv6 mtu
- #### DETAILED STEPS

Command or Action

Purpose

Step 1

enable Example:

Device> enable

Enables privileged EXEC mode.

Step 2

configure terminal Example:

Device# configure terminal

Enters global configuration mode.

Step 3

ipv6 flowset Example:

Device(config)# ipv6 flowset

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

Reference: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_basic/configuration/15-mt/ip6b-15-mt-book/ip6-mtu-path-disc.html

QUESTION 71

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: (none)

Explanation

Explanation/Reference:

Explanation:

Table 5 show ip flow export Field Descriptions

Field

Description

Exporting flows to 10.1.1.1 (1000) and 10.2.1.1

Specifies the export destinations and ports. The ports are in parentheses.

Exporting using source IP address 10.3.1.1

Specifies the source address or interface.

Version 5 flow records

Specifies the version of the flow.

11 flows exported in 8 udp datagrams

The total number of export packets sent, and the total number of flows contained within them.

0 flows failed due to lack of export packet

No memory was available to create an export packet.

0 export packets were sent up to process level

The packet could not be processed by CEF or by fast switching, possibly because another feature requires running on the packet.

0 export packets were dropped due to no fib

0 export packets were dropped due to adjacency issues Indicates that CEF was unable to switch the packet or forward it up to the process level. 0 export packets were dropped due to fragmentation failures 0 export packets were dropped due to encapsulation fixup failures Indicates that the packet was dropped because of problems constructing the IP packet.

0 export packets were dropped enqueueing for the RP

0 export packets were dropped due to IPC rate limiting Indicates that there was a problem transferring the export packet between the RP and the line card.

0 export packets were dropped due to output drops

Indicates that the send queue was full while the packet was being transmitted.

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

QUESTION 72

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 ip nat 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 ip nat 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 ip nat 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: (none)

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
```

```
ip nat 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
```

```
ip nat 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  
ip nat outside
```

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

```
ip nat pool ovrlD 172.16.10.1 172.16.10.1 prefix 24  
!
```

!--- Defines a NAT pool named ovrlD with a range of a single IP !--- address, 172.16.10.1.

```
ip nat inside source list 7 pool ovrlD 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 ovrlD. !--- 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 "ovrlD" only has a range of one address. The keyword overload used in the ip nat inside source list 7 pool ovrlD 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 73

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: (none)

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 74

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: (none)

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:

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

QUESTION 75

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: (none)

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 ip cef' 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)# ip cef
```

```
Router(config)#
```

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

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 76

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: (none)

Explanation

Explanation/Reference:

Explanation:

show flow exporter exporter-name Example:

Device# show flow exporter FLOW_EXPORTER-1

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

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

QUESTION 77

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: (none)

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 78

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: (none)

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 79

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: (none)

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

1:1 translation

1:N translation

No conservation of IPv4 address

Conserves IPv4 address

Assures end-to-end address transparency and scalability Uses address overloading, hence lacks in end-to-end address transparency No state or bindings created on the translation

State or bindings are created on every unique translation Requires IPv4-translatable IPv6 addresses assignment (mandatory requirement) No requirement on the nature of IPv6 address assignment Requires either manual or DHCPv6 based address assignment for IPv6 hosts Free to choose any mode of IPv6 address assignment viz. Manual, DHCPv6, SLAAC Reference: http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/enterprise-ipv6-solution/white_paper_c11-676277.html

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

QUESTION 80

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: (none)

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:

QUESTION 81

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: (none)

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 82

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: (none)

Explanation

Explanation/Reference:

Explanation:

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

Syntax Description

`ip-address`

IP address of the workstation to which you want to send the 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