CISCO 640-811 Interconnecting Cisco Networking Devices Exam (ICND)

Practice Test

Version 1.3

QUESTION NO: 1

What is the purpose of the OSPF router ID in a DR/BDR election?

A. It is used with the OSPF priority values to determine which OSPF router will become the DR or BDR in a point-to-point network.
B. It is used with the OSPF priority values to determine which interface will be used to form a neighbor relationship with another OSPF router.
C. It is used to determine which interfaces will send Hello packets to neighboring OSPF routers.
D. It is used with the OSPF priority values to determine which router will become the DR or BDR in a multiaccess network.

Answer: D

Explanation:
The router ID is the highest IP address or the highest IP address among loopback addresses (if one is configured) on the Cisco router or can be configured manually by "router-id x.x.x.x". Once the router ID is chosen, it will not be changed unless the OSPF process is reset(clear ip ospf process xx) or the router is reloaded. The IP address of router ID doesn't need to be reachable, but it is used to determine which will router will become the DR or BDR in a multi-access network.

QUESTION NO: 2

Which wild card mask will enable a network administrator to permit access to the Internet for only hosts that are assigned an address in the range of 192.168.8.0 through 192.168.15.255?

A. 0.0.0.0
B. 0.0.3.255
C. 0.0.7.255
D. 0.0.255.255
E. 0.0.0.255

Answer: C

Explanation:
Wildcard masks are used with access lists to specify an individual host, a network, or a certain range of a network or networks. To understand a wildcard mask, you need to understand what a block size is; block sizes are used to specify a range of addresses. Some of the different block sizes available are 64, 32, 16, 8, and 4.
When you need to specify a range of addresses, you choose the next-largest block size for your needs. For example, if you need to specify 34 networks, you need a block size of 64. If you want to specify 18 hosts, you need a block size of 32. If you only specify 2 networks, then a block size of 4 would work. You use wildcards with the host or network address to tell the router a range of available addresses to filter. To specify a host, the address would look like this: 172.16.30.5

0.0.0.0

The four zeros represent each octet of the address. Whenever a zero is present, it means that octet in the address must match exactly. To specify that an octet can be any value, the value of 255 is used. As an example, here's how a /24 subnet is specified with a wildcard:

172.16.30.0 0.0.0.255 This tells the router to match up the first three octets exactly, but the fourth octet can be any value.

Now, that was the easy part. What if you want to specify only a small range of subnets? This is where the block sizes come in. You have to specify the range of values in a block size. In other words, you can't choose to specify 20 networks. You can only specify the exact amount as the block size value. For example, the range would either have to be 16 or 32, but not 20.

Let's say that you want to block access to part of network that is in the range from 172.16.8.0 through 172.16.15.0. That is a block size of 8. Your network number would be 172.16.8.0, and the wildcard would be 0.0.7.255. Whoa! What is that?!? The 7.255 is what the router uses to determine the block size. The network and wildcard tell the router to start at 172.16.8.0 and go up a block size of eight addresses to network 172.16.15.0.

QUESTION NO: 3

Why has the network shown in the exhibit failed to converge?

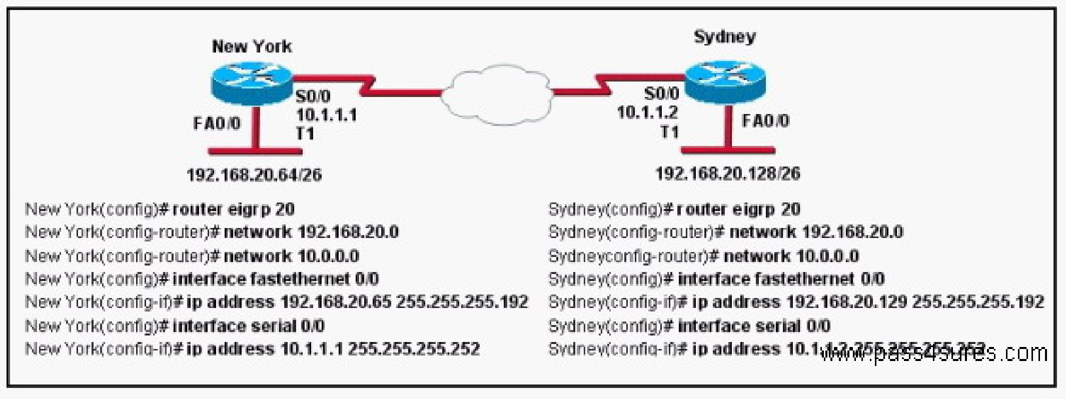

```
New York(config)# router eigrp 20
New York(config-router)# network 192.168.20.0
New York(config-router)# network 10.0.0.0
New York(config)# interface fastethernet 0/0
New York(config-if)# ip address 192.168.20.65 255.255.255.192
New York(config)# interface serial 0/0
New York(config-if)# ip address 10.1.1.1 255.255.255.252
```

```
Sydney(config)# router eigrp 20
Sydney(config-router)# network 192.168.20.0
Sydneyconfig-router)# network 10.0.0.0
Sydney(config)# interface fastethernet 0/0
Sydney(config-if)# ip address 192.168.20.129 255.255.255.192
Sydney(config)# interface serial 0/0
Sydney(config-if)# ip address 10.1.1.2 255.255.255.252
```

A. The autonomous system number has not been properly configured.
B. The no auto-summary command needs to be applied to the routers.
C. The network numbers have not been properly configured on the routers.
D. The bandwidth values have not been properly configured on the serial interfaces.
E. The subnet masks for the network numbers have not been properly configured.

Answer: B

Explanation:

To restore the default behavior of automatic summarization of subnet routes into network-level routes, use the auto-summary command in router configuration mode. To disable this function and transmit subprefix routing information across classful network boundaries, use the no form of this command.

auto-summary

no auto-summary

QUESTION NO: 4

Refer to the exhibit. Switch-1 needs to send data to a host with a MAC address of 00b0.d056.efa4. What will Switch-1 do with this data?

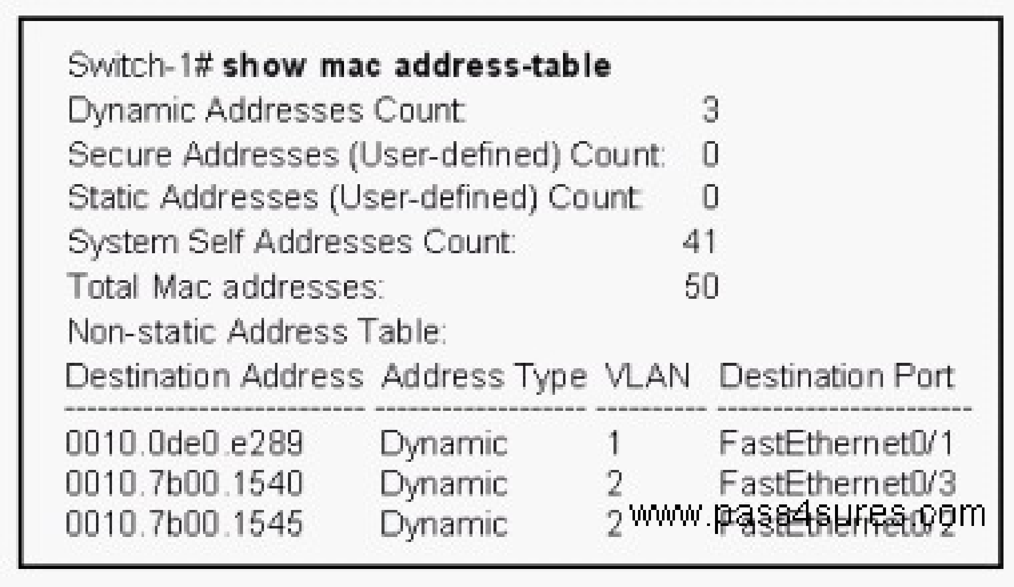

```
Switch-1# show mac address-table
Dynamic Addresses Count:                    3
Secure Addresses (User-defined) Count:      0
Static Addresses (User-defined) Count:      0
System Self Addresses Count:               41
Total Mac addresses:                       50
Non-static Address Table:
Destination Address  Address Type  VLAN   Destination Port
-------------------  ------------  -----  ----------------
0010.0de0.e289       Dynamic       1      FastEthernet0/1
0010.7b00.1540       Dynamic       2      FastEthernet0/3
0010.7b00.1545       Dynamic       2      FastEthernet0/2
```

A. Switch-1 will flood the data out all of its ports except the port from which the data originated.
B. Switch-1 will forward the data to its default gateway.
C. Switch-1 will send an ARP request out all its ports except the port from which the data originated.
D. Switch-1 will drop the data because it does not have an entry for that MAC address.

Answer: A

Explanation:

Switches learn the MAC addresses of PCs or workstations that are connected to their switch ports by examining the source address of frames that are received on that port.

Machines may have been removed from a port, turned off, or moved to another port on the same switch or a different switch.

This could cause confusion in frame forwarding.

The MAC address entry is automatically discarded or aged out after 300 seconds

If there is not MAC address of destination host in MAC table, switch sends broadcast to all ports except the source to findout the destination host.

In output there is no MAC address of give host so switch floods to all ports except source port.

QUESTION NO: 5

Host 1 is trying to communicate with Host 2. The e0 interface on Router C is down. Which of the following are true? (Choose two.)

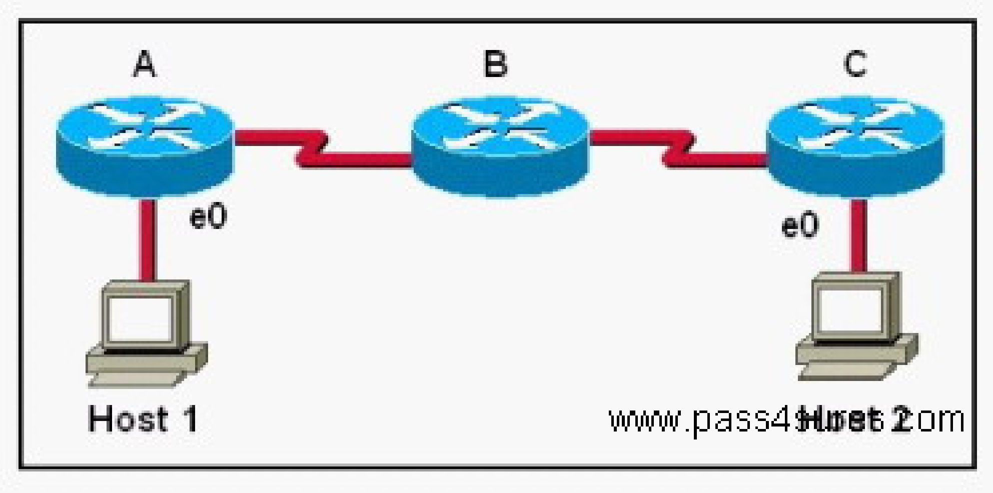

A. Router C will send a Source Quench message type.
B. Router C will use ICMP to inform Router B that Host 2 cannot be reached.
C. Router C will use ICMP to inform Host 1, Router A, and Router B that Host 2 cannot be reached.
D. Router C will use ICMP to inform Host 1 that Host 2 cannot be reached.
E. Router C will send a Destination Unreachable message type.
F. Router C will send a Router Selection message type.

Answer: D,E

Explanation:
When a packet reaches a router that is destined for a network that is not in the routing table or for a network that is down, the router will send an ICMP destination unreachable message back to the sender. This informs the sending station that the packet could not be forwarded to the destination, and this information will be sent to the sending station, not to the router.

QUESTION NO: 6 DRAG DROP

Drag Drop

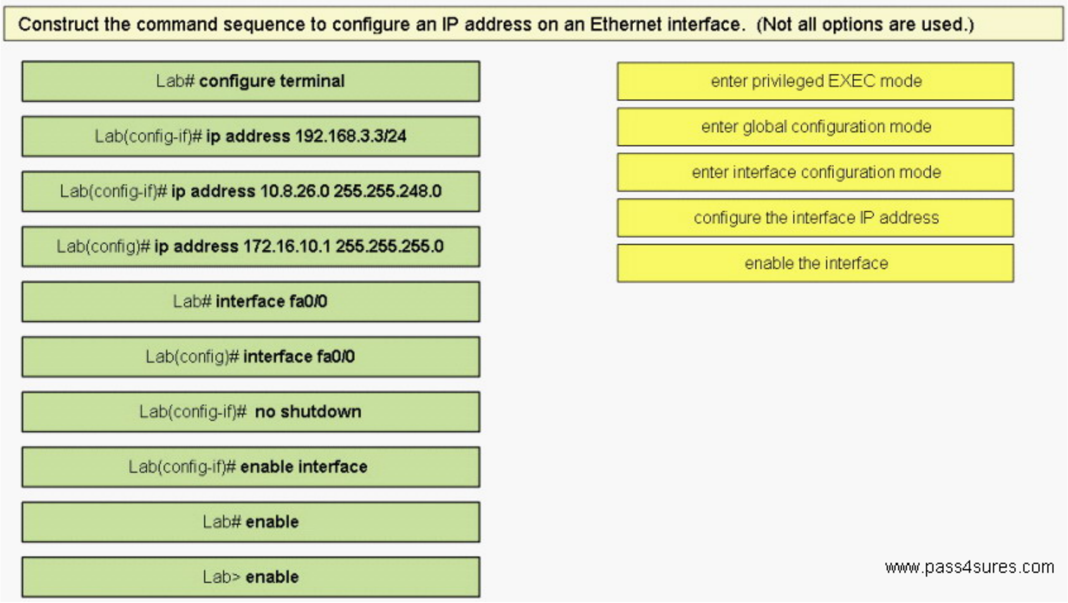

Construct the command sequence to configure an IP address on an Ethernet interface. (Not all options are used.)

Lab# configure terminal

Lab(config-if)# ip address 192.168.3.3/24

Lab(config-if)# ip address 10.8.26.0 255.255.248.0

Lab(config)# ip address 172.16.10.1 255.255.255.0

Lab# interface fa0/0

Lab(config)# interface fa0/0

Lab(config-if)# no shutdown

Lab(config-if)# enable interface

Lab# enable

Lab> enable

enter privileged EXEC mode

enter global configuration mode

enter interface configuration mode

configure the interface IP address

enable the interface

www.pass4sures.com

Answer:

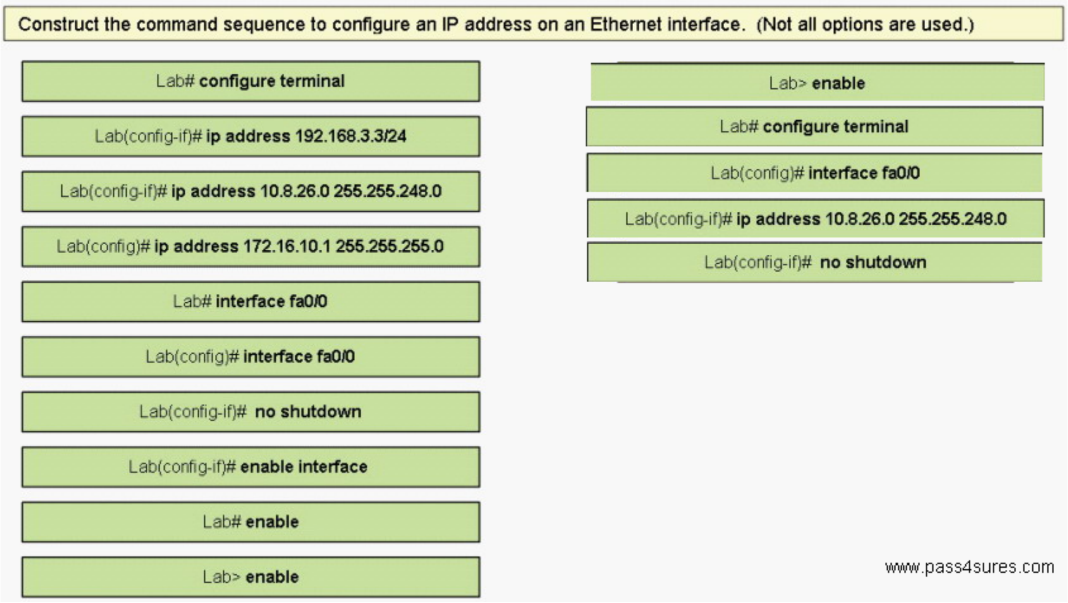

Construct the command sequence to configure an IP address on an Ethernet interface. (Not all options are used.)

Lab# configure terminal

Lab(config-if)# ip address 192.168.3.3/24

Lab(config-if)# ip address 10.8.26.0 255.255.248.0

Lab(config)# ip address 172.16.10.1 255.255.255.0

Lab# interface fa0/0

Lab(config)# interface fa0/0

Lab(config-if)# no shutdown

Lab(config-if)# enable interface

Lab# enable

Lab> enable

Lab> enable

Lab# configure terminal

Lab(config)# interface fa0/0

Lab(config-if)# ip address 10.8.26.0 255.255.248.0

Lab(config-if)# no shutdown

www.pass4sures.com

QUESTION NO: 7

Which of the following IP addresses fall into the CIDR block of 115.64.4.0/22? (Choose three.)

A. 115.64.3.255
B. 115.64.7.64
C. 115.64.5.128
D. 115.64.6.255
E. 115.64.8.32
F. 115.64.12.128

Answer: B,C,D

Explanation:

22 Bits for Network is used
Network ID=256-254=4
First Subnet is 115.64.4
Second Subnet is 115.64.8
So B,C,E Belongs to 115.64.4.0/22 subnet.

QUESTION NO: 8

A company has the following addressing scheme requirements:

-uses a Class B IP address
-currently has 60 subnets
-has a maximum of 1000 computers on any network segment
-needs to leave the fewest unused addresses in each subnet
-uses RIP v1

Which subnet mask is appropriate to use in this network?

A. 255.255.255.128
B. 255.255.252.0
C. 255.255.255.248
D. 255.255.240.0
E. 255.255.248.0
F. 255.255.255.0

Answer: B

Explanation:

NO EXPLANATION

QUESTION NO: 9

Refer to the exhibit. Given the partial configuration shown in the exhibit, why do internal workstations on the 192.168.1.0 network fail to access the Internet?

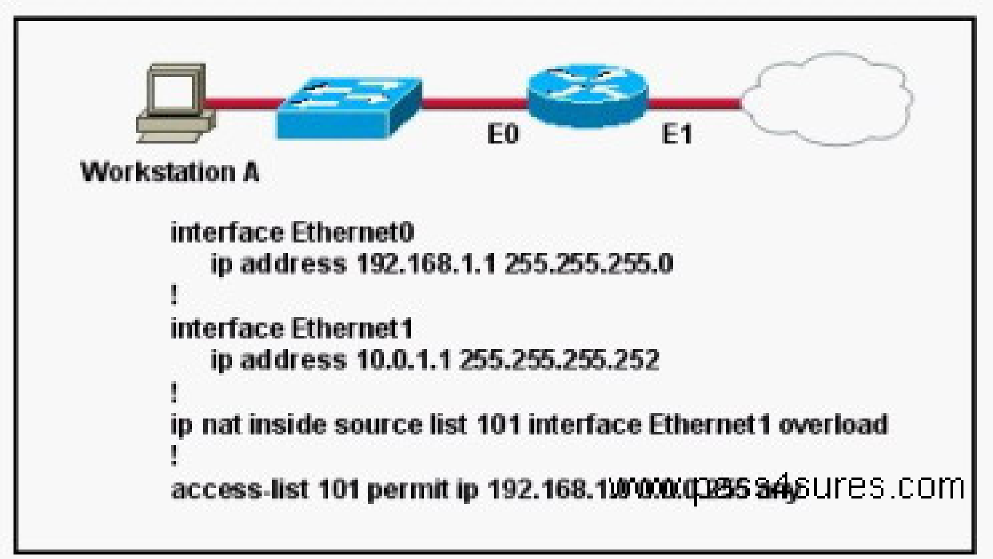

```
interface Ethernet0
    ip address 192.168.1.1 255.255.255.0
!
interface Ethernet1
    ip address 10.0.1.1 255.255.255.252
!
ip nat inside source list 101 interface Ethernet1 overload
!
access-list 101 permit ip 192.168.1.0 0.0.0.255 any
```

A. NAT has not been applied to the inside and outside interfaces.

B. The wrong interface is overloaded.

C. A NAT pool has not been defined.

D. The access list has not been applied to the proper interface to allow traffic out of the internal network.

Answer: A

Explanation:

NO EXPLANATION

QUESTION NO: 10

Refer to the exhibit. A technician has configured the FastEthernet 0/1 interface on Sw11 as an access link in VLAN 1. Based on the output from the show vlan brief command issued on Sw12, what will be the result of making this change on Sw11?

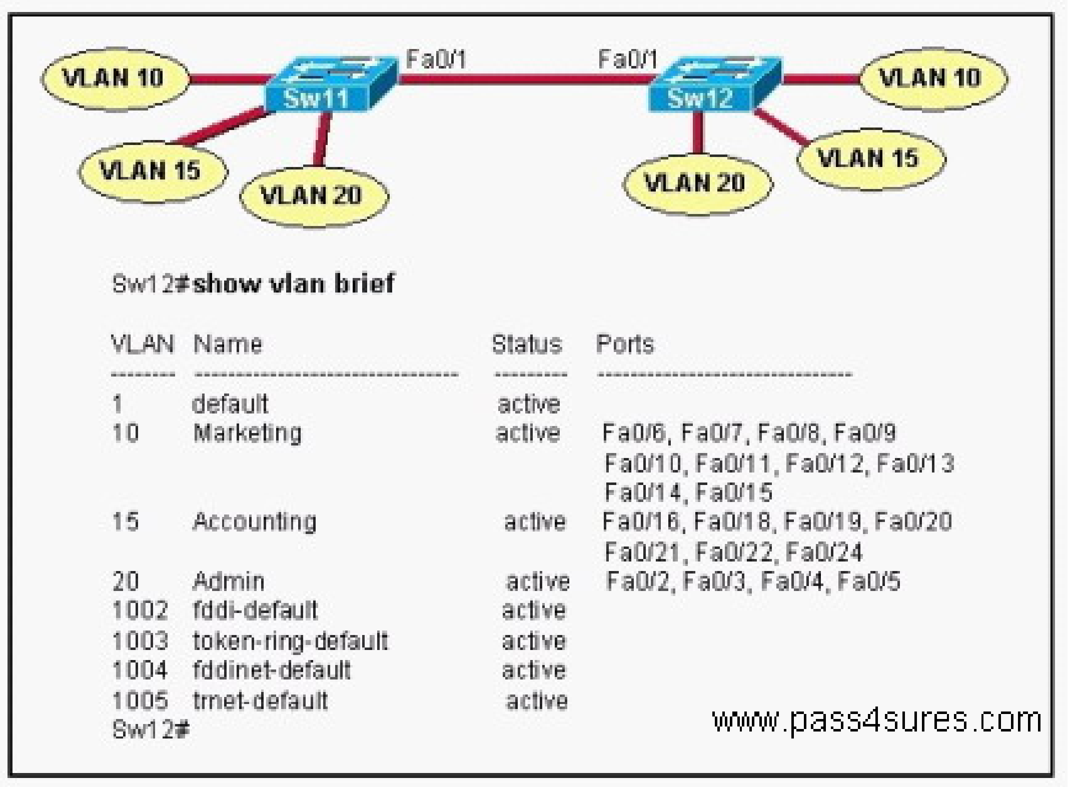

```
Sw12#show vlan brief

VLAN  Name                              Status    Ports
----  --------------------------------  --------  -------------------------------
1     default                           active
10    Marketing                         active    Fa0/6, Fa0/7, Fa0/8, Fa0/9
                                                  Fa0/10, Fa0/11, Fa0/12, Fa0/13
                                                  Fa0/14, Fa0/15
15    Accounting                        active    Fa0/16, Fa0/18, Fa0/19, Fa0/20
                                                  Fa0/21, Fa0/22, Fa0/24
20    Admin                             active    Fa0/2, Fa0/3, Fa0/4, Fa0/5
1002  fddi-default                      active
1003  token-ring-default                active
1004  fddinet-default                   active
1005  trnet-default                     active
Sw12#
```

A. Only the hosts in VLAN 10 and VLAN 15 on the two switches will be able to communicate with each other.

B. Hosts will not be able to communicate between the two switches.

C. The hosts in all VLANs on the two switches will be able to communicate with each other.

D. Only the hosts in VLAN 1 on the two switches will be able to communicate with each other.

Answer: B

Explanation:

NO EXPLANATION

QUESTION NO: 11

In the route highlighted in the graphic, what does the number 782 represent?

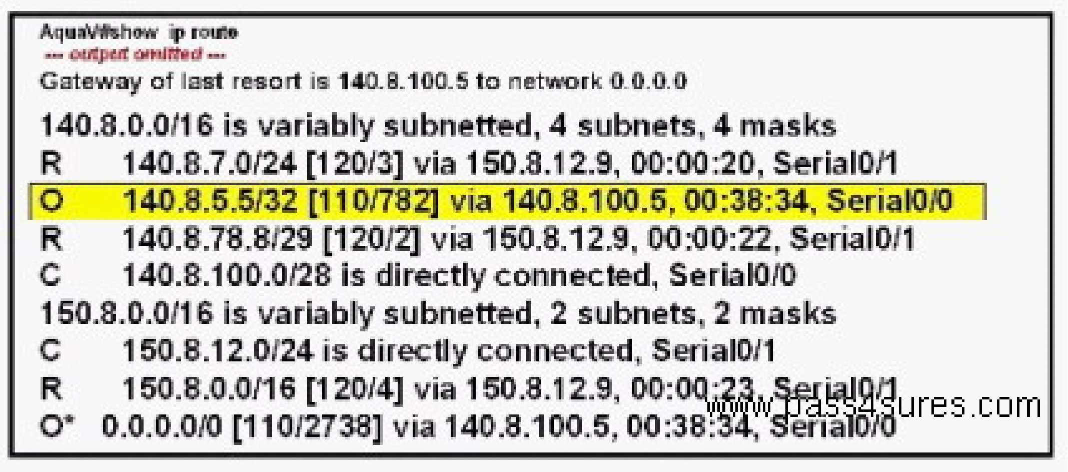

```
AquaW#show ip route
--- output omitted ---
Gateway of last resort is 140.8.100.5 to network 0.0.0.0
140.8.0.0/16 is variably subnetted, 4 subnets, 4 masks
R      140.8.7.0/24 [120/3] via 150.8.12.9, 00:00:20, Serial0/1
O      140.8.5.5/32 [110/782] via 140.8.100.5, 00:38:34, Serial0/0
R      140.8.78.8/29 [120/2] via 150.8.12.9, 00:00:22, Serial0/1
C      140.8.100.0/28 is directly connected, Serial0/0
150.8.0.0/16 is variably subnetted, 2 subnets, 2 masks
C      150.8.12.0/24 is directly connected, Serial0/1
R      150.8.0.0/16 [120/4] via 150.8.12.9, 00:00:23, Serial0/1
O*   0.0.0.0/0 [110/2738] via 140.8.100.5, 00:38:34, Serial0/0
```

A. cost of the route
B. administrative distance
C. hop count
D. delay to the destination

Answer: A

Explanation:
NO EXPLANATION

QUESTION NO: 12

Which statements describe two of the benefits of VLAN Trunking Protocol? (Choose two.)

A. VTP simplifies switch administration by allowing switches to automatically share VLAN configuration information.
B. VTP enhances security by preventing unauthorized hosts from connecting to the VTP domain.
C. VTP allows a single switch port to carry information to more than one VLAN.
D. VTP allows routing between VLANs.
E. VTP helps to limit configuration errors by keeping VLAN naming consistent across the VTP domain.
F. VTP allows physically redundant links while preventing switching loops.

Answer: A,E

Explanation:
The role of the VLAN Trunking Protocol (VTP) is to maintain VLAN configuration consistency across the entire network. VTP is a messaging protocol that uses Layer 2 trunk frames to manage the addition, deletion, and renaming of VLANs on a network-wide basis from a centralized switch that is in the VTP server mode. VTP is responsible for synchronizing VLAN information within a VTP domain. This reduces the need to configure the same VLAN information on each switch.
VTP minimizes the possible configuration inconsistencies that arise when changes are made. These inconsistencies can result in security violations, because VLANs can crossconnect when duplicate names are used. They also could become internally disconnected when they are mapped from one LAN type to another, for example, Ethernet to ATM LANE ELANs or FDDI 802.10 VLANs. VTP provides a mapping scheme that enables seamless trunking within a network employing mixed-media technologies.
VTP provides the following benefits: VLAN configuration consistency across the network Mapping scheme that allows a VLAN to be trunked over mixed media Accurate tracking and monitoring of VLANs Dynamic reporting of added VLANs across the network Plug-and-play configuration when adding new VLANs