

300-101: Cisco CCNP 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

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

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

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

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

Explanation

Explanation/Reference:

Explanation:

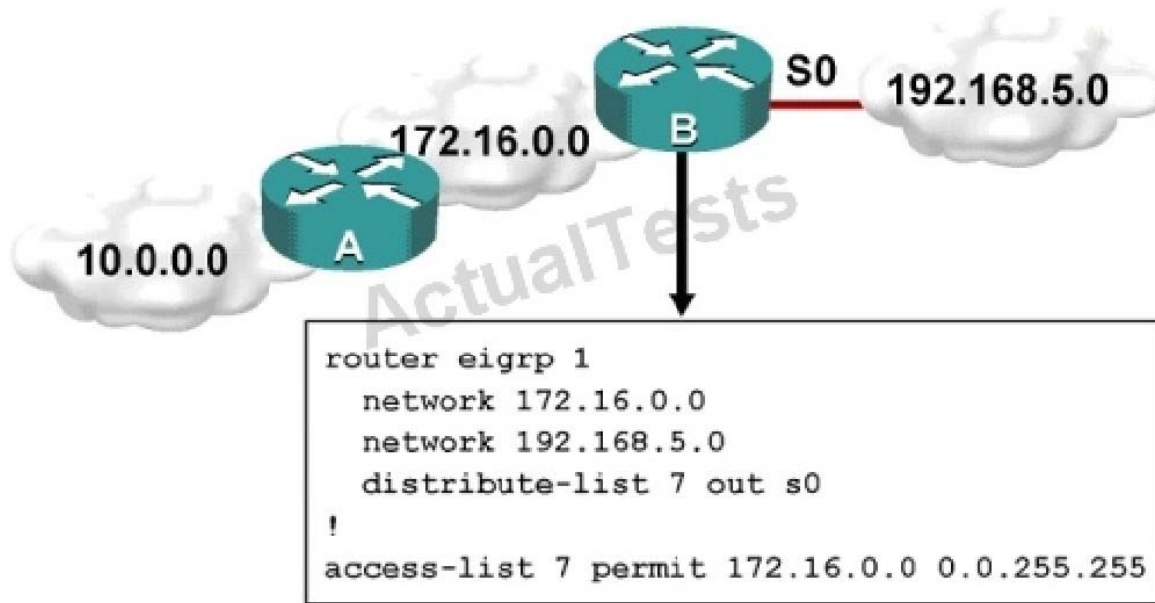
Dual stack means that devices are able to run IPv4 and IPv6 in parallel. It allows hosts to simultaneously reach IPv4 and IPv6 content, so it offers a very flexible coexistence strategy. For sessions that support IPv6, IPv6 is used on a dual stack endpoint. If both endpoints support IPv4 only, then IPv4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

Refer to the exhibit.



Which one statement is true?

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

Correct Answer: E

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.

- 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

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 +

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

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.

QUESTION ,

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

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.