

Cisco 642-902



Implementing Cisco IP Routing (ROUTE)


Version: 12.7

Topic 1, Implement an EIGRP based solution, given a network design and a set of requirements

QUESTION NO: 1 CORRECT TEXT

((EIGRP OSPF Redistribution Sim))

In this question you need to redistribute between OSPF and EIGRP such that 172.16.100.1 is reachable from router R1.



R2 is an ASBR for EIGRP 100 and OSPF AREA 24

R3 is an ASBR for EIGRP 100 and OSPF AREA 34

NotE. There are TWO separate areas on TWO separate ASBRS thus you need to do redistribution on R2 and R3

R1 is ONLY in EIGRP 100, and is THE ONLY router you can ping from. R4 has a loopback interface that must be pinged from R1.

R4 is running OSPF and has redundant links to EIGRP network over R3 router.

NotE. You should ping from R1 to 172.16.100.1 network to make sure everything is working correctly.

Answer: First we need to find out 5 K-Values used for EIGRP (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 its metric unit is 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
```

(Notice: In fact, these parameters are just used for reference and we can use other parameters with no problem. Also, a candidate said that the simulator didn't accept the Bandwidth of 1544; in that case, we can use a lower value, like 128.)

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
```

QUESTION NO: 2 CORRECT TEXT

(EIGRP Stub Sim)

By increasing the first distant office, JS manufactures has extended their business. They configured the remote office router (R3) from which they can reach all Corporate subnets. In order to raise network stableness and lower the memory usage and CPU utilization to R3, JS manufactures makes use of route summarization together with the EIGRP Stub Routing feature. Another network engineer is responsible for the implementing of this solution. However, in the process of configuring EIGRP stub routing connectivity with the remote network devices off of R3 has been missing.



Presently JS has configured EIGRP on all routers in the network R2, R3, and R4. Your duty is to find and solve the connectivity failure problem with the remote office router R3. You should then configure route summarization only to the distant office router R3 to complete the task after the problem has been solved. The success of pings from R4 to the R3 LAN interface proves that the fault has been corrected and the R3 IP routing table only contains two 10.0.0.0 subnets.

Answer: First we have to figure out why R3 and R4 can not communicate with each other.

Explanation:

Use the show runningconfig command on router R3

```
R3# show running-config
```

```
<output omitted>
```

```
!
```

```
router eigrp 123
```

```
network 10.0.0.0
```

```
network 172.16.0.0
```

```
no auto-summary
```

```
eigrp stub receive-only
```

```
!
```

```
<output omitted>
```

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

command and replace it with the eigrp stub command:

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

Now R3 will send updates containing its connected and summary routes to other routers. Notice that the eigrp stub command equals to the eigrp stub connected summary because the connected and summary options are enabled by default. Next we will configure router R3 so that it has only 2 subnets of 10.0.0.0 network. Use the show ip route command on R3 to view its routing table

```
R3#show ip route
```

```
Router3# show ip route
```

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

```
D    10.0.0.0/8 is a summary, 02:04:27, Null0
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, Serial0/0
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, 00:00:02, Serial0/0
D    10.2.8.0/24 [90/161280] via 10.2.3.4, 00:00:02, Serial0/0
D    10.2.9.0/24 [90/161280] via 10.2.3.4, 00:00:02, Serial0/0
```

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

```
D    172.16.0.0/16 is a summary, 02:04:27, 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 summarize the networks. Therefore, to make the routing table of R3 has only 2 subnets we have to summarize the other subnets into one subnet. In conclusion, we will use the ip summary-address eigrp 123 10.0.0.0 255.0.0.0 at the interface s0/0 of R4 to summarize.

```
R4>enable
```

```
R4#configure terminal
```

```
R4(config)#interface s0/0
```

```
R4(config-if)#ip summary-address eigrp 123 10.0.0.0 255.0.0.0
```

```
Router3# show ip route
```

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

```
D    10.0.0.0/8 is a summary, 00:06:22, Null0
```

```
C    10.2.3.0/24 is directly connected, Serial0/0
```

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

```
D    172.16.0.0/16 is a summary, 00:06:22, Null0
```

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

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

```
Router3# show ip route
```

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

```
C    10.2.3.0/24 is directly connected, Serial0/0
```

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

```
D    172.16.0.0/16 is a summary, 00:18:43, Null0
```

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

But in your real exam, if you don't see the line "10.0.0.0/8 is a summary,....Null0" then you can summarize using the network 10.2.0.0/16. This summarization is better because all the pings can work well. Finally don't forget to use the copy running-config startup-config command on routers R3 and R4 to save the configurations.

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

```
R4#copy running-config startup-config
```

QUESTION NO: 3

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

- A. EIGRP supports five generic packet types, including Hello, Database Description (DBD), Linkstate Request (LSR), Link-State Update (LSU), and LSAck.
- B. EIGRP sends periodic hello packets to the multicast IP address 224.0.0.10.
- C. EIGRP will not form a neighbor relationship with another peer when their AS number and K values, either or both are mismatched.
- D. EIGRP sends periodic hello packets to the multicast IP address 224.0.0.9.
- E. EIGRP will form a neighbor relationship with another peer even when their K values are mismatched.
- F. EIGRP supports five generic packet types, including Hello, Update, Query, Reply, and ACK packets.

Answer: B,C,F

Reference: <http://www.ietf.org/id/draft-savage-eigrp-00.txt> (see eigrp packets)

QUESTION NO: 4

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

- A. EIGRP will find the interface with the lowest MAC address
- B. The route will transition to the active state
- C. The route will transition to the passive state
- D. EIGRP will automatically use the route with the lowest feasible distance(FD)
- E. EIGRP will automatically use the route with the lowest advertised distance(AD)


Answer: B

Explanation: Explanation

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

QUESTION NO: 5

Refer to the exhibit.



Routers R1 and R2 have established a neighbor relationship and are exchanging routing

information. The network design requires that R1 receive routing updates from R2, but not advertise any routes to R2. Which configuration command sequence will successfully accomplish this task?

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

Answer: C

Explanation:

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

QUESTION NO: 6

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

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

Answer: A

Explanation:

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

QUESTION NO: 7

Refer to the exhibit.

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

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

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

Answer: D


Explanation:

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

The fields to note in this output are as follows:

QUESTION NO: 8

Refer to the exhibit.



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

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

Answer: A

Explanation:

Unlike the ip default-gateway command, you can use ip default-network when ip routing is enabled on the Cisco router. When you configure ip default-network the router considers routes to that network for installation as the gateway of last resort on the router.

For every network configured with ip default-network, if a router has a route to that network, that route is flagged as a candidate default route.


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

Reference:

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

QUESTION NO: 9

Refer to the exhibit.



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

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

Answer: E

Explanation:

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

By default, split horizon is enabled on all interfaces.

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

QUESTION NO: 10

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

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

Answer: A,E

Explanation:

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

Example:

EIGRP

```
router eigrp 1
```

```
redistribute ospf 1 metric 1544 5 255 1 1500
```

```
redistribute rip metric 1544 5 255 1 1500
```

```
network 15.0.0.0
```

RIP

```
router rip
version 2
redistribute eigrp 1 metric 2
redistribute ospf 1 metric 3
network 16.0.0.0
```

QUESTION NO: 11

When troubleshooting an EIGRP connectivity problem, you notice that two connected EIGRP routers are not becoming EIGRP neighbors. A ping between the two routers was successful. What is the next thing that should be checked?

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

Answer: D

Explanation:


The point of this question is about the condition of establish EIGRP neighbor.

You can use these ways to troubleshoot the EIGRP connectivity problem.

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

QUESTION NO: 12

Refer to the exhibit.



You are the network administrator of the Route.com company. You have been tasked to implement a hub and spoke EIGRP topology over Frame Relay to provide connectivity between the networks at headquarters and all 300 spokes.

Before you begin the actual implementation, which three pieces of information are more important to know than the others? (Choose three.)


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

Answer: A,B,E

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

QUESTION NO: 13

Refer to the exhibit.



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

Which is the most important thing that you can quickly verify first to resolve this issue?

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

Answer: E

Explanation: Stub routing is commonly used in a hub and spoke network topology. In a hub and spoke network, one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is adjacent only to

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

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

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

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

QUESTION NO: 14

Refer to the exhibit.



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

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

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

Answer: B,F

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

QUESTION NO: 15

Refer to the exhibit.



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

What will be the effect?

- A. EIGRP will over utilize the 512 kb/s link.


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

Answer: C

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

QUESTION NO: 16

Refer to the exhibit.



ROUTE.com has just implemented this EIGRP network. A network administrator came to you for advice while trying to implement load balancing across part of their EIGRP network.

If the variance value is configured as 2 on all routers and all other metric and K values are configured to their default values, traffic from the Internet to the data center will be load balanced across how many paths?

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

Answer: C

Explanation: Explanation

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

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

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

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

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

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


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

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

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

QUESTION NO: 17

Refer to the exhibit.



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

Which configuration change do you recommend?

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

Answer: A


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

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

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

QUESTION NO: 18

Refer to the exhibit.



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

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

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

Answer: B

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

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

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

QUESTION NO: 19

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

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

Answer: B

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

QUESTION NO: 20

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

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

Answer: C

Explanation: Explanation

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

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

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

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

QUESTION NO: 21

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

- A.** EIGRP would be appropriate if LMI type ANSI is NOT used.


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

Answer: D

Explanation: Explanation

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

For example, in our Frame Relay topology a Hub is connected with 4 Spoke routers. The main Frame Relay interface on Hub router is 512Kbps which is not enough to use for 6 links of 128 Kbps (= 768 Kbps).



The solution here is we can use $512 / 6 = 85$ Kbps on each subinterface of Hub by using "bandwidth 85 command. For example:

Hub(config)#interface Serial0/0.1 point-to-pointHub(config-subif)#bandwidth 85

Also on Spoke routers we need to set this value. For example on Spoke1:

Spoke1(config)#interface Serial0/1.0 point-to-pointSpoke1(config-subif)#bandwidth 85

Notice that by default, EIGRP limits itself to use no more than 50% of the configured interface bandwidth. In this case EIGRP will not use more than 42.5 Kbps (50% of 85 Kbps).

(For more information about implementing EIGRP over Frame Relay, please read http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094063.shtml)

QUESTION NO: 22

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

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

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

Answer: C

Explanation: Explanation


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

redistributing from a routing protocol into another:

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

QUESTION NO: 23

Refer to the exhibit.



Which three statements are true? (Choose three.)

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

Answer: A,B,D

Explanation: Explanation


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

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

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

QUESTION NO: 24

Refer to the exhibit.



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

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

Answer: E

Explanation: Explanation

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

QUESTION NO: 25

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

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

Answer: D

Explanation: Explanation

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

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

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

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

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

QUESTION NO: 26

You have implemented mutual route redistribution between OSPF and EIGRP on a border router.

When checking the routing table on one of the OSPF routers within the OSPF routing domain, you are seeing some, but not all of the expected routes.

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

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

Answer: D,E

Explanation: When troubleshooting EIGRP routes being distributed into OSPF, the first item to check is that the missing EIGRP routes are in fact on the router performing the redistribution (border router). If not, the routes cannot be redistributed into OSPF and troubleshooting should occur within the EIGRP AS.

The second step is to verify that the subnets keyword has been appended to the redistribute eigrp statement under the OSPF routing process. The question mentions that the networks being redistributed are using a /28 mask, with implies that the networks are actually subnets not assigned to classful boundaries.

OSPF will only redistribute classful networks by default. The keyword "subnets" allows subnets to also be redistributed into OSPF (not merely classful routes).

QUESTION NO: 27

Refer to the exhibit.

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

All EIGRP load balancing parameters are set to their defaults. You want to use all the routes in the EIGRP topology for IP load balancing. Which two EIGRP subcommands would you use to

accomplish this goal? (Choose two.)

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

Answer: C,E

Explanation: Explanation

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

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

QUESTION NO: 28

Refer to the exhibit.

R1#show ip route

```

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

```


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

- A. router eigrp 190
redistribute static
!
ip route 0.0.0.0 0.0.0.0 Null0
- B. ip default-network 20.0.0.0
- C. router eigrp 190
redistribute static
!
ip route 0.0.0.0 255.255.255.255 Null0
- D. ip default-network 20.20.20.0

Answer: A

Explanation:

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

QUESTION NO: 29

Refer to the exhibit.

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

A partial routing configuration is shown. Complete the configuration so that only the default-

network is redistributed from EIGRP 190 into EIGRP 212. Which ACL statement completes the configuration correctly?

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

Answer: C

Explanation: Explanation

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


Note:

any equals **0.0.0.0 255.255.255.255**

host 0.0.0.0 equals **0.0.0.0 0.0.0.0**

QUESTION NO: 30

Refer to the exhibit.



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

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

Answer: A,C,F

Reference:

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

QUESTION NO: 31

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



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


Answer: D

Explanation:

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

QUESTION NO: 32

Refer to the exhibit.



In a redundant hub-and-spoke deployment using EIGRP, what feature can be used to ensure that routers C through F are not used as transit routers for data traveling from router B to network 10.1.1.0? Select the best response

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

Answer: B

Explanation:

Explanation

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

QUESTION NO: 33

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

- A. EIGRP
- B. OSPF
- C. RIPv1
- D. RIPv2
- E. BGP
- F. IS-IS

Answer: A


Explanation:

Unlike most internal routing protocols, EIGRP has a really cool feature that allows you to share the load of your traffic across multiple links and not just links that have the same cost values. EIGRP allows you to make full use of your redundant links that could be in place just to have for back up but you are paying out a lot of money just to sit there and do nothing. EIGRP makes it easy for us the network engineers to make this happen. Before jumping in the in's and out lets run through a few things first when it comes to EIGRP Load Sharing, also refereed to as Load Balancing sometimes.

Reference: <http://ericleahy.com/index.php/eigrp-equal-and-unequal-cost-load-sharing/>

QUESTION NO: 34

Refer to the exhibit.



ROUTE Enterprises has many stub networks in their enterprise network, such as router B and its associated network. EIGRP is to be implemented on router A so that neither the prefix for the S0/0/0 interface nor the prefixes from router B appear in the routing tables for the router in the enterprise network. Which action will accomplish this goal?

- A. Declare router B a stub router using the `eigrp stub` command.
- B. Use the `passive-interface` command for interface Serial0/0/0.
- C. Use a mask with the `network` command to exclude interface Serial0/0/0.
- D. Implement a distribute list to exclude the link prefix from the routing updates.

Answer: C

Explanation: Explanation

If we declare router B a stub router then the routers in Enterprise Network still learn about the network for S0/0/0 interface and the network behind router B.

If we use the `passive-interface` command on s0/0/0 interface then router A & B can not become neighbors because they don't exchange hello messages -> A can not send traffic to the network behind B .


Theoretically, we can use a distribute list to exclude both the link prefix and the prefix from router B but it is not efficient because:

+ We have many stub networks so we will need a "long" distribute list.+ We declare networks in stub routers (like router B) while filter them out at router A -> it is a waste.

Not totally sure about answer C because if we "use a mask with the `network` command to exclude interface Serial0/0/0 then router A and B can not become neighbors and the situation is same as answer B. But from many discussions about this question, maybe C is the best answer.

QUESTION NO: 35

Refer to the exhibit.



Which statement about dynamic routing protocols for this network is true?

- A. No dynamic interior routing protocol can summarize as shown.
- B. Unless configured otherwise, EIGRP would automatically summarize the prefixes as shown in the exhibit.
- C. With this IP addressing scheme, EIGRP can be manually configured to summarize prefixes at the specified summarization points.
- D. The IP address design lends itself to OSPF. Each summarizing router would be an ABR, summarizing to the next area in the address hierarchy.

Answer: C

Explanation:

Summarization may be manually applied at any point in the network. You can configure manual summarization on any router interface. Consider summarization for both upstream and downstream neighbors. Upstream neighbors should receive a consolidated route, and downstream neighbors can receive a default route.

QUESTION NO: 36

After implementing EIGRP on your network, you issue the `show ip eigrp traffic` command on router

C. The following output is shown:

RouterC#show ip eigrp traffic

IF-EIGRP Traffic Statistics for process 1

Hellos sent/received: 481/444

Updates sent/received: 41/32

Queries sent/received: 5/1

Replies sent/received: 1/4

Acks sent/received: 21/25

Input queue high water mark 2, 0 drops

SIA-Queries sent/received: 0/0

SIA-Replies sent/received: 0/0

Approximately 25 minutes later, you issue the same command again. The following output is shown:

RouterC#show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 1

Hellos sent/received: 1057/1020

Updates sent/received: 41/32

Queries sent/received: 5/1

Replies sent/received: 1/4

Acks sent/received: 21/25

Input queue high water mark 2, 0 drops

SIA-Queries sent/received: 0/0

SIA-Replies sent/received: 0/0

Approximately 25 minutes later, you issue the same command a third time. The following output is shown:

RouterC#show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 1

Hello sent/received: 1754/1717

Updates sent/received: 41/32

Queries sent/received: 5/1

Replies sent/received: 1/4

Acks sent/received: 21/25

Input queue high water mark 2, 0 drops

SIA-Queries sent/received: 0/0

SIA-Replies sent/received: 0/0

What can you conclude about this network?

- A. The network has been stable for at least the last 45 minutes.
- B. There is a flapping link or interface, and router C knows an alternate path to the network.
- C. There is a flapping link or interface, and router A does not know an alternate path to the network.
- D. EIGRP is not working correctly on router C.
- E. There is not enough information to make a determination.

Answer: A

Explanation: Explanation

In three times using the command, the "Queries sent/received" & "Replies sent/received" are still the same -> the network is stable.

QUESTION NO: 37

After implementing EIGRP on your network, you issue the show ip eigrp traffic command on router C. The following output is shown:

RouterC#show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 1

Hello sent/received: 2112/2076

Updates sent/received: 47/38

Queries sent/received: 5/3

Replies sent/received: 3/4

Acks sent/received: 29/33

Input queue high water mark 2, 0 drops

SIA-Queries sent/received: 0/0

SIA-Replies sent/received: 0/0

Moments later, you issue the same command a second time and the following output is shown:

RouterC#show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 1

Hellos sent/received: 2139/2104

Updates sent/received: 50/39

Queries sent/received: 5/4

Replies sent/received: 4/4

Acks sent/received: 31/37

Input queue high water mark 2, 0 drops

SIA-Queries sent/received: 0/0

SIA-Replies sent/received: 0/0

Moments later, you issue the same command a third time and the following output is shown:

RouterC#show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 1

Hellos sent/received: 2162/2126

Updates sent/received: 53/42

Queries sent/received: 5/5

Replies sent/received: 5/4

Acks sent/received: 35/41

Input queue high water mark 2, 0 drops

SIA-Queries sent/received: 0/0

SIA-Replies sent/received: 0/0

What information can you determine about this network?

- A. The network is stable.
- B. There is a flapping link or interface, and router C knows an alternate path to the network.
- C. There is a flapping link or interface, and router C does not know an alternate path to the network.
- D. EIGRP is not working correctly on router C.
- E. There is not enough information to make a determination.


Answer: B

Explanation: Explanation

We notice that the "Queries received" number is increased so router C has been asked for a route. The "Replies sent" number is also increased -> router C knows an alternate path to the network.

QUESTION NO: 38

Refer to the exhibit.



EIGRP has been configured on router D. Router C is performing mutual redistribution between EIGRP and OSPF. While verifying that the redistribution is functioning properly, you discover that while router C has all of the EIGRP routes in its routing table, router A does not have any routes from the EIGRP domain. What on router C may be the cause of the problem?

- A. The no auto-summary command needs to be added under router eigrp 1.
- B. The subnets keyword was not included in the redistribute command under router ospf 1.
- C. The metric specified for the redistributed EIGRP routes is too large; making the EIGRP routes

unreachable by router A.

D. The default-information originate command needs to be added under router ospf 1.

E. The administrative distance of either OSPF or EIGRP must be changed so that EIGRP has a higher administrative distance than OSPF.


Answer: B

Explanation: Explanation

If we don't use the "subnets" keyword when redistributing routes learned from another routing process into OSPF, only classful routes will be redistributed. This is an important thing to remember when redistributing into OSPF.

QUESTION NO: 39

During the redistribution process configured on RTA, some of the EIGRP routes, such as 10.1.1.0/24 and 10.2.2.0/24, are not being redistributed into the OSPF routing domain.



Which two items could be a solution to this problem? (Choose two.)

- A.** Change the metric-type to 2 in the redistribute command.
- B.** Configure the redistribute command under router eigrp 1 instead.
- C.** Change the EIGRP AS number from 100 to 1 in the redistribute command.
- D.** Add the subnets option to the redistribute command.
- E.** Change the metric to an EIGRP compatible metric value (bandwidth, delay, reliability, load, MTUs) in the redistribute command.

Answer: C,D

Explanation:

In this example, the router is configured for EIGRP AS 1, but EIGRP AS 100 is being redistributed

into OSPF so the EIGRP AS needs to be changed from 100 to 1.

The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword, only networks that are not subnetted are redistributed by OSPF.

Example:

```
RTA(config)#router ospf 109RTA(config-router)#redistribute rip subnetsRTA(config-router)#network 130.10.62.0 0.0.0.255 area 0RTA(config-router)#network 130.10.63.0 0.0.0.255 area 0
```

The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword, only networks that are not subnetted are redistributed by OSPF.

QUESTION NO: 40

Given the accompanying output, which additional command is needed to redistribute IGRP into EIGRP?

```
Router eigrp 123
```

```
Network 10.10.10.0
```

```
No auto-summary
```

```
!
```

```
Router igrp 123
```

```
Network 172.16.0.0
```

```
Network 172.17.0.0
```

- A. Under the router igrp mode add redistribute eigrp 123
- B. Under the router eigrp mode add redistribute igrp 123
- C. Under the router igrp mode add redistribute igrp 123 subnets
- D. None, EIGRP and IGRP are automatically redistributed in this instance.

Answer: D

Explanation:

The point of this question is redistribute IGRP into EIGRP.

When redistributing IGRP into EIGRP, there is a feature that they are automatically redistributed if they have same autonomous system number; in opposite, they need to manually redistributed if they have different autonomous system number..

QUESTION NO: 41

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

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

Answer: B

Explanation:

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

Example:

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

Router# show ip eigrp traffic

IP-EIGRP Traffic Statistics for process 77

Hellos sent/received. 218/205

Updates sent/received. 7/23

Queries sent/received. 2/0

Replies sent/received. 0/2

Acks sent/received. 21/14

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

QUESTION NO: 42

Which three statements are true about EIGRP operation? (Choose three.)

- A. When summarization is configured, the router will also create a route to null 0.
- B. The summary route remains in the route table, even if there are no more specific routes to the network.
- C. Summarization is configured on a per-interface level.
- D. The maximum metric for the specific routes is used as the metric for the summary route.

E. Automatic summarization across major network boundaries is enabled by default.

Answer: A,C,E

Reference: <http://astorinonetworks.com/2011/07/20/summary-routes-to-null0-the-protocols-that-love-them/>

QUESTION NO: 43

Which two statements about the EIGRP DUAL process are correct? (Choose two.)

- A. An EIGRP route will go active if there are no successors or feasible successors in the EIGRP topology table.
- B. An EIGRP route will go passive if there are no successors in the EIGRP topology table.
- C. DUAL will trigger an EIGRP query process while placing the flapping routes in the holddown state.
- D. A feasible successor in the EIGRP topology table can become the successor only after all the query requests have been replied to.
- E. The stuck in active state is caused when the wait for the query replies have timed out.
- F. EIGRP queries are sent during the loading state in the EIGRP neighbor establishment process.

Answer: A,E

Reference: <http://routemyworld.com/category/routing-protocols/eigrp/>

QUESTION NO: 44

What are three key concepts that apply when configuring the EIGRP stub routing feature in a hub and spoke network? (Choose three.)

- A. A hub router prevents routes from being advertised to the remote router.
- B. Only remote routers are configured as stubs.
- C. Stub routers are not queried for routes.
- D. Spoke routers connected to hub routers answer the route queries for the stub router.
- E. A stub router should have only EIGRP hub routers as neighbors.
- F. EIGRP stub routing should be used on hub routers only.


Answer: B,C,E

Reference: <http://astorinonetworks.com/2011/06/14/eigrp-stub-routing/>

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

QUESTION NO: 45

Refer to the exhibit. EIGRP is configured with the default configuration on all routers. Autosummarization is enabled on routers R2 and R3, but it is disabled on router R1. Which two EIGRP routes will be seen in the routing table of router R3? (Choose two.)



- A. 10.0.0.0/8
- B. 10.10.0.0/16
- C. 10.10.10.0/24
- D. 172.16.0.0/16
- E. 172.16.0.0/24
- F. 172.16.10.0/24

Answer: C,D

Explanation: Explanation

EIGRP performs an auto-summarization each time it crosses a border between two different major networks. In this case all different networks are in different major networks so EIGRP will perform auto-summarization when it exits an interface. But R1 has been configured with “no auto-summary” command so EIGRP will not summarize on S0 interface of R1. So the routing table of R2 will have the network 10.10.10.0/24 (not be summarized).

When exiting S1 interface of R2, EIGRP summarizes network 172.16.10.0/24 into the major 172.16.0.0/16 network but it does not summarize network 10.10.10.0/24 because it is not directly connected with this network. Therefore in the routing table of R3 there will have:

+ Network 10.10.10.0/24 (not summarized)+ Network 172.16.0.0/16 (summarized)

-> C and D are correct.

Note: I simulated this question on GNS3, you can see the final outputs of the “show ip route” commands on these routers (I connected these routers via FastEthernet, not Serial interfaces so

the outputs are slightly different but the main points are not changed).

```
R1#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

    172.16.0.0/24 is subnetted, 1 subnets
C       172.16.10.0 is directly connected, FastEthernet0/0
D       192.168.20.0/24 [90/30720] via 172.16.10.2, 00:02:15, FastEthernet0/0
    10.0.0.0/24 is subnetted, 1 subnets
C       10.10.10.0 is directly connected, Loopback0
```

```
R2#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

    172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C       172.16.10.0/24 is directly connected, FastEthernet0/0
D       172.16.0.0/16 is a summary, 00:02:33, Null0
C       192.168.20.0/24 is directly connected, FastEthernet0/1
    10.0.0.0/24 is subnetted, 1 subnets
D       10.10.10.0 [90/156160] via 172.16.10.1, 00:02:36, FastEthernet0/0
```

```
R3#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route


Gateway of last resort is not set

D       172.16.0.0/16 [90/30720] via 192.168.20.2, 00:03:38, FastEthernet0/0
C       192.168.20.0/24 is directly connected, FastEthernet0/0
    10.0.0.0/24 is subnetted, 1 subnets
D       10.10.10.0 [90/158720] via 192.168.20.2, 00:03:38, FastEthernet0/0
```

QUESTION NO: 46

Refer to the exhibit and the partial configuration on router R2. On router R4 all RIP routes are

redistributed into the OSPF domain. A second redistribution is configured on router R2 using a route map. Based on the configuration on router R2, which EIGRP external routes will be present in the routing table of R1?



- A. the routes originating from the RIP routing domain
- B. the routes originating from the OSPF stub area
- C. all OSPF inter and intra-area routes
- D. all routes originating from RIP and OSPF routing domains

Answer: A


Explanation: Explanation

R2 sees the routes from RIP domain as external routes while it sees the routes from OSPF Stub Area as internal routes. From the output we learn that the "route-type external" is redistributed from OSPF to EIGRP (via route-map ABC) so we will see the routes from the RIP domain (external) in the routing table of R1.

In the case we want to redistribute routes from OSPF Stub Area (Area 1) to EIGRP we need to use the "match route-type internal" command instead.

QUESTION NO: 47

Refer to the exhibit. EIGRP has been configured on routers R1 and R2. However, R1 does not show R2 as a neighbor and does not accept routing updates from R2. What could be the cause of the problem?



- A. The no auto-summary command has not been issued under the EIGRP process on both routers.
- B. Interface E0 on router R1 has not been configured with a secondary IP address of 10.1.2.1/24.
- C. EIGRP cannot exchange routing updates with a neighbor's router interface that is configured with two IP addresses.
- D. EIGRP cannot form neighbor relationship and exchange routing updates with a secondary address.

Answer: D

Explanation: Explanation

EIGRP updates always use the primary IP address of the outgoing interface as the source address. In this case R2 will use the 10.1.2.2/24 address, which is not in the same subnet of R1, to send EIGRP update to R1. Therefore R1 does not accept this update and generates the "not on common subnet" error message.

Answer D is a bit unclear. It should state that "EIGRP cannot form neighbor relationship and exchange routing updates if the two primary addresses on two routers are not in the same subnet".

Notice that although R1 does not accept R2 as its EIGRP neighbors but R2 accepts R1 as its


EIGRP neighbor accepts R1 hello packets..

For more information about this problem, please read

http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a0080093f09.shtml.

QUESTION NO: 48

Refer to the exhibit.



EIGRP had converged in AS 1 when the link between router R1 and R2 went down. The console on router R2 generated the following messages:

*Mar 20 12:12:06: %DUAL-5-NBRCHANGE. IP-EIGRP 1: Neighbor 10.1.4.3 (Serial0) is down: stuck in active

*Mar 20 12:15:23: %DUAL-3-SIA. Route 10.1.1.0/24 stuck-in-active state in IP-EIGRP 1. Cleaning up The network administrator issued the show ip eigrp topology active command on R2 to check the status of the EIGRP network. Which statement best describes the reason for the error messages?

- A. Incorrect bandwidth configuration on router R3 prevents R2 from establishing neighbor adjacency.
- B. Incorrect bandwidth configuration on router R5 prevents R2 from establishing neighbor adjacency.
- C. Router R3 did not reply to the query about network 10.1.1.0/24 sent by router R2 .
- D. Router R5 did not reply to the query about network 10.1.1.0/24 sent by router R2 .

Answer: C

Explanation: Explanation

When the link between R1 and R2 is down, R2 loses its successor for the network 10.1.1.0/24. R2 checks its topology table for a feasible successor but it can't find one. So R2 goes active on the that route to find a new successor by sending queries out to its neighbors (R3 and R5) requesting a path to the lost route. Both R3 and R5 also go "active" for the that route. But R5 doesn't have any neighbor to ask besides R2 so it will send an "unreachable message" to indicate it has no alternative path for that route and has no other neighbor to query. R3 also checks its EIRGP topology table for a feasible successor but it has none, too. Unlike R5, R3 has a neighbor (R4) so it continues to query this router.

Now suppose there is a problem on the link between R3 and R4 so R4 never receives the query from R3 and of course, R3 also never receives a reply back from R4. Therefore, R3 can't reply back to R2. After about 3 minutes, the "Stuck in active" (SIA) timer on R2 expires and R2 marks the route 10.1.1.0/24 as "stuck in active" route.

The output line "via 10.1.3.3 (Infinity/Infinity), r, Serial0, serno 1232 indicates R2 has sent a query to 10.1.3.3 and is waiting for a reply (the lowercase r).

(Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008010f016.shtml)

QUESTION NO: 49

Which EIGRP packet statement is true?

- A. On high-speed links, hello packets are broadcast every 5 seconds for neighbor discovery.
- B. On low-speed links, hello packets are broadcast every 15 seconds for neighbor discovery.
- C. Reply packets are multicast to IP address 224.0.0.10 using RTP.
- D. Update packets route reliable change information only to the affected routers.
- E. Reply packets are used to send routing updates.

Answer: D


Reference:

http://docwiki.cisco.com/wiki/Enhanced_Interior_Gateway_Routing_Protocol#EIGRP_Packet_Types

QUESTION NO: 50

Refer to the exhibit. EIGRP has been configured on all routers in the network. The command

metric weights 0 0 1 0 0 has been added to the EIGRP process so that only the delay metric is used in the path calculations. Which router will R1 select as the successor and feasible successor for Network A?



- A. R4 becomes the successor for Network A and will be placed in the routing table. R2 becomes the feasible successor for Network A.
- B. R4 becomes the successor for Network A and will be included in the routing table. No feasible successor will be selected as the advertised distance from R2 is higher than the feasible distance.
- C. R2 becomes the successor and will be placed in the routing table. R4 becomes the feasible successor for Network A.
- D. R2 becomes the successor and will be placed in the routing table. No feasible successor will be selected as the reported distance from R4 is lower than the feasible distance.

Answer: B

Explanation:

The point of the question is DUAL of EIGRP.

$FD = 5 + 10 + 5 = 20$

R4 is successor

No FS, because no other router's AD is lower the FD

QUESTION NO: 51

During a redistribution of routes from OSPF into EIGRP, an administrator notices that none of the OSPF routes are showing in EIGRP. What are two possible causes? (Choose two.)

- A. incorrect distribute lists have been configured
- B. missing ip classless command
- C. CEF not enabled
- D. no default metric configured for EIGRP

Answer: A,D


Explanation: Explanation

An incorrect distribute list can filter out updates therefore none of the OSPF routes are showing in EIGRP.

The default metric when redistributing into EIGRP is infinite so we must specify a seed metric for EIGRP to work with.

QUESTION NO: 52

You are a network technician, study the exhibit carefully.



What must be done on router A in order to make EIGRP work effectively in a Frame Relay multipoint environment?

- A. Issue the command bandwidth 56 on the physical interface.
- B. Issue the command bandwidth 56 on each subinterface.
- C. Issue the command bandwidth 224 on each subinterface.
- D. Issue the command bandwidth 224 on the physical interface.

Answer: D

Explanation: Explanation

In Frame Relay, all neighbors share the same bandwidth, regardless of the actual CIR of each individual PVC. In this case the CIR of each PVC is the same so we can find the bandwidth of the

main interface (multipoint connection interface) by $56 \times 4 = 224$.

Notice that if the bandwidth on each PVC is not equal then we get the lowest bandwidth to multiply.

QUESTION NO: 53

Based on the exhibited output, which three statements are true? (Choose three.)

```
R1#show ip route

Gateway of last resort is 10.1.1.2 to network 0.0.0.0

C    1.0.0.0/8 is directly connected, Loopback0
  172.17.0.0/24 is subnetted, 1 subnets
D    172.17.1.0 [90/25632000] via 10.1.1.2, 00:05:20, Serial0/0
  172.16.0.0/24 is subnetted, 1 subnets
D    172.16.1.0 [90/23072000] via 10.1.1.2, 00:05:20, Serial0/0
      [90/20640000] via 10.1.1.3, 00:00:13, Serial0/0
D    172.19.0.0/16 [90/391248640] via 10.1.1.3, 00:05:20, Serial0/0
D    172.22.0.0/16 [90/20640000] via 10.1.1.3, 00:05:21, Serial0/0
D EX 172.25.0.0/16 [170/32032000] via 10.1.1.2, 00:00:10, Serial0/0
  10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
D    10.2.0.0/16 is a summary, 00:06:18, Null0
C    10.2.1.0/24 is directly connected, FastEthernet0/0
C    10.1.1.0/24 is directly connected, Serial0/0
D*EX 0.0.0.0/0 [170/20514560] via 10.1.1.2, 00:00:11, Serial0/0
R1#
```

- A. R1 is configured with the variance command.
- B. The route to 10.2.0.0/16 was redistributed into EIGRP.
- C. A default route has been redistributed into the EIGRP autonomous system.
- D. R1 is configured with the ip summary-address command.
- E. The router at 10.1.1.2 is configured with the ip default-network 0.0.0.0 command.
- F. R1 is sourcing an external EIGRP route from Null0.

Answer: A,C,D

Explanation: Explanation

From the routing table above, we see that network 172.16.1. can be reached via 2 unequal paths (with FD of 23072000 & 20640000) so surely R1 has been configured with the "variance" command

By configuring a default route and redistribute it into EIGRP you will get the line "D *EX 0.0.0.0/0 ..." line in the routing table of that router

From the line "10.2.0.0/16 is a summary, 00:16:18, Null0 we know that this network has been summarized with the "ip summaray-address" command (notice that 10.2.0.0 is not the major network of net

QUESTION NO: 54

Examine the exhibit carefully.

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

EIGRP is configured on all routers in the network. What conclusion can be derived from the show ip eigrp topology output provided?

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

Answer: D

Explanation:

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

The fields to note in this output are as follows:

QUESTION NO: 55

Which three statements are true about EIGRP route summarization? (Choose three.)

- A.** Manual route summarization is configured in router configuration mode when the router is configured for EIGRP routing.
- B.** Manual route summarization is configured on the interface.
- C.** When manual summarization is configured, the summary route will use the metric of the largest specific metric of the summary routes.
- D.** The ip summary-address eigrp command generates a default route with an administrative distance of 90.
- E.** The ip summary-address eigrp command generates a default route with an administrative distance of 5.
- F.** When manual summarization is configured, the router immediately creates a route that points to null0 interface.

Answer: B,E,F

Explanation:


The purpose of route summarization is small routing tables, smaller updates. On major network boundaries, subnetworks are summarized to a single classful network and automatic route summarization is enabled by default. Manual route summarization can be configured on per interface basis. When summarization is configured on an interface, the router immediately creates a route pointing to null0.

Route summarization works in conjunction with the ip summary-address eigrp interface configuration command, in which additional summarization can be performed. If automatic summarization is in effect, there usually is no need to configure network level summaries using the ip summary-address eigrp command. You can configure a summary aggregate address for a specified interface. If there are any more specific routes in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes

Reference: http://www.cisco.com/en/US/docs/ios/12_0/np1/configuration/guide/1ceigrp.html

QUESTION NO: 56

Refer to the exhibit.



Routers R1 and R2 are running EIGRP and have converged. On the basis of the information that is presented, which statement is true?

- A. All outgoing routing updates from router R1 to router R2 will be suppressed, but the inbound updates will continue to be received.
- B. All incoming routing updates from R2 will be suppressed, but the outgoing updates will continue to be sent.
- C. Both outgoing and incoming routing updates on R1 will be stopped because of the passive-interface Serial0/0 configuration statement.
- D. Both outgoing and incoming routing updates on R1 will be permitted because the distribute-list 20 out Serial0/0 command cannot be used with association with the outgoing interface.

Answer: C

Explanation:

You can use the passive-interface command to control the advertisement of routing information. The command enables the suppression of routing updates over some interfaces while it allows updates to be exchanged normally over other interfaces.

With most routing protocols, the passive-interface command restricts outgoing advertisements only. However, when used with Enhanced Interior Gateway Routing Protocol (EIGRP), the effect is slightly different. With EIGRP running on a network, the passive-interface command stops both outgoing and incoming routing updates, since the effect of the command causes the router to stop sending and receiving hello packets over an interface.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0a.shtml

QUESTION NO: 57

What does the default value of the EIGRP variance command of 1 mean?

- A. Load balancing is disabled on this router.
- B. The router performs equal-cost load balancing.
- C. Only the path that is the feasible successor should be used.
- D. The router only performs equal-cost load balancing on all paths that have a metric greater than 1.

Answer: B


Explanation: Explanations

The point of the question is about the balance configuration of EIGRP.

If variance is 1, it means that it support Equal cost path.

QUESTION NO: 58

Refer to the exhibit.



R1 and R2 have been configured to share routing information via EIGRP. What will be the result of the configuration section shown for R2?

- A. Any routes learned by R2 from the interface tied to the 172.16.0.0 network will not be advertised to neighbors on the 192.168.2.0 network.
- B. Only routes learned by R2 from the interface tied to the 172.16.0.0 network will be advertised to neighbors on the 192.168.2.0 network.


- C.** Only the 172.16.0.0 network will be advertised to neighbors on the 192.168.2.0 network.
D. All networks, except the 172.16.0.0 network will be advertised to neighbors on the 192.168.2.0 network.

Answer: C

Explanation:

QUESTION NO: 59

Refer to the exhibit.



Which two statements are true? (Choose two.)

- A.** The eigrp stub command prevents queries from being sent from R2 to R1.
B. The eigrp stub command will automatically enable summarization of routes on R2.
C. The eigrp stub command prevents all routes except a default route from being advertised to R1.
D. Router R1 will advertise connected and summary routes only.
E. Router R1 will advertise connected and static routes. The sending of summary routes will not be permitted.
F. Router R1 is configured as a receive-only neighbor and will not send any connected, static or summary routes.

Answer: A,D

Explanation: Explanation


The command "eigrp stub" turns R1 into a stub router so R2 will never send any query to R1 because R2 knows that a stub router will only route packets for networks it has explicitly advertised.

The command "eigrp stub" is same as "eigrp stub connected summary" command because connected and summarized routes are advertised by default.

Note: Because the network 192.168.50.0 is not advertised by "network" statement, it is necessary to redistribute connected route with the "redistribute connected" command.

QUESTION NO: 60

Refer to the exhibit.



The routing protocols EIGRP and OSPF have been configured as indicated in the exhibit. Given the partial configuration of router R2, which network will be present in the routing table of R4?

- A. Network A
- B. Network B
- C. Network A and Network B
- D. neither Network A nor Network B

Answer: B

Explanation:

In this exhibit the OSPF domain is redistributed into the EIGRP 100 domain so Network B will present into Router R-4. However, the Network A network will not be seen on router R-4 (The bottom router which is improperly labeled Network B) because EIGRP 50 was not redistributed into EIGRP 100.

QUESTION NO: 61

When an EIGRP topology change is detected, what is the correct order of events when there is a FS?

- A. The neighbor adjacency is deleted.
The feasible route is used.
DUAL is notified.
Remove all topology entries learned from that neighbor.
- B. DUAL is notified.
Remove all topology entries learned from that neighbor.
The neighbor adjacency is deleted.
Routes enter the Active state and the feasible route is used.
- C. The neighbor adjacency is deleted.
Routes enter the Active state and the feasible route is used.
DUAL is notified.
Remove all topology entries learned from that neighbor.
- D. DUAL is notified.
The neighbor adjacency is deleted.
Remove all topology entries learned from that neighbor.
The feasible route is used.

Answer: D

Explanation:


If a packet is not received before the expiration of the hold time, the neighbor adjacency is deleted, and all topology table entries learned from that neighbor are removed, as if the neighbor had sent an update stating that all the routes are unreachable. If the neighbor is a successor for any destination networks, those networks are removed from the routing table, and alternative paths, if available, are computed. This lets the routes quickly reconverge if an alternative feasible route is

available.

QUESTION NO: 62

Refer to the exhibit. The network administrator is trying to configure mutual redistribution between EIGRP and OSPF. Autosummarization in EIGRP 100 AS is disabled. After adding OSPF configuration to router B1, the network administrator checked the routing table of router B2, but none of the EIGRP routes appeared there.

To redistribute the EIGRP AS 100 routes into OSPF, which command should be added, or edited, on router B1 under router ospf 10?



```
B2# show ip route
  10.0.0.0/24 is subnetted, 1 subnets
C    10.10.18.0 is directly connected, Ethernet1/0
  172.16.0.0/24 is subnetted, 1 subnets
C    172.16.18.0 is directly connected, FastEthernet0/1
C    192.168.20.0/24 is directly connected, FastEthernet0/0
```

- A. redistribute eigrp 100 metric-type 1
- B. redistribute eigrp 100 subnets
- C. no auto-summary 10.0.0.0 255.0.0.0
- D. area 0 range 10.10.0.0 255.255.0.0

Answer: B

Explanation: Explanation

When redistributing into OSPF without keyword "subnets", only classful networks will be redistributed. Classful networks here mean networks with the default major subnet masks (for example 10.0.0.0/8; 180.1.0.0/16; 200.200.200.0/24...).

In fact, the routing table on the exhibit above is not totally correct. The network 192.168.110.0/24 will be redistributed and shown in the routing table of B2 even if the keyword "subnets" is not used

because it belongs to class C with the default subnet mask of class C.

To make all the networks, including subnets appear in the routing table of B2 we must use keyword "subnets" when redistributing into OSPF. This is also an important thing to remember when redistributing into OSPF.

QUESTION NO: 63

Which of the following are methods EIGRP uses to initially populate (seed) its EIGRP topology table, before learning topology data from neighbors? (Choose two.)

- A. By adding all subnets listed by the show ip route connected command
- B. By adding the subnets of working interfaces over which static neighbors have been defined
- C. By adding subnets redistributed on the local router from another routing source
- D. By adding all subnets listed by the show ip route static command

Answer: B,C

Explanation:

Other than the two listed correct answers, the local router also adds connected routes for which the network command matches the corresponding interfaces, so it may not add all connected routes. Also, EIGRP does not add static routes to the EIGRP topology table, unless those routes are redistributed.

QUESTION NO: 64

An engineer has added the following configuration snippet to an implementation planning document. The configuration will be added to Router R1, whose Fa0/0 interface connects to a LAN to which Routers R2 and R3 also connect. R2 and R3 are already EIGRP neighbors with each other. Assuming the snippet shows all commands on R1 related to EIGRP authentication, which answer lists an appropriate comment to be made during the implementation plan peer review?

key chain fred

key 3

key-string whehew

interface fa0/0

ip authentication key-chain eigrp 9 fred

- A. The configuration is missing one authentication-related configuration command.
- B. The configuration is missing two authentication-related configuration commands.
- C. Authentication type 9 is not supported; type 5 should be used instead.
- D. The key numbers must begin with key 1, so change the key 3 command to key 1.

Answer: A

Explanation:

The configuration requires the ip authentication mode eigrp asn md5 command, which is currently missing. This command enables MD5-style authentication, rather than the default of no authentication. Adding this one command completes the configuration. Any valid key numbers can be used. Also, the 9 in the ip authentication key-chain eigrp 9 fred command refers to the EIGRP ASN, not an authentication type.

QUESTION NO: 65

Which of the following settings could prevent two potential EIGRP neighbors from becoming neighbors? (Choose two answers.)

- A. The interface used by one router to connect to the other router is passive in the EIGRP process.
- B. Duplicate EIGRP router IDs
- C. Mismatched Hold Timers.
- D. IP addresses of 10.1.1.1/24 and 10.2.2.2/24, respectively.

Answer: A,D

Reference: <http://smitley.net/?p=167> (see 'configuration settings that could prevent neighbor relationships')

QUESTION NO: 66

Based on the need to limit processing and bandwidth utilization due to dynamic routing protocol operation, the following routing requirements have been specified for your network.

- Partial and incremental routing updates
- Only the devices affected by a topology change perform route recomputation
- Route recomputation only occurs for routes that were affected

Which dynamic routing protocol should be deployed in your network to best meet these requirements?

- A. BGP
- B. OSPF
- C. IS-IS
- D. EIGRP
- E. RIPv2

Answer: D

Explanation:


The bandwidth utilization issue has been addressed by implementing partial and incremental updates. Therefore, only when a topology change occurs does routing information get sent. Regarding processor utilization, the feasible successor technology greatly reduces the total processor utilization of an AS by requiring only the routers that were affected by a topology change to perform the route recomputation. Furthermore, the route recomputation only occurs for routes that were affected. Only those data structures are accessed and used. This greatly reduces search time in complex data structures.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f07.shtml (See frequently asked questions)

QUESTION NO: 67

Refer to the exhibits. Router B should advertise the network connected to the E0/0/0 interface to router A and block all other network advertisements. The IP routing table on router A indicates that it is not receiving this prefix from router B. What is the probable cause of the problem?



- A. An access list on router B is causing the 192.168.3.16/28 network to be denied.
- B. An access list on router B is causing the 192.168.3.32/28 network to be denied.
- C. The distribute list on router B is referencing a numbered access list that does not exist on router B.
- D. The distribute list on router B is referencing the wrong interface.

Answer: A

Explanation: Explanation


This is an unclear question. The question says "Router B should advertise the network connected to the E0/0/0 interface to router A and block all other network advertisements. The IP routing table on router A indicates that it is not receiving this prefix from router B." That means the network 192.168.3.16/28 (including the IP 192.168.3.21/28) is not received on router A -> A is the most suitable answer.

Note: Distribute list are used to filter routing updates and they are based on access lists.

QUESTION NO: 68

Refer to the exhibit. Router B and router C are performing mutual redistribution between OSPF and EIGRP, and their default metrics are configured the same. Router D has equal cost paths to networks where both paths are not really equal cost. For example, network 172.16.54.0 shows

equal cost through both router B and router C, though in reality the cost is greater using router C. Other routers, though not shown, are connected to the 172.16.54.0 and 172.16.55.0 networks, and the same issues exist to those routers and the networks connected to them. What can be done so that data will be routed along the most optimal path in the network?



- A. Redistribute connected interfaces on router B and router C.
- B. Set the maximum number of equal cost paths to 1 in all routers.
- C. When redistributing EIGRP into OSPF, set the external metric type to type E1.
- D. Adjust the default metrics in router B and router C so that the values are different in each router.
- E. None of these solutions will fix the problem. Migrate to a single dynamic routing protocol.

Answer: E

Explanation: Explanation

From the output, we learn that all the External OSPF routes have metrics of 100 (the second parameters in [110/100]). This is not the default metric of OSPF Type 2 External route (the default value is 20) so the metrics of redistributed routes have been modified. Maybe when redistributing into OSPF, the "metric" in the "redistribute" command or the "default-metric" command was used on router B & C to assign the metric of these routes. Something like this:

```
router ospf 1 redistribute eigrp 1 metric 100 subnets
```

or

```
router ospf 1.....default-metric 100
```

Therefore even if we use the metric type E1 the problem still exists because the link B-D & C-D seems to have the same metric -> the total metrics remains the same .

We can use route-map and set different metrics for each networks but some unshown networks will have the same issues -

QUESTION NO: 69

Based on the exhibited output,

```
R1# show ip eigrp topology
IP-EIGRP Topology Table for process 200
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - Reply status

P 192.168.1.64/28 1 successors, FD is 281600
   via Connected, Ethernet
  P 192.168.1.48/28, 1 successors, FD is 40512000
     via Connected, Serial1
P 192.168.1.48/28, 1 successors, FD is 40537600
   via 192.168.1.66 (40537600/40512000), Ethernet0
   via 192.168.1.17 (41024000/40512000), Serial0
   via 192.168.1.33 (41024000/40512000), Serial1
P 192.168.1.16/28 1 successors, FD is 40512000
   via Connected, Serial0
```

Which three statements are true? (Choose three.)

- A. R1 is in AS 200.
- B. R1 will load balance between three paths to reach the 192.168.1.48/28 prefix because all three paths have the same advertised distance (AD) of 40512000.
- C. The best path for R1 to reach the 192.168.1.48/28 prefix is via 192.168.1.66.
- D. 40512000 is the advertised distance (AD) via 192.168.1.66 to reach the 192.168.1.48/28 prefix.

- E. All the routes are in the passive mode because these routes are in the hold-down state.
- F. All the routes are in the passive mode because R1 is in the query process for those routes.

Answer: A,C,D

Explanation:

It can be determined that AS 200 is used, from the fact that the IS-IS process ID is labeled as 200. The best path to reach the network 192.168.1.48/28 is the first one displayed in the routing table. This can be further demonstrated by the fact that the metric is less than the alternative route, via serial 0. Finally, the AD can be found by viewing the second number within the parentheses, which in this case is 40512000.

QUESTION NO: 70

Study the exhibit carefully.

```
R1#show running-config
<Output omitted>
!
router eigrp 100
 network 172.16.0.0
 distribute-list prefix TEST out
 auto-summary
 no eigrp log-neighbor-changes
!
ip prefix-list TEST seq 5 permit 172.16.1.0/26
!
<Output omitted>
```

Router R1 is connected to networks 172.16.1.0 /26 and 172.16.1.64 /27. Based on the partial output in the exhibit, which description is correct?

- A. Router R1 should be reconfigured with an ACL instead of an ip prefix-list command.
- B. Router R1 will advertise both routes.
- C. Router R1 will deny the 172.16.1.0/27 route while permitting the 172.16.1.0/26 route to be advertised.
- D. Router R1 will deny the 172.16.1.0/26 route while permitting the 172.16.1.64/27 route to be advertised.

Answer: C

Explanation: Explanation

Prefix lists are configured with permit or deny keywords to either permit or deny the prefix based on the matching condition. A prefix list consists of an IP address and a bit mask. The IP address can be a classful network, a subnet, or a single host route. The bit mask is entered as a number from 1 to 32.

Prefix lists are configured to match an exact prefix length or a prefix range. The ge and le keywords are used to specify a range of the prefix lengths to match, providing more flexible configuration than can be configured with just the network/length argument. The prefix list is processed using an exact match when neither ge nor le keyword is entered.

Therefore in this case the exact 172.16.1.0/26 network is permitted while other networks are denied.

(Reference:

http://www.cisco.com/en/US/docs/ios/12_3t/ip_route/command/reference/ip2_i2gt.html)

QUESTION NO: 71

Refer to the exhibit.

```
Core1#show ip eigrp topology all-links
IP EIGRP Topology table for AS(65001) / ID (172.17.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.17.3.128/25, 2 successors, FD is 30720, serno 9
   via 172.17.10.2 (30720/28160), FastEthernet0/1
   via 172.17.3.2 (30720/28160), FastEthernet0/3
P 10.140.0.0/24, 1 successors, FD is 156160, serno 16
   via 172.17.3.2 (156160/128256), FastEthernet0/3
   via 172.17.10.2 (157720/155160), FastEthernet0/1
P 172.17.10.0/24, 1 successors, FD is 28160, serno 1
   via Connected, FastEthernet0/1
P 172.17.0.0/30, 1 successors, FD is 20514560, serno 15
   via 172.17.1.1 (20514560/205122000), FastEthernet0/2
   via 172.17.10.2 (20516120/20513560), FastEthernet0/1
P 172.17.1.0/24, 1 successors, FD is 28160, serno 2
   via Connected, FastEthernet0/2
P 172.17.2.0/24, 1 successors, FD is 30720, serno 8
   via 172.17.10.2 (30720/28160), FastEthernet0/1
   via 172.17.3.2 (33280/30720), FastEthernet0/3
P 172.17.3.0/25, 1 successors, FD is 28160, serno 3
   via Connected, FastEthernet0/3
Core1#
```

BigBids Incorporated is a worldwide auction provider. The network uses EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's question.


Which two networks does the Core1 device have feasible successors for? (Choose two)

- A. 172.17.0.0/30
- B. 172.17.1.0/24
- C. 172.17.2.0/24
- D. 172.17.3.0/25
- E. 172.17.3.128/25
- F. 10.140.0.0/24

Answer: A,F

Explanation: Explanation

To understand the output of the "show ip eigrp topology all-links command" command, let's analyze an entry (we choose the second entry because it is better for demonstration than the first one)



```
P 10.140.0.0/24, 1 successors, FD is 156160, serno 16
Feasible Distance -> (156160/128256), FastEthernet0/3
via 172.17.10.2 (157720/155160) *Advertised Distance
```

The first line tells us there is only **1 successor** for the path to 10.140.0.0/24 network but there are 2 lines below. So we can deduce that one line is used for successor and the other is used for another route to that network. Each of these two lines has 2 parameters: the first one ("156160 or "157720) is the Feasible Distance (FD) and the second ("128256 or "155160) is the Advertised Distance (AD) of that route.

The next thing we want to know is: if the route via 172.17.10.2 (the last line) would become the feasible successor for the 10.140.0.0/24 network. To figure out, we have to compare the Advertised Distance of that route with the Feasible Distance of the successor's route, if $AD < FD$ then it will become the feasible successor. In this case, because $AD (155160) < FD (156160)$ so it will become the feasible successor. Therefore we can conclude the network 10.140.0.0/24 has 1 feasible successor.

Because the question asks about feasible successor so we just need to focus on entries which have more paths than the number of successor. In this case, we find 3 entries that are in blue boxes because they have only 1 successor but has 2 paths, so the last path can be the feasible successor.

By comparing the value of AD (of that route) with the FD (of successor's route) we figure out there are 2 entries will have the feasible successor: the first and the second entry. The third entry has $AD = FD (30720)$ so we eliminate it.

QUESTION NO: 72

Refer to the exhibit.

```
Core1#show ip eigrp topology all-links
IP EIGRP Topology table for AS(65001) / ID (172.17.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.17.3.128/25, 2 successors, FD is 30720, serno 9
  via 172.17.10.2 (30720/28160), FastEthernet0/1
  via 172.17.3.2 (30720/28160), FastEthernet0/3
P 10.140.0.0/24, 1 successors, FD is 156160, serno 16
  via 172.17.3.2 (156160/128256), FastEthernet0/3
  via 172.17.10.2 (157720/155160), FastEthernet0/1
P 172.17.10.0/24, 1 successors, FD is 28160, serno 1
  via Connected, FastEthernet0/1
P 172.17.0.0/30, 1 successors, FD is 20514560, serno 15
  via 172.17.1.1 (20514560/205122000), FastEthernet0/2
  via 172.17.10.2 (20516120/20513560), FastEthernet0/1
P 172.17.1.0/24, 1 successors, FD is 28160, serno 2
  via Connected, FastEthernet0/2
P 172.17.2.0/24, 1 successors, FD is 30720, serno 8
  via 172.17.10.2 (30720/28160), FastEthernet0/1
  via 172.17.3.2 (33280/30720), FastEthernet0/3
P 172.17.3.0/25, 1 successors, FD is 28160, serno 3
  via Connected, FastEthernet0/3
Core1#
```

BigBids Incorporated is a worldwide auction provider. The network uses EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's question.

Which three EIGRP routes will be installed for the 172.17.3.128/25 and 172.17.2.0/24 networks? (Choose three)

- A. 172.17.3.128/25 [90/28160] via 172.17.1.2, 01:26:35, FastEthernet0/2
- B. 172.17.3.128/25 [90/30720] via 172.17.3.2, 01:26:35, FastEthernet0/3
- C. 172.17.3.128/25 [90/30720] via 172.17.10.2, 01:26:35, FastEthernet0/1
- D. 172.17.2.0/24 [90/30720] via 172.17.10.2, 02:10:11, FastEthernet0/1
- E. 172.17.2.0/24 [90/28160] via 172.17.10.2, 02:10:11, FastEthernet0/1
- F. 172.17.2.0/24 [90/33280] via 172.17.3.2, 02:10:11, FastEthernet0/3

Answer: B,C,D

Explanation: Explanation

First indicate the positions of these networks:

Network 172.17.3.128/25 has 2 successors, therefore the two paths below are both successors.

Network 172.17.2.0/24 has only 1 successor, therefore the path lies right under it is the successor.

QUESTION NO: 73

Refer to the exhibit.

Core1#show ip eigrp topology all-links

IP EIGRP Topology table for AS(65001) / ID (172.17.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
r - reply Status, s - sia Status

P 172.17.3.128/25, 2 successors, FD is 30720, serno 9

via 172.17.10.2 (30720/28160), FastEthernet0/1

via 172.17.3.2 (30720/28160), FastEthernet0/3

P 10.140.0.0/24, 1 successors, FD is 156160, serno 16

via 172.17.3.2 (156160/128256), FastEthernet0/3

via 172.17.10.2 (157720/155160), FastEthernet0/1

P 172.17.10.0/24, 1 successors, FD is 28160, serno 1

via Connected, FastEthernet0/1

P 172.17.0.0/30, 1 successors, FD is 20514560, serno 15

via 172.17.1.1 (20514560/205122000), FastEthernet0/2

via 172.17.10.2 (20516120/20513560), FastEthernet0/1

P 172.17.1.0/24, 1 successors, FD is 28160, serno 2

via Connected, FastEthernet0/2

P 172.17.2.0/24, 1 successors, FD is 30720, serno 8

via 172.17.10.2 (30720/28160), FastEthernet0/1

via 172.17.3.2 (33280/30720), FastEthernet0/3

P 172.17.3.0/25, 1 successors, FD is 28160, serno 3

via Connected, FastEthernet0/3

Core1#

BigBids Incorporated is a worldwide auction provider. The network uses EIGRP as its routing protocol throughout the corporation. The network administrator does not understand the convergence of EIGRP. Using the output of the show ip eigrp topology all-links command, answer the administrator's question.

Which three networks is the router at 172.17.10.2 directly connected to? (Choose three)

- A. 172.17.0.0/30
- B. 172.17.1.0/24
- C. 172.17.2.0/24
- D. 172.17.3.0/25
- E. 172.17.3.128/25
- F. 172.17.10.0/24

Answer: C,E,F

Explanation:

```
Core1#show ip eigrp topology all-links
IP EIGRP Topology table for AS(65001) / ID (172.17.10.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 172.17.3.128/25, 2 successors, FD is 30720, serno 9
   via 172.17.10.2 (30720/28160), FastEthernet0/1
   via 172.17.3.2 (30720/28160), FastEthernet0/3
P 10.140.0.0/24, 1 successors, FD is 156160, serno 16
   via 172.17.3.2 (156160/128256), FastEthernet0/3
   via 172.17.10.2 (157720/155160), FastEthernet0/1
P 172.17.10.0/24, 1 successors, FD is 28160, serno 1
   via Connected, FastEthernet0/1
P 172.17.0.0/30, 1 successors, FD is 20514560, serno 15
   via 172.17.1.1 (20514560/205122000), FastEthernet0/2
   via 172.17.10.2 (20516120/20513560), FastEthernet0/1
P 172.17.1.0/24, 1 successors, FD is 28160, serno 2
   via Connected, FastEthernet0/2
P 172.17.2.0/24, 1 successors, FD is 30720, serno 8
   via 172.17.10.2 (30720/28160), FastEthernet0/1
   via 172.17.3.2 (33280/30720), FastEthernet0/3
P 172.17.3.0/25, 1 successors, FD is 28160, serno 3
   via Connected, FastEthernet0/3
Core1#
```

First, we should notice about the entry in the orange box, it shows that the network 172.17.10.0/24 is directly connected with this router and has a FD of 28160. So we can guess the networks that directly connected with router at 172.17.10.2 will be shown with an AD of 28160. From that, we find out 3 networks which are directly connected to the router at 172.17.10.2 (they are green underlined). The network 172.17.10.0/24 is surely directly connected to the router at 172.17.10.2 (in fact it is the network that links the router at 172.17.10.2 with Core1 router).

QUESTION NO: 74

Which two statements are true about EIGRP manual summarization? (Choose two.)

- A.** Manual summarization is configured on a per interface basis.
- B.** Manual summaries can be configured with the classful mask only.
- C.** When manual summarization is configured, autosummarization is automatically disabled by default.
- D.** The summary address is assigned an administrative distance of 10 by default.
- E.** The summary address is entered into the routing table and is shown to be sourced from the Null0 interface.

Answer: A,E

Explanation:

You can configure a summary addresses on a per-interface basis. You need to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary or if you want to use summary addresses on a ASA with automatic route summarization disabled. If any more specific routes are in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

QUESTION NO: 75

Refer to the exhibit.



Implementation Plan

1. Establish a traffic throughput baseline.
2. Configure variance on R1 and R4.
3. Use traceroute to validate load balancing has been activated.
4. Establish a new traffic throughput baseline.
5. Compare the new and old baselines and verify that load balancing is implemented as desired.

ROUTE.com is planning to implement EIGRP load balancing for traffic between hosts on the 172.16.10.0/24 and 172.16.20.0/24 networks. You have been asked to review the implementation plan for this project. Which statement about the plan is true?

- A. It is complete as written.
- B. It should include a task to configure multipath to equal a value of 2 on R1 and R4.
- C. It should use a ping instead of a traceroute to validate that load balancing has been activated.
- D. It should contain a task that documents the changes made to the configurations.

Answer: A

Explanation:

This implementation plan is complete because it has all the requirements for an EIGRP load balancing process.

QUESTION NO: 76 DRAG DROP

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

Answer:

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

redistribution	bandwidth management
bandwidth management	authentication
authentication	redistribution
stubs	stubs

Explanation:

bandwidth management
authentication
redistribution
stubs

Redistribution - integrating two merging companies

Bandwidth management - low speed WAN links

Authentication - WAN link to external supplier

stubs - 256 kbps/CIR FR hub and spokes.

QUESTION NO: 77

A network administrator is managing a hub-and-spoke network with EIGRP routing that has been enabled. The hub router is trying to query a remote router. However, delays are occurring that are caused by certain paths being stuck in active (SIA). How should the administrator configure EIGRP in order to limit the scope of the query range and prevent SIA from occurring?

- A. Configure the hub router with a scope limit of 1.
- B. Configure the remote router with a scope limit of 1.
- C. Configure the hub to indicate that the remote router is a stub router.
- D. Configure the hub and remote router as stub routers.
- E. Configure the remote router as a stub router.
- F. Disable the SIA feature of EIGRP on the remote router.

Answer: E

Explanation:

Configuring a router as a stub also helps the rest of the network. Queries are responded to much quicker and convergence happens much faster. Sometimes a query can cause delays that result in the path being SIA. If the stub configuration is applied, the router responds to queries as inaccessible, thus limiting the scope of the query range and preventing SIA from occurring.

QUESTION NO: 78

What administrative distance is given to EIGRP summary routes?

- A. 0
- B. 1
- C. 5
- D. 90
- E. 95
- F. 170

Answer: C

Reference: http://www.cisco.com/en/US/docs/ios/iproute_eigrp/command/reference/ire_i1.html
(See usage guidelines)

QUESTION NO: 79

What are two possible causes for EIGRP Stuck-In-Active routers? (Choose Two)

- A. Some query or reply packets are lost between the routers.
- B. The neighboring router starts receiving route updates from this router.
- C. A failure causes traffic on a link between two neighboring routers to flow in only one direction (unidirectional link).
- D. The neighboring router stops receiving ACK packets from this router.

Answer: A,C**Explanation:**

Generally, a route shown as Active is going to be there for a very short period of time by the time you repeat the command, hopefully that Active route has gone Passive. Sometimes that doesn't happen, though, and the route becomes SIA - Stuck In Active.

A route becomes SIA when a query goes unanswered for so long that the neighbor relationship is reset. From experience, I can tell you that troubleshooting SIA routes is more of an art form than a

science, but there are four main reasons a route becomes SIA:

The link is unidirectional, so the query can't possibly be answered. The queried router's resources are unavailable, generally due to high CPU utilization. The queried router's memory is corrupt or otherwise unable to allow the router to answer the query.

The link between the two routers is of low quality, allowing just enough packets through to keep the neighbor relationship intact, but not good enough to allow the replies through.

To sum it up, routes generally become SIA when a neighbor either doesn't answer a query, or either the query or reply took a wrong turn somewhere. I told you it wasn't the easiest thing to troubleshoot!

QUESTION NO: 80

When configuring EIGRP to run across a 56 Kbps serial PPP link, what command do you need to put under the serial interface ensure proper convergence of EIGRP routes?

- A. bandwidth 56
- B. bandwidth 56000
- C. ip bandwidth-percent eigrp 1 56
- D. ip bandwidth-percent eigrp 1 56000

Answer: A


Explanation:

When configuring serial links using EIGRP it is important to configure the bandwidth setting on the interface. If the bandwidth setting is not changed for these interfaces EIGRP assumes the default bandwidth on the link instead of the true bandwidth. If the link is slower, the router may not be able to converge, routing updates might become lost, or suboptimal path selection may result.

Router(config-if)#bandwidth kilobits The value, kilobits, indicates the intended bandwidth in kilobits per second. For generic serial interfaces, such as PPP or HDLC, set the bandwidth to the line speed.

QUESTION NO: 81

Refer to the Exhibit.



Routers in the Diagram are configured with EIGRP. If RTB and RTC fail, which action will RTA take with respect to the HQ network?


- A. RTA will automatically route packets via RTD to the HQ network.
- B. RTA will place the route via RTD into the hold down state.
- C. RTA will go into the active state for all routers.
- D. RTA will go into the active state for the route to HQ network.

Answer: D

Explanation:

When RTB and RTC fails, RTA will go into active state for the HQ network route.

QUESTION NO: 82 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.

Answer: Here are the step by Step Solution for this:

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

```
R1#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 NO: 83

Exhibit:



Refer to the topology diagram R3 is redistributing the EIGRP routes into OSPF. What will the EIGRP routes appear in the routing table of R1?

- A. O
- B. O IA
- C. E2
- D. D
- E. D EX

Answer: C

Reference: http://kowon.dongseo.ac.kr/~nok60/lecture/CCNP1-V50/Labs/CCNP1v50_L5-2.pdf
(page 10 and 11)

QUESTION NO: 84

In EIGRP, when the IP default-network command is configured on a router, what is generated in the router's configuration?

- A. A static route
- B. A directly connected route
- C. An EIGRP route
- D. A default route


Answer: A

Explanation:

When you configure the `ip default-network` command and specify a subnet, a static route (the `ip route` command) is generated in the router's configuration; however, the IOS does not display a message to indicate that this has been done. The entry appears as a static route in the routing table of the router where the command is configured. This can be confusing when you want to remove the default network; the configuration must be removed with the `no ip route` command, not with the `no ip default-network` command.

QUESTION NO: 85

Refer to the exhibit.



Which router configuration command can be given that will restrict router RTB from sharing its routing information with router RTA?

- A. the `eigrp stub` command on router RTA
- B. the `eigrp stub` command on router RTB
- C. the `eigrp stub connected` command on router RTA
- D. the `eigrp stub connected` command on router RTB
- E. the `eigrp stub receive-only` command on router RTA
- F. the `eigrp stub receive-only` command on router RTB

Answer: F

Explanation:

This is a hub and spoke network, so EIGRP stub receive-only command on RTB will restrict the router from sharing its routing information with RTA.

QUESTION NO: 86

The following command was issued on R2

```
R2#sho ip rou
<output omitted>
C    10.1.1.0 is directly connected, FastEthernet0/0
D    172.16.0.0/16 [90/156160] via 10.1.1.1, 00:07:48, FastEthernet0/0
D EX 192.168.1.0/24 [170/1308160] via 10.1.1.1, 00:00:11, FastEthernet0/0
```

Given the above output, which statement is true?

- A. 192.168.1.0 is a static route.
- B. 192.168.1.0 is a summarized route.
- C. 192.168.1.0 is a redistributed route into EIGRP.
- D. 192.168.1.0 is equal path load balancing with 172.16.1.0.

Answer: C

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009487e.shtml
(administrative distance, second para)

QUESTION NO: 87

Which two types of routes will be advertised with the EIGRP configuration as shown?

(Choose two.)

```
router eigrp 100
```

```
network 10.0.0.0
```

eigrp stub

- A. static
- B. receive-only
- C. summary
- D. stub
- E. connected
- F. dynamic

Answer: C,E

Explanation:

The Enhanced Interior Gateway Routing Protocol (EIGRP) Stub Routing feature improves network stability, reduces resource utilization, and simplifies stub router configuration.

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

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

Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router will depend on the distribution router to send the proper updates to all peers. Router(config-router)#eigrp stub [receive-only | connected] static | summary] :Configures a remote router as an EIGRP stub router.

QUESTION NO: 88

Refer to the exhibit.

```
Router# show ip route
```

```
...
```

```
C    10.1.3.0/24 is directly connected, Serial2
```

```
D    10.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1
```

```
D    10.0.0.0/8 is a summary, 00:23:20, Null0
```

```
C    10.1.1.0/24 is directly connected, Serial1
```

```
S    192.168.20.0/24 is directly connected, Ethernet0
```

What happens when the router stops receiving advertisements for the 10.1.2.0/24 network?

- A. The summary route will be removed from the table.
- B. The summary route will remain in the table.
- C. The more specific routes will be advertised from the table.
- D. 10.1.2.0/24 will still be advertised but packets destined for it will be dropped when they reach this router.

Answer: B

Explanation:


If you look very closely at the routing table output, we can conclude that R1 has "auto-summary" enabled under the EIGRP routing process.

D 10.0.0.0/8 is a summary, 00:23:20, Null 0

Anytime the "auto-summary" is enabled under the routing process the router will install a summary route to "null 0" as long as the router has one or more subnets within the "classful" network. In the case above, you have two directly connected interfaces (Serial1 & Serial2) that are within the "classful" network. Therefore, regardless of whether you learn a route via EIGRP that is in the "classful" network, R1 will still install this summary route to "null 0".

QUESTION NO: 89

Refer to the exhibit.



On all routers in the network, EIGRP has been configured for load balancing across the three links. However, traffic destined for Network B from R1 is only load balanced over paths R1-R2-R5 and R1-R3-R5. What is the cause of the problem?

- A. EIGRP will not select more than two links for unequal cost path load balancing.
- B. Because the path has a different link type, EIGRP will not select path R1-R4-R5 for load balancing.
- C. Because Router R4 is not a feasible successor, EIGRP will not select path R1-R4-R5 for load balancing.
- D. EIGRP will not select path R1-R4-R5 for load balancing unless the value of the variance parameter is increased.

Answer: C

Explanation:

Since R4 is not configured as a feasible successor, EIGRP will not select that path for load balancing. IN EIGRP, you need to configure feasible successor to enable load balancing on the path.

QUESTION NO: 90

Identify three characteristics of EIGRP feasible successors? (Choose three.)

- A. A feasible successor is selected by comparing the advertised distance of a non-successor route to the feasible distance of the best route.
- B. If the advertised distance of the non-successor route is less than the feasible distance of best route, then that route is identified as a feasible successor.
- C. If the successor becomes unavailable, then the feasible successor can be used immediately without recalculating for a lost route.
- D. The feasible successor can be found in the routing table.

E. Traffic will be load balanced between feasible successors with the same advertised distance.

Answer: A,B,C

Reference: <http://packetlife.net/blog/2010/aug/9/eigrp-feasible-successor-routes/>

QUESTION NO: 91

Which are three features of EIGRP? (Choose three)


- A. Support VLSM and discontiguous subnets
- B. Link-state protocol
- C. Partial routing updates
- D. External Administrative distance is 100
- E. Fast convergence.
- F. Used by other vendors than Cisco

Answer: A,C,E

Reference: http://cisco.jjc.edu/cnt205/ch2/2_1_1/index.html (See first three bullets)

QUESTION NO: 92

Refer to the exhibit.



Implementation Plan

1. Establish a traffic throughput baseline
2. Configure variance on R1 and R4
3. Use traceroute to validate load balancing has been activate
4. Establish a new traffic throghput baseline
5. Compare the new and old baselines and verify that load balancing is implemented as desired.

ROUTE.com is planning to implement load balancing for traffic between hosts on the 172.16.10.0/24 and 172.16.20.0/24 networks. You have been asked to review the implementation plan for this project. Which statement about the plan is true?


- A. It is complete as written.
- B. It should include a task to configure EIGRP multipath equal to 2 on R1 and R4.
- C. It should include a task to implement OSPF because it handles unequal cost load balancing most efficiently using variance.
- D. It should include a task that establishes a baseline before and after the configuration has been changed.

Answer: A

Explanation:

QUESTION NO: 93

Refer to exhibit.



RouterA (DR) failed, and after 10 minutes it came back. Which two statements are true? (Choose two)

- A. RouterA is a DR
- B. RouterA is a BDR
- C. RouterA is a DROTHER
- D. RouterB is a DR
- E. RouterB is a BDR
- F. RouterC is a DROTHER

Answer: C,D

Reference: <http://packetlife.net/blog/2011/jun/2/ospf-designated-router-election/>

QUESTION NO: 94

Which of the below mentioned conditions form a neighbor relation in EIGRP? (Choose Three)

- A. Hello or ACK received
- B. AS number match
- C. Hello timer match
- D. Identical metric(k values)
- E. Dead Timer Match
- F. Network Time Match


Answer: A,B,D

Explanation:

EIGRP routers create a neighbor relationship by exchanging hello packets. Certain fields in the hello packets must match before an EIGRP neighbor relationship is established:

QUESTION NO: 95

Refer to exhibit.



The exhibit shows the R1 topology table to reach 192.168.1.0/24 network.

Which route(s) will be installed in routing table of R1 to reach network 192.168.1.0/24 after configuring R1 with the following command?

Router(config-router)# variance 2

- A. R2 only
- B. R2 and R3

- C. R2 and R4
- D. R2, R3 and R4

Answer: A

Explanation:

Every routing protocol supports equal cost path load balancing. In addition, Interior Gateway Routing Protocol (IGRP) and EIGRP also support unequal cost path load balancing. Use the **variance *n*** command in order to instruct the router to include routes with a metric of less than ***n*** times the minimum metric route for that destination. The variable ***n*** can take a value between 1 and 128. The default is 1, which means equal cost load balancing. Traffic is also distributed among the links with unequal costs, proportionately, with respect to the metric.

QUESTION NO: 96

Refer to the following.

Router # sh ip route eigrp

13.0.0.0/8 is variably subnetted, 2 subnets, 2 masks

D 13.0.0.0/8 is a summary, 00:00:32, Null0

What happens to packets that are forwarded from the 13.0.0.0/8 network to the Null0 interface?

- A. Flagged
- B. Accepted
- C. Summarized
- D. Dropped

Answer: D

Explanation:

When an EIGRP router summarizes, it automatically builds a route to null0 for the summarized route. The route to null0 prevents packets that do not match a specific entry in the routing table from following a default route. (The route to null0 causes the packet to be dropped).

QUESTION NO: 97

In which state do DR and BDR establish adjacency with each OSPF router in the network?

- A. Init State
- B. Exstart State
- C. Exchange State
- D. Loading State

Answer: B

Explanation:

QUESTION NO: 98

A stub area is typically created using what kind of topology?

- A. Broadcast
- B. Point-to-point
- C. Hub and spoke
- D. Full Mesh

Answer: C

Explanation:

A stub area is typically created when you have a hub-and-spoke topology, with the spoke being the stub area, such as a branch office. In this case, the branch office does not need to know about every network at the headquarters site; instead, it can use a default route to get there.

QUESTION NO: 99

A network administrator is troubleshooting an EIGRP connection between RouterA, IP address 10.1.2.1, and RouterB, IP address 10.1.2.2.

```
RouterA# debug eigrp packets
...
01:39:13: EIGRP: Received HELLO on Serial0/0 nbr 10.1.2.2
01:39:13: AS 100, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0 peerQ un/rely 0/0
01:39:13:      K-value mismatch
```

Given the debug output on RouterA, which two statements are true?

- A. RouterA received a hello packet with mismatched autonomous system numbers.

- B. RouterA received a hello packet with mismatched hello timers.
- C. RouterA received a hello packet with mismatched authentication parameters.
- D. RouterA received a hello packet with mismatched metric-calculation mechanisms.
- E. RouterA will form an adjacency with RouterB.
- F. RouterA will not form an adjacency with RouterB.


Answer: D,F

Explanation:

If the k-value mismatch occurs, Router A will never form an adjacency with Router B since it is one of the basic requirements of adjacency. If you see the exhibit, Router A received HELLO packet with a mismatched metric.

QUESTION NO: 100

Refer to the exhibit.



Network administrators have set up a hub and spoke topology with redundant connections using EIGRP. However, they are concerned that a network outage between Router R1 and Router R2 will cause traffic from the 10.1.1.x network to the 10.1.2.x network to traverse the remote office links and overwhelm them. What command should be used to configure the spoke routers as EIGRP stub routers that will not advertise connected networks, static routes, or summary addresses?

- A. eigrp stub
- B. eigrp stub receive-only
- C. eigrp stub connected static
- D. no eigrp stub connected static
- E. No additional command is needed beyond a default EIGRP configuration.

Answer: B

Explanation:

A router that is configured as a stub with the **eigrp stub** command shares connected and summary routing information with all neighbor routers by default. Four optional keywords can be used with the **eigrp stub** command to modify this behavior:

- receive-only
- connected
- static
- summary

This section provides configuration examples for all forms of the **eigrp stub** command. The **eigrp stub** command can be modified with several options, and these options can be used in any combination except for the **receive-only** keyword. The **receive-only** keyword will restrict the router from sharing any of its routes with any other router in that EIGRP autonomous system, and the **receive-only** keyword will not permit any other option to be specified because it prevents any type of route from being sent. The three other optional keywords (**connected**, **static**, and **summary**) can be used in any combination but cannot be used with the **receive-only** keyword. If any of these three keywords is used individually with the **eigrp stub** command, connected and summary routes will not be sent automatically.

The **connected** keyword will permit the EIGRP Stub Routing feature to send connected routes. If the connected routes are not covered by a network statement, it may be necessary to redistribute connected routes with the **redistribute connected** command under the EIGRP process. This option is enabled by default.

The **static** keyword will permit the EIGRP Stub Routing feature to send static routes. Without the configuration of this option, EIGRP will not send any static routes, including internal static routes that normally would be automatically redistributed. It will still be necessary to redistribute static routes with the **redistribute static** command.

The **summary** keyword will permit the EIGRP Stub Routing feature to send summary routes. Summary routes can be created manually with the **summary address** command or automatically at a major network border router with the **auto-summary command** enabled. This option is enabled by default.

In the following example, the **eigrp stub** command is used to configure the router as a stub that advertises connected and summary routes.

```
router eigrp 1
```

eigrp stub

In the following example, the **eigrp stub** command is issued with the **connected** and **static** keywords to configure the router as a stub that advertises connected and static routes (sending summary routes will not be permitted):

```
router eigrp 1
network 10.0.0.0
eigrp stub connected static
```

In the following example, the **eigrp stub** command is issued with the **receive-only** keyword to configure the router as a receive-only neighbor (Connected, summary, and static routes will not be sent):

```
router eigrp 1
network 10.0.0.0 eigrp
eigrp stub receive-only
```

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a0080087026.html

QUESTION NO: 101

Which configuration command is used to enable EIGRP unequal-cost path load balancing?

- A. maximum-paths
- B. distance
- C. metric
- D. variance
- E. default-metric

Answer: D


Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009437d.shtml (See traffic sharing

<http://networkninja.co.za/page/2/?s=botha&cat=plus-5-results>

QUESTION NO: 102

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.

Answer: E

Explanation: In this example, access-list 7 is used to specify which networks will be advertised out the interface on the distribute list. This ACL only allows for the 172.16.0.0/16 to be advertised out, so 10.0.0.0/8 will indeed be filtered and will not be advertised to any upstream routers.

QUESTION NO: 103

If the primary path goes down, what will EIGRP use to reach a destination?

- A. administrative distance
- B. advertised successor
- C. successor
- D. feasible successor

Answer: D

Explanation: The key to this question is the four terminology about DUAL.

Enhanced Interior Gateway Routing Protocol (EIGRP) is an enhanced distance-vector protocol based on the diffusing update algorithm (DUAL). It is capable of (conservatively) finding all loop-free paths to any given destination based on route advertisements from neighbors. The neighbor (or neighbors) with the best path to a destination is called the successor. The remaining neighbors with loop-free paths to the destination are called feasible successors. To reduce traffic load on the network, EIGRP maintains neighbor relationships and exchanges routing information only as needed, using a query process to find alternate paths when all loop-free paths to a destination have failed.

QUESTION NO: 104

A network administrator would like to configure an EIGRP router as a stub router that advertises directly connected and summary routes only. What command must the administrator issue to accomplish this?

- A. `eigrp stub`
- B. `eigrp stub connected`
- C. `eigrp stub summary`
- D. `eigrp stub connected static`
- E. `eigrp stub receive-only`

Answer: A

Reference: http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/eigrpstb.html#wp1036215

QUESTION NO: 105

Which three statements are true regarding EIGRP? (Choose three.)

- A. By default, EIGRP performs auto-summarization across classful network boundaries.
- B. EIGRP uses an area hierarchy to increase network scalability.
- C. To speed convergence, EIGRP attempts to maintain a successor and feasible successor path for each destination.
- D. EIGRP uses hellos to establish neighbor relationships.
- E. By default, EIGRP uses the Dijkstra algorithm to determine the best path to a destination network based on bandwidth and delay.


Answer: A,C,D

Reference:

<http://ptgmedia.pearsoncmg.com/images/1587131706/samplechapter/1587131706content.pdf>
(See page 66 which is the 2nd page of this pdf)

QUESTION NO: 106

Refer to the graphic.



What optional EIGRP configurations will be required in order to achieve full connectivity within AS 100?

- A. Use the EIGRP no auto-summary command on R1 and R2.
- B. Use the EIGRP no auto-summary command on R3 and R4.
- C. Use the passive interface on the R1 and R2 interface that connects to the 10.1.1.0/24 and 10.1.2.0/24 subnet respectively.
- D. Use the passive interface command between the R3 and R1 connection and between the R3 and R2 connection.
- E. Use the variance command on R3.

Answer: A

Explanation: EIGRP automatically summarizes routes at the classful boundary, the boundary where the network address ends as defined by class-based addressing. In most cases, auto summarization is a good thing, keeping the routing tables as compact as possible. In the presence of discontinuous subnetworks, automatic summarization must be disabled for routing to work properly. To turn off auto-summarization, use the following command:
Router(config-router)#**no auto-summary**

QUESTION NO: 107

Which two statements are EIGRP characteristics? (Choose two.)

- A. Updates are sent as multicast.
- B. Updates are sent as broadcast.
- C. Metric values are represented in a 32-bit format for granularity.
- D. LSAs are sent to adjacent neighbors.

Answer: A,C

Explanation:

EIGRP uses the same algorithm for metric calculation as IGRP, but represents values in 32-bit format to give additional granularity. EIGRP supports unequal metric load balancing, which allows administrators to better distribute traffic flow in their networks. Multicast and unicast: EIGRP uses multicast and unicast, rather than broadcast. The multicast address used for EIGRP is 224.0.0.10.

QUESTION NO: 108

What is the purpose of the `eigrp stub` configuration command?

- A. to increase scalability by limiting the EIGRP query range
- B. to reduce the size of the routing table by blocking the D EX (External EIGRP) routes into the EIGRP stub router
- C. to reduce the convergence time by enabling the EIGRP stub router to propagate the EIGRP queries from the EIGRP hub router
- D. to reduce the convergence time by enabling the EIGRP stub router to also perform query requests to the EIGRP hub router

Answer: A

Explanation:

Complex, redundant EIGRP networks can cause scalability problems. The best solution to this is to provide a means within the context of the EIGRP protocol itself to control traffic flows and limit query depth. The EIGRP Stub Router functionality in Cisco IOS Software Release 12.0(7)T can achieve this solution.

The Enhanced Interior Gateway Routing Protocol (EIGRP) Stub Routing feature improves network stability, reduces resource utilization, and simplifies stub router configuration.

Stub routing is commonly used in a hub and spoke network topology. In a hub and spoke network, one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is adjacent only to one or more

distribution routers. The only route for IP traffic to follow into the remote router is through a distribution router. This type of configuration is commonly used in WAN topologies where the distribution router is directly connected to a WAN. The distribution router can be connected to many more remote routers. Often, the distribution router will be connected to 100 or more remote routers. In a hub and spoke topology, the remote router must forward all non local traffic to a distribution router, so it becomes unnecessary for the remote router to hold a complete routing table. Generally, the distribution router need not send anything more than a default route to the remote router.

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

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

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a0080087026.html

QUESTION NO: 109

Which is the most effective technique to contain EIGRP queries?

- A. route summarization
- B. configuring route filters
- C. using a hierarchical addressing scheme
- D. establishing separate autonomous systems

Answer: A

Explanation:

The most basic SIA routes occur when it simply takes too long for a query to reach the other end of the network and for a reply to travel back. One of the most effective techniques for containing EIGRP queries is to use route summarization or stub networks.

Stub networks are configured in a hub and spoke network and is configured only on the stub


router. The command is

```
router eigrp 100
eigrp stub
```

The result of this configuration is that the stub routers will send updates about routes they have to the hub router but the hub router will never query the stub router for updates in the event of a route being lost.

QUESTION NO: 110

Refer to the exhibit.



```
R1#show running-config
<output omitted>
key chain troy
key 1
  key-string 0987654321
key 2
  key-string 1234576890
<output omitted>
interface FastEthernet0/0
  ip address 10.10.10.1 255.255.255.0
  ip authentication mode eigrp 1 md5
  ip authentication key-chain eigrp 1 troy
<output omitted>
router eigrp 1
  network 10.0.0.0
  no auto-summary
```

```
R2#show running-config
<output omitted>
key chain albany
  key 1
    key-string 0987654321
  key 2
    key-string 1234576890
<output omitted>
interface GigabitEthernet0/0
  ip address 10.10.10.2 255.255.255.0
  ip authentication mode eigrp 1 md5
  ip authentication key-chain eigrp 1 albany
!
interface GigabitEthernet0/1
  ip address 172.16.1.1 255.255.255.0
  ip authentication mode eigrp 1 md5
  ip authentication key-chain eigrp 1 albany
<output omitted>
router eigrp 1
  network 10.0.0.0
  network 172.16.0.0
  no auto-summary
```

```
R3#show running-config
<output omitted>
key chain schenectady
key 1
  key-string 0987654321
key 2
  key-string 1234576890
<output omitted>
interface GigabitEthernet0/1
  ip address 172.16.1.2 255.255.255.0
  ip authentication mode eigrp 1 md5
  ip authentication key-chain eigrp 1 schenectady
<output omitted>
router eigrp 1
  network 172.16.0.0
  no auto-summary
```

EIGRP is configured on all routers in the network. On the basis of the output provided, which statement is true?

- A. Because the key chain names do not match, router R1 will not be able to ping routers R2 and R3 .
- B. Because the key strings do not match, router R1 will not be able to ping routers R2 and R3.
- C. Because authentication is misconfigured on interfaces Gi0/0 and Gi0/1 on router R2, router R1 will not be able to ping routers R2 and R3.
- D. Because autosummarization needs to be turned on for EIGRP on all routers, router R1 will not be able to ping routers R2 and R3.
- E. Router R1 will be able to ping routers R2 and R3.

Answer: E

Explanation: Here we see that all of the routers have correctly included the proper networks in the EIGRP process, and the authentication is also correct so all networks would be reachable from R1. Even though the name of the authentication keys are different, the actual keys are identical so authentication will work.

QUESTION NO: 111 DRAG DROP

Click and drag the command on the left to the associated task on the right.

show ip eigrp neighbor	confirm what EIGRP is learning
show ip eigrp interface	confirm what is actually being used
show ip eigrp topology	view route information sources
show ip route eigrp	verify the routing of specific networks

Answer:

Click and drag the command on the left to the associated task on the right.

show ip eigrp neighbor	show ip eigrp topology
show ip eigrp interface	show ip route eigrp
show ip eigrp topology	show ip eigrp neighbor
show ip route eigrp	show ip eigrp interface

Explanation:

show ip eigrp topology
show ip route eigrp
show ip eigrp neighbor
show ip eigrp interface

QUESTION NO: 112

An engineer is trying to summarize the following networks using the "ip summary-address eigrp" command:

10.8.88.0/25

10.8.89.48/29

10.8.64.96/27

Which network and subnet mask below would be the smallest EIGRP summary address to include all three subnets?

- A. 10.8.64.0 255.255.224.0
- B. 10.8.64.0 255.255.128.0
- C. 10.8.64.0 255.255.192.0
- D. 10.8.0.0 255.255.192.0

Answer: A

Explanation: 10.8.64.0 255.255.224.0 will include networks 10.8.64.0 - 10.8.96.255, so this would be the smallest summary route that would include the 3 routes listed in the question.

QUESTION NO: 113

EIGRP performs automatic summarization at network boundaries. What administrative distance is given to these EIGRP summary routes?

- A. 1
- B. 90
- C. 95
- D. 0
- E. 5
- F. 170
- G. 255
- H. None of the other alternatives apply

Answer: E

Explanation:


By default, EIGRP summary routes are given an administrative distance value of 5.

Topic 2, Implement a multi-area OSPF Network, given a network design and a set of requirements

QUESTION NO: 114 CORRECT TEXT

(OSPF Sim)

OSPF is configured on routers Amani and Lynaic. Amani's S0/0 interface and Lynaic's S0/1 interface are in Area 0. Lynaic's Loopback0 interface is in Area 2.



Your task is to configure the following:

Portland's S0/0 interface in Area 1

Amani's S0/1 interface in Area 1

Use the appropriate mask such that ONLY Portland's S0/0 and Amani's S0/1 could be in Area 1. Area 1 should not receive any external or inter-area routes (except the default route).

Answer: First, we configure Portland's S0/0 interface so that it belongs to Area 1. So, we have to find out which subnetwork the IP address 192.168.4.5/30 (the IP of interface S0/0 of Portland) belongs to. This address belongs to a subnetwork which has:

Increment: 4 (/30 = 255.255.255.252 or 1111 1111.1111 1111.1111 1111.1111 1100)

Network address: 192.168.4.4 (because $4 = 4 * 1$ and $4 < 5$)

Broadcast address: 192.168.4.7 (because $7 = 4 + 4 - 1$) (It is not necessary to find out the broadcast address but we should know it)

The question requires that only Portland's S0/0 and Amani's S0/1 could be in Area 1, therefore we must use a wildcard of 0.0.0.3 (this wildcard is equivalent with a subnet mask of /30) so that there are only 2 IP addresses can participate in area 1 (they are 192.168.4.5 & 192.168.4.6). The full command we use here is `network 192.168.4.4 0.0.0.3 area 1`. The question also requires that "Area 1 should not receive any external or inter-area routes (except the default route)". Recall that if we don't want the router to receive external routes, we have to stop LSA Type 5. And if we don't want to receive inter-area routes, we have to stop LSA Type 3 and Type 4. Therefore we have to configure area 1 as a totally stubby area. For your information, here is the definition of a totally stubby area. "Totally stubby area - This area does not accept summary LSAs from other areas

(types 3 or 4) or external summary LSAs (Type 5). Types 3,4 and 5 LSAs are replaced by the Area Border Router(ABR) with a default router. Totally stubby areas protect internal routers by minimizing the routing table and summarizing everything outside the area with a default route." (CCNP BSCI Official Exam Certification Guide, Fourth Edition) In conclusion, we have to configure area 1 as a totally stubby area. We do that by configuring Portland as stub and configuring Amani (ABR router) as a stub + "no-summary"suffix. + Configure Portland router as a stub:

```
Portland#configure terminal
```

```
Portland(config)#router ospf 1
```

Allow network 192.168.4.4/30 to join Area 1, notice that you have to convert subnet mask into wildcard mask:

```
Portland(config-router)#network 192.168.4.4 0.0.0.3 area 1
```

Configure Portland as a stub:

```
Portland(config-router)#area 1 stub
```

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

```
Portland#copy running-config startup-config
```

+ Configure Amani router as a "totally stub":

```
Amani#configure terminal
```

```
Amani(config)#router ospf 1
```

```
Amani(config-router)#network 192.168.4.4 0.0.0.3 area 1
```

Make area 1 become a totally stubby area, notice that we can only use this command on ABR router:


```
Amani(config-router)#area 1 stub no-summary
```

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

```
Amani#copy running-config startup-config
```

QUESTION NO: 115 CORRECT TEXT

Acme is a small export company 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 OSPF 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 ping from R1 to R4 loopback address.



Answer: 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:

Explanation:

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

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

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

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

```
R2>enable
```

```
R2#configure terminal
```

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

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

Save the configuration:

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

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

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

```
R3>enable
```

```
R3#configure terminal
```

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

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

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

Save the configuration:

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

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

You should check the configuration of R4, too.

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

QUESTION NO: 116 CORRECT TEXT


Scenario

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 OSPF 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.

Topology



Answer: On router R2:

Config t

Router ospf 1

Redistribute eigrp 100 metric 100 subnets

distance eigrp 90 105


Router eigrp 100

Redistribute ospf 1 metric 100000 10 200 10 1500

Then from R1 you should be able to ping 172.16.1.100

Explanation: Since R2 router has the potential to be the shortest path to OSPF network from loopback 101. So configure the router R2 to have redistributable eigrp.

QUESTION NO: 117



The company and the company network have both been growing rapidly. Multiple adds, moves and changes have been applied to the network. Your boss has asked you to troubleshoot a recent OSPF synchronization problem that has arisen. There have been synchronization problems at separate locations in the OSPF area 0. There have been reported link failures during the rapid growth of the company network. You are required to resolve the OSPF problem. OSPF must be able to converge when the network changes.

Refer to the information above to answer the following question

Examine the following excerpt from the "show ip ospf" command on D1:

Area BACKBONE (0)

Number of interfaces in the this area is 1

Area has no authentication

SPF algorithm last executed 00:00:31.280 ago

SPF algorithm executed 5 times

Area ranges are

Number of LSA 13. Checksum Sum 0x16F0FD

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 16

Number of interfaces in this area is 2

Area has message digest authentication

SPF algorithm last executed 00:00:34.928 ago

SPF algorithm executed 7 times

Area ranges are

Number of LSA 5. Checksum Sum 0x02FCD3

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

Based on the information shown above, what is most likely causing the different missing routes throughout the network?


- A. Area 16 is configured with authentication.
- B. Area 16 has been configured to use the same interfaces as Area 0.
- C. Area 0 and Area 32 have been configured with mismatched LSA numbers.
- D. Area 16 has been configured as a total stub network
- E. Area 16 has been configured as a stub network
- F. Area 0 is discontinuous.
- G. None of the above

Answer: F

Explanation:

From the topology, we see D1 has 2 interfaces belong to Area 0, that are interfaces Fa0/1 & Fa0/2 but the output says there is only one interface in Area 0 (Number of interfaces in the this area is 1). Therefore we can deduce that a link in area 0 was down and area 0 is discontinuous.

QUESTION NO: 118



The company and the company network have both been growing rapidly. Multiple adds, moves and changes have been applied to the network. Your boss has asked you to troubleshoot a recent OSPF synchronization problem that has arisen. There have been synchronization problems at separate locations in the OSPF area 0. There have been reported link failures during the rapid growth of the company network. You are required to resolve the OSPF problem. OSPF must be able to converge when the network changes.

Refer to the information above to answer the following question.

Which configuration command on D1 (with a similar command on D2) will provide an immediate solution to the missing route problem?


- A. no area 16 stub
- B. no area 16 authentication message-digest
- C. area 16 virtual-link 8.187.175.82
- D. area 16 virtual-link 172.16.4.2
- E. no area 16 stub no-summary
- F. network 172.16.0.0.0.0.255.255 area 16
- G. None of the above

Answer: C

Explanation: Explanation

To fix this problem immediately without changing the topology we need to create virtual link between D1 & D2.

QUESTION NO: 119



The company and the company network have both been growing rapidly. Multiple adds, moves and changes have been applied to the network. Your boss has asked you to troubleshoot a recent OSPF synchronization problem that has arisen. There have been synchronization problems at separate locations in the OSPF area 0. There have been reported link failures during the rapid growth of the company network. You are required to resolve the OSPF problem. OSPF must be able to converge when the network changes.

Refer to the information above to answer the following question.

The log of d1 reports the following:

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to down
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down
%OSPF-5-ADJCHG: Process 1, Nbr 8.187.175.82/32 on FastEthernet0/1 from FULL to DOWN, Neighbor Down: Interface down or detached
```

This event was anticipated due to maintenance; however, it resulted in excessive lost routes. Which route should be the only one removed from the routing tables of the routers?


- A. 8.187.175.82/32
- B. 10.138.43.0/30
- C. 10.206.180.0/30
- D. 4.249.113.59/32
- E. 10.201.0.0/30
- F. None of the above

Answer: E

Explanation:

From the log we learn that the link of Interface Fa0/1 has been down. This link belongs to network 10.201.0.0/30 so we just need to remove this route from the routing table.

QUESTION NO: 120



The company and the company network have both been growing rapidly. Multiple adds, moves and changes have been applied to the network. Your boss has asked you to troubleshoot a recent OSPF synchronization problem that has arisen. There have been synchronization problems at separate locations in the OSPF area 0. There have been reported link failures during the rapid growth of the company network. You are required to resolve the OSPF problem. OSPF must be able to converge when the network changes.

Refer to the information above to answer the following question.

The R2 router has lost connectivity to R1. The following is R1s current route table:

```
172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
O IA 172.16.230.0/24 [110/21] via 10.138.43.1, 00:00:03, FastEthernet0/0
O IA 172.16.209.0/24 [110/22] via 10.138.43.1, 00:00:03, FastEthernet0/0
O IA 172.16.1.0/30 [110/21] via 10.138.43.1, 00:00:03, FastEthernet0/0
10.0.0.0/30 is subnetted, 1 subnets
C 10.138.43.0 is directly connected, FastEthernet0/0
```

Which expected route is missing from R1s route table based on the topology during the maintenance period?

- A. o 172.16.0.0 [110/2] via 10.138.43.1, 00:00:09, FastEthernet0/0
- B. o IA 9.152.105.122 [110/3] via 10.138.43.1, 00:00:09, FastEthernet0/0
- C. o IA 10.138.0.0 [110/3] via 10.138.43.1, 00:00:09, FastEthernet0/0
- D. o IA 10.249.0.0 [110/2] via 10.138.43.1, 00:00:09, FastEthernet0/0
- E. o IA 4.249.113.59 [110/2] via 10.138.43.1, 00:00:09, FastEthernet0/0
- F. o 8.187.175.82 [110/3] via 10.138.43.1, 00:00:09, FastEthernet0/0

Answer: F

Explanation: 0 8.287.175.82 [110/3] is not to be seen in the R1 current route table which means it has lost connectivity via 10.138.43.1. Also 10.0.0.0/30 has only 1 subnet.

QUESTION NO: 121

An administrator types in the command `router ospf 1` and receives the error message:

"OSPF process 1 cannot start." (Output is omitted.) What should be done to correctly set up OSPF?

- A. Ensure that an interface has been configured with an IP address.
- B. Ensure that an interface has been configured with an IP address and is up.
- C. Ensure that IP classless is enabled.
- D. Ensure that the interfaces can ping their directly connected neighbors.

Answer: B

Explanation:


A loopback interface can override the OSPF router ID. If a loopback interface exists, the router ID is the highest IP address on any active loopback interface.

The OSPF router-id command can be used to override the OSPF router ID.

Using a loopback interface or a router-id command is recommended for stability.

QUESTION NO: 122

During a recent OSPF election among three routers, RTA was elected the DR and RTB was elected the BDR, as seen in the graphic. Assume that RTA fails, and that RTB takes the place of the DR while RTC becomes the new BDR. What will happen when RTA comes back online?



- A. RTA will take the place of DR immediately upon establishing its adjacencies.
- B. RTA will take the place of DR only if RTB fails.
- C. RTA will take the place of DR only if both RTB and RTC fail.
- D. A new election will take place establishing an all new DR and BDR based on configured priority levels and MAC addresses.

Answer: C

Explanation:

If a router with a higher priority value gets added to the network, it does not preempt the DR and BDR. The only time a DR and BDR changes is if one of them is out of service. If the DR is out of service, the BDR becomes the DR, and a new BDR is selected. If the BDR is out of service, a new

BDR is elected. In a multi-access network, the router that is powered on first will generally become the DR, since the DR/BDR process is not pre-emptive.


Reference:: CCNP Self-Study Second Edition P.243

QUESTION NO: 123

Refer to the exhibit. During the process of configuring a virtual link to connect area 2 with the backbone area, the network administrator received this console message on R3:

*Mar 1 00:25:01.084: %OSPF-4-ERRRCV: Received invalid packet: mismatch area ID, from backbone area must be virtual link but not found from 20.20.20.1, Serial 0

How should the virtual link be configured on the OSPF routers to establish full connectivity between the areas?



- A. R1(config-router)# area 1 virtual-link 30.30.30.3
R3(config-router)# area 1 virtual-link 20.20.20.1
- B. R1(config-router)# area 1 virtual-link 20.20.20.2
R3(config-router)# area 1 virtual-link 30.30.30.2
- C. R1(config-router)# area 0 virtual-link 1.1.1.1
R3(config-router)# area 2 virtual-link 3.3.3.3
- D. R1(config-router)# area 1 virtual-link 3.3.3.3
R3(config-router)# area 1 virtual-link 1.1.1.1
- E. R1(config-router)# area 1 virtual-link 2.2.2.2
R3(config-router)# area 1 virtual-link 2.2.2.2

Answer: D

Explanation:

When designing a multi-area OSPF network, all areas should be connected to the backbone area. However, there may be instances when an area will need to cross another area to reach the backbone area like area 2 in this case. A virtual link has the following two requirements:

- + It must be established between two routers that share a common area and are both ABRs. +
- One of these two routers must be connected to the backbone.


In this case, two routers that satisfy the above requirements are R1 and R3. The syntax for creating a virtual link across an area is:

Area virtual-link

The area-id is the number of the transit area, in this example Area 1 and neighbor-router-id is the IP address of the highest loopback interface configured or can be manually set on the neighboring router.

QUESTION NO: 124

OSPF is configured over a Frame Relay network as shown in the exhibit. All PVCs are active. However, R1 and R3 fail to see all OSPF routes in their routing tables. The show ip ospf neighbor command executed on R2 shows the state of the neighbors. What should be done to fix the problem?



- A. The ip ospf network non-broadcast command should be configured on each Frame Relay interface.
- B. The ip ospf network broadcast command should be configured on each Frame Relay interface.
- C. The neighbor command should be configured under the OSPF routing process on all routers.
- D. The ip ospf priority value on the hub router should be set to 0.
- E. The ip ospf priority value on the spoke routers should be set to 0.

Answer: E

Explanation:

In an NBMA network topology, neighbors are not discovered automatically. OSPF tries to elect a DR and a BDR due to the multi-access nature of the network, but the election fails since neighbors are not discovered because NBMA environment doesn't forward broadcast and multicast packets. Neighbors must be configured manually to overcome these problems.

Also, additional configuration is necessary in a hub and spoke topology to make sure that the hub routers, which have connectivity with every other spoke router, are elected as the DR and BDR.

You must set the spoke interfaces to an OSPF priority of zero, this ensures that the spokes will not become the DR or BDR.

QUESTION NO: 125

Refer to the exhibit.

```
O I 2001:0DB8:0:0:7::/64 [110/20]
  via FE80::A8BB:CCFF:FE00:6F00, FastEthernet0/0
O 1 2001:0DB8:0:0:8::/64 [110 / 100]
  via FE80::A8BB:CCFF:FE00:6F00, FastEthernet0/0
O I 2001:0DB8:0:0:9::/64 [110/40]
  via FE80::A8BB:CCFF:FE00:6F00, FastEthernet0/0
```

In the show ipv6 route output, what would the metric be for a summary route that summarizes all three OSPFv3 routes displayed?

- A. 20
- B. 40
- C. 100
- D. 120
- E. 140
- F. 160


Answer: C

Explanation:

The cost of the summarized routes is the highest cost of the routes being summarized. In fact, in the old RFC 1583 standard, the cost of the summary route was the cost of the lowest metric. But when OSPF was updated in RFC 2178 and RFC 2328, the summary route should have the same cost as the highest-cost summarized route. In this case, the highest-cost is 100 according to the second entry.

QUESTION NO: 126

Study the exhibit below carefully.



In order to summarize all routes from area 0 to area 1, what must be configured on the router?

- A. area 0 range 172.16.96.0 255.255.224.0
- B. area 1 range 172.16.96.0 255.255.224.0
- C. area 1 range 172.16.96.0 255.255.0.0
- D. area 0 range 172.16.96.0 255.255.255.0

Answer: A

Explanation:

This identifies area 0 as the area containing the range of networks to be summarized. The networks will be summarized into area 1. The ABR R1 summarizes the range of subnets from

172.16.32.0 to 172.16.63.0 into one range: 172.16.32.0 255.255.224.0.

QUESTION NO: 127

Refer to the exhibit.

```
router eigrp 1
 redistribute ospf 1 route-map ospf-to-eigrp
 default-metric 20000 2000 255 1 1500
!
!
route-map ospf-to-eigrp deny 10
 match tag 6
 match route-type external type-2
!
route-map ospf-to-eigrp permit 20
 match ip address prefix-list pfx
 set metric 40000 1000 255 1 1500
!
route-map ospf-to-eigrp permit 30
 set tag 8
```

Which three statements accurately describe the result of applying the exhibited route map?
(Choose three.)

- A. All routes that do not match clauses 10 and 20 of the route map are redistributed with their tags set to 8.
- B. The map prohibits the redistribution of all type 2 external OSPF routes.
- C. The map permits the redistribution of all type 1 external OSPF routes.
- D. The map prohibits the redistribution of all type 2 external OSPF routes with tag 6 set.

Answer: A,C,D

Explanation:

The route-map command is used to configure policy routing, which is often a complicated task. A route map is defined using the syntax shown in the figure.

Syntax:

RouterA(Config)#route-map map-tag [permit | deny] <Sequence Number>

RouterA(Config-map-router)#

The map-tag is the name, or ID, of the route map. This map-tag can be set to something easily

recognizable name. The route-map command changes the mode on the router to the route-map configuration mode, from there conditions can be configured for the route map.

In this example, the first route-map entry (10) will deny all type 2 routes with a tag of 6 (both must be true).

The second route map (20) will set the metrics for all routes matching the pfx prefix list. Choice A is incorrect because the metrics shown do not need to match. These metrics will be assigned to the routes that match.

The last route-map (30) will set all other routes that did not already match 10 or 20 with a tag of 8.

QUESTION NO: 128

Which three restrictions apply to OSPF stub areas? (Choose three.)

- A. No virtual links are allowed.
- B. The area cannot be a backbone area.
- C. Redistribution is not allowed unless the packet is changed to a type 7 packet.
- D. The area has no more than 10 routers.
- E. No Autonomous System Boundary Routers are allowed.
- F. Interarea routes are suppressed.

Answer: A,B,E

Reference: http://www.juniper.net/techpubs/en_US/junos13.1/topics/concept/ospf-routing-understanding-ospf-areas-overview.html (see stub areas)

QUESTION NO: 129

Refer to the output.

```
Routing Process "ospfv3 1" with ID 172.16.3.3
It is an autonomous system boundary router
Redistributing External Routes from,
static
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msecs
Retransmission pacing timer 66 msecs
Numbers of external LSA 1. Checksum sum 0x218D
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Area 1
Number of interfaces in this area is 2
SPF algorithm executed 9 times
Number of LSA 15. Checksum Sum 0x67581
Number of DCbitless LSA 0
Number of indication LSA 0
Number of DoNotAge LSA 0
Flood list length 0
```

What IOS command produces this output?

Select the best response

- A. show ip ospf
- B. show ip ospf interface
- C. show ipv6 ospf interface
- D. show ipv6 ospf

Answer: D

Reference:

http://www.cisco.com/en/US/docs/ios/ipv6/command/reference/ipv6_15.html#wp2439467

QUESTION NO: 130

How is authentication handled with OSPFv3?


- A. OSPFv3 for IPv6 authentication is supported by IPv6 IPsec.
- B. OSPFv3 for IPv6 authentication is supported by MD5 authentication.
- C. OSPFv3 for IPv6 authentication is supported by IPv4 IPsec.
- D. OSPFv3 for IPv6 authentication is supported by SHA-1 authentication.

Answer: A

Reference: http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_ospf/configuration/xe-3s/ip6-route-ospfv3-auth-ipsec-xe.html

QUESTION NO: 131

Refer to the exhibit.



Routers R1 and R2 have been configured to operate with OSPF. Routers R1 and R3 have been configured to operate with RIP. After configuring the redistribution between OSPF and RIP on R1, no OSPF routes are distributed into RIP. What should be done to correct this problem?

- A. The redistribution command should be reentered with the match route-type parameter included.
- B. The redistribution command should be reentered with the route-map map-tag parameter included.
- C. The redistribution command should be reentered with the metric metric-value parameter included.
- D. Routes will first need to be distributed into another protocol, and then into RIP.

Answer: C

Explanation: Explanation

Notice that RIP metric is based on hop count only, and the maximum valid metric is 15. Anything above 15 is considered infinite. By default, when no metric is assigned when redistributing from EIGRP, OSPF, IS-IS, BGP into RIP, the default metric will be infinite. Therefore we must define a metric that is understandable to the receiving protocol. Usually, we should use a small value (like 1, 2, 3) so that after redistributing, that route can be advertised through many routers (because the limit is 15).

QUESTION NO: 132

By default, which statement is correct regarding the redistribution of routes from other routing protocols into OSPF?

- A. They will appear in the OSPF routing table as type E1 routes.
- B. They will appear in the OSPF routing table as type E2 routes.
- C. Summarized routes are not accepted.
- D. All imported routes will be automatically summarized when possible.
- E. Only routes with lower administrative distances will be imported.

Answer: B

Explanation: Explanation

Type E1 external routes calculate the cost by adding the external cost to the internal cost of each link that the packet crosses while the external cost of E2 packet routes is always the external cost only. E2 is useful if you do not want internal routing to determine the path. E1 is useful when internal routing should be included in path selection. E2 is the default external metric when redistributing routes from other routing protocols into OSPF.

QUESTION NO: 133

When implementing OSPFv3, which statement describes the configuration of OSPF areas?

- A. In interface configuration mode, the OSPFv3 area ID combination assigns interfaces to OSPFv3 areas.
- B. In router configuration mode, the network wildcard area ID combination assigns networks to OSPFv3 areas.
- C. In interface configuration mode, the IPv6 OSPF process area ID combination assigns interfaces to OSPFv3 areas.
- D. In router configuration mode, the IPv6 OSPF interface area ID combination assigns interfaces to OSPFv3 areas.


Answer: C

Reference:

http://www.hh.se/download/18.4cf286ee134f03ddb7b800015/1326882212358/Chapter3_VT2012.pdf (slide 42)

QUESTION NO: 134

Refer to the exhibit. OSPF is configured on all routers in the network. On the basis of the show ip ospf neighbor output, what prevents R1 from establishing a full adjacency with R2?



- A. Router R1 will only establish full adjacency with the DR and BDR on broadcast multiaccess networks.
- B. Router R2 has been elected as a DR for the broadcast multiaccess network in OSPF area
- C. Routers R1 and R2 are configured as stub routers for OSPF area 1 and OSPF area 2.
- D. Router R1 and R2 are configured for a virtual link between OSPF area 1 and OSPF area 2.
- E. The Hello parameters on routers R1 and R2 do not match.


Answer: A

Explanation: Explanation

From the output, we learn that R4 is the DR and R3 is the BDR so other routers will only establish full adjacency with these routers. All other routers have the two-way adjacency established.

QUESTION NO: 135

Refer to the exhibit.



On the basis of the configuration provided, how are the Hello packets sent by R2 handled by R5 in OSPF area 5?

- A. The Hello packets will be exchanged and adjacency will be established between routers R2 and R5.
- B. The Hello packets will be exchanged but the routers R2 and R5 will become neighbors only.
- C. The Hello packets will be dropped and no adjacency will be established between routers R2 and R5.
- D. The Hello packets will be dropped but the routers R2 and R5 will become neighbors.

Answer: C

Explanation:

The point of this question is the conditions of OSPF establish adjacency relationship.

For ospf, the optional capabilities must set the same between neighbors, but from the exhibit, R5 was configured as a stub area while R2 in area 0 is a normal area. So there will be no adjacent relationship established between routers R2 and R5.

QUESTION NO: 136

Which statement is true about OSPF Network LSAs?

- A.** They are originated by every router in the OPSF network. They include all routers on the link, interfaces, the cost of the link, and any known neighbor on the link.
- B.** They are originated by the DR on every multi-access network. They include all attached routers including the DR itself.
- C.** They are originated by Area Border Routers and are sent into a single area to advertise destinations outside that area.
- D.** They are originated by Area Border Router and are sent into a single area to advertise an Autonomous System Border Router.

Answer: B

Explanation:

The point of this question is OSPF Network LSAs

The feature of OSPF Network LSAs is that they are generated by DR, and DR only exist on multi-access network, the use of OSPF Network LSAs is that it list all neighbors around and send it to every router which run OSPF.

Incorrect Answer: OSPF Network LSAs are not originated by Area Border Routers.

QUESTION NO: 137

Refer to the exhibit.

```
NProuter#debug ip ospf events
```

```
OSPF events debugging is on
```

```
NProuter #
```

```
00:02:03: OSPF: Rcv hello from 172.16.1.1 area 0 from Serial0/0 10.1.1.1
```

```
00:02:03: OSPF: Mismatched hello parameters from 10.1.1.1
```

```
00:02:03: OSPF: Dead R 120 C 10, Hello R 30 C 30
```

```
00:02:26: OSPF: Rcv hellofrom 192.168.1.2 area 0 from Serial0/0 10.1.1.2
```

```
00:02:26: OSPF: Mismatched hello parameters from 10.1.1.2
```

```
00:02:26: OSPF: Dead R 120 C 10, Hello R 30 C 30
```

You are the network administrator responsible for the NProuter, the 10.1.1.1 router, and the

10.1.1.2 router. What can you determine about the OSPF operations from the debug output?

- A. The NProuter has two OSPF neighbors in the "Full" adjacency state.
- B. The NProuter serial0/0 interface has the OSPF dead timer set to 10 seconds.
- C. The NProuter serial0/0 interface has been configured with an OSPF network type of "point-to-point".
- D. The 10.1.1.1 and 10.1.1.2 routers are not using the default OSPF dead and hello timers setting.
- E. The "Mismatched" error is caused by the expiration of the OSPF timers.

Answer: B

Explanation: Explanation

First we should understand clearly about the line

Dead R 120 C 10, Hello R 30 C 30

The "R" here means "Received" and "C" means "Configured". In other words, "Dead R" is the Dead Timer Received from the neighbor and the "Dead C" is the Dead Timer of the local router. Therefore in this case "Dead R 120 C 10" means the Death Timer of the neighbor is 120 seconds while the local Dead Timer is 10 seconds, which causes a mismatch. Also we can learn that the local OSPF dead timer is set to 10 seconds.

For your information, by default, OSPF uses a 10-second hello timer and 40-second hold timer on broadcast and point-to-point links, and a 30-second hello timer and 120-second hold timer for all other network types.

QUESTION NO: 138

You have just completed an OSPF implementation. While executing your verification plan, you determine that R1 is not able to establish full OSPF adjacency with R2. The show ip ospf neighbor command output on R1 shows that R2 is stuck in the INIT state.

What could be the cause of this problem?

- A. DR and BDR election errors between R1 and R2.
- B. The R2 router has not received the OSPF hello packets from the R1 router.
- C. Mismatched interface maximum transmission unit (MTU) configuration between the R1 and R2.
- D. Mismatched OSPF hello interval configuration between the R1 and R2.
- E. Corrupted LSAs exchanges between the R1 and R2.

Answer: B

Explanation: Explanation


When a router receives an OSPF Hello from a neighbor, it sends the Hello packet by including that neighbor's router ID in the Hello packet. If the neighbor does not receive this packet (means that it

doesn't see itself in this packet), it will be stuck in INIT state. INIT state can be understood as a one-way Hello. An example of a router stuck in INIT state is shown below:

R1# show ip ospf neighbor fa0/0					
Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.2	1	INIT/-	00:00:35	10.1.1.2	Fast Ethernet0/0

QUESTION NO: 139

Refer to the exhibit. You have completed an OSPF implementation, and you are verifying OSPF operation. You notice that router A and router B are stuck in the two-way state. From the show ip ospf interface command output, what is the cause of this issue?



- A. All OSPF implementations must have at least one interface in area 0.
- B. You are attempting to run in the broadcast mode over an NBMA interface.
- C. Both routers are configured to function as a BDR; therefore, there is no DR router.
- D. Someone has changed the OSPF router ID; therefore you must clear the OSPF process.
- E. The OSPF priority is set to 0 on both routers; therefore neither can become the DR.

Answer: E

Explanation: Explanation

When OSPF adjacency is formed, a router goes through several state changes before it becomes fully adjacent with its neighbor. The states are Down, Attempt, Init, 2-Way, Exstart, Exchange, Loading, and Full.

An OSPF neighbor reaches the 2-way state when bidirectional communication is established (each router has seen the other's hello packet). This is the beginning of an OSPF adjacency. On broadcast media and non-broadcast multiaccess networks, the DR and BDR are elected in this state. But the priority on both routers are 0 so no DR and BDR are elected -> These routers stay in the 2-way state.

(Reference and a good resource of OSPF Neighbor states:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0e.shtml)

QUESTION NO: 140

You have completed an OSPF implementation, and you are verifying OSPF operation. During this verification, you notice that the OSPF route of 172.16.10.0 is repeatedly appearing and disappearing from the routing table. Further investigation finds that the OSPF CPU utilization is very high and the routers are constantly performing SPF calculations. You determine that 172.16.20.2 is the source of the 172.16.10.0 route. Using the show ip ospf database router 172.16.20.1 command, you notice that when this show command is performed repeatedly, the contents of the LSA change every few seconds.

What could be the cause of this problem?

- A. OSPF authentication errors between some of the routers.
- B. Two routers have the same OSPF router ID.
- C. Issues with mistuned OSPF timers.
- D. OSPF LSA pacing issues between some of the routers.
- E. OSPF neighbor adjacency problems between some of the routers.

Answer: B

Explanation:

When two routers use the same router ID in an OSPF domain, routing possibly does not work

correctly. Cisco bug IDs CSCdr61598 and CSCdu08678 enhance the detection and reporting mechanisms of duplicate router IDs. Access the Bug Toolkit (registered customers only) in order to view additional information about these Cisco bug ID.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080117102.shtml

QUESTION NO: 141

When an OSPF design is planned, which implementation can help a router not have memory resource issues?

- A.** Have a backbone area (area 0) with 40 routers and use default routes to reach external destinations.
- B.** Have a backbone area (area 0) with 4 routers and 30,000 external routes injected into OSPF.
- C.** Have less OSPF areas to reduce the need for interarea route summarizations.
- D.** Have multiple OSPF processes on each OSPF router. Example, router ospf 1, router ospf 2.

Answer: A

Explanation:

Memory issues usually come up when too many external routes are injected in the OSPF domain. A backbone area with 40 routers and a default route to the outside world would have less memory issues compared with a backbone area with 4 routers and 33,000 external routes being injected into OSPF. Router memory could also be conserved by using a good OSPF design.

Summarization at the area border routers and use of stub areas could further minimize the number of routes exchanged.

The total memory used by OSPF is the sum of the memory used in the routing table (show ip route summary) and the memory used in the LSDB. The following numbers are a "rule of thumb" estimate. Each entry in the routing table will consume between approximately 200 and 280 bytes plus 44 bytes per extra path. Each LSA will consume a 100 byte overhead plus the size of the actual LSA, possibly another 60 to 100 bytes (For router links, this depends on the number of interfaces on the router). These amounts should be added to memory already used by other processes and by the IOS itself.

If you really want to know the exact number, you can do a show memory with and without OSPF being turned on. The difference in the processor memory used would be the answer.

QUESTION NO: 142

The maximum number of routers per OSPF area typically depends on which three factors? (Choose three.)

- A. the kind of OSPF areas being implemented
- B. the number of external LSAs in the network
- C. the number of DRs and BDRs in the areas
- D. the number of virtual links in the areas
- E. how well the areas can be summarized
- F. the use of LSA filters

Answer: A,B,E

Reference:

https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=6&cad=rja&ved=0CFwQFjAF&url=http%3A%2F%2Ffaculty.valenciacollege.edu%2Fwyousif%2FCCNP%2FSemester5%2FPr esentations%2FMAOSPF_P2.ppt&ei=VUurUbmOA9OThgeDhYDoCg&usg=AFQjCNE5mLCAUIW Czou_vUX_DGhOOwcYxw&sig2=_7fgBDpXZCFi0Tay60wYmw&bvm=bv.47244034,d.ZG4 (Slide 85)

QUESTION NO: 143

You are troubleshooting an OSPF problem where external routes are not showing up in the OSPF database. Which two options are valid checks that should be performed first to verify proper OSPF operation? (Choose two.)

- A. Are the ASBRs trying to redistribute the external routes into a totally stubby area?
- B. Are the ABRs configured with stubby areas?
- C. Is the subnets keyword being used with the redistribution command?
- D. Is backbone area (area 0) contiguous?
- E. Is the CPU utilization of the routers high?

Answer: A,C

Explanation: Explanation

A totally stubby stubby area cannot have an ASBR so it will discard this type of LSA (LSA Type 5)
-> A is a valid check.

Each stubby area needs an ABR to communicate with other areas so it is normal -> B is not a valid check.

When pulling routes into OSPF, we need to use the keyword "subnets" so that subnets will be redistributed too. For example, if we redistribute these EIGRP routes into OSPF:

+ 10.0.0.0/8+ 10.10.0.0/16+ 10.10.1.0/24

without the keyword "subnets"

router ospf 1 redistribute eigrp 1

Then only 10.0.0.0/8 network will be redistributed because other routes are not classful routes, they are subnets. To redistribute subnets we must use the keyword "subnets"

router ospf 1 redistribute eigrp 1 subnets

-> C is a valid check.

We don't need to care if area 0 is contiguous or not -> D is not a valid check.

CPU utilization cannot be the cause for this problem -> E is not a valid check.

QUESTION NO: 144

When verifying the OSPF link state database, which type of LSAs should you expect to see within the different OSPF area types? (Choose three.)

- A. All OSPF routers in stubby areas can have type 3 LSAs in their database.
- B. All OSPF routers in stubby areas can have type 7 LSAs in their database.
- C. All OSPF routers in totally stubby areas can have type 3 LSAs in their database.
- D. All OSPF routers in totally stubby areas can have type 7 LSAs in their database.
- E. All OSPF routers in NSSA areas can have type 3 LSAs in their database.
- F. All OSPF routers in NSSA areas can have type 7 LSAs in their database.

Answer: A,E,F

Explanation: Explanation

Below summarizes the LSA Types allowed and not allowed in area types:

Area Type	Type 1 & 2 (within area)	Type 3 (from other areas)	Type 4	Type 5	Type 7
Standard & backbone	Yes	Yes	Yes	Yes	No
Stub	Yes	Yes	No	No	No
Totally stubby	Yes	No	No	No	No
NSSA	Yes	Yes	No	No	Yes
Totally stubby NSSA	Yes	No	No	No	Yes

Popular LSA Types are listed below:

LSA Type	Description	Details
1	Router LSA	Generated by all routers in an area to describe their directly attached links
2	Network LSA	Advertised by the DR of the broadcast network (does not cross ABR)
3	Summary LSA	Advertised by the ABR of originating area
4	Summary LSA	Generated by the ABR of the originating area to advertise an ASBR to all other areas in the autonomous system
5	AS external LSA	Used by the ASBR to advertise networks from other autonomous systems
7	Defined for NSSAs	Generated by an ASBR inside a Not-so-stubby area (NSSA) to describe routes redistributed into the NSSA

QUESTION NO: 145

When verifying OSPF virtual link problems, which is an important item to check on the two transit OSPF routers?

- A. OSPF process ID
- B. OSPF router ID
- C. OSPF network type
- D. OSPF memory usage
- E. OSPF CPU utilization
- F. OSPF stub area configurations


Answer: B

Explanation: Explanation

The OSPF router IDs of the two transit OSPF routers are used to form the virtual link (with the area area-id virtual-link neighbor-router-id command) so it is an important item to check -> B is correct.

QUESTION NO: 146

Refer to the exhibit.



A network administrator wants to reduce the number of OSPF routes advertised from Area 1 into Area 2. As the router configuration specialist, what two things would you do to accomplish this goal? (Choose two.)


- A. Enter the configuration on router B1.
- B. Enter the configuration on router B4.
- C. On the same router, enter the Summary-address 10.10.1.0 255.255.255.128 subcommand.
- D. On the same router, enter the Area 1 range 10.10.1.0 255.255.255.128 subcommand.
- E. On the same router, enter the Area 2 range 10.10.1.0 255.255.255.128 subcommand.

Answer: A,D

Explanation:

QUESTION NO: 147

Refer to the exhibit.



A company would prefer all Internet-bound OSPF routed traffic to use ISP ABC with ISP DEF as a backup. As the network consultant, what three configuration changes should you make? (Choose three.)

- A. The default-information originate command should be configured on router B1 and B4.
- B. The default-information originate command should be configured on router B2 and B3.
- C. If the metric value for ISP ABC is set at the default, the ISP DEF metric value should be set to 1.

- D. If the metric value for ISP ABC is set at the default, the ISP DEF metric value should be set to 25.
- E. The metric type value should be set to type 1.
- F. The metric type value should be set to type 2.

Answer: B,D,F

Explanation: Explanation

Routers B2 & B3 need to advertise a default route to the Internet for “inside” OSPF routers so we should use the “default-information originate” command with a default route (something like “ip route 0.0.0.0 0.0.0.0”) pointing to the ISP router.

If no metric is specified, OSPF puts a default value of 20 when redistributing routes from all protocols except BGP routes (BGP routes get a metric of 1). We use ISP DEF as a backup so its metric value should be set to a higher value than 20.

There are two types of external routes: external type 1 and external type 2. The difference between the two is in the way the cost (metric) of the route is being calculated: + The cost of a type 2 route is always the external cost, irrespective of the interior cost to reach that route. + Type 1 cost is the addition of the external cost and the internal cost used to reach that route.

-> We should configure the type 2 external route to make sure the ISP ABC is always referred over ISP DEF because internal routing does not determine the path.

Note: E2 is the default external metric, but E1 is preferred over E2 if two equal-cost paths exist.

QUESTION NO: 148

The administrator wants to verify the current state of the OSPF database loading process. Which show command should the administrator use?

- A. show ip ospf [process-id] interface
- B. show ip ospf neighbor
- C. show ip ospf [process-id]
- D. show ip ospf [process-id area-id] database

Answer: B

Explanation: Explanation


The “show ip ospf neighbor” command can be used to view the current state of the OSPF database loading process. In the output below we can see router 2.2.2.2 is in 2way state, router 3.3.3.3 is elected as the BDR & router 4.4.4.4 is the BR.


```
R1# show ip ospf neighbor fa0/0
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.2	1	2WAY/DR0THER	00:00:35	10.1.1.2	Fast Ethernet0/0
3.3.3.3	1	FULL/BDR	00:00:38	10.1.1.3	Fast Ethernet0/0
4.4.4.4	1	FULL/BR	00:00:34	10.1.1.4	Fast Ethernet0/0

QUESTION NO: 149

Refer to the exhibit. OSPF is running throughout the network. You want to minimize the propagation of LSAs into and out of Area 1. Which OSPF feature would best achieve this goal?



- A. stub
- B. totally stubby
- C. NSSA
- D. totally NSSA

Answer: D

Explanation: Explanation

We need to redistribute RIP from R1 to Area 1 so Area 3 cannot be a stub or totally stubby area. To minimize the propagation of LSAs into and out of Area 1 we should configure it as a totally NSSA. Notice that a NSSA allows LSA Type 3 & 7 while a Totally NSSA only allows LSA Type 7

Note: Both Totally Stubby Area & Totally Stubby NSSA do not accept external AS routes or inter-area routes (LSA Types 3, 4 and 5). They recognize only intra-area routes and the default route 0.0.0.0. The main difference between them is Totally Stubby NSSA accepts routes from other AS while Totally Stubby Area does not.

Below summarizes the LSA Types allowed and not allowed in area types:

Area Type	Type 1 & 2 (within area)	Type 3 (from other areas)	Type 4	Type 5	Type 7
Standard & backbone	Yes	Yes	Yes	Yes	No
Stub	Yes	Yes	No	No	No
Totally stubby	Yes	No	No	No	No
NSSA	Yes	Yes	No	No	Yes
Totally (stubby) NSSA	Yes	No	No	No	Yes

QUESTION NO: 150

Refer to the exhibit.

```

ASBR
router ospf 123
 redistribute eigrp 1 route-map eigrp-to-ospf
 !
 route-map eigrp-to-ospf permit 10
  match ip address prefix-list private
  set tag 255
 route-map eigrp-to-ospf permit 20
 !
 ip prefix-list private permit 10.0.0.0/8 ge 8 le 30
 ip prefix-list private permit 172.16.0.0/11 ge 11 le 30
 ip prefix-list private permit 192.168.0.0/16 ge 16 le 30

Area 0 ABR
router ospf 123
 distribute-list route-map private-filter in
 !
 route-map private-filter deny 10
  match tag 255
 route-map private-filter permit 20

```

The partial configuration for an OSPF ASBR and an Area 0 ABR is shown. Assume the OSPF configurations throughout the network are operable. Which statement about these configurations is true?

- A. The ASBR route-maps are basically useless, because there are no deny prefix-lists.
- B. LSA Type 5s will not be received by the ABR from the ASBR.

- C.** The OSPF backbone will not learn any RFC 1918 addresses.
D. The matched prefix-list addresses will be given a metric of 255, which is essentially unreachable.

Answer: C

Explanation: Explanation

The ASBR accepts RFC 1918 addresses and set these networks to "tag 255 but when advertising into Area 0, the ABR Area 0 filters out these networks because they match "tag 255 so the OSPF backbone will not learn any RFC 1918 addresses.

Note that if you use an ACL in a route-map **deny** clause, routes that are permitted by the ACL are not redistributed.

All the networks with "tag 255 are blocked by the clause 10 while all other networks are permitted by the clause 20 of the route-map (if a **match** command is not present, all routes match the clause).

Note:

RFC 1918 addresses include:

+ Class A: 10.0.0.0 – 10.255.255.255 (10/8 prefix)+ Class B: 172.16.0.0 – 172.31.255.255 (172.16/12 prefix)+ Class C: 192.168.0.0 – 192.168.255.255 (192.168/16 prefix)

QUESTION NO: 151

Refer to the exhibit.

```
Router eigrp 10
  Network 10.0.0.0
Router ospf 100
  redistribute eigrp 10 subnets
  network 172.16.0.0 0.0.255.255 area 1
.....
A partial routing table
  C    10.10.10.16/28 is directly connected, Ethernet1/0
  D    10.10.10.64/26 [90/284160] via 10.10.24.2, 01:54:39, Ethernet1/0
  C    172.16.10.0/24 is directly connected, FastEthernet0/1
```

Examine the partial configuration and the routing table excerpt. Which routes would be redistributed into OSPF area 1?

- A.** 10.10.10.16/28 only
B. 10.10.10.16/28 and 10.10.10.64/26

- C. 10.10.10.16/28, 10.10.10.64/26, and 172.16.10.0/24
- D. 10.10.10.64/26 only

Answer: D

Explanation:

QUESTION NO: 152

Which two statements about route redistribution when implementing OSPF are true? (Choose two.)

- A. Routes learned using any IP routing protocol can only be redistributed into non IP routing protocols.
- B. OSPF can import routes learned using EIGRP, RIP, and IS-IS.
- C. OSPF routes cannot be exported into EIGRP, RIP, and IS-IS.
- D. At the interdomain level, OSPF cannot import routes learned using BGP.
- E. OSPF routes can be exported into BGP.

Answer: B,E

Reference:

http://www.cisco.com/en/US/docs/switches/lan/catalyst3750x_3560x/software/release/12.2_55_se/configuration/guide/swiprout.html (See configuring OSPF)

QUESTION NO: 153

Which two routing interface parameters are supported in OSPF implementations? (Choose two.)

- A. retransmit-interval
- B. dead-interval
- C. stub area
- D. virtual link
- E. NSSA area

Answer: A,B

Explanation: Explanation

When OSPF sends an advertisement to an adjacent router, it expects to receive an acknowledgment from that neighbor. If no acknowledgment is received, the router will retransmit the advertisement to its neighbor. The retransmit-interval timer controls the number of seconds between retransmissions. To edit the retransmit-interval, use the "**ip ospf retransmit-interval seconds**" in interface configuration mode.

Dead-interval is the number of seconds without hello packets before an adjacency is declared down. To edit the dead-interval, use the "**ip ospf dead-interval seconds**" in interface configuration mode.

Other answers are not correct because they are not interface parameters.

QUESTION NO: 154

RouterA#

~~~~~

!

router ospf 1

log-adjacency-changes

network 10.0.0.0 0.255.255.255 area 1

network 172.16.1.0 0.0.0.255 area 1

!

~~~~~

RouterB#

~~~~~

router ospf 1

log-adjacency-changes

network 10.0.0.0 0.255.255.255 area 2

network 172.16.2.0 0.0.0.255 area 2

!

~~~~~

RouterC#

~~~~~

!

router ospf 1


log-adjacency-changes

network 10.0.0.0 0.255.255.255 area 0

!

~~~~~

Refer to the exhibits. You are verifying your OSPF implementation, and it does not seem to be functioning properly. What can you conclude from the exhibit and the show running-configuration command output?



- A. The OSPF areas are not configured correctly.
- B. The wildcard masks for the 10.x.x.x networks are incorrect.
- C. The 172.16.x.x networks need to be connected to area 0 using virtual links.
- D. The 172.16.x.x networks are discontinuous. OSPF is automatically summarizing them to 172.16.0.0/16 and data is being "black holed."
- E. There is not enough information to make a determination.

Answer: A

Explanation: Explanation

The E0/0 & E0/1 interfaces of router C belong to area 0 while E0/0 of router A belongs to area 1; E0/0 of router B belongs to area 2 -> it is not correct. Both E0/0 interfaces of router A & B should be in area 0 -> A is not correct.

QUESTION NO: 155

OSPF is enabled on router A. You execute the following command on router A and receive the

accompanying output:

```
RouterA#ping 224.0.0.5 repeat 1
```

Type escape sequence to abort.

Sending 1, 100-byte ICMP Echos to 224.0.0.5, timeout is 2 seconds:

```
Reply to request 0 from 10.100.100.1, 4 ms
```

10.100.100.1 is the IP address of a loopback interface on router A.

What can you conclude about router A?

- A.** Only the router A loopback interface is participating in the OSPF routing process.
- B.** None of the router A interfaces are participating in the OSPF routing process.
- C.** Router A is using the loopback interface IP address as its OSPF router ID.
- D.** Router A does not have any reachable OSPF neighbors.


Answer: D

Explanation: Explanation

The 224.0.0.5 address is the multicast for OSPF routers. Therefore when you ping to this address all OSPF routers will reply but from the output we learn that only the local router responded -> Router A does not have any reachable OSPF neighbors.

QUESTION NO: 156

Refer to the exhibit.



What is the effect of the OSPF configuration on router B?

- A. All interfaces will be in area 0.
- B. The router will be an ABR with s1/0 in area 0 and f0/0 and 0/1 in area 1.
- C. The router will be an ABR with s1/0 in area 0 and f0/0 and 0/1 in area 2.
- D. The router will be an ABR with s1/0 in area 0, f0/0 in area 1, and f0/1 in area 2.


Answer: D

Reference:

https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=4&cad=rja&ved=0CEIQFjAD&url=https%3A%2F%2Flearningnetwork.cisco.com%2Fservlet%2FJiveServlet%2Fdownload%2F4907-2332%2FCase%2520..&ei=sU-rUeGvFlanhAfR-IHwDQ&usg=AFQjCNH4zqER_WNMmrHEY2VhuMzDTy7D1xQ&sig2=v1xIJWspSrBE6szoQAP8cg&bvm=bv.47244034,d.ZG4

QUESTION NO: 157

Refer to the exhibit. Two routers are connected by Frame Relay and are running OSPF between them. Each router has been configured with the appropriate network statements under router ospf 1, but the routers are not forming an adjacency. Which of the following three commands could be configured on each router to correct this problem? (Choose three.)



A.

RouterC(config-if)#ip ospf network broadcast
RouterD(config-if)#ip ospf network broadcast

B.

RouterC(config-if)#ip ospf network point-to-point
RouterD(config-if)#ip ospf network point-to-point

C.

RouterC(config-router)#neighbor 10.100.100.4
RouterD(config-router)#neighbor 10.100.100.3

D.

RouterC(config-router)#neighbor 10.255.255.254
RouterD(config-router)#neighbor 10.255.255.253

Answer: A,B,D

Explanation: Explanation

By default, Frame Relay is classified as a non-broadcast network, meaning it doesn't send any broadcasts/multicasts like RIP, OSPF or EIGRP updates across the network (hello packets of OSPF are multicast to 224.0.0.5). Hence, in NBMA network, the neighbors are not discovered automatically; they must be configured manually. There are two ways to simulate a broadcast model on an NBMA network:

- + Define the network type as broadcast with the "ip ospf network broadcast" interface sub-command
- + Configure the neighbor statements under router ospf mode (though configuring the neighbor statement on one end is sufficient to form adjacency, it is a good practice to have it configured on both the ends)

Besides these two ways, another way for making OSPF work with Frame Relay is configuring the network as a Point-to-Point network (or Point-to-Multipoint, notice that OSPF treats Point-to-Multipoint network as a collective of point-to-point links).

Some information about Point-to-Multipoint (or Point-to-Point) network:

Note: Point-to-Multipoint networks do not maintain a DR/BDR relationship.

QUESTION NO: 158

Your network has a mixture of Fast Ethernet and Gigabit Ethernet links. What needs to be done to ensure optimal data routing when using OSPF?

- A. Nothing. OSPF will determine the most optimal path for routing data by default.
- B. Adjust the hello and dead timers for more rapid detection of link failures.
- C. Increase the reference-bandwidth used to calculate the interface default metrics, on all routers in your network.
- D. Set the priority values on every broadcast interface to ensure that the designated and backup designated routers are the routers with the most processor and memory resources.

Answer: C

Explanation: Explanation

The default formula to calculate the cost for the OSPF metric is $(108/BW)$. Therefore when using default reference bandwidth (100M) to calculate ospf, Gi and Fa interfaces could have the same cost (1). This problem can be fixed by configuring reference bandwidth to 1000M.

Default Bandwidth (100M)	Configured Bandwidth (1000M)
Fa = $100M / 100M = 1$	Fa = $1000M / 100M = 10$
Gi = $100M / 1000M = 0.1 \rightarrow$ round up to 1	Gi = $1000M / 1000M = 1$

This is how to configure reference bandwidth to 1000Mbps:

Router(config)#router ospf 1Router(config-router)#auto-cost reference-bandwidth 1000

QUESTION NO: 159

To make OSPF area 4 a totally stubby area, which two things need to be done? (Choose two.)

- A. Apply the area 4 stub command to all routers in the area.
- B. On the ABR, use the area 4 stub command with the no-summary keyword.
- C. On the ABR, specify a default cost for the area with the area default-cost command.
- D. On the ABR, use the default-information originates command to inject a default route into area 4.

E. Use the auto-cost command on each router in the area to automatically determine the cost to other OSPF areas.

Answer: A,B

Explanation: Explanation

To make an area "totally stubby" we must apply the "**area area-id stub no-summary**" command on the ABR and **area area-id stub** commands to all other routers in that area.

Note: The ABR in a totally stubby area does not create Type 3 summary LSA. It only creates a default route to outside destinations.

QUESTION NO: 160

Refer to the exhibit. Router C was configured so that it could form an adjacency with three OSPF neighbors, one connected to each of its three physical interfaces.

```
RouterC#show ip interface brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
Ethernet0/0	10.1.0.1	YES	NVRAM	up	up
Serial0/0	10.255.255.253	YES	NVRAM	up	up
Ethernet0/1	10.2.0.1	YES	NVRAM	up	up
Loopbac k0	10.100.100.3	YES	NVRAM	up	up

```

RouterC#show ip ospf
Routing Process "ospf 1" with ID 10.100.100.3
Supports only single TOS(TOS0) routes
Supports opaque LSA
It is an area border router
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
Number of external LSA 0. Checksum Sum 0x000000
Number of opaque AS LSA 0. Checksum Sum 0x000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 2. 1 normal 1 stub 0 nssa
External flood list length 0
Area BACKBONE(0)
  Number of interfaces in this area is 3
  Area has no authentication
  SPF algorithm executed 10 times
  Area ranges are
  Number of LSA 14. Checksum Sum 0x053D21
  Number of opaque link LSA 0. Checksum Sum 0x000000
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0
Area 4
  Number of interfaces in this area is 1
  It is a stub area, no summary LSA in this area
    generates stub default route with cost 1
  Area has no authentication
  SPF algorithm executed 9 times
  Area ranges are
  Number of LSA 4. Checksum Sum 0x01F200
  Number of opaque link LSA 0. Checksum Sum 0x000000
  Number of DCbitless LSA 0
  Number of indication LSA 0
  Number of DoNotAge LSA 0
  Flood list length 0

```

Which statement is correct about router C?

- A. It is configured and functioning correctly as an OSPF internal router.
- B. It is configured and functioning correctly as an ABR attached to stub area 4.
- C. It is configured and functioning correctly as an ASBR attached to external area 4.
- D. It is configured and functioning correctly as an ABR attached to totally stubby area 4.
- E. It is not configured correctly to function as specified.

Answer: D

Explanation: Explanation

From the output of the "show ip ospf" command we notice 2 lines:

+ It is an area border router+ It is a stub area, no summary LSA in this area (in Area 4)

Therefore we can conclude RouterC is an ABR and it is attached to a totally stubby area (with no summary LSA advertised).

QUESTION NO: 161

Refer to the exhibits. What can be done to fix the problem?

RouterA#debug ip ospf events


OSPF events debugging is on

RouterA#

04:43:16: OSPF. Rcv pkt from 10.3.0.2, Ethernet0/1, area 0.0.0.1 mismatch area 0.0.0.2 in the header

04:43:19: OSPF. Rcv hello from 10.100.100.3 area 0 from Ethernet0/0 10.1.0.1

04:43:19: OSPF. End of hello processing



- A. Change router B E0/1 interface to area 0.
- B. Change router A interface E0/1 to area 0.0.0.2.
- C. Configure the E0/1 interfaces of router A and router B to be in area 0.
- D. Shut down the E0/1 interfaces in router A and router B as OSPF does not allow "back doors" between areas.
- E. Remove the E0/1 interfaces in router A and router B from the OSPF process. Use static routes to route data directly from router A to router B to avoid passing data through router C in area 0.

Answer: C


Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml (See

topic: the backbone and area 0)

QUESTION NO: 162

Refer to the exhibit.



Router B is performing bidirectional redistribution between EIGRP and OSPF. The network 10.100.1.0/24 should not be reachable from the 10.100.9.0/24 network. However, it needs to be reachable from any network within the EIGRP domain. All other networks should be seen in both domains. Which change to router B would accomplish these goals?

- A. Under the EIGRP process, insert the distribute-list block_net_1 out ospf 1 command.
- B. Under the OSPF process, insert the distribute-list block_net_1 in serial1/0 command.
- C. Under the EIGRP process, insert the distribute-list block_net_1 in serial1/0 command.
- D. Under the OSPF process, insert the distribute-list block_net_1 out eigrp 1 command.


Answer: D

Explanation: Explanation

The network 10.100.1.0/24 should not be reachable from network 10.100.9.0/24 -> we need to filter updates **from EIGRP to OSPF** so that the routing table of router C doesn't have network 10.100.1.0/24 -> we need to filter it under OSPF process and the direction should be out of router B ->.

QUESTION NO: 163

Refer to the exhibit. ROUTE.com is planning to implement a new secure OSPF network to support traffic between clients on the 172.16.10.0/24 network and the file server on the 172.16.20.0/24 network. You have been asked to review the implementation plan for the OSPF project. Which statement about the plan is true?

Implementation Plan

1. Enable OSPF process 1 on all routers.
2. Enable area 0 on R2, R3, and R4.
3. Enable area 10 on R1, R2 and R3.
4. Enable area 20 on R4 and R5.
5. Verify that all routers contain a complete routing table.
6. Verify that the clients can successfully access the file server.
7. Document configuration changes.

- A. It is complete as written.
- B. It should include a task that shuts down all unused interfaces.
- C. It should include tasks that enable and verify OSPF authentication.
- D. It should include a task that establishes a file transfer baseline before and after the configuration is changed.

Answer: C

Explanation: Explanation


The complete implementation plan should be

1. Enable OSPF process 1 on all routers
2. Enable area 0 on R2, R3 and R4
3. Enable area 10 on R1, R2 and R3
4. Enable area 20 on R4 and R5
5. Enable and verify OSPF authentication
6. Verify that all routers contain a complete routing table
7. Verify that the clients can successfully access

the server8. Document configuration changes

QUESTION NO: 164

Refer to the exhibit. ROUTE.com is planning to implement a secure OSPF network to support traffic between clients on the 172.16.10.0/24 network and the file server on the 172.16.20.0/24 network. You have been asked to review the implementation and verification plans for this OSPF project. Which statement about the plan is true?



Clients

Verification Plan

1. On R1 and R5, verify end-to-end connectivity between prefixes 172.16.10.0 and 172.16.20.0 using the extended ping command.
2. Verify that all routers have established a full neighbor relationship with the appropriate neighbors.
3. Verify the proper authentication method is active for each neighbor relationship.
4. Verify that each prefix is assigned to the appropriate OSPF area.
5. Verify the appropriate routes and summaries are in the routing table of each router.
6. Verify end-to-end connectivity between the clients and the file server.

- A. It is complete as written.
- B. It should include a task that verifies that the interarea routes are using the proper MED.
- C. It should include a task that verifies that load sharing is active on R1 and R4.
- D. It should include a task that verifies that all redundant links will become active when the primary links are shut down.


Answer: D

Reference:

http://www.cisco.com/en/US/docs/routers/access/as5350/software/feature/guide/pull_rlm.html

QUESTION NO: 165

Refer to the exhibit. Which two Cisco IOS commands on R2 would verify its OSPF neighbor relationships? (Choose two.)



- A. show ip ospf
- B. show ip ospf interface
- C. show ip ospf neighbor
- D. show ip ospf database
- E. show ip ospf statistics
- F. show running-config | begin router ospf

Answer: B,C

Explanation: Explanation

The **show ip ospf interface** command shows us information about the neighbor count and adjacent neighbor count:

```

RouterA# show ip ospf int s1/0
Serial 1/0 is up, line protocol is up
  Internet Address 10.100.3.253/30, Area 1
  Process ID 1, Router ID 1.1.1.1, Network Type BROADCAST, Cost:64
  Transmit Delay is 1 sec, State DROTHER, Priority 0
  No designated router on this network
  No backup designated router on this network
  Old designated Router (ID) 2.2.2.2, Interface address 10.100.3.254
  Flush timer for old DR LSA due in 00:01:12
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    oob-resync timeout 40
    Hello due in 00:00:01
  Supports Link-Local Signaling (LLS)
  Index 1/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 3
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)

```

The **show ip ospf neighbor** command shows us the role of each neighbor (DR, BDR, DROTHER).

```
Dev-1#sh ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.200.200.13	1	FULL/BDR	00:00:33	10.1.1.3	Ethernet 0/0

```
Dev-3#sh ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.31.1.1	2	FULL/DR	00:00:31	10.1.1.1	Ethernet 0/0

QUESTION NO: 166

Which three statements about OSPF areas are true? (Choose three.)

- A. Areas introduce a boundary on the link-state updates.
- B. Areas are logical definitions specific to any given router.
- C. All routers within an area have the exact link-state database.
- D. The calculation of the Dijkstra algorithm on a router is limited to changes within an area.

E. The area designated router will always have a priority of 0.

Answer: A,C,D

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml
(areas and border routers)

QUESTION NO: 167

Given the following partial configuration for Router A:

```
interface serial 0  
  
ip address 10.1.1.1 255.255.255.0  
  
encapsulation frame-relay  
  
ip ospf network point-to-multipoint  
  
router ospf 7  
  
network 10.1.1.0 0.0.0.255 area 0
```

Which two statements are correct? (Choose two.)

- A. DR/BDR elections do not take place.
- B. The area 0 NBMA cloud is configured as more than one subnet.
- C. The router is restricted to a hub and spoke topology.
- D. OSPF neighbor statements are not necessary.

Answer: A,D

Explanation:

In an OSPF Point-to-Multipoint environment, DR/BDR elections do not take place.

The neighbor command became somewhat obsolete with the introduction of the capability to configure other network modes for the interface, regardless of the underlying physical topology.

Reference: Building Scalable Cisco Networks (Cisco Press) page 130 and 181

Point-to-Multipoint Network:

Point-to-multipoint is a single interface that connects to multiple destinations. The underlying network treats the network as a series of point-to-point circuits. It replicates LSA packets for each circuit. OSPF traffic is sent as multicast. There is no DR or BDR election. This technology uses one IP subnet for all endpoints on the network.

By default, the network is considered to be a series of point-to-point interfaces. There is no need to specify neighbors, because the neighbors will see each other and simply become adjacent, with no need for the election of a DR or a BDR. Point-to-multipoint does not try to reduce adjacencies using a DR. Instead, it accepts the extra overhead of having a full set of adjacencies for the sake of stability. Point-to-multipoint forms an adjacency automatically along any PVC, which causes more overhead but is more resilient than NBMA.

QUESTION NO: 168

Which command displays the number of times that the OSPF Shortest Path First (SPF) algorithm has been executed?

- A. show ip protocol
- B. show ip ospf
- C. show ip ospf database
- D. show ip ospf interface

Answer: B

Explanation:

The following table describes the output of the "show ip ospf" command and their meanings: Table 52 show ip ospf Field Descriptions

Field	Description
Routing process "ospf 201" with ID 10.0.0.1	Process ID and OSPF router ID.
Supports...	Number of types of service supported (Type 0 only).
SPF schedule delay	Delay time of SPF calculations.
Minimum LSA interval	Minimum interval between link-state advertisements.
LSA group pacing timer	Configured LSA group pacing timer (in seconds).
Interface flood pacing timer	Configured LSA flood pacing timer (in milliseconds).
Retransmission pacing timer	Configured LSA retransmission pacing timer (in milliseconds).
Number of...	Number and type of link-state advertisements that have been received.
Number of external LSA	Number of external link-state advertisements.
Number of opaque AS LSA	Number of opaque link-state advertisements.
Number of DCbitless external and opaque AS LSA	Number of demand circuit external and opaque link-state advertisements.
Number of DoNotAge external and opaque AS LSA	Number of do not age external and opaque link-state advertisements.
Number of areas in this router is	Number of areas configured for the router.
External flood list length	External flood list length.

Reference:http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/iprrp_r/ip2_s3g.htm#wp1036469

QUESTION NO: 169

The Dev-1 and Dev-3 routers are OSPF neighbors over the Ethernet 0/0 connection. Based on the show ip ospf neighbor output from the Dev-1 and Dev-3 routers, which statement is true?

Dev-1#sh ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
10.200.200.13	1	FULL/BDR	00:00:33	10.1.1.3	Ethernet0/0
Dev-3#sh ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
172.31.1.1	2	FULL/DR	00:00:31	10.1.1.1	Ethernet0/0

- A. Dev-1 is the DR because it has a higher OSPF router priority.
- B. Dev-1 is the DR because it has a lower OSPF router ID.
- C. Dev-3 is the DR because it has a higher OSPF router priority.
- D. Dev-3 is the DR because it has a lower OSPF router ID.
- E. Both Dev-1 and Dev-3 are using the default OSPF router priority.

Answer: A

Explanation:

The output shown above displays information about the neighbors, so from Dev-3 we see that the priority of Dev-1 is 2, and that Dev-1 is the DR. From Dev-1 we can see that its neighbor (Dev-3) is the BDR and has an OSPF priority of 1.

QUESTION NO: 170

When other routing protocol routes are being redistributed into OSPF, what is one of the most common problems?

- A. missing the tag option in the redistribute command.
- B. missing the subnet option in the redistribute command.
- C. missing the metric option in the redistribute command.
- D. misconfiguring the metric-type option in the redistribute command to type-1.
- E. misconfiguring the metric-type option in the redistribute command to type-2.

Answer: B

Explanation:

The point of this question is about the redistribution of OSPF.


To redistribute routes from one routing domain into another routing domain, use the redistribute command in router configuration mode. redistribute protocol [process-id] {level-1 | level-1-2 | level-2} [as-number] [metric {metric-value | transparent}] [metric-type type-value] [match {internal | external 1 |

external 2}} [tag tag-value] [route-map map-tag] [subnets]

The subnets keyword tells OSPF to redistribute all subnet routes. Without the subnets keyword, only networks that are not subnetted are redistributed by OSPF.

QUESTION NO: 171

Refer to the diagram. Which OSPF configuration command is required to configure Area 3 as a totally stubby area?



- A. On the ABR
router ospf 10
area 3 nssa
- B. On the ABR
router ospf 10
area 3 stub
- C. On all Area 3 routers
router ospf 10
area 3 stub no-summary
- D. On all Area 3 routers
router ospf 10
area 3 stub
- E. On the ABR
router ospf 10
area 3 stub no-summary

Answer: E

Explanation: Explanation

To make an area "totally stubby" we must apply the "**area area-id stub no-summary**" command

on the ABR (the middle router in this case) and **area area-id stub** commands to all other routers in that area.

QUESTION NO: 172

Which show command will display only the Type 5 LSAs in the OSPF topology database?

- A. show ip ospf database external
- B. show ip ospf database nssa-external
- C. show ip route ospf
- D. show ip ospf database summary
- E. show ip route

Answer: A

Explanation: Explanation

The "show ip ospf database external" command displays information only about external LSAs (Type 5 LSAs).

Below is an example of the "show ip ospf database external" command

Notice the line "LS Type: AS External Link", which means LSA Type 5. For your information, the "Link State ID: 143.105.0.0 indicates the network being advertised; the "Advertising Router: 10.187.70.6 indicates the router that originated this LSA.

Note:

+ The "show ip ospf database summary" command displays information only about the summary LSAs. + The "show ip ospf database nssa-external" command displays information only about the not so stubby area (NSSA) external LSAs.

```
Router# show ip ospf database external
OSPF Router with id(192.168.239.66) (Autonomous system 300)

    Displaying AS External Link States

LS age: 280
Options: (No TOS-capability)
LS Type: AS External Link
Link State ID: 143.105.0.0 (External Network Number)
Advertising Router: 10.187.70.6
LS Seq Number: 80000AFD
Checksum: 0xC3A
Length: 36
Network Mask: 255.255.0.0
    Metric Type: 2 (Larger than any link state path)
    TOS: 0
    Metric: 1
    Forward Address: 0.0.0.0
    External Route Tag: 0
```

QUESTION NO: 173

Which three are advantages to creating multiple areas in OSPF? (Choose three.)

- A. less frequent SPF calculations
- B. fewer hello packets
- C. smaller routing tables
- D. reduced LSU overhead
- E. fewer adjacencies needed

Answer: A,C,D

Explanation: Explanation

OSPF routers within an area only need to know about other routers within their own area, not outside their area, and all OSPF routers within a given area share the same link state database. This keeps the routing tables small enough to prevent processing bottlenecks from occurring. Also SPF only needs to calculate paths to routers within that area.

If a router receives an LSA with old information then it will send a LSU to the sender to update the sender with the newer information. The Link State Update (LSU) holds the LSAs. Instead of sending multiple LSUs the ABR / ASBR summarizes a route and sends only one LSU-.

Note: The LSA has a 30 minute timer that causes the router to send an LSU to everyone on the network once it ages out.

QUESTION NO: 174

What are two Cisco IOS commands that can be used to view neighbor adjacencies? (Choose two.)

- A. show ip ospf database
- B. show ip ospf neighbors
- C. show ip ospf protocols
- D. show ip ospf interfaces

Answer: B,D

Explanation: Explanation

The output of these commands are shown below:

```
R1# show ip ospf neighbor fa0/0
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.2	1	INIT/-	00:00:35	10.1.1.2	Fast Ethernet0/0


```
RouterA# show ip ospf int s1/0
Serial 1/0 is up, line protocol is up
Internet Address 10.100.3.253/30, Area 1
Process ID 1, Router ID 1.1.1.1, Network Type BROADCAST, Cost:64
Transmit Delay is 1 sec, State DROTHER, Priority 0
No designated router on this network
No backup designated router on this network
Old designated Router (ID) 2.2.2.2, Interface address 10.100.3.254
Flush timer for old DR LSA due in 00:01:12
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  oob-resync timeout 40
  Hello due in 00:00:01
Supports Link-Local Signaling (LLS)
Index 1/2, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 1, maximum is 3
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 1, Adjacent neighbor count is 1
Suppress hello for 0 neighbor(s)
```

Notice that for the “show ip ospf interface” command, the “Neighbor Count” is the number of OSPF

neighbors discovered on this interface while the "Adjacent neighbor count" is the number of routers running OSPF that are fully adjacent with this router. Adjacent means that their databases are fully synchronized. In this example, this router has one neighbor on its Ethernet0 interface.

QUESTION NO: 175

Refer to the exhibit. All routers have simultaneously been reloaded, and the DR election has concluded as expected. Which state is RTC in?



RTD# show ip ospf neighbor					
Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.2	1	output omitted	00:00:37	192.168.1.2	Ethernet0
192.168.1.1	1	output omitted	00:00:38	192.168.1.1	Ethernet0
192.168.1.3	1	output omitted	00:00:32	192.168.1.3	Ethernet0

- A. 2WAY/DROTHER
- B. 2WAY/BDR
- C. 2WAY/DR
- D. FULL/DROTHER
- E. FULL/BDR
- F. FULL/DR

Answer: E

Explanation:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094059.shtml

All the routers are allowed to take part in the DR/BDR election because all the "Pri"s (Priority) are

1 and they are equal so with default parameters, the router with highest IP address will become DR (RTD in this case) and the router with second highest IP address (RTC) will become BDR.

QUESTION NO: 176

A network administrator is troubleshooting a redistribution of RIP routes into OSPF. Given the exhibited configuration commands, which statement is true?

```
router rip
  network 10.0.0.0
!
router ospf 5
  network 172.10.0.0 0.0.255.255 area 0
  redistribute rip
```

- A. Redistributed routes will be tagged as external type 1 (E1) with a metric of 30.
- B. Redistributed routes will be tagged as external type 2 (E2) with a metric of 20.
- C. Redistributed routes will maintain their original RIP routing metric.
- D. Redistributed routes will have a default metric of 0 and will be treated as unreachable and not advertised.
- E. Redistributed routes will have a default metric of 0 but will not be treated as reachable and will be advertised.

Answer: B

Explanation: Explanation

By default, all routes redistributed into OSPF will be tagged as external type 2 (E2) with a metric of 20, except for BGP routes (with a metric of 1).

Note: The cost of a type 2 route is always the external cost, irrespective of the interior cost to reach that route. A type 1 cost is the addition of the external cost and the internal cost used to reach that route.

QUESTION NO: 177

Which two statements are true of the OSPF link-state routing protocol? (Choose two.)

- A.** Using the Bellman-Ford algorithm, each OSPF router independently calculates its best paths to all destinations in the network.
- B.** Using the DUAL algorithm, each OSPF router independently calculates its best paths to all destinations in the network.
- C.** OSPF sends summaries of individual link-state entries every 30 minutes to ensure LSDB synchronization.
- D.** OSPF sends triggered updates when a network change occurs.
- E.** OSPF sends updates every 10 seconds.
- F.** When a link changes state, the router that detected the change creates a link-state advertisement (LSA) and propagates it to all OSPF devices using the 224.0.0.6 multicast address.

Answer: C,D


Explanation:

The point of this question is the basis of OSPF.

Incorrect answer A. OSPF send hello packets every 10 seconds, not the updates, OSPF sends triggered updates when a network change occurs. For OSPF, D Rather use the multicast address 224.0.0.6 to send packets to DR and BDR, only DR and BDR can get the information from this multicast address.

QUESTION NO: 178

Refer to the exhibit.



Which statement is true about the configuration?

- A. RTA will not establish an OSPF adjacency with RTB.
- B. RTA will not accept OSPF hello packets from RTB.
- C. RTA will send OSPF hello packets, but will not send OSPF updates.
- D. RTA will send OSPF updates, but will not establish an adjacency with RTB.

Answer: A

Explanation: Explanation


Passive-interface command is used to disable sending updates out from a specific interface. For OSPF it prevents Hello packets from being sent out or received through the interface. This will not allow to create a neighbor adjacency and prevent a router from learning prefixes from those neighbors.

For RIP, configuring an interface as passive will still allow the interface to listen to incoming routing updates for other RIP neighbors, but the interface will no longer send them.

Note: The passive-interface command is used in router configuration mode, not interface mode.

QUESTION NO: 179

Refer to the exhibit.



Which statement is true?

- A. Router RTA is directly connected to interface 192.168.45.1.

- B. Neighbor 192.168.45.1 has changed its OSPF priority number.
- C. Router RTA and neighbor 192.168.45.2 are exchanging OSPF LSAs.
- D. Router RTA is the BDR.


Answer: D

Explanation: Explanation

There are only 3 routers in this segment. From the output we learn that RTB (10.0.0.2) is the DR and RTC (10.0.0.1) is not the DR or BDR so we can deduce RTA is the BDR.

QUESTION NO: 180

Refer to the exhibit.



OSPF has been configured on all routers in the network. However, router R1 does not receive a default route to router R2 as intended. Which configuration change would ensure that R1 would receive a default route from R2?

- A. Add the area 1 stub command on routers R1.
- B. Add the always keyword to the default-information originate configuration command on router R2.

- C.** Remove the default information originate configuration command from router R2 and place it on router R1.
- D.** Add the ip route 5.0.0.0 255.255.255.0 0.0.0.0 command to router R2.

Answer: B


Explanation:

The point of this question is about advertise a default route into an ospf area.

There are two ways to advertise a default route into an ospf area. The first is to advertise 0.0.0.0 into the OSPF domain, provided the advertising router already has a default route. The second is to advertise 0.0.0.0 regardless of whether the advertising router already has a default route. The second method can be accomplished by adding the keyword always to the default-information originate command. In this example, the "always" keyword would be required in order to propagate this default route.

QUESTION NO: 181

Refer to the exhibit.



OSPF has been configured on all routers in the network and Area 1 has been configured as a NSSA. Which statement is true about the NSSA Area 1?

- A.** Redistributed RIP and IGRP routes will appear in Area 1. They will be advertised via type 5 LSAs.
- B.** Only redistributed RIP routes will appear in Area 1. They will be advertised via type 7 LSAs.
- C.** Only redistributed IGRP routes will appear in Area 1. They will be advertised via type 7 LSAs.
- D.** No redistributed routes can appear in Area 1, only summary routes.

Answer: C

Explanation: Explanation

Same explanation of Question 6. Please notice that the IGRP routes are redistributed to NSSA area 1 via LSA Type 7. The NSSA ABR of area 1 will convert it into a LSA Type 5 before flooding to area 0 & area 2.

QUESTION NO: 182

According to RFC 2328, what is the stateful order in which an OSPF router transitions to a full adjacency with a neighbor router?

- A. Down, Init, 2-Way, Exstart, Exchange, Loading, and Full
- B. Down, Init, 2-Way, Exchange, Exstart, Loading, and Full
- C. Down, 2-Way, Init, Loading, Exstart, Exchange, and Full
- D. Down, 2-Way, Init, Exchange, Exstart, Loading, and Full
- E. Down, Init, 2-Way, Loading, Exstart, Exchange, and Full
- F. Down, 2-Way, Init, Exstart, Exchange, Loading, and Full

Answer: A

Explanation: Explanation

When OSPF adjacency is formed, a router goes through several state changes before it becomes fully adjacent with its neighbor. The states are Down -> Attempt (optional) -> Init -> 2-Way -> Exstart -> Exchange -> Loading -> Full. Short descriptions about these states are listed below:

Down: no information (hellos) has been received from this neighbor.

Attempt: only valid for manually configured neighbors in an NBMA environment. In Attempt state, the router sends unicast hello packets every poll interval to the neighbor, from which hellos have not been received within the dead interval.

Init: specifies that the router has received a hello packet from its neighbor, but the receiving router's ID was not included in the hello packet

2-Way: indicates bi-directional communication has been established between two routers.

Exstart: Once the DR and BDR are elected, the actual process of exchanging link state information can start between the routers and their DR and BDR.

Exchange: OSPF routers exchange database descriptor (DBD) packets

Loading: In this state, the actual exchange of link state information occurs

Full: routers are fully adjacent with each other

(Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0e.shtml)

QUESTION NO: 183

You have been tasked with setting up OSPF on an existing company router using IPv6.

Which command enables OSPF for IPv6 on a router?

- A. `ipv6 router ospf process-id`
- B. `router ospf process-id`
- C. `router ospf ipv6 process-id`
- D. `ipv6 ospf process-id area-id`
- E. None of the above

Answer: A

Explanation:

To enable Open Shortest Path First (OSPF) for IPv6 router configuration mode, use the `ipv6 router ospf` command in global configuration mode.

`ipv6 router ospf process-id`

Syntax Description

process-id

Internal identification. It is locally assigned and can be a positive integer from 1 to 65535. The number used here is the number assigned administratively when enabling the OSPF for IPv6 routing process.

Reference:

http://www.cisco.com/en/US/products/ps6922/products_command_reference_chapter09186a00801d6615.html#wp2027608

QUESTION NO: 184

One of the most important characteristics of OSPF is multiple areas. Which statement best describes why this feature is such an important enhancement to earlier routing protocols?

- A. The network domain, when divided into areas, allows for the use of both IANA classful addressing and private addressing.
- B. The use of multiple areas allows for the use of prioritization.
- C. All computation is kept within the area, with minimum communication between the areas, allowing the network to scale to larger sizes.
- D. It is easier to implement security.

Answer: C

Reference: <http://www.ciscopress.com/articles/article.asp?p=31577> (OSPF within an area)

QUESTION NO: 185

When learning a new route, if a LSA received is not found in the topological database, what will an internal OSPF router do?

- A.** The sequence numbers are checked, and if the LSA is valid it is entered into the topology database.
- B.** The LSA is placed in the topological database and an acknowledgement is sent to the transmitting router.
- C.** The LSA is dropped and a message is sent to the transmitting router.
- D.** The LSA is flooded immediately out of all the OSPF interfaces, except the interface from which the LSA was received.

Answer: D

Explanation:

Internal routers flood the area with new LSA's, and ABR's create and flood each area they are in with LSA's. The sequencing happens at the 2-way/exstart stage, when exchanging the DBD.

QUESTION NO: 186

Which statement is correct based upon the following output from the show command on RT1?

```
RT1# show ipv6 ospf interface
FastEthernet0/0 is up, line protocol is up
Link Local Address FE80::218:B9FF:FE12:2CD1, Interface ID 4
Area 0, Process ID 1, Instance ID 0, Router ID 10.1.1.1
Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State Drother, Priority 1
Designated Router (ID) 10.1.3.1, local address FE80::218:B9FF:FECD:BEF0
Backup Designated router (ID) 10.1.2.1, local address FE80::218:B9FF:FE92:28D8
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:04
Index 1/3/3, flood queue lenght 0
Next 0x(0)/0x0(0)/0x0(0)
Last flood scan length is 2, maximum is 2
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 2, Adjacent neighbor count is 2
  Adjacent with neighbor 10.1.3.1 (Designated Router)
  Adjacent with neighbor 10.1.2.1 (Backup Designated Router)
Suppress help for 0 neighbor(s)
```


Choose one from below:

- A. OSPFv3 establishes neighbor adjacencies by using Link-local addresses.
- B. OSPFv3 establishes neighbor adjacencies by using IPv4 addresses.
- C. OSPFv3 establishes neighbor adjacencies by using global IPv6 addresses.
- D. RT1 owns a subnet mask of 64 bits.

Answer: A


Explanation:

The point of this question is the OSPFv3.

As shown information, we notice that the router with both ipv6 address and ipv4 address and both of them are in ospf process, so in OSPFv3, we can find that, the DR BDR address are all link local address, so in OSPFv3 it establish neighbor with link local address.

QUESTION NO: 187

Refer to the exhibit.



OSPF is configured on all routers in the network. Area 5 is configured as an NSSA area. The RIPv2 routes are redistributed into the OSPF domain on router R5. What two types of LSAs will be originated by router R5? (Choose two.)

- A. type 1 Router LSA
- B. type 2 Network LSA
- C. type 3 Network Summary LSA
- D. type 4 ASBR Summary LSA
- E. type 5 AS External LSA
- F. type 7 NSSA External LSA

Answer: A,F

Explanation:

The point of the question is the LSA types in different kinds of areas.

NSSA area has 4 kinds of LSA : 1, 2 , 3 ,7

There is no LSA3 because it is produced by ABR to advertise routing entries between OSPF areas. There is no LSA2 because there are serial interface but no fast Ethernet interface, no DR elected, no LSA2. LSA2 is produced by DR.

QUESTION NO: 188

A network administrator has enabled OSPF across an NBMA network and has issued the command `ip ospf network nonbroadcast`. Given those facts, which two statements are true? (Choose two.)

- A. DR and BDR elections will occur.
- B. DR and BDR elections will not occur.
- C. All routers must be configured in a fully meshed topology with all other routers.
- D. The neighbor command is required to build adjacencies.
- E. Interfaces will automatically detect and build adjacencies with neighbor routers.

Answer: A,D

Explanation:

Even if there is only one router, broadcast multiaccess networks elect a DR and a BDR to serve as focal points for routing information. In contrast, point-to-point OSPF networks do not elect a DR because they can never include more than two nodes.

Another type of OSPF network, Nonbroadcast Multiaccess (NBMA), can include more than two nodes. Therefore, NBMA will try to elect a DR and a BDR. Common NBMA implementations include Frame Relay, X.25, and SMDS. NBMA networks follow rules at Layer 2 that prevent the delivery of broadcasts and multicasts.

A non-broadcast environment requires that all OSPF neighbors be manually configured. This is the default setting for physical interfaces with Frame Relay encapsulation, as well as for their point-to-multipoint subinterfaces. By manually configuring each neighbor, OSPF knows exactly which neighbors need to participate and which neighbor is identified as the DR. Also, communication between neighbors is done via unicast instead of multicast. This configuration also requires a full mesh and has the same weakness as the broadcast environment.

For non-broadcast networks the default Hello interval is 30 seconds and the Dead interval is four

times the Hello interval, 120 seconds. Non-broadcast multi-access networks do elect a DR and BDR, due to their multi-access nature. In order to set which router you want as the DR, you must set the priority in the neighbor statement to elect the neighbor as the DR. In order to manually configure who your neighbors are, the following command must be entered in router configuration mode for the selected OSPF process:

`neighbor ip_address`


ip_address = the ip address of the neighbor.

If you would like to set the priority of this router to become the DR, you will need to append the priority of the neighbor:

`neighbor ip_address priority value`

QUESTION NO: 189

Refer to the topology diagram.



R3 is redistributing the EIGRP routes into OSPF. What will the EIGRP routes appear as in the routing table of R1?

- A. O
- B. O IA
- C. O E2
- D. D
- E. D EX

Answer: C

Reference: <http://netpxel.com/wiki2/doku.php/ospf-eigrp-route-redistribution>

QUESTION NO: 190

Refer to the exhibit.

```
R3#show ip ospf neighbor
Neighbor ID    Pri   State           Dead Time   Address        Interface
192.168.0.4    1     FULL/DR         00:01:46   116.16.34.4   Serial1/0.1
192.168.0.5    1     FULL/BDR        00:01:46   116.16.35.5   Serial2/0.1
R3#
R3#show ip route ospf
140.140.0.0/32 is subnetted, 3 subnets
O IA   140.140.1.1 [110/65] via 116.16.34.4, 01:21:23, Serial1/0.1
O IA   140.140.3.1 [110/65] via 116.16.34.4, 01:21:23, Serial1/0.1
O IA   140.140.2.1 [110/65] via 116.16.34.4, 01:21:23, Serial1/0.1
192.168.0.0/32 is subnetted, 2 subnets
O IA   192.168.0.4 [110/65] via 116.16.34.4, 01:21:23, Serial1/0.1
```

Based on the command output, what is one reason why no routes from the OSPF neighbor 192.168.0.5 are installed in the IP routing table?

- A. R3 will only install routes from the neighbor with the lowest priority (Pri). If routes have the same priority, routes from the neighbor with the lowest IP address are used.
- B. R3 did not receive any LSAs from 192.168.0.5.
- C. Routes from backup designated routers are never installed in the IP routing table.
- D. 192.168.0.5 is a redundant link to 192.168.0.4, and load balancing is not enabled.


Answer: B

Explanation: Explanation

R3 may not receive any LSAs from neighbor 192.168.0.5 because a distribute-list blocks it. But notice that the LSAs are not filtered out in the LSDB since all routers in an OSPF area must be the same (synchronized).

QUESTION NO: 191

Refer to the exhibit.



What additional commands should be used to configure OSPF area 5 as a Totally Stubby area?

- A. area 0 stub on routers R4 and R5
- B. area 5 stub on routers R4 and R5
- C. area 5 stub no-summary on routers R4 and R5
- D. area 0 stub no-summary on router R4 and area 5 stub no-summary on router R5
- E. area 5 stub no-summary on router R4 and area 5 stub on router R5

Answer: E

Explanation: Explanation

To define a totally stub area, use the **area area-id stub no-summary** command on the ABR (in OSPF router configuration) and the **area area-id stub** on the totally stub router. The ABR will inject a default route into the area so routers in this type of area only see routing information local to their area, plus a default route pointing to the ABR, from which they can reach all other areas and all other networks

QUESTION NO: 192

Refer to the partial configurations in the exhibit.

```
Router1#show run

**** output omitted ****

interface serial1/1
  ipv6 address 2001:410:FFFF: 1::64/64
  ipv6 ospf 100 area 0
!
interface serial2/0
  ipv6 address 3FFF:B00:FFFF:1::2/64
  ipv6 ospf 100 area 0
!
ipv6 router ospf
  router-id 10.1.1.3
```

What address is utilized for DR and BDR identification on Router1?

- A. the configured router-id address
- B. an arbitrarily generated internal address
- C. the sequential 1/1 address
- D. the sequential 2/0 address

Answer: A

Explanation:

The point of this question is DR and BDR election of OSPFv3.

OSPF test the title of the DR and BDR election.

DR and BDR election there are two standards:

1. Interface priority, the default is 1.
2. Router-ID. manually specify> Highest loopback interface> Highest active IP address

This title, there is no priority list, it is the default, manually specify the Router-ID. 10.1.1.3, so hand-Router-id specified in the election has become the standard DR and BDR.

The selection of DR and BDR is the same between OSPFv2 and OSPFv3.

QUESTION NO: 193

You are developing a verification plan for an upcoming OSPF implementation. Part of this plan is to verify the status of type 3 LSAs within the network. Which routers should you verify first to ensure that the configurations are correct for generating type 3 LSAs?

- A. Internal routers within the backbone area (area 0)
- B. Internal routers within the NSSAs
- C. Internal routers within the stubby areas
- D. ASBRs
- E. ABRs
- F. DRs and BDRs

Answer: E

Explanation: Explanation

Type 3 LSA (Summary LSA) is advertised by the ABR of originating area to advertise network from other areas so we should check the ABRs first.

QUESTION NO: 194

Refer to the exhibit. Based on this command output, what can we conclude about R3?

```
R3#show ip ospf database

      OSPF Router with ID (172.16.1.1) (Process ID 123)
      Router Link States (Area 0.0.0.0)
  Link ID  ADV Router  Age      Seq#       Checksum Link count
  172.16.1.1 172.16.1.1   128      0x80000002 0x00D748 1
  192.168.0.4 192.168.0.4  128      0x8000000c 0x000EA9 1
      Net Link States (Area 0.0.0.0)
  Link ID  ADV Router  Age      Seq#       Checksum
  116.16.34.4 192.168.0.4  128      0x80000003 0x00CC36
      Summary Net Link States (Area 0.0.0.0)
  Link ID  ADV Router  Age      Seq#       Checksum
  116.16.35.0 172.16.1.1   245      0x80000001 0x00E4C0
  140.140.1.1 192.168.0.4  728      0x80000001 0x00A013
  140.140.2.1 192.168.0.4  728      0x80000001 0x00951D
  140.140.3.1 192.168.0.4  728      0x80000001 0x008A27
  192.168.0.4 192.168.0.4  728      0x80000001 0x0095CB
      Router Link States (Area 2)
  Link ID  ADV Router  Age      Seq#       Checksum Link count
  172.16.1.1 172.16.1.1   127      0x80000002 0x00F725 1
  192.168.0.5 192.168.0.5  128      0x8000000E 0x00258B 1
      Net Link States (Area 2)
  Link ID  ADV Router  Age      Seq#       Checksum
  116.16.35.5 192.168.0.5  128      0x80000003 0x00BB43
      Summary Net Link States (Area 2)
  Link ID  ADV Router  Age      Seq#       Checksum
  116.16.34.0 172.16.1.1   245      0x80000001 0x00EFB6
  140.140.1.1 172.16.1.1   124      0x80000001 0x00A77A
  140.140.2.1 172.16.1.1   124      0x80000001 0x009C84
  140.140.3.1 172.16.1.1   124      0x80000001 0x00918E
  192.168.0.4 172.16.1.1   124      0x80000001 0x009C33
```

- A. R3 is an ABR.
- B. R3 is not connected to the backbone.
- C. R3 has four neighbors.
- D. R3 is the DR for area 2.

Answer: A

Explanation: Explanation

From the output we learned that this router is getting Type 3 LSAs (Summary Net Link States) for both Area 0 and Area 3 -> It is an ABR between Area 0 & Area 3;

From the LSA Type 1 we learn that R3 only has 1 neighbor, which is 192.168.0.4

To find the DR for an area, we need to look at the LSA Type 2 of that area. In this case we find out 116.16.35.5 is the DR for area 2, not R3.

To help you understand more clearly about the "show ip ospf database" command, we want to explain more about the fields in the output:

```
R3#show ip ospf database
```

OSPF Router with ID (172.16.1.1) (Process ID 123)					
Router Link states (Area 0.0.0.0) LSA Type 1 (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	Link count
172.16.1.1	172.16.1.1	128	0x80000002	0x00D748	1
192.168.0.4	192.168.0.4	128	0x8000000C	0x000EA9	1
Net Link states (Area 0.0.0.0) LSA Type 2 (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	
116.16.34.4	192.168.0.4	128	0x80000003	0x00CC36	
Summary Net Link States (Area 0.0.0.0) LSA Type 3 (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	
116.16.34.0	172.16.1.1	245	0x80000001	0x00B4C0	
140.140.1.1	192.168.0.4	728	0x80000001	0x00A013	
140.140.2.1	192.168.0.4	728	0x80000001	0x00951D	
140.140.3.1	192.168.0.4	728	0x80000001	0x008A27	
192.168.0.4	192.168.0.4	728	0x80000001	0x0095CB	
Router Link States (Area 2) LSA Type 1 (Area 2)					
Link ID	ADV Router	Age	Seq#	Checksum	Link count
172.16.1.1	172.16.1.1	127	0x80000002	0x00F725	1
192.168.0.5	192.168.0.4	128	0x8000000B	0x00258B	1
Net Link States (Area 2) LSA Type 2 (Area 2)					
Link ID	ADV Router	Age	Seq#	Checksum	
116.16.35.5	192.168.0.5	128	0x80000003	0x00BB43	
Summary Net Link States (Area 2) LSA Type 3 (Area 2)					
Link ID	ADV Router	Age	Seq#	Checksum	
116.16.34.0	172.16.1.1	245	0x80000001	0x00EFB6	
140.140.1.1	172.16.1.1	124	0x80000001	0x00A77A	
140.140.2.1	172.16.1.1	124	0x80000001	0x009C84	
140.140.3.1	172.16.1.1	124	0x80000001	0x00918E	
192.168.0.4	172.16.1.1	124	0x80000001	0x009C33	

+ **Link ID** is the OSPF Router-ID of a router in the area for LSA Type 1 & 2 but it is can be the Router-ID or the network address for LSA Type 3, 5 & 7.+ The **ADV Router** is the ID of the router that sent the LSA (Advertising Router) into the area+ **Age**: The max age of the link state+ **Seq#** and **Checksum**: these fields are used to verify link-state integrity.

For example, from the Router Link States (Area 0.0.0.0):

OSPF Router with ID (172.16.1.1) (Process ID 123)					
Router Link states (Area 0.0.0.0) LSA Type 1 (Area 0)					
Link ID	ADV Router	Age	Seq#	Checksum	Link count
172.16.1.1	172.16.1.1	128	0x80000002	0x00D748	1
192.168.0.4	192.168.0.4	128	0x8000000C	0x000EA9	1

The first "Link ID" 172.16.1.1 is the Router-ID of the local router R3 (because it is the same as

“OSPF Router with ID”). This “Link ID” is, of course, advertised by itself so the ADV Router has the same value.

The second “Link ID” is the Router-ID of 192.168.0.4, which is advertised by 192.168.0.4 so it is directly connected to R3.

Notice that these 2 routers belong to Area 0.

From the Net Link States:

Net Link states (Area 0.0.0.0)				
Link ID	ADV Router	Age	Seq#	Checksum
116.16.34.4	192.168.0.4	128	0x80000003	0x00CC36

We learn that the “Link ID” 116.16.34.4 is advertised from the neighbor 192.168.0.4, which is directly connected to R3. This router (116.16.34.4) also belongs to Area 0 and it is the DR of that segment.

The Summary Net Link States gives us information about LSA Type 3 (advertised by the ABR of area 0. Recall that ABRs generate a Type 3 LSA for each subnet in one area, and advertises each Type 3 LSA into the other areas)


Summary Net Link States (Area 0.0.0.0)				
Link ID	ADV Router	Age	Seq#	Checksum
116.16.34.0	172.16.1.1	245	0x80000001	0x00B4C0
140.140.1.1	192.168.0.4	728	0x80000001	0x00A013
140.140.2.1	192.168.0.4	728	0x80000001	0x00951D
140.140.3.1	192.168.0.4	728	0x80000001	0x008A27
192.168.0.4	192.168.0.4	728	0x80000001	0x0095CB

From the output we learn that the subnet 116.16.35.0 is advertised via 172.16.1.1 and four routers (in fact, interfaces) are advertised by 192.168.0.4. Notice that these routers (interfaces) can belong to other areas.

QUESTION NO: 195

Refer to the exhibit. A new TAC engineer comes to you for advice. The engineer wants to configure RIPv2- OSPF two-way redistribution while avoiding routing loops.

Which two additions to the router B1 configuration should the engineer make? (Choose two.)



A. access-list 40 deny 172.16.1.0 0.0.0.255

access-list 40 permit any

router rip

redistribute ospf 100 metric 5

distribute-list 40 out ospf 100

B. ip prefix-list rip_routes permit 172.16.1.16/25 ge 26 le 28 route-map redis-ospf deny 10

match ip address prefix-list rip_routes

router rip

redistribute ospf 10 route-map redis-ospf subnets

C. ip prefix-list rip-to-ospf permit 10.1.1.8/25 ge 26 le 28 route-map redis-rip deny 20

match ip address prefix-list rip-to-ospf

router ospf 100

redistribute rip route-map redis-rip subnets

D. access-list 15 permit 10.1.1.0 0.0.0.63

access-list 15 permit any

route-map redis-rip deny 10

match ip address 15

route-map redis-rip permit 20

router ospf 100

redistribute rip route-map redis-rip subnets


Answer: A,D

Explanation: Explanation

B1 is not the only router that redistributes between RIP & OSPF. The "small" router below B1 can be configured for this task too so B1 can try to redistribute networks advertised by that "small" router again. Therefore it is necessary to filter out networks that have been advertised by the "small" router. For example, we need to prevent network 172.16.1.0/24 from advertised back into RIPv2 or network 10.1.1.0/26 from advertised back into OSPF. Notice that all networks in OSPF domain (including 10.1.1.8/30, 10.1.1.12/30, 10.1.1.48/28, 10.1.1.32/28) can be summarized as 10.1.1.0/26 and all networks in RIP domain (including 172.16.1.24/30, 172.16.1.20/30, 172.16.1.32/28, 172.16.1.48/28) can be summarized as 172.16.1.0/24 -> answers A & D are

correct.

In answer B, the command "ip prefix-list rip_routes permit 172.16.1.16/25 ge 26 le 28 means:
+ First check the first 25 bits of the address -> this will allow addresses from 172.16.1.0 to 172.16.1.127

172.16.1.16/25 = 1010 1100.0001 0000.0000 0001.0001 0000

 1010 1100.0001 0000.0000 0001.0000 0000 = 172.16.1.0
 1010 1100.0001 0000.0000 0001.0111 1111 = 172.16.1.127
 variable

+ If those match then check the subnet mask, which in this case can be GREATER THAN or EQUAL to 26 bits & LESS THAN or EQUAL to 28 bits -> meaning that /26, /27, /28 subnet masks would match.

For example, networks 172.16.1.0/26; 172.16.1.16/28 would match (but notice networks 172.16.1.0/25; 172.16.1.128/26 wouldn't).

In the "ip prefix-list rip_routes permit 172.16.1.16/25 ge 26 le 28, the prefix-list "rip_routes" only covers networks 172.16.1.32/28 & 172.16.1.48/28 but can't cover networks 172.16.1.24/30 & 172.16.1.20/30. Also, the OSPF process in the "redistribute" command should be 100, not 10. Same problem as answer B, the prefix-list in answer C can't cover networks 10.1.1.8/30 & 10.1.1.12/30 .

QUESTION NO: 196

what is the best order for choosing the Router-ID in ospf?

- A. high loopbak interface address,high physical interface address, Router ID command
- B. high loopbak interface address, Router ID command, high physical interface address
- C. Router ID command, high loopbak interface address, high physical interface address ||
- D. Router ID command, high physical interface address, high loopbak interface address
- E. high physical interface address, high loopbak interface address, Router ID command
- F. high physical interface address, Router ID command, high loopbak interface address

Answer: C

Explanation: The *Router ID (RID)* is an IP address used to identify the router in the OSPF database (LSDB). Cisco chooses the Router ID by using the highest IP address of all configured loopback interfaces. If no loopback interfaces are configured with addresses, OSPF will choose the highest IP address of all active physical interfaces

QUESTION NO: 197 DRAG DROP

Click and drag the associated set of OSPF LEAs on the left to the corresponding area type on the right where this set of LEAs may be seen.

LSA 1, 2, 3, 4, 5	Stub
LSA 1, 2, 3	NSSA
LSA 1, 2	Backbone or transit
LSA 1, 2, 3, 7	Totally NSSA
LSA 1, 2, 7	Totally Stubby

Answer:

Click and drag the associated set of OSPF LEAs on the left to the corresponding area type on the right where this set of LEAs may be seen.


LSA 1, 2, 3, 4, 5	LSA 1, 2, 3
LSA 1, 2, 3	LSA 1, 2, 3, 7
LSA 1, 2	LSA 1, 2, 3, 4, 5
LSA 1, 2, 3, 7	LSA 1, 2, 7
LSA 1, 2, 7	LSA 1, 2

Explanation:

LSA 1, 2, 3, 4, 5 - Backbone or transit
 LSA 1, 2, 3 - Stub
 LSA 1, 2 - Totally Stubby
 LSA 1, 2, 3, 7 - NSSA
 LSA 1, 2, 7 - Totally NSSA

QUESTION NO: 198

Refer to the exhibit. ROUTE.com is planning to implement a secure OSPF network to support traffic between clients on the 172.16.10.0/24 network and the file server on the 172.16.20./24 network. You have been asked to review the implementation and verification plans. Which statement about the plan is true?



1. On R1 and R5, verify end-to-end connectivity between prefixes 172.16.10.0 and 172.16.20.0 using the extended ping command.
2. Verify that all routers have established a full neighbor relationship with the appropriate neighbors.
3. Verify the proper authentication method is active for each neighbor relationship.
4. Verify that each prefix is assigned to the appropriate OSPF area.
5. Verify the appropriate routes and summaries are in the routing table of each router.
6. Verify that all redundant links will become active when the primary links are shut

- A.** It is complete as written.
- B.** It should include a task that verifies that the interarea routes are using the proper MED.
- C.** The plan should include a task that verifies that load sharing is active on the appropriate links.
- D.** The plan should include a task that verifies end-to-end connectivity between the clients and the file server.

Answer: D

Explanation:

The plan for secure OSPF network to support traffic between network and file server requires a task that verifies end-to-end connectivity between the clients and the file server.

QUESTION NO: 199

What is the sequence, from first to last, that OSPF will follow when choosing a router ID?

- A.** 1. the highest IP address on any active loopback interface; 2. The router-id command results; 3. The highest IP address of any active physical interface.
- B.** 1. the highest IP address on any active loopback interface; 2. The highest IP address of any active Physical interface; 3. The router-id command result.
- C.** 1. The highest IP address of any active Physical interface; 2. the highest IP address on any active loopback interface; 3. The router-id command results.

- D.** 1. The highest IP address of any active Physical interface; 2. The router-id command result; 3. The highest IP address of any active physical interface.
- E.** 1. The router-id command results; 2. The highest IP address of any active Physical interface; 3. The highest IP address of any active physical interface.
- F.** 1. the router-id command result; 2. the highest IP address on any active loopback interface; 3. the highest IP address of any active physical interface

Answer: F

Explanation:

1. The router chooses the numerically highest IP address on any of its loopback interfaces.
2. If no loopback interfaces are configured with IP addresses, the router chooses the numerically highest IP address on any of its physical interfaces. The interface from which the Router ID is taken does not have to be running OSPF.

Using addresses associated with loopback interfaces has two advantages:

The loopback interface is more stable than any physical interface. It is active when the router boots up, and it only fails if the entire router fails.

The network administrator has more leeway in assigning predictable or recognizable addresses as the Router IDs.

Cisco's OSPF will continue to use a Router ID learned from a physical interface even if the interface subsequently fails or is deleted. Therefore, the stability of a loopback interface is only a minor advantage. The primary benefit is the ability to control the Router ID.

The OSPF router begins a neighbor relationship by advertising its Router ID in Hello packets.

QUESTION NO: 200 DRAG DROP

Click the resources on the left that you need to create an implementation plan for an OSPF project and drag them to the target zone on the right.

local preference value	Required Resources
update query boundaries	
summarization boundaries	
OSPF process ID that will be used	
authentication type that will be used	
OSPF areas and associated prefixes	
load sharing method that will be used	
amount of bandwidth that will be allocated to OSPF	

Answer:

Click the resources on the left that you need to create an implementation plan for an OSPF project and drag them to the target zone on the right.

Available Resources	Required Resources
local preference value	summarization boundaries
update query boundaries	OSPF process ID that will be used
summarization boundaries	authentication type that will be used
OSPF process ID that will be used	OSPF areas and associated prefixes
authentication type that will be used	
OSPF areas and associated prefixes	
load sharing method that will be used	
amount of bandwidth that will be allocated to OSPF	

Explanation:

Required Resources
summarization boundaries
OSPF process ID that will be used
authentication type that will be used
OSPF areas and associated prefixes

QUESTION NO: 201 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 routes Chicago and New York. The S0/0 interface on Chicago and the S0/1 interface on New York are in area 0. The loopback0 on New York is in Area 1. However, they cannot ping from the serial interface of the satellite router to the loopback interface of the New York router. You have been asked to complete the implementation to allow this ping.

ROUTE.com's corporate implementation guideline requires:

- The OSPF process ID for all routers must be 24.
- The routing protocol for all each interface must be enabled under the routing process.
- The routing protocol must be enabled for each interface using the most specific wild card mask possible.
- The serial link between Seattle and Chicago must be in OSPF area 8.
- OSPF area 8 must not receive any inter-area or external routers.

Network information

Seattle

S0/0 192.168.34.5/30 – link between Seattle and Chicago

Secret Password: Cisco

Chicago

S0/0 192.168.40.9/30 – link between Chicago and New York


S0/0 192.168.34.6/30 – link between Chicago and New York

Secret Password: Cisco

Name : Seattle
S0/0 : 192.168.38.5/30
Secret Password : cisco

Name : Chicago
S0/0 : 192.168.76.9/30
S0/1 : 192.168.38.6/30
Secret Password : cisco

Name : Orlando
S0/1 : 192.168.76.10/30
Loopback0 : 172.16.101.101/32



Answer: Here is the solution:

On Seattle Router:

Router ospf 24

network 192.168.34.5 0.0.0.0 area 8

area 8 stub

On Chicago Router:

Router ospf 24

area 8 stub no-summary


network 192.168.34.6 0.0.0.0 area 8

Explanation:

You have to configure area 8 stub on seattle router and no-summary on Chicago router. This will allow the ping to go through.

QUESTION NO: 202

Refer to the exhibit.



OSPF is running throughout the network and RIP is being redistributed into OSPF in R2. You want to minimize the propagation of LSAs into and out of Area 1.

Which OSPF feature would be best to achieve this goal?

- A. stub
- B. totally stubby
- C. NSSA
- D. totally NSSA

Answer: D

Explanation:

We need to redistribute RIP from R1 to Area 1 so Area 3 cannot be a stub or totally stubby area. To minimize the propagation of LSAs into and out of Area 1 we should configure it as a totally NSSA. Notice that a NSSA allows LSA Type 3 & 7 while a Totally NSSA only allows LSA Type 7. Note: Both Totally Stubby Area & Totally Stubby NSSA do not accept external AS routes or inter-area routes (LSA Types 3, 4 and 5). They recognize only intra-area routes and the default route 0.0.0.0. The main difference between them is Totally Stubby NSSA accepts routes from other AS while Totally Stubby Area does not.


Below summarizes the LSA Types allowed and not allowed in area types:

Area Type

Type 1 & 2 (within area)
Type 3 (from other areas)
Type 4
Type 5
Type 7
Standard & backbone
Yes
Yes
Yes
Yes
No
Stub
Yes
Yes
No
No
No
Totally stubby
Yes
No
No
No
No
NSSA
Yes
Yes
No
No
Yes
Totally (stubby) NSSA
Yes
No
No
No
Yes

QUESTION NO: 203

Refer to the exhibit.



What must be configured on router RTA to summarize all routes from area 0 to area 1?

- A. area 0 range 172.16.96.0 255.255.224.0
- B. area 0 range 172.16.96.0 255.255.255.0
- C. area 1 range 172.16.96.0 255.255.224.0
- D. area 1 range 172.16.96.0 255.255.0.0
- E. summary-address 172.16.96.0 255.255.224.0
- F. summary-address 172.16.96.0 0.0.63.255

Answer: A

Explanation:

Recall that summarization is the consolidation of multiple routes into one single, supernet advertisement. See Module 2 for more details on this. Proper summarization requires contiguous, sequential, addressing. 200.10.0.0, 200.10.1.0, 200.10.2.0, and so on are examples of contiguous addressing. OSPF routers can be manually configured to advertise a supernet route, which is different from an LSA summary route.

Route summarization directly affects the amount of bandwidth, CPU, and memory resources that are consumed by the OSPF process. With summarization, if a network link fails or flaps, the topology change will not be propagated into the backbone, and other areas by way of the backbone. As discussed in previous modules, route summarization protects routers from needless routing table recalculations. Because the SPF calculation places a significant demand on a router CPU, proper summarization is an important part of OSPF configuration.

OSPF supports the following two types of summarization:

To configure an ABR to summarize routes for a specific area before injecting them into a different area, use the following syntax:

Router(config-router)#**area** *area-id* **range** *address mask*

To configure an ASBR to summarize external routes before injecting them into the OSPF domain, use the following syntax:

Router(config-router)#**summary-address** *address mask*

Use the following commands to configure RTA for external router summarization as shown in the Figure.

RTA(config)#**router ospf 1**RTA(config-router)#**summary-address 200.9.0.0 255.255.0.0**


Once configured, RTA will send only a single summary route, 200.9.0.0/16, into the OSPF domain. Because RTB sits on the border between Area 0 and Area 1, it should be configured to perform interarea summarization, shown as follows:

RTB(config)#**router ospf 1**RTB(config-router)#**area 1 range 192.168.16.0 255.255.252.0**

Notice that the **area 1 range** command in this example specifies the area containing the range to be summarized before being injected into Area 0.

QUESTION NO: 204

Refer to the exhibit.



All routers have simultaneously been reloaded and the DR election has concluded as expected. What state is RTB in?

- A. 2WAY/DROTHER
- B. 2WAY/BDR

- C. 2WAY/DR
- D. FULL/DROTHER
- E. FULL/BDR
- F. FULL/DR

Answer: D


Explanation: The point of this question is about the select principles of DR and BDR.

DR and BDR election is done via the Hello protocol. Hello packets are exchanged via IP multicast packets (Appendix B) on each segment. The router with the highest OSPF priority on a segment will become the DR for that segment. The same process is repeated for the BDR. In case of a tie, the router with the highest RID will win. The default for the interface OSPF priority is one.

Remember that the DR and BDR concepts are per multiaccess segment. Setting the ospf priority on an interface is done using the ip ospf priority <value> interface command. In this case, all routers have the same priority, but RTD has the highest RID, and RTC was followed, so RTC was the BDR.

QUESTION NO: 205

Given the above OSPF network, which command will RTB use to summarize routes for the 192.168.16.0/22 supernet before injecting them into Area 0?



- A. area 10 range 192.168.16.0 255.255.252.0

- B. summary-address 192.168.16.0 255.255.252.0
- C. ip summary-address ospf 101 192.168.16.0 255.255.252.0
- D. area 0 range 192.168.16.0 255.255.252.0
- E. ip summary-address area 0 192.168.16.0 255.255.252.0

Answer: A

Explanation:

The **area range** command is used only with Area Border Routers (ABRs) which is router RT2 in this example. It is used to consolidate or summarize routes for an area. The result is that a single summary route is advertised to other areas by the ABR. Routing information is condensed at area boundaries. External to the area, a single route is advertised for each address range. Only Choice A specifies the correct syntax and route summarization network mask.

QUESTION NO: 206

To create an NSSA totally stubby area in Area 1, what commands should be configured on the NSSA ABR?

- A. router ospf 1
area 1 nssa
- B. router ospf 1
area 1 nssa no-summary
- C. router ospf 1
area 1 nssa no-redistribution
- D. router ospf 1
area 1 nssa default-information originate
- E. router ospf 1
area 1 nssa default-information originate metric-type 2

Answer: B

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094a88.shtml (See the topic: defining an NSSA totally stub area)

QUESTION NO: 207

How is the configuration of a totally stubby area different from that of a stub area?

- A. The totally stubby area requires the no-summary command on the ABR.

- B. The totally stubby area requires the totally stubby command on the ABR.
- C. The no-summary command should be included on all routers within the totally stubby area.
- D. The totally stubby command should be included on all routers within the totally stubby area.
- E. The totally stubby area requires the no-summary command on the ASBR.

Answer: A

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094a74.shtml

QUESTION NO: 208

When the passive-interface command is used on a router, which two routing protocols will continue to receive routing updates on an interface that is configured as passive? (Choose two.)

- A. OSPF
- B. RIP
- C. RIPv2
- D. EIGRP

Answer: B,C

Explanation:

You can use the **passive-interface** command to control the advertisement of routing information. The command enables the suppression of routing updates over some interfaces while it allows updates to be exchanged normally over other interfaces.

With most routing protocols, such as RIP version 1 and 2, the **passive-interface** command restricts outgoing advertisements only. However, when used with Enhanced Interior Gateway Routing Protocol (EIGRP) and OSPF, the effect is slightly different. The use of the **passive-interface** command in EIGRP suppresses the exchange of hello packets between two routers, which results in the loss of their neighbor relationship. This stops not only routing updates from being advertised, but it also suppresses incoming routing updates.

The same is true for OSPF. To stop routers from becoming OSPF neighbors on a particular interface, issue the **passive-interface** command at the interface.


References:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0a.shtml

http://www.cisco.com/en/US/tech/tk365/technologies_q_and_a_item09186a0080094704.shtml

QUESTION NO: 209

Refer to the exhibit.



Which statement is true?

- A.** RTA will redistribute the RIP routes into the NSSA as type 7 LSAs. RTB will translate the type 7 LSAs into type 5 LSAs and flood them throughout the OSPF backbone.
- B.** RTA will redistribute the RIP routes into the NSSA as type 7 LSAs. RTB will flood the type 7 LSAs throughout the backbone.
- C.** RTA will redistribute the RIP routes into the NSSA as type 5 LSAs. RTB will flood the type 5 LSAs throughout the backbone.
- D.** RTA will redistribute the RIP routes into the NSSA as type 5 LSAs. RTB will translate the type 5 LSAs into type 7 LSAs and flood them throughout the OSPF backbone.
- E.** RTA will not redistribute the RIP routes into the NSSA.


Answer: A

Reference:

https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=15&ved=0CEUQFjAEOAo&url=http%3A%2F%2Ffaculty.valenciacollege.edu%2Fwyousif%2FCCNP%2FSemester5%2FPresentations%2FMAOSPF_P2.ppt&ei=YaGsUcL9G4mthQez4oD4Cw&usg=AFQjCNE5mLCAUIWCzou_vUX_DGhOOwcYxw&sig2=uaCCrMUuvgeL-Zwbf_dCkg&bvm=bv.47244034,d.ZG4&cad=rja (see slide 25, 26 and 27)

QUESTION NO: 210

Refer to the exhibit.



Given the exhibited router output, which command sequence can be added to R1 to generate a default route into the OSPF domain?

- A. default-router
- B. ip default-network
- C. default-information originate always
- D. ip default-gateway

Answer: C

Explanation:

QUESTION NO: 211

Look at the following configuration below and select the statements that are true.

RD2(config)#router ospf 100

RD2(config-router)# distance 180 10.1.10.1 0.0.0.0 match-me

RD2(config-router)#!

RD2(config-router)#ip access-list standard match-me

RD2(config-std-nacl)# permit host 172.16.0.0

Choose two.

- A. OSPF will have an AD of 180 for routes from 10.1.10.1 in the 172.16.0.0/16 range.
- B. The neighbor RID is 10.1.10.1
- C. OSPF will have an AD of 180 for all routes.
- D. The local RID is 10.1.10.1

Answer: A,B

Explanation:

If you see the configuration, the OSPF will be having an AD of 180 routers from 10.1.10.1 in the range permitted by the host. Since the distance 180 is configured the neighbor RID will be 10.1.10.1.

QUESTION NO: 212

Which of the following must match for two directly connected routers running OSPF to establish a neighbor adjacency? (Select 2 choices.)

- A. area IDs
- B. process IDs
- C. router IDs
- D. hello timers

Answer: A,D

Explanation:

Hello packets are sent every 10 seconds by default. In order for OSPF routers to establish neighbor adjacencies and exchange routing information successfully, the hello interval needs to match all OSPF routers in the OSPF area.

The area number is a number from 0-255, declared at the end of the network command after the wildcard bits. Routers in the same area will exchange routing information or Link State Updates or LSUs

QUESTION NO: 213

Which of the following is an OSPF configuration parameter that is used on an ABR, but not on an internal router?

- A. A virtual link to area 0.
- B. OSPF summarization command.

- C. default-cost extension to the area command.
- D. no-summary extension to the area stub command.
- E. None of the other alternatives apply

Answer: D

Explanation:


The no-summary extension of the area stub command is used only for ABRs connected to totally stubby areas. It prevents an ABR from sending summary link advertisements into the stub area. This option is used for creating a totally stubby area.

Topic 3, Implement an eBGP based solution, given a network design and a set of requirements

QUESTION NO: 214 CORRECT TEXT

(Policy Based Routing Sim)

Company Acne has two links which can take it to the Internet. The company policy demands that you use web traffic to be forwarded only to Frame Relay link and other traffic can go through any links that are available. Use BGP attributes to solve the requested action.



Answer: 1) Create an Access list that catches the HTTP traffic:


```
BorderRouter#access-list 101 permit tcp any any eq www
```

Note that the server was not directly connected to the Border Router. There were a lot of EIGRP routes on it. In the real exam you do not know the exact IP address of the server in the EIGRP network so we have to use the source as "any" to catch all the source addresses.

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

```
BorderRouter(config)#route-map pbr permit 10
```

```
BorderRouter(config-route-map)#match ip address 101
```

```
BorderRouter(config-route-map)#set ip next-hop 10.1.101.1
```

```
BorderRouter(config-route-map)#exit
```

```
BorderRouter(config)#route-map pbr permit 20
```

(Notice. the route-map pbr permit 20 line allows other traffic than HTTP to be routed. Otherwise, other traffic will be dropped)

3) Apply the route-map on the interface to the server in the EIGRP Network:

```
BorderRouter(config-route-map)#exit
```

```
BorderRouter(config)#int fa0/0
```

```
BorderRouter(config-if)#ip policy route-map pbr
```

```
BorderRouter(config-if)#exit
```

```
BorderRouter(config)#exit
```

4) There is a "Host for Testing", click on this host to open a box in which there is a button named "Generate HTTP traffic". Click on this button to generate some packets for HTTP traffic. Jump back to the BorderRouter and type the command "show route-map".

```
BorderRouter#show route-map
```

In the output you will see the line "Policy routing matches: 9 packets...". It means that the route-map we configured is working properly.


Explanation:

First the access list was implemented so that the traffic can be conditioned. Then the router map is configured to set the next hop address to ISP1 and the rest of the traffic is permitted to pass.

Since there is EIGRP network, route map is applied on the interface to the server in the EIGRP network.

QUESTION NO: 215

Refer to the exhibit.



```
RTR# debug ip bgp
```

```
*Mar 1 03:27:00.879: %TCP-6-BADAUTH: No MD5 digest from 10.0.0.1:179 to 10.0.0.2:11032
*Mar 1 03:27:00.883: %TCP-6-BADAUTH: No MD5 digest from 10.0.0.1:179 to 10.0.0.2:11032
*Mar 1 03:27:04.311: %TCP-6-BADAUTH: No MD5 digest from 10.0.0.1:11022 to 10.0.0.2:179
*Mar 1 03:27:12.315: %TCP-6-BADAUTH: No MD5 digest from 10.0.0.1:11022 to 10.0.0.2:179
*Mar 1 03:28:39.659: %TCP-6-BADAUTH: Invalid MD5 digest from 10.0.0.5:11026 to 10.0.0.6:179
*Mar 1 03:28:41.659: %TCP-6-BADAUTH: Invalid MD5 digest from 10.0.0.5:11026 to 10.0.0.6:179
```

Router RTR is attempting to establish BGP neighbor relationships with routers RT1 and RT3. On the basis of the information that is presented in the exhibit, which two statements are true? (Choose two)

- A. RTR has a BGP password set but neighbor 10.0.0.1 does not
- B. RTR has a BGP password set but neighbor 10.0.0.5 does not
- C. RTR has a BGP password set but neighbor 10.0.0.1 has an incorrect password set
- D. RTR has a BGP password set but neighbor 10.0.0.5 has an incorrect password set
- E. Neighbor 10.0.0.1 has a BGP password set but RTR does not
- F. Neighbor 10.0.0.5 has a BGP password set but RTR does not

Answer: A,D


Explanation:

You can configure MD5 authentication between two BGP peers, MD5 authentication must be configured with the same password on both BGP peers; otherwise, the connection between them will not be made. If a router has a password configured for a peer, but the other peer does not, a message "No MD5 digest from..." will appear on the console while the routers attempt to establish a Multicast Source Discovery Protocol (MSDP) session between them. Therefore A is correct because RT1 (with an ip address of 10.0.0.1) is not configured with a password.

Similarly, if the two routers have different passwords configured, a message "Invalid MD5 digest from..." will appear on the screen.

QUESTION NO: 216

Refer to the exhibit diagram and configuration. RTB is summarizing its networks from AS 64100 with the aggregate-address command. However, the show ip route command on RTA reveals the RTB individual networks as well as its summary route. Which option would ensure that only the summary route would appear in the routing table of RTA?



- A. Delete the four network statements and leave only the aggregate-address statement in the BGP configuration.
- B. Add the keyword summary-only to the aggregate-address command.
- C. Add a static route with a prefix of 192.168.24.0 255.255.252.0 pointing to the null0 interface.
- D. Create a route map permitting only the summary address.


Answer: B

Explanation:

When the aggregate-address command is used within BGP routing, the aggregated address is advertised, along with the more specific routes. The exception to this rule is through the use of the summary-only command. The "summary-only" keyword suppresses the more specific routes and announces only the summarized route.

QUESTION NO: 217

Refer to the exhibit.



The neighbor 10.1.1.1 weight 200 BGP configuration command has been configured on router A. What will be the result of this configuration?

- A. Router A will prefer the path through router B for network 172.20.0.0
- B. Router A will prefer the path through router C for network 172.20.0.0
- C. Packets from router D will prefer the path through router B for networks advertised by router A
- D. Packets from router D will prefer the path through router C for networks advertised by router A

Answer: A

Explanation:

The weight attribute is a special Cisco attribute that is used in the path selection process when there is more than one route to the same destination. The higher the weight value, the better the path. The default weight is 0. Therefore, by configuring weight 200 to the neighbor 10.1.1.1, Router-A will prefer the path through B for network 172.20.0.0 then the path through Router-C. The weight attribute is local to the router and not propagated to other routers. In this case the weight is local to Router-A so it has no effect on the decision of transferring packets from Router-D.

QUESTION NO: 218

Based on the show ip bgp summary output,

```

R1# show ip bgp summary
BGP table version is 8, main routing table version 8
4 network entries (8/12 paths) using 832 bytes of memory
5 BGP path attribute entries using 576 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
2 received paths for inbound soft reconfiguration

Neighbor      V    AS  MsgRcvd MsgSent  TblVer  InQ   OutQ   Up/Down   State/PfxRcd
10.1.1.1      4    50001   80      81       8     0     0    04:15:10       2
10.2.2.2      4    50002   79      81       0     0     0    00:00:15   Active
10.3.3.3      4    50003   82      82       0     0     0    02:00:00    Idle

```

Which two statements are true? (Choose two.)

- A. The BGP session to the 10.1.1.1 neighbor is established.
- B. The BGP session to the 10.2.2.2 neighbor is established.
- C. The BGP session to the 10.3.3.3 neighbor is established.
- D. The router is attempting to establish a BGP peering session with the 10.1.1.1 neighbor.
- E. The BGP session to the 10.3.3.3 neighbor is established, but the router has not received any BGP routing updates from the 10.3.3.3 neighbor.
- F. The router is attempting to establish a BGP peering session with the 10.2.2.2 neighbor.

Answer: A,F

Explanation:

The main point of this question is the "State/PfxRcd" column, which shows the BGP states. Below is the list of BGP states in order, from startup to peering:

- 1 – Idle:** the initial state of a BGP connection. In this state, the BGP speaker is waiting for a BGP start event, generally either the establishment of a TCP connection or the re-establishment of a previous connection. Once the connection is established, BGP moves to the next state.
- 2 – Connect:** In this state, BGP is waiting for the TCP connection to be formed. If the TCP connection completes, BGP will move to the OpenSent stage; if the connection cannot complete, BGP goes to Active
- 3 – Active:** In the Active state, the BGP speaker is attempting to initiate a TCP session with the BGP speaker it wants to peer with. If this can be done, the BGP state goes to OpenSent state.
- 4 – OpenSent:** the BGP speaker is waiting to receive an OPEN message from the remote BGP speaker
- 5 – OpenConfirm:** Once the BGP speaker receives the OPEN message and no error is detected, the BGP speaker sends a KEEPALIVE message to the remote BGP speaker
- 6 – Established:** All of the neighbor negotiations are complete. You will see a number (2 in this case), which tells us the number of prefixes the router has received from a neighbor or peer group.

QUESTION NO: 219

Which command displays the IBGP and EBGP neighbors that are configured?

- A. show ip bgp
- B. show ip bgp paths
- C. show ip bgp peers
- D. show ip bgp summary

Answer: D

Explanation:

The picture below shows the output of the show ip bgp summary

```
R1#show ip bgp summary
BGP router identifier 172.12.123.1, local AS number 100
BGP table version is 1, main routing table version 1

Neighbor      V     AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down State/PfxRcd
3.3.3.3        4     300      0       0        0    0    0 never    Active
```

Notice that the “show ip bgp” command to display BGP topology database. Below is the output of the “show ip bgp” command:

```
RouterA# show ip bgp
BGP table version is 14, local router ID is 172.31.11.1
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop        Metric LocPrf Weight Path
*> 10.1.0.0/24     0.0.0.0          0         32768 i
* i               10.1.0.2          0        100      0 i
*> 10.1.1.0/24     0.0.0.0          0         32768 i
*>i10.1.2.0/24     10.1.0.2          0        100      0 i
*> 10.97.97.0/24   172.31.1.3        0         0 64998 64997 i
*                 172.31.11.4        0         0 64999 64997 i
* i               172.31.11.4        0        100      0 64999 64997 i
*> 10.254.0.0/24   172.31.1.3        0         0 64998 i
*                 172.31.11.4        0         0 64999 64998 i
* i               172.31.1.3        0        100      0 64998 i
r> 172.31.1.0/24   172.31.1.3        0         0 64998 i
r                 172.31.11.4        0         0 64999 64998 i
r i               172.31.1.3        0        100      0 64998 i
*> 172.31.2.0/24   172.31.1.3        0         0 64998 i
<output omitted>
```

QUESTION NO: 220

BGP contains two paths to a destination. Assuming both routes were originated locally and have an equal weight. What will be the next determining factor in choosing the best path?

- A. lowest MED
- B. highest local preference
- C. lowest neighbor IP address
- D. lowest origin code

E. shortest AS-path

Answer: B

Explanation:

Memorizing the BGP decision process steps is very useful and you should remember them. The table below lists the complete path selection process:

1. Weight (Bigger is better)
2. Local preference (Bigger is better)
3. Self originated (Locally injected is better than iBGP/eBGP learned)
4. AS-Path (Smaller is better)
5. Origin (Prefer ORIGIN code I over E, and E over ?)
6. MED (Smaller is better)
7. External (Prefer eBGP over iBGP)
8. IGP cost (Smaller is better)
9. EBGP Peering (Older is better)
10. RID (Lower is better)

QUESTION NO: 221

Refer to the exhibit.

```
hostname RAR1
!
<output omitted>
!
router bgp 100
 neighbor 172.16.1.2 remote-as 200
 neighbor 172.16.1.2 distribute-list 101 in
!
access-list 101 permit ip 10.10.0.0 0.0.0.0 255.255.224.0 0.0.0.0
```

Which statement is true?

- A. Router RAR1 will accept only route 10.10.0.0/19 from its BGP neighbor.
- B. Router RAR1 will send only route 10.10.0.0/19 to its BGP neighbor.
- C. Only traffic with a destination from 10.10.0.0/19 will be permitted.
- D. Only traffic going to 10.10.0.0/19 will be permitted.

Answer: A

Explanation:

Wild card mask for /19 is 0.0.31.255. We can see that there is distributed list inbound used to filter routers from BGP neighbor. Distribute list is linked with extend ACL No. 101, which statement permits network 10.10.0.0/19 and denies all other networks, that comes from implicit deny rule at

the end of any access list.

QUESTION NO: 222

What are the two reasons for the appearance of 0.0.0.0 as the next hop for a network in the show ip bgp command output? (Choose two.)

- A. The network was originated via redistribution of an interior gateway protocol into BGP.
- B. The network was defined by a static route.
- C. The network was learned via IBGP.
- D. The network was learned via EBGP.
- E. The network was originated via a network or aggregate command.

Answer: A,E

Explanation: Explanation

"The network was originated via a network or aggregate command" has the next hop of 0.0.0.0. In short, the router on which you use the "network" or "aggregate" command will set the next hop of 0.0.0.0 for that route.

About redistribution please refer to the Understanding Redistribution of OSPF Routes into BGP: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800943c5.shtml you will see after the network 2.0.0.0 has been redistributed in BGP it has the next hop of 0.0.0.0.

QUESTION NO: 223

Refer to the exhibit.

```
router bgp 100
 neighbor internal peer-group
 neighbor internal remote-as 100
 neighbor internal update-source loopback 0
 neighbor internal route-map set-med out
 neighbor internal filter-list 1 out
 neighbor internal filter-list 1 in
 neighbor 171.69.232.53 peer-group internal
 neighbor 171.69.232.54 peer-group internal
 neighbor 171.69.232.55 peer-group internal
 neighbor 171.69.232.55 filter-list 3 in
```

Which two statements are true about the partial configuration that is provided? (Choose two.)

- A. All the configured neighbors are in autonomous system 100.
- B. The peer group shortens the IBGP configuration.
- C. The peer group shortens the EBGP configuration.
- D. Only the outgoing filters are applied to BGP updates.
- E. Three AS-path filters are applied to each BGP neighbor.

Answer: A,B

Explanation: Explanation

This is an IBGP peer group because the AS numbers in “router bgp {AS number}” and “neighbor internal remote-as {AS number}” are the same.

A BGP peer group reduces the load on system resources by allowing the routing table to be checked only once, and updates to be replicated to all peer group members instead of being done individually for each peer in the peer group. In addition, a BGP peer group also simplifies the BGP configuration.

This is the process of creating a peer-group (used the output above):

neighbor internal peer-group	Create a peer-group (named internal)
neighbor internal remote-as 100 neighbor internal update-source loopback 0 neighbor internal route-map set-med out neighbor internal filter-list 1 out neighbor internal filter-list 2 in	Configure needed commands for the peer-group
neighbor 171.69.232.53 peer-group internal neighbor 171.69.232.54 peer-group internal neighbor 171.69.232.55 peer-group internal	Assign BGP neighbor into a peer group

Just one thing to notice is the last command “neighbor 171.69.232.55 filter-list 3 in” indicates the filter-list 3 will be applied for neighbor 171.69.232.55 while other neighbors will be applied filter-list 2 as the inbound filter-list (all neighbors use outbound filter-list 1).

QUESTION NO: 224

Refer to the exhibit.

```
Router#show ip bgp
```

```
BGP table version is 5, local router ID is 10.0.33.34
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.1.0.0	0.0.0.0	0		32768	?
* 10.2.0.0	10.0.33.35	10		0	35 ?
*>	0.0.0.0	0		32768	?
* 10.0.0.0	10.0.33.35	10		0	35 ?
*>	0.0.0.0	0		32768	?
*> 192.168.0.0/16	10.0.33.35	10		0	35 ?

Which two statements are correct? (Choose two.)

- A. All six routes will be installed in the routing table.
- B. Two routes will be installed in the routing table.
- C. Four routes will be installed in the routing table.
- D. All the routes were redistributed into BGP from an IGP.
- E. All the routes were originated by BGP with the network command.

Answer: C,D

Explanation: Explanation

Only the valid & best routes (represented by *>) will be installed into the routing table.

All the routes were redistributed into BGP from an IGP so we will see a next hop of 0.0.0.0. For more information about this type of redistribution please read

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800943c5.shtml

QUESTION NO: 225

A router has two paths to reach another network in a different autonomous system. Neither route was generated by the local router and both routes have the same default weight and local preference values. Which statement is true about how BGP would select the best path?

- A. If the command `bgp always-compare-med` has been given, then the router will prefer the route with the highest MED.
- B. The router will prefer the route with the lower MED.
- C. The router will prefer the shortest autonomous system path.
- D. To influence one route to be preferred, its default local preference value will be changed via the use of the command `bgp default local-preference 50`.

Answer: C

Explanation:

In the Route selection decision process, if the weight, local preference & route originated are the same then the shortest AS path will be chosen.

The full Route selection decision process is listed below:


1. Prefer highest weight (local to router)
2. Prefer highest local preference (global within AS)
3. Prefer route originated by the local router (next hop = 0.0.0.0)
4. Prefer shortest AS path
5. Prefer lowest origin code (IGP < EGP)
6. Prefer lowest MED (exchanged between autonomous systems)
7. Prefer EBGP path over IBGP path
8. Prefer the path through the closest IGP neighbor (IGP cost)
9. Prefer oldest route for EBGP paths
10. Prefer the path with the lowest neighbor BGP router ID
11. Prefer the path with the lowest neighbor IP address

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094431.shtml

QUESTION NO: 226

Refer to the exhibit.



BGP table version is 12, local router ID is 192.168.100.1

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
* 172.16.0.0	192.168.1.2	75		0	64512 i
*>	192.168.1.6	50		0	64512 i
*> 192.168.1.0/30	0.0.0.0	0		32768	i
*> 192.168.1.4/30	0.0.0.0	0		32768	i
*> 192.168.100.0	0.0.0.0	0		32768	i

On the basis of the information in the exhibit, which two statements are true? (Choose two.)

- A. The output was generated by entering the show ip bgp command on the ISP router.
- B. The output was generated by entering the show ip bgp command on the SanJose1 router.
- C. The serial 0/0/1 interface on the ISP router has been configured with the set metric 50 command.
- D. The serial 0/0/1 interface on the ISP router has been configured with the set metric 75 command.
- E. When traffic is sent from the ISP to autonomous system 64512, the traffic will be forwarded to SanJose1 because of the lower MED value of SanJose1.
- F. When traffic is sent from the ISP to autonomous system 64512, the traffic will be forwarded to SanJose2 because of the higher MED value of SanJose2.


Answer: A,D

Explanation:

The “show ip route bgp” command will display any BGP-learned routes that make it into the IP routing table, the command “show ip bgp” is required to display the contents of the actual BGP routing table. This output was seen on ISP because the local router ID is 192.168.100.1 (ISP). Since we know that this output must have been seen by ISP, we know the serial 0/0/1 interface has been configured with a metric of 75, as this is the metric to the peer with IP address 192.168.1.2 (the other side of the serial 0/0/1 interface).

QUESTION NO: 227

Refer to the exhibit.



On the basis of the configuration that is provided, how would the BGP updates that come from router R1 be replicated inside autonomous system 65200?

- A.** All BGP updates that are received on router R2 will be sent to routers R3 and R4. Routers R3 and R4 will then forward those BGP updates to router R5.
- B.** All BGP updates that are received on router R2 will not be sent to routers R3 and R4.
- C.** All BGP updates that are received on router R2 will be sent directly to router R5.
- D.** None of the BGP updates that are received on router R2 will ever be received by router R5.

Answer: D

Explanation: Explanation

All BGP updates that are received on router R2 will be sent to routers R3 and R4 but R3 & R4 will **not** forward those BGP updates to R5. This is called the BGP split-horizon rule (which states that a route learned from one IBGP neighbor will not be advertised to another IBGP neighbor) .


The BGP updates received on router R2 will be sent to R3 and R4 without violating the BGP split-horizon rule because R2 receives updates from an EBGP (R1), not IBGP.

From the configuration of R2, we learn that R2 did not establish neighbor relationship with R5 so they are not neighbors -> no BGP updates will be sent from R2 to R5.

The BGP split-horizon rule prevents updates received on R2 from being sent to R5 is correct.

QUESTION NO: 228

Refer to the exhibit.



Router RT-1 chooses one path to network 198.133.219.0/24. Indicate the reason Router RT-1 chooses this "best" path.

- A. In making its decision about the best path, RT-1 gives precedence to the origin code.
- B. In making its decision about the best path, RT-1 gives precedence to the BGP MED values.
- C. IP address 128.107.2.2 is lower than 128.107.255.2.
- D. In making its decision about the best path, RT-1 prefers the IGP metrics.
- E. RT-1 prefers internal BGP routes.
- F. IP address 128.107.254.2 is lower than 128.107.255.2.

Answer: A


Explanation:

After BGP receives updates about different destinations from different autonomous systems, it chooses the best path to reach a specific destination.

The following process summarizes how BGP chooses the best route on a Cisco router.

QUESTION NO: 229

Refer to the exhibit.



Router RT-1 and router RT-2 both advertise network 131.25.0.0/16 to router RT-3 via internal BGP. What is the reason that router RT-3 chose router RT-1 as its best path to network 131.25.0.0/16.

- A. It advertises the best AS-path.
- B. It advertises the best origin code.
- C. It advertises the best MED.
- D. It advertises the best local preference.
- E. It has a better router ID.
- F. It advertises a lower autonomous system.

Answer: E

Explanation:

After BGP receives updates about different destinations from different autonomous systems, it chooses the best path to reach a specific destination.

The following process summarizes how BGP chooses the best route on a Cisco router.

In this case, since everything else is equal the router ID will be used, and the specific router ID that is chosen is 165.1.1.1 (lowest one), which is RT-3.

QUESTION NO: 230

During BGP configuration on a router that has peered with other BGP speakers, the BGP command `aggregate-address 172.32.0.0 255.255.252.0` is issued. However, the peers do not receive this aggregate network in BGP advertisements. Also, the router does not have this aggregate network in its BGP table. Which option indicates a possible reason this command did not cause the router to advertise the aggregate network to its peers?

- A. Interface NULL 0 is likely shutdown.
- B. The BGP command `no synchronization` is missing.
- C. The BGP command `no auto-summary` is missing.
- D. Subnets of 172.32.0.0/22 do not exist in the BGP table.
- E. The IGP running on this router does not have network 172.32.0.0/22 installed.
- F. The next hop IP address must be a loopback address.

Answer: D

Explanation:


BGP allows the aggregation of specific routes into one route using the "aggregate-address address mask" command. Aggregation applies to routes that exist in the BGP routing table. This is in contrast to the network command, which applies to the routes that exists in IP routing table. Aggregation can be performed if at least one or more of the specific routes of the aggregate address exists in the BGP routing table.

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800945ff.shtml

QUESTION NO: 231

Refer to the exhibit.



It is desired to set up a BGP neighbor relationship between routers R1 and R4. BGP packets between them could travel through R2 or R3. What is the simplest configuration that will allow for failover?

- A. Configure BGP neighbor relationships between all interfaces on R1 and R4
- B. Install a direct connection between R1 and R4.
- C. Configure loopback interfaces on R1 and R4 to provide the update source address for BGP packets.
- D. Configure only one neighbor relationship between R1's 192.168.1.2 interface and R4's 172.16.10.2 interface.

Answer: C

Explanation:

Assume R4 only uses this configuration:

```
R4(config)#router bgp 65000R4(config-router)#neighbor 192.168.1.2 remote-as 65000
```

Then if R1 sends BGP packets to R4 via R3, the source IP address of the packets is 192.168.2.2. But R4 does not recognize this IP address because 192.168.2.2 is not configured in the "neighbor" command of R4. Therefore the IBGP session between R1 & R4 is not established -> these

packets will be dropped.

The simplest configuration to allow R1 and R4 communicate via both R2 & R3 is to use a loopback interface address rather than a physical interface address as the source IP address for all BGP packets. To do this, use the command:

R4(config-router)#neighbor 1.1.1.1 update-source loopback0

(In which 1.1.1.1 is the loopback interface of R1). In practical, we should establish neighborship with the loopback interface rather than the physical interface because if the physical interface goes down, the neighborship would be lost while a loopback interface never goes down.

Also, when configuring "1.1.1.1 as the neighbor, you must configure on R1 the "neighbor 4.4.4.4 update-source **loopback0**" command so that the source IP address of packets sent from R1 (loopback0 – 1.1.1.1) will be matched with the neighbor command configured on R4.

QUESTION NO: 232

Which BGP option is required when load sharing over multiple equal-bandwidth parallel links from a single CE router to a single ISP router over eBGP?

- A. eBGP Multipath
- B. eBGP Multihop
- C. BGP Synchronization
- D. Public AS numbers

Answer: B

Explanation: Explanation

The eBGP multihop allows a neighbor connection between two external peers that do not have direct connection. The multihop is only for eBGP and not for iBGP. For example, in the topology below router A wants to establish neighbor relationship with the loopback0 of router B (to allow load balancing), which does not have direct connection so it must use "ebgp-multihop"



For your reference, the full configurations of both router A & B are shown below:

```
A# int loopback 0
ip address 10.10.10.2 255.255.255.255
router bgp 1
neighbor 10.10.10.1 remote-as 2
neighbor 10.10.10.1 ebgp-multihop
neighbor 10.10.10.1 update-source loopback 0
network 10.10.10.2 mask 255.255.255.255
ip route 10.10.10.1 255.255.255.255 172.16.10.1
ip route 10.10.10.1 255.255.255.255 172.16.10.6
```

```
B# int loopback 0
ip address 10.10.10.1 255.255.255.255
router bgp 2
neighbor 10.10.10.2 remote-as 1
neighbor 10.10.10.2 ebgp-multihop
neighbor 10.10.10.2 update-source loopback 0
network 10.10.10.1 mask 255.255.255.255
ip route 10.10.10.2 255.255.255.255 172.16.10.2
ip route 10.10.10.2 255.255.255.255 172.16.10.5
```

Note: If router B wants to establish neighbor relationship with the directly connected interface of router A, it only needs these commands:

```
B#
router bgp 2
neighbor 172.16.10.2 remote-as 1
```


But notice the traffic from router B would be sent to 172.16.10.2 interface only and load balancing would not take place.

(Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800c95bb.shtml)

QUESTION NO: 233

Refer to the exhibit.



The customer has reported an issue about eBGP loading. Currently the AS 100 eBGP links have an average outbound load of 65% and 20% respectively. On further investigation, traffic from 10.10.24.0 accounts for 45%, and 10.10.25.0 and 10.10.32.0 accounts for 20% each of the outbound load. The customer wants to spread the load between the two eBGP links more evenly. The BGP attributes are currently set at their default values. If you are located at AS 100 and want to influence how AS 100 sends traffic to AS 200, what BGP attribute would you configure to cause AS 100 outbound traffic to load the eBGP links more evenly?

- A. On router A, set the default local-preference to 50.
- B. On router B, set the default metric to 150.
- C. On router B, configure a route map for 10.10.25.0/24 that sets the local preference to 150.
- D. On router B, set the default local-preference to 150.

Answer: C


Explanation:

To make the eBGP links more evenly we should use the link B-F for network 10.10.25.0/24 so that the total traffic going through B-F link is about 40%. In this case we should apply a route map on B to set the local preference of 10.10.25.0/24 to a higher value than 100. But notice that we must use a second clause to permit other traffic if not they will be filtered out.

Note: The default value for local preference is 100. A path with higher local preference is preferred.

QUESTION NO: 234

Refer to the exhibit.



There are two eBGP links between AS100 and AS200. Currently the router A to router E link has an average inbound load of 65% and the router B to router F link has an average inbound load of 20%. After further investigation, it is found that traffic to the 10.10.1.16/28 subnet is using 45% of the bandwidth on the A to router E link and traffic to the 10.10.1.32/28 subnet is using 20% of the bandwidth on the A to router E link. Traffic to the 10.10.1.48/28 subnet is using 20% of the bandwidth on the B to router F link. If you want to influence how AS200 sends traffic to AS100, which two commands would you configure in AS100 to influence AS200 to use the eBGP links more evenly? (Choose two.)


- A. neighbor 192.168.30.2 route-map as_50 out
- B. neighbor 192.168.20.2 route-map as_50 out
- C. route-map as_50 permit 10
match ip address 50
set metric 150
access-list 50 permit 10.10.1.16 0.0.0.240
- D. route-map as_50 permit 10
match ip address 50
set metric 150
access-list 50 permit 10.10.1.32 0.0.0.15

Answer: A,D

Explanation:

QUESTION NO: 235

Refer to the exhibit.



AS 65500 is running OSPF as its IGP and has two eBGP connections to the Internet provider AS 65501. The provider requires that you advertise all of the prefixes in AS 65500 to the provider as a single 192.168.12.0/22 route. Currently only Routers A and B are able to reach the Internet. Which of the following additions to the configuration will solve this issue?

- A. RouterA(config)#ip cef
- B. RouterA(config)# ip route 192.168.12.0 255.255.252.0 null 0
- C. RouterA(config-router)#ebgp multihop 1
- D. RouterA(config-router)#redistribute ospf 1
- E. RouterA(config-router)#neighbor 192.168.14.253 next-hop-self
- F. RouterA(config-router)#neighbor 192.168.14.253 local-as 65500

Answer: B

Explanation:

The synchronization rule states "A BGP router should not use, or advertise to an external neighbor, a route learned by IBGP, unless that route is local or is learned from the IGP". Notice that IGP here can be a static route.

In this case, unless there is an entry about network 192.168.12.0 in the routing table of RouterA,

RouterA will not advertise this network to its EBGp -> B is correct (even if this static route points to Null0).

Note: Although this question states that OSPF is being used as IGP but for some reasons, network 192.168.12.0/22 is not advertised to RouterA -> RouterA does not have this route in its routing table.

QUESTION NO: 236

Which BGP command provides the router ID, local preference, next hop, and BGP path in its output?

- A. show ip route bgp
- B. show ip bgp
- C. show ip bgp neighbors
- D. show ip bgp summary

Answer: B

Explanation: Explanation

The “show ip bgp” command is used to display entries in the BGP routing table. An example of the “show ip bgp” output is shown below:

```
R4#show ip bgp
BGP table version is 2, local router ID is 34.34.34.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 1.1.1.0/24        34.34.34.3              0  23  1  i
```

In the output you can see the router ID, local preference, next hop, and BGP path.

QUESTION NO: 237

Refer to the exhibit.

```
Router# show ip bgp neighbors 172.16.254.3
```

```
BGP neighbor is 172.16.254.3, remote AS 150, internal link  
BGP version 4, remote router ID 172.16.254.3  
BGP state = Active, table version 0  
Last read 00:00:06, hold time is 180, keepalive interval is 60 seconds  
Default minimum time between advertisement runs is 5 seconds  
Received 3 messages, 0 notifications, 0 in queue  
Connections established 1; dropped 1  
Last user reset 00:00:54, due to User reset request  
No active TCP connection
```

For the accompanying router output, which of the following statements describes the state that neighbor 172.16.254.3 is in?

- A. BGP can exchange routing information in this state.
- B. The router will not accept connections from the peer.
- C. The router is listening on its server port for connection requests from the peer.
- D. The router has sent out an active TCP connection request to the peer.

Answer: C

Explanation: Explanation

The BGP state in the output is "Active", which means BGP speaker is attempting to initiate a TCP session with the BGP speaker it wants to peer with. If this can be done, the BGP state goes to OpenSent state.

QUESTION NO: 238

Above is the output from show ip bgp neighbors command.

```
BGP neighbor is 172.16.254.3, remote AS 150, internal link
1. BGP version 4, remote router ID 172.16.254.3
2. BGP state = Established, up for 19:24:07
3. Last read 00:00:06, hold time is 180, keepalive interval is 60 seconds
4. Neighbor capabilities:
5. Route refresh:advertised and received(new)
6. Address family IPv4 Unicast:advertised and received
7. Graceful Restart Capability:advertised and received
8. Remote Restart timer is 120 seconds
9. Address families preserved by peer
10. IPv4 Unicast
11. Received 4231 messages, 0 notifications, 0 in queue
12. Sent 4167 messages, 0 notifications, 0 in queue
13. Default minimum time between advertisement runs is 5 seconds
14. For address family:IPv4 Unicast
15. BGP table version 159559, neighbor version 159559
16. Index 90, Offset 11, Mask 0x4
17. Route refresh request:received 0, sent 0
18. 10031 accepted prefixes consume 441364 bytes
19. Prefix advertised 29403, suppressed 0, withdrawn 9801
20. Number of NLRI in the update sent:max 242, min 0
21. Connections established 2; dropped 1
22. Last reset 19:26:54, due to NSF peer closed the session
23. Connection state is ESTAB, I/O status:1, unread inout bytes:0
24. Local host:150.254.254.2, Local port:11005
25. Foreign host:172.16.254.3, Foreign port:179
```

What is line 21 stating about the BGP connection?

- A. the number of neighbors that the router has
- B. the number of times the router has established a TCP connection
- C. the number of consecutive TCP connections to the specified remote neighbor
- D. the number of total TCP connections that the router has

Answer: B

Explanation: According to

http://www.cisco.com/en/US/docs/ios/12_3/iproute/command/reference/ip2_s2g.html#wp1040913, "Connections established" is the number of times a TCP and BGP connection have been successfully established while "dropped" is the number of times that a valid session has failed or been taken down.

QUESTION NO: 239

The 192.168.0.0 network is not being propagated throughout the network. Observe the BGP configuration commands from the advertising router. What is the reason the 192.168.0.0 route is not being advertised?

```
router bgp 65111  
  
neighbor 172.16.1.1 remote-as 65111  
  
neighbor 172.16.2.1 remote-as 65112  
  
network 192.168.0.0  
  
network 10.0.0.0  
  
!  
  
ip route 192.168.0.0 255.255.0.0 null0
```

- A. The network 192.168.0.0 statement is missing mask 255.255.0.0
- B. The network 192.168.0.0 statement is missing mask 0.0.255.255.
- C. The network 10.0.0.0 statement is missing mask 255.0.0.0.
- D. The network 10.0.0.0 statement is missing mask 0.255.255.255.
- E. The auto-summary configuration is missing.

Answer: A

Explanation: Explanation

The “network” statement in other routing protocols (EIGRP, OSPF, RIP...) is used to enable routing protocol on the interfaces within that “network” statement. But in BGP, the function of a network statement is to tell the router to search the IP routing table for a particular network, and if that network is found, originate it into the BGP database. But notice that you must have an exact match in the IP routing table to appear the network in the BGP routing table (in this case we don’t see the auto-summary command so we suppose it is disabled in this case). For example: + network 10.10.10.0/8 will appear in BGP if network 10.10.10.0/8 appears in the IP routing table. + network 10.10.10.0/24 will appear in BGP if network 10.10.10.0/24 appears in the IP routing table. Therefore, in this question the static route “ip route 192.168.0.0 255.255.0.0 null0 was used to put a route to 192.168.0.0/16 into the routing table (although it points to Null0 but this command really makes that route appears in the routing table). But the “network 192.168.0.0 statement tells the router to lookup network 192.168.0.0/24 (if the network statement under BGP-mode does not specify a subnet mask, the default subnet mask of that class will be used). The router only finds network 192.168.0.0/16 -> The network 192.168.0.0 is not being propagated throughout the network because of the mismatch of the subnet mask.

Just for your information, in fact we have to suppose there is no entry of the network 192.168.0.0/24 exist in the routing table except the static route “ip route 192.168.0.0 255.255.0.0 null0. If such an entry exists (for example, a directly connected entry like “C 192.168.0.0/24 is directly connected”) then the router still advertises it with the “network 192.168.0.0 (without mask

255.255.0.0) command.

QUESTION NO: 240

Which BGP attribute is used by BGP to prevent routing loops?

- A. AS-path
- B. next-hop
- C. MED
- D. weight
- E. local preference
- F. origin

Answer: A

Explanation: Explanation

The AS-PATH attribute is used to prevent BGP routing loops. When receiving an BGP advertisement, the router checks the AS-PATH attribute, if it see it's own AS number in the AS-Path then it is a routing loop so the router will not install this route in its BGP table.

QUESTION NO: 241

Which BGP path attribute is Cisco proprietary?

- A. weight
- B. MED
- C. local preference
- D. origin
- E. next-hop
- F. AS-path

Answer: A

Explanation: Explanation

The weight attribute is a Cisco proprietary technology. This attribute is local to a router and is not advertised to neighboring routers. If the router learns about more than one route to the same destination, the route with the highest weight will be preferred and will be installed in the routing table. By default, the value of weight is 0 and the range is from 0 to 65535.

The route with the highest weight will be installed in the IP routing table.

QUESTION NO: 242

Which attribute must exist in the BGP update packet?

- A. LOCAL_PREF
- B. AGGREGATOR
- C. AS_Path
- D. Weight

Answer: C

Explanation: Explanation

Origin, AS_Path, Next_Hop are well-known mandatory BGP attributes that all BGP Updates must include.

Note:

There are 4 BGP attribute types:

+ **Well-known Mandatory**: recognized by all implementations of BGP and must appear in a BGP update message. If missing, a notification error will be generated.+ **Well-Known Discretionary**: recognized by all implementations of BGP but may not be sent in the BGP update message (include LOCAL_PREF, ATOMIC_AGGREGATOR).+ **Optional Transitive**: may or may not be recognized by all BGP implementations. Because the attribute is transitive, BGP accepts and advertises the attribute even if it is not recognized (include Community attribute).+ **Optional Nontransitive**: may or may not be recognized by all BGP implementations. Whether or not the receiving BGP router recognizes the attribute, it is nontransitive and is not passed along to other BGP peers (include MED).

QUESTION NO: 243

Which BGP attribute will not be advertised in routing updates to its neighboring routers?

Select the best response


- A. weight
- B. local preference
- C. origin
- D. AS_path
- E. next hop

Answer: A

Explanation: The weight attribute is a Cisco proprietary technology. This attribute is local to a router and is not advertised to neighboring routers. If the router learns about more than one route to the same destination, the route with the highest weight will be preferred and will be installed in the routing table. By default, the value of weight is 0 and the range is from 0 to 65535. The route with the highest weight will be installed in the IP routing table.

QUESTION NO: 244

Observe the diagram.



RTC is the hub router and RTA and RTB are the spokes. There are no virtual circuits between the spoke locations. What is needed to successfully route traffic to the 11.11.11.0/24 network from RTA?


- A. The neighbor 10.10.10.1 next-hop-self command on RTA.
- B. The neighbor 10.10.10.1 next-hop-self command on RTB.
- C. The neighbor 10.10.10.1 next-hop-self command on RTC.
- D. Nothing is required. This is the default behavior on this topology.

Answer: C

Explanation:

The following example illustrates the issue:

BGP Next Hop (NBMA)



If the common media as you see in the shaded area above is a frame relay or any NBMA cloud then the exact behavior will occur as if we were connected via Ethernet. RTC will advertise 180.20.0.0 to RTA with a next hop of 170.10.20.3.

The problem is that RTA does not have a direct PVC to RTD, and cannot reach the next hop. In this case routing will fail.

In order to remedy this situation a command called next-hop-self is created.

The next-hop-self Command

Because of certain situations with the next hop as we saw in the previous example, a command called next-hop-self is created. The syntax is:

```
neighbor {ip-address|peer-group-name} next-hop-self
```

The next-hop-self command allows us to force BGP to use a specified IP address as the next hop rather than letting the protocol choose the next hop.

In the previous example, the following configuration solves our problem:

RTC#

```
router bgp 300
```

```
neighbor 170.10.20.1 remote-as 100
```

```
neighbor 170.10.20.1 next-hop-self
```

RTC advertises 180.20.0.0 with a next hop = 170.10.20.2

Reference: <http://www.cisco.com/warp/public/459/bgp-toc.html#bgpnexthop>

QUESTION NO: 245

Which statement is true about IBGP routers?

- A. They must be fully meshed.
- B. They can be in a different AS.
- C. They must be directly connected.
- D. They do not need to be directly connected.

Answer: D

Explanation:

When two routers establish a TCP enabled BGP connection, they are called neighbors or peers. Each router running BGP is called a BGP speaker. Peer routers exchange multiple messages to open and confirm the connection parameters, such as the version of BGP to be used. If there are any disagreements between the peers, notification errors are sent and the connection fails. To be a neighbor of BGP no need to be directly connected.

QUESTION NO: 246

Which statement is true about EBGP?

- A. An internal routing protocol can be used to reach an EBGP neighbor.
- B. The next hop does not change when BGP updates are exchanged between EBGP neighbors.
- C. A static route can be used to form an adjacency between neighbors.
- D. EBGP requires a full mesh.


Answer: C

Explanation:

When BGP is running between routers in different autonomous systems, it is called External BGP (EBGP). When BGP is running between routers in the same AS, it is called Internal BGP (IBGP). BGP allows the path that packets take to be manipulated by the AS, as described in this module. It is important to understand how BGP works to avoid creating problems for your AS as a result of running BGP.

QUESTION NO: 247

Refer to the exhibit.



By default, when RTB passes BGP advertisements from RTA about network 192.168.2.0 to RTC, what address will be listed as the next-hop address?

Select the best response

- A. 192.168.1.49
- B. 10.1.1.1
- C. 10.1.1.2
- D. 192.168.1.50

Answer: D

Explanation:

The next hop attribute is a well-known mandatory attribute, type code 3. In terms of an IGP, such as RIP, the "next hop" to reach a route is the IP address of the router that has announced the route.

The next hop concept with BGP is more complex and takes one of the following three forms:

QUESTION NO: 248

Which two statements are true about IBGP neighbor relationships? (Choose two.)

- A. An EGP or static routing is required between IBGP neighbors.
- B. A full-mesh IBGP requires that neighbor relationships be established between all BGP enabled routers in the autonomous system.
- C. IBGP neighbors must be in different autonomous systems.
- D. The BGP split-horizon rule specifies that routes learned via EBGp are never propagated to other IBGP peers.

E. The BGP split horizon rule specifies that routes learned via IBGP are never propagated to other IBGP peers.

Answer: B,E

Reference: <http://packetlife.net/wiki/bgp/> (see confederations)

QUESTION NO: 249

A router is running BGP and receives more than one route for a particular prefix. Assume all the routes for this prefix have the same attributes. Which three path features would be reasons be for the router to ignore some of the routes and not consider them as candidates for the best path? (Choose three.)

- A.** paths that are marked as synchronized in the show ip bgp output
- B.** paths that are marked as not synchronized in the show ip bgp output
- C.** paths for which the NEXT_HOP is accessible
- D.** paths for which the NEXT_HOP is inaccessible
- E.** paths from an external BGP (eBGP) neighbor if the local autonomous system (AS) appears in the AS_PATH
- F.** paths from an internal BGP (iBGP) neighbor if the local autonomous system (AS) appears in the AS_PATH


Answer: B,D,E

Explanation: Explanation

Only synchronized routes (answer B) with no AS loops (answer E) and a valid next hop (answer D) will be considered as candidates for the best path route selection decision process.

QUESTION NO: 250

Refer to the exhibit. Autonomous systems 200 and 300 have EBGP sessions established with their directly connected routers in autonomous system 100. IGP has been configured on all routers in autonomous system 100 and they successfully exchange routing updates. Traffic originated in autonomous system 200 cannot reach the destination autonomous system 300. What configuration should be done on the routers in autonomous system 100 in order for the traffic coming from autonomous system 200 to be forwarded to autonomous system 300?




- A. IBGP session must be established between routers R1 and R3, and the synchronization must be turned on.
- B. IBGP session must be established between routers R1 and R3, and the synchronization must be turned off.
- C. IBGP session must be established between routers R1 R2 and R2 R3, and the synchronization must be turned on.
- D. IBGP session must be established between routers R1 R2 and R2 R3, and the synchronization must be turned off.
- E. IBGP speakers within autonomous 100 must be fully meshed, and the synchronization must be turned on.
- F. IBGP speakers within autonomous 100 must be fully meshed, and the synchronization must be turned off.

Answer: F

Explanation: Explanation

The synchronization rule states that if an AS provides transit service to another AS, BGP should not advertise a route until all of the routers within the AS have learned about the route via an IGP. To understand why this rule exists, let's take an example if this rule is not there.




Suppose Rt-A wants RT-B to access its local LAN 1.1.1.0, so it advertises this network through R1. R1 and R3 are running IBGP so R1 sends this update to R3 through R2 (using the next-hop-self to use its own interface's IP address). In turn, R3 announces to RT-B that it can reach 1.1.1.0 via R3.

Now Rt-B really wants to send traffic to 1.1.1.0 so it will send to R3. R3 does a look up and sees that the network can be reachable via R1. It then does a lookup for R1's IP address and sees that it is reachable via R2 -> so it forwards packets to R2. But R2, running IGP (like OSPF), does not find an entry for 1.1.1.0 so R2 drops all the packets for that network – a black-hole is created!

That is why the BGP synchronization rule is born. With this rule, when R3 receives an advertisement for 1.1.1.0 from R1, it adds that route to its BGP table and before sending advertisement to RT-B, it first checks its IGP routing table to see whether an entry exists for that route. In this example, R3's IGP routing table does not know how to reach 1.1.1.0 so R3 will not advertise this network to RT-B. This route is only advertised to RT-B when IGP makes an entry in the routing table for 1.1.1.0.

Well, now you understand the importance of BGP Synchronization rule but now I wish to explain why this rule causes trouble in fully-meshed IBGP!



Synchronization prevents fully-meshed IBGP from working properly. Because no IGP is running so R3 cannot advertise any route to R1 even if no black-hole exists in this topology.

Note: A “fully-meshed” can be a physical fully-meshed topology or a topology where all routers in the same AS established IBGP connections with each other (although they do not need to be directly connected). So in the topology above, the connection between R1 & R3 is represented by a dashed line, which means it can be physically connected or not (but an IBGP connection must be established on both routers).

Therefore if all routers in AS 100 is fully-meshed, the synchronization rule must be turned off.

QUESTION NO: 251

Refer to the exhibit.

```
R2# show ip bgp 6.6.6.0
BGP routing table entry for 6.6.6.0/24, version 2
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to non peer-group peer:
    10.10.23.3
    30
    10.10.23.3 from 10.10.23.3 (6.6.6.1)
      Origin IGP, metric 0, localpref 125, valid, external
      Community: 100:250
```

Which statement is true about the 6.6.6.0/24 prefix?

- A. Route 6.6.6.0/24 is learned by an IBGP peer.
- B. The command neighbor send-community is configured on BGP neighbor 10.10.23.3.
- C. The route 10.10.23.3 is not being advertised to other BGP neighbors.
- D. If another path advertises the 6.6.6.0/24 path and has the default local preference, that path is more preferred.

Answer: B

Explanation: Explanation

By default, the community attributes are not advertised to BGP neighbors. But in the output we see the "Community: 100:250 which means the command "neighbor ... send-community" was used to send community attributes of the local router to the neighbor.

QUESTION NO: 252

If a metric is not specified for routes that are redistributed into OSPF, the default metric that is assigned to the routes is 20, except for redistributed BGP routes. What is the metric that is assigned to redistributed BGP routes?

- A. 0
- B. 1
- C. 10
- D. 200

Answer: B

Explanation:

If a metric is not specified, OSPF puts a default value of 20 when redistributing routes from all protocols except Border Gateway Protocol (BGP) routes, which get a metric of 1.

Reference:

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

QUESTION NO: 253

Which one of the following statements about BGP is FALSE?

- A. BGP uses TCP port 179.
- B. BGP ensures reliability of updates by using the reliable transport services of TCP.
- C. The network command with the mask option never installs a prefix into the BGP table unless there is a matching prefix exists in the IP route table.

- D. A TCP connection is required before exchanging updates.
- E. BGP uses notification and the update messages to establish and maintain the BGP neighbor relationship.

Answer: E

Explanation: Explanation

An underlying connection between two BGP speakers is established before any routing information is exchanged. This connection takes place on TCP port 179.

Unlike other routing protocols, the router must be manually configured with the neighbor information on both sides of the connection (which is FALSE in this question)

QUESTION NO: 254

Which option is true regarding the synchronization rule?

- A. Do not use or internally advertise a route until the route is learned from a source other than BGP.
- B. Do not use or advertise a route until the route is learned from a BGP peer.
- C. Do not use or advertise routes marked PARTIAL.
- D. Wait until a CONFIRM message is received before using routes from BGP neighbors.

Answer: A

Explanation: Explanation

The complete synchronization rule is "A BGP router should not use, or advertise to an external neighbor, a route learned by IBGP, unless that route is local or is learned from the IGP."

With the default of synchronization disabled, BGP can use and advertise to external BGP neighbors routes learned from an IBGP neighbor that are not present in the local routing table

The "synchronization" here means "synchronization between iBGP with its IGP (such as OSPF, EIGRP...)

You can disable synchronization if one of the following conditions is true: Your AS does not pass traffic from one AS to another AS. All the transit routers in your AS run BGP.

Note: BGP synchronization is disabled by default in Cisco IOS Software Release 12.2(8)T and later.

For more information about BGP Synchronization please read the explanation of Question 5 in this page. Also another good resource is:

http://docwiki.cisco.com/wiki/Internetworking_Case_Studies_-_Using_the_Border_Gateway_Protocol_for_Interdomain_Routing#Synchronization

QUESTION NO: 255

Study the configuration presented in the exhibit carefully.

```
router bgp 50001

network 10.0.0.0
network 172.16.0.0
neighbor 10.1.1.1 remote-as 5003
neighbor 10.1.1.1 route-map test out
!
access-list 1 permit 10.0.0.0
access-list 2 permit any
!
route-map test permit 10
match ip address 1
set metric 200
!
route-map test permit 20
match ip address 2
```

What is the objective of the route map named test?

- A. marks all prefixes received from the 10.1.1.1 neighbor with a MED of 200
- B. marks the 10.0.0.0/8 prefix received from the 10.1.1.1 neighbor with a MED of 200
- C. marks the 10.0.0.0/8 prefix advertised to the 10.1.1.1 neighbor with a MED of 200
- D. marks all prefixes advertised to the 10.1.1.1 neighbor with a MED of 200

Answer: C

Explanation:

The Multiple-exit-discriminator (MED) attribute is an optional nontransitive attribute, type code 4. MED informs external neighbors about the preferred path into an AS that has multiple entry points.

A lower MED is preferred over a higher MED.


Unlike Local Preference, the MED attribute is exchanged between autonomous systems, but a MED attribute that comes into an AS does not leave the AS. When an update enters the AS with a certain MED value, that value is used for decision making within the AS. When BGP forwards the routing update to another AS, the MED is reset to zero. This is true unless the outgoing MED is set to a specific value.

Example:

```
RTB(config)#route-map setmedout permit 10
RTB(config-route-map)#set metric 50
RTB(config)#router bgp 400
```

QUESTION NO: 256

Refer to the exhibit.



Routers RTA and RTB are running BGP but the session is not active. What command needs to be added to establish the BGP session?

A. `ip route 10.10.10.1 255.255.255.255 s0/0`

```
ip route 10.10.10.1 255.255.255.255 s0/1
```

- B.** no synchronization
- C.** network 10.10.10.0
- D.** neighbor 10.10.10.1 next-hop-self


Answer: A

Explanation:

When BGP is running between routers in different autonomous systems, it is called External BGP (EBGP). When BGP is running between routers in the same AS, it is called Internal BGP (IBGP). BGP allows the path that packets take to be manipulated by the AS, as described in this module. It is important to understand how BGP works to avoid creating problems for your AS as a result of running BGP. A static route can be used to form an adjacency between EBGP neighbors.

QUESTION NO: 257

Refer to the exhibit.



Network 10.0.0.0/8 is being advertised to autonomous system 65550 via both external links. Which statement about the preferred path to the 10.0.0.0/8 network is true?

- A. Router R1 will be preferred because its neighbor has the higher autonomous system number.
- B. Router R1 will be preferred because it has the lower neighbor IP address.
- C. Router R1 will be preferred because it has a lower local preference.
- D. Router R2 will be preferred because its neighbor has a lower autonomous system number.
- E. Router R2 will be preferred because it has the higher neighbor IP address.
- F. Router R2 will be preferred because it has a higher local preference.

Answer: F

Explanation:

The preferred path to 10.0.0.0/8 network is R2 because it has a higher local preference.

The following process summarizes how BGP chooses the best route on a Cisco router.

In this example, since the weights remained the same (default) value the next thing that is looked at is the highest local preference.

QUESTION NO: 258

What technique should be used on BGP edge routers to prevent a multi-homed autonomous system from becoming a transit system?

- A. Advertise with a high MED value all networks that are discovered via external BGP.
- B. Set the no-export community attribute on all networks that are advertised externally.
- C. Remove the AS-Path information on all routes in the BGP table prior to advertising externally.
- D. Set the origin code to incomplete for all networks that are discovered via external BGP.
- E. Only advertise networks externally if they have been discovered via internal BGP.
- F. Use an outgoing distribution list to filter all networks not originating from inside the autonomous system.

Answer: F

Explanation:

The minimum configuration that guarantees you won't become a transit AS is shown in the following example:

```
router bgp 65000neighbor 10.1.1.1 filter-list 1 out!ip as-path access-list 1 permit ^$
```

In this example, the outgoing filter list says that all traffic not sourced from the local AS, should not be advertised. This will prevent the announcement of routes that originated from other Autonomous Systems, and prevent traffic destined for networks in another AS from going through your AS. This is useful real-world information to know and if you configure your network with BGP in a multi-homed environment, odds are good that you will want to implement this kind of filter.

QUESTION NO: 259

Which two methods advertise internal networks to external ISPs via BGP? (Choose two.)

- A. using aggregate routes
- B. disabling synchronization
- C. forcing the next-hop address
- D. defining routes via the network statement

Answer: A,D

Explanation:

In BGP, the network command tells the BGP process what locally learned networks to advertise. The networks can be connected routes, static routes, or routes learned by way of a dynamic routing protocol, such as RIP. These networks must also exist in the routing table of the local router or they will not be sent out in updates. The mask keyword can be used with the network command to specify individual subnets. Routes learned by the BGP process are propagated by default but are often filtered by a routing policy.

```
Router(config-router)#network network-number [mask network-mask]
```

QUESTION NO: 260

Why should iBGP sessions be fully meshed within a Transit AS?

- A. BGP requires redundant TCP sessions between iBGP peers.
- B. A full mesh allows for optimal routing within the Transit AS.
- C. Routes learned via iBGP are never propagated to other eBGP peers.
- D. Routes learned via iBGP are never propagated to other iBGP peers.
- E. Routes learned via eBGP are never propagated to other iBGP peers.

Answer: D

Explanation:

Any two routers that have formed a TCP connection in order to exchange BGP routing information are called peers or neighbors. It is important to remember that the BGP peers will never become established unless there is IP connectivity between the two peers.

BGP does not advertise routes learned by way of iBGP peers to other iBGP peers. If BGP did, BGP routing inside the AS would present a dangerous potential for routing loops. For iBGP routers to learn about all BGP routes inside the AS, they must connect to every other iBGP router in a full iBGP mesh. This full mesh needs to be only logical, not physical. In other words, as long as the iBGP peers can connect to each other using TCP/IP, a logical full mesh can be created even if the routers are not directly connected

QUESTION NO: 261

When the BGP path selection process is being performed on a Cisco router, which BGP attribute is used first when determining the best path?

- A. local preference
- B. MED
- C. weight
- D. origin
- E. next-hop
- F. AS-path

Answer: C

Explanation: Explanation

Weight is the first attribute BGP uses in the route selection process. Route with a higher weight is preferred when multiple routes exist to the same destination.

QUESTION NO: 262

Which two conditions can cause BGP neighbor establishment to fail? (Choose two.)

- A. There is an access list blocking all TCP traffic between the two BGP neighbors.
- B. The IBGP neighbor is not directly connected.
- C. BGP synchronization is enabled in a transit autonomous system with fully-meshed IBGP neighbors.
- D. The BGP update interval is different between the two BGP neighbors.
- E. The BGP neighbor is referencing an incorrect autonomous system number in its neighbor statement.

Answer: A,E

Explanation:

BGP uses TCP port 179 to establish and maintain neighbor relationships, so any access lists or firewalls must permit this port for BGP to function.

By default, EBGP multi-hop is not enabled, so the EBGP peer must be directly connected in order for the local BGP router to know how to reach the EBGP peer.

BGP uses neighbor statements that specify the AS number of the BGP peer. If it is the same number as the local BGP router process, then the BGP router knows that IBGP is used. If the AS

number for the specified peer is different, then EBGP is used. Either way, the specified neighbor must be configured correctly, or the BGP peers will not become neighbors.

QUESTION NO: 263

Refer to the exhibit.

```
!
router bgp 65001
  neighbor 172.16.1.1 remote-as 65001
  neighbor 172.26.1.1 remote-as 65555
  neighbor 172.26.1.1 route-map local_pref in
  neighbor 172.30.1.1 remote-as 65510
!
route-map local_pref permit 10
  match ip address 65
  set local-preference 200
!
access-list 65 permit 192.168.1.0 0.0.0.255
!
```

Routing updates for the 192.168.1.0 network are being received from all three neighbors. Which statement is correct regarding the result of the configuration shown?

- A. The router will prefer the next hop of 172.16.1.1 for packets destined for the 192.168.1.0 network.
- B. The router will prefer the next hop of 172.26.1.1 for packets destined for the 192.168.1.0 network.
- C. The router will advertise the 192.168.1.0 network only to 172.30.1.1.
- D. The router will advertise the 192.168.1.0 network only to 172.26.1.1.
- E. The router will prefer the next hop of 172.26.1.1 for packets except those destined for the

Answer: B

Explanation: Explanation

The local-preference of the next hop 172.26.1.1 is set to 200, which is higher than the default value (100) so this path is preferred for packets destined to the 192.168.1.0 network.

QUESTION NO: 264 DRAG DROP

Drag and drop question. Drag each item to its proper location.

Place the BGP attributes in the correct order used for determining a route.

originate route	1st
AS-Path	2nd
weight	3rd
local preference	4th
MED	5th

Answer:

Place the BGP attributes in the correct order used for determining a route.

originate route	weight
AS-Path	local preference
weight	AS-Path
local preference	originate route
MED	MED

Explanation:

Place the BGP attributes in the correct order used for determining a route.

originate route	weight
AS-Path	local preference
weight	AS-Path
local preference	originate route
MED	MED

The point of this question is bgp 11 routing principles.

BGP 11 routing principles:


- 1, The weight is higher well.
- 2, The local priority is higher well.
- 3, As soon as gets down jumps is, or is the straight company.
- 4, The shortest AS way.
- 5, The route origins (i: network to be best, E. egp issues us this time well?: Obtains what through the route heavy distribution is worst)
- 6, MED (a generation of value) smallest better
- 7, EBGP learns well compared to IBGP
- 8, As soon as gets down jumps nearer way.
- 9, What studies most early is better.
- 10, What lowest neighbor's router-ID sends is better.
- 11, What the lowest neighbor IP address sends is better.

QUESTION NO: 265 DRAG DROP

Place the BGP commands to the proper locations

Select From here	Place here
Path selection values	show ip bgp summary
Memory Usage	show ip route bgp
AD of BGP	show ip bgp
Notification, Update	show ip bgp neighbor

Answer:

**Explanation:**

1. show ip bgp: path selection values
2. show ip bgp summary: Memory usage
3. show ip route bgp: AD of BGP
4. show ip bgp neighbor: Notification, update...

QUESTION NO: 266

Which BGP feature should be used to avoid high memory utilization on a router?

- A. soft-reconfiguration
- B. route refresh
- C. BGP communities
- D. full-mesh BGP peering

Answer: B

Explanation: Explanation

BGP routers have enormous routing tables so it uses much memory to proceed these routes. When a BGP policy is changed, the BGP session needs to be reset for the policy to take effect. But the resetting results in route churn and route flapping. There are two ways to clear a BGP session without resetting the TCP session between them (this is often called "soft reset"):

Soft-reconfiguration: stores all received (inbound) routing policy updates without modification in a table so that when a new filter is applied, the router will use this table to calculate the changes without resetting the TCP session between the two BGP peers . This is a memory-intensive (high memory utilization) method and is not recommended.

Route-refresh: allows a BGP router to request a remote peer resend its BGP Adj-RIB-Out. This allows the BGP router to reapply the inbound policy. The route-refresh capability requires no extra memory on the local router

QUESTION NO: 267 DRAG DROP

Click and drag the BGP attribute characterization on the left to the correct BGP attribute on the right.

is propagated throughout the local autonomous system	MED Attribute
is not advertised to neighboring routers	
used for one router with multiple exit points out of the autonomous system	Local Preference Attribute
is propagated between autonomous systems	
	Weight Attribute

Answer:

Click and drag the BGP attribute characterization on the left to the correct BGP attribute on the right.

is propagated throughout the local autonomous system	MED Attribute
is not advertised to neighboring routers	is propagated between autonomous systems
used for one router with multiple exit points out of the autonomous system	Local Preference Attribute
is propagated between autonomous systems	is propagated throughout the local autonomous system
	Weight Attribute
	is not advertised to neighboring routers
	used for one router with multiple exit points out of the autonomous system

Explanation:

MED Attribute

is propagated between autonomous systems

Local Preference Attribute

is propagated throughout the local autonomous system

Weight Attribute

is not advertised to neighboring routers

used for one router with multiple exit points out of the autonomous system

QUESTION NO: 268

Which Ipv4-mapped Ipv6 address is equivalent to IPv6 address ::ffff:AC11:AC11?

- A. ::ffff:10.12.10.12
- B. ::ffff:10.14.10.14
- C. ::ffff:44.49.44.49
- D. ::ffff:161.193.161.193
- E. ::ffff:172.17.172.17
- F. ::ffff:193.11.193.11

Answer: E

Reference: http://www.tcpipguide.com/free/t_IPv6IPv4AddressEmbedding-2.htm

QUESTION NO: 269

A problem was reported that the 10.10.10.0 /24 prefix was not injected into the local BGP table on RouterA. The following information is available from RouterA.

```
router bgp 65001
```

```
network 10.0.0.0
```

```
neighbor 172.16.1.1 remote-as 65002
```

```
no auto-summary
```

```
show ip route
```

```
o 10.10.10.0 /24 [110/11] via 192.168.1.1 2d00h, Ethernet0/0
```

Why is this prefix not in the local BGP table?

- A. This route is not a BGP learned route.
- B. The network command is wrong.
- C. The 172.16.1.1 neighbor is down.
- D. The prefix 10.10.10.0 /24 is not a 'connected' route.

Answer: B

Explanation:

Network command is wrong because to advertise any network through bgp, the prefix and network mask must be exact as per the IP routing table. So to advertise this ospf route of 10.10.10.0/24, the network command in bgp should be : network 10.10.10.0 mask 255.255.255.0

QUESTION NO: 270

Based on the P1R3 show ip bgp output, which statement is true?

```
P1R3#show ip bgp
BGP table version is 16, local router ID is 10.200.200.13
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i192.168.11.0	10.200.200.12	0	100	101	64998 65222 65223 i
* i	10.200.200.11	0	200	0	64998 65222 65223 i
*>i192.168.12.0	10.200.200.12	0	100	101	64998 65222 65223 i
* i	10.200.200.11	0	200	0	64998 65222 65223 I

- A.** The best path to reach the 192.168.11.0 prefix is via 10.200.200.11.
- B.** The best path to reach the 192.168.11.0 prefix is via 10.200.200.12.
- C.** The best path to reach the 192.168.11.0 prefix is via both 10.200.200.11 and 10.200.200.12; BGP will automatically load balance between the two.
- D.** The 192.168.11.0 and 192.168.12.0 prefixes were learned via EBGP from the 10.200.200.11 and 10.200.200.12 EBGP neighbors.

Answer: B

Explanation:

The best path to any given destination is noted by the ">" in the IP BGP table. In this case, the best path to 192.168.11.0 is via next hop 10.200.200.12 due to the fact that the weight is higher (101) than the path via the alternative next hop. Weight is a cisco proprietary method for path determination and the weight value is used above all other values.

QUESTION NO: 271

If no metric is specified for the routes being redistributed into IS-IS, what metric value is assigned to the routes?

- A.** 20
- B.** 10
- C.** 0
- D.** 1

Answer: C

Explanation:

The IS-IS metric must be between 1 and 63. There is no default-metric option in IS-IS—you should define a metric for each protocol. If no metric is specified for the routes being redistributed into IS-IS, a metric value of 0 is used by default.

Topic 4, Implement an IPv6 based solution, given a network design and a set of requirements

QUESTION NO: 272

Which two reductions are the correct reductions of the IPv6 address

2001:0d02:0000:0000:0014:0000:0000:0095? (Choose two)

- A. 2001:0d02:::0014:::0095
- B. 2001:d02::14::95
- C. 2001:d02:0:0:14::95
- D. 2001:d02::14:0:0:95

Answer: C,D

Explanation:

We can't use triple colons (:::) in IPv6 presentation. Also We can't use double colons (::) twice. You can use it only once in any address because if two double colons are placed in the same address, there will be no way to identify the size of each block of 0s. Remember the following techniques to shorten an IPv6 address:

- Omit leading 0s in the address field, so :0000 can be compressed to just :0 and :0d02 can be compressed to :d02 (but :1d00 can not be compressed to :1d)

- Use double colons (::), but just once, to represent a contiguous block of 0s, so

2001:0d02:0000:0000:0014:0000:0000:0095 can be compressed to 2001:0d02::14:0:0:95 or 2001:0d02:0:0:14::95

QUESTION NO: 273

What is the IPv6 address FF02::2 used for?

- A. all hosts in a local segment
- B. all routers in a local segment
- C. all hosts in a particular multicast group
- D. all routers in an autonomous system

Answer: B

Explanation:

To identify all nodes for the node-local and link-local scopes, the following multicast addresses are defined:

- FF01::1 (node-local scope all-nodes address)
- FF02::1 (link-local scope all-nodes address)

To identify all routers for the node-local, link-local, and site-local scopes, the following multicast addresses are defined:


- FF01::2 (node-local scope all-routers address)

- FF02::2 (link-local scope all-routers address)
- FF05::2 (site-local scope all-routers address)

Reference: <http://technet2.microsoft.com/windowsserver/en/library/8c07faf2-35b3-4f2a-920f-d61ad76581ad1033.mspx?mfr=true>

QUESTION NO: 274

Refer to the exhibit.



Routers R1 and R2 are IPv6 BGP peers that have been configured to support a neighbor relationship over an IPv4 internet work. Which three neighbor IP addresses are valid choices to use in the highlighted section of the exhibit? (Choose three.)

- A. ::0A43:0002
- B. 0A43:0002::
- C. ::10.67.0.2
- D. 10.67.0.2::
- E. 0:0:0:0:0:0:10.67.0.2
- F. 10.67.0.2:0:0:0:0:0:0

Answer: A,C,E

Explanation:


The automatic tunneling mechanism uses a special type of IPv6 address, termed an "IPv4-compatible" address. An IPv4-compatible address is identified by an all-zeros 96-bit prefix, and holds an IPv4 address in the low-order 32-bits. IPv4-compatible addresses are structured as follows:

0:0:0:0:0:0	IPv4 address
96 bits	32 bits

Therefore, an IPv4 address of 10.67.0.2 will be written as ::10.67.0.2 or 0:0:0:0:0:0:10.67.0.2 or ::0A43:0002 (with 10[decimal] = 0A[hexa] ; 67[decimal] = 43[hexa] ; 0[hexa] = 0[decimal] ; 2[hexa] = 2[decimal])

QUESTION NO: 275

Refer to the exhibit.



The 6to4 overlay tunnel configuration has been applied on each router to join isolated IPv6 networks over a IPv4 network. Which statements regarding the 6to4 overlay tunnel is true?

- A. The least significant 32 bits in the address referenced by the ipv6 route 2002::/16 Tunnel0 command will correspond to the interface E0/0 IPv4 address
- B. The least significant 32 bits in the address referenced by the ipv6 route 2002::/16 Tunnel0 command will correspond to the IPv4 address assigned to the tunnel source
- C. The configuration is invalid since the tunnel source command must be configured with an IPv6 address
- D. This is actually a configuration example of an IPv4-compatible tunnel and not a 6to4 tunnel
- E. This is actually a configuration example of an ISATAP overlay tunnel and not a 6to4 tunnel

Answer: B

Explanation:

6to4 tunnels use IPv6 addresses that concatenate 2002::/16 with the 32-bit IPv4 address of the edge router, creating a 48-bit prefix. The tunnel interface on R1 has an IPv6 prefix of 2002:4065:4001:1::/64, where 4065:4001 is the hexadecimal equivalent of 64.101.64.1, the IPv4 address of its interface in the IPv4 network. The tunnel interface on R2 has an IPv6 prefix of 2002:4065:4101:1::/64, where 4065:4101 is the hexadecimal equivalent of 64.101.65.1, the IPv4 address of its interface in the IPv4 network.

When R1 receives a packet with IPv6 destination address of 2002:4065:4101:1:: (from the left IPv6 network, for example) R1 will:

* Take the IPv6 destination address of that packet (2002:4065:4101:1::) and convert it into an IPv4 address. In this case, the IPv4 address is 40.65.41.01 in hexa, which is 64.101.65.1 in decimal format.* R1 encapsulates the IPv6 packet in an IPv4 packet with a destination address of 64.101.65.1; the packet is routed normally through the IPv4 network to R2* R2 receives the IPv4 packet, decapsulates and routes it normally to its final IPv6 destination.

QUESTION NO: 276

What happens when an IPv6 enabled router running 6to4 must send a packet to a remote destination and the next hop is the address of 2002::/16?

- A. The IPv6 packet has its header removed and replaced with an IPv4 header
- B. The IPv6 packet is encapsulated in an IPv4 packet using an IPv4 protocol type of 41
- C. The IPv6 packet is dropped because that destination is unable to route IPv6 packets
- D. The packet is tagged with an IPv6 header and the IPv6 prefix is included

Answer: B

Explanation:

6to4 and Teredo are dynamic tunneling techniques used by desktop operating systems to help their users gain access to the IPv6 Internet. These techniques tunnel the IPv6 packets within IPv4 packets. The 6to4 method places the IPv6 packets within IPv4 protocol 41 packets. The Teredo

method places the IPv6 packets within IPv4 packets with a UDP 3544 header.

Reference: <http://www.networkworld.com/community/node/47345>

QUESTION NO: 277

What are three IPv6 transition mechanisms? (Choose three)

- A. 6to4 tunneling
- B. VPN tunneling
- C. GRE tunneling
- D. ISATAP tunneling
- E. PPP tunneling
- F. Teredo tunneling

Answer: A,D,F

Explanation:

Below is a summary of IPv6 transition technologies:

6 to 4 tunneling: This mechanism allows IPv6 sites to communicate with each other over the IPv4 network without explicit tunnel setup. The main advantage of this technology is that it requires no end-node reconfiguration and minimal router configuration but it is not intended as a permanent solution.

ISATAP tunneling (Intra-Site Automatic Tunnel Addressing Protocol): is a mechanism for transmitting IPv6 packets over IPv4 network. The word "automatic" means that once an ISATAP server/router has been set up, only the clients must be configured to connect to it.

Teredo tunneling: This mechanism tunnels IPv6 datagrams within IPv4 UDP datagrams, allowing private IPv4 address and IPv4 NAT traversal to be used.

In fact, GRE tunneling is also a IPv6 transition mechanism but is not mentioned in BSCI so we shouldn't choose it (there are 4 types of IPv6 transition mechanisms mentioned in BSCI; they are manual, 6-to-4, Teredo and ISATAP).

QUESTION NO: 278

What are two rules for compacting IPv6 addresses? (Choose two.)

- A. Every 16-bit segment segment that consists of all zeroes can be represented with a single colon.
- B. The trailing zeroes in any 16-bit segment do not have to be written.

- C. The leading zeroes in any 16-bit segment do not have to be written.
- D. Any single, continuous string of one or more 16-bit segments that consists of all zeroes can be represented with a double colon.
- E. The maximum number of times a double colon can replace a 16-bit segment that consists of all zeroes is two.
- F. Two zeroes in the middle of any 16-bit segment do not have to be written.

Answer: C,D

Reference: <http://www.9tut.com/ipv6-tutorial>

QUESTION NO: 279

What is the difference between the IPv6 addresses ::/0 and ::/128?

- A. ::/0 is the default route, and ::/128 is the unspecified address.
- B. ::/0 is the unicast address, and ::/128 is the anycast address.
- C. ::/0 is the anycast address, and ::/128 is the multicast address.
- D. ::/0 is the unicast address, and ::/128 is the multicast address.
- E. ::/0 is the unspecified address, and ::/128 is the multicast address.
- F. ::/0 is the anycast address, and ::/128 is the default address.

Answer: A

Reference: <http://tools.ietf.org/html/rfc5156>

QUESTION NO: 280

Using the rules for IPv6 addressing, how can the address 2031:0000:240F:0000:0000:09C0:123A:121B be rewritten?

- A. 2031:0:240F::09C0:123A:121B
- B. 2031::240F::09C0:123A:121B
- C. 2031::240F:9C0::123A:121B
- D. 2031::240F:::09C0:123A:121B

Answer: A

Explanation:

Leading zeros can be truncated. For example "0000" can be just written as empty. In the above example :: indicates that it has multiple 0s in that location. Typically the IPv6 format can be written

down in three ways 1) compressed, 2) uncompressed and 3) fully uncompressed as shown below. All of the following are the same:

QUESTION NO: 281

Refer to the exhibit.

```
Router1#  
interface S1/1  
  ipv6 address  
  2001:410:FFFF:1::1/64  
  ipv6 ospf 100 area 0  
  
interface S2/0  
  ipv6 address  
  3FFF:B00:FFFF:1::2/64  
  ipv6 ospf 100 area 1  
  
  ipv6 router ospf 100  
    router-id 10.1.1.3  
  
Router2#  
interface S3/0  
  ipv6 address  
  3FFE:B00:FFFF:1::1/64  
  ipv6 ospf 100 area 1  
  
  ipv6 router ospf 100  
    router-id 10.1.1.4
```

- A. Interface authentication must be configured.
- B. The routing processes must be configured with an area ID.
- C. IP unicast routing must be enabled.
- D. IPv4 addresses must be applied to the interfaces.

Answer: C

Reference:

http://www.cisco.com/en/US/docs/switches/lan/catalyst3560/software/release/12.2_55_se/configuration/guide/swiprout.html

QUESTION NO: 282

Which functionality is required within an IP router that is situated at the boundary of an IPv4 network and an IPv6 network to allow communication between IPv6-only and IPv4-only nodes?

- A. Autoconfiguration
- B. Automatic 6to4 Tunnel
- C. Automatic 6to4 Relay
- D. Network Address Translator-Protocol Translator (NAT-PT)
- E. Intrasite Automatic Tunnel Address Protocol (ISATAP)

Answer: D

Reference:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps8802/ps6969/ps1835/prod_white_paper09186a008011ff51.html (Introduction, see 4th para)

QUESTION NO: 283

During the IPv6 auto configuration, what does the device append to the 64-bit prefix that it receives from the router to create its IPv6 address?

- A. a pseudorandom generated number
- B. its locally configured IPv4 address
- C. the DHCP-supplied device ID
- D. its MAC address

Answer: D

Explanation: Explanation

The automatic configuration is a great feature of IPv6. Imagine you have to manually configure an IPv6 address with 128-bit long, what a pain! With this feature, it is no longer necessary to configure each host manually. But notice that host only autonomously configures its own Link-local address (the IP address used on a LAN). The Link-local address can be created automatically using a link-local prefix of FE80::/10 and a 64-bit interface identifier (based on 48-bit MAC

address).

For example, if your MAC address is 00:12:34:56:78:9a, your 64-bit interface identifier is 0012:34FF:FE56:789a (16-bit FFFE is inserted in the middle). And notice that the notation has been changed because IPv6 addresses require 16-bit pieces to be separated by ":".

Then, according to the RFC 3513 we need to invert the **Universal/Local** bit ("U/L" bit) in the 7th position of the first octet (start counting from 0). The "u" bit is set to 1 to indicate Universal, and it is set to zero (0) to indicate local scope. In this case we set this bit to 1 because the MAC address is universally unique. Thus the result is: 0212:34FF:FE56:789a.

Finally, add the link-local prefix FE80 to create the full IPv6 address:

FE80:0:0:0:0212:34FF:FE56:789a (or FE80::212:34FF:FE56:789a, in short form).

Note: The reason for inverting the "U/L" bit is to allow ignoring it for short values in the manual configuration case. For example, you can manually assign the short address fc80::1 instead of the long fc80:0:0:0:0200::1.

QUESTION NO: 284

Refer to the exhibit.

R1#show ipv6 neighbor

IPv6 Address	Age	Link-layer Addr	State	Interface
FE80::21E:79FF:FEAB:3141	2	001e.79ab.3141	STALE	Gi0/1

Which statement about this neighbor of R1 is true?

- A. OSPFv3 adjacency has been lost, which causes the neighbor to be considered Stale.
- B. Aggregate global addresses are always used between IPv6 neighbors.
- C. OSPFv3 adjacency will not work between link-local addresses.
- D. R1 used ICMP to learn about this neighbor.

Answer: D

Explanation:

ICMP is a ping echo. IPv6 uses ICMP to learn about its neighbor.

QUESTION NO: 285

Which three are characteristics of IPv6? (Choose three.)

- A. An IPv6 address is 128 bits long.
- B. An IPv6 header is 20 bits long.
- C. An IPv6 header contains the next header field.
- D. An IPv6 header contains the protocol field.
- E. IPv6 routers send RA messages.
- F. An IPv6 header contains the header checksum field.

Answer: A,C,E

Explanation:

All IPv6 addresses are 128 bits long to accommodate a far larger number of stations than what was possible with the 32 bit IPv4 addresses.

The following displays the IPv6 header field in detail:

IPv6 header:

00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31																								
Version				Traffic Class								Flow Label																																											
Payload Length																Next Header								Hop Limit																															
Source address ...																																																							
Destination address ...																																																							
Data ...																																																							

Version. 4 bits. IPv6 version number.

Traffic Class. 8 bits. Internet traffic priority delivery value.

Flow Label. 20 bits. Used for specifying special router handling from source to destination(s) for a sequence of packets.

Payload Length. 16 bits unsigned. Specifies the length of the data in the packet. When cleared to zero, the option is a hop-by-hop Jumbo payload.

Next Header. 8 bits. Specifies the next encapsulated protocol. The values are compatible with those specified for the IPv4 protocol field.

Hop Limit. 8 bits unsigned. For each router that forwards the packet, the hop limit is decremented by 1. When the hop limit field reaches zero, the packet is discarded. This replaces the TTL field in the IPv4 header that was originally intended to be used as a time based hop limit.

Source address. 16 bytes. The IPv6 address of the sending node.

Destination address. 16 bytes. The IPv6 address of the destination node.

Reference: <http://www.networksorcery.com/enp/protocol/ipv6.htm>

QUESTION NO: 286

When an IPv6 enabled host boots, it sends a router solicitation (RS) message. An IPv6 router responds with a router advertisement (RA). Which two items are contained in the RA?

(Choose two.)

- A. IPv6 address for the host
- B. lifetime of the prefix
- C. prefixes for the link
- D. keepalive timers
- E. request for the local host IP address
- F. any route advertisements it has received

Answer: B,C

Explanation:

In IP Version 6, Router Advertisements have the following attributes:

- Routers advertise periodically
 - Max. time between advertisements can be in the range from 4 and 1800 seconds
 - The advertisement has a lifetime (= 0 if not a default router)
- Advertisement contains one or more prefixes
 - Prefixes have a lifetime
- Preferred lifetime
- Valid lifetime
- Specifies if stateful or stateless autoconfiguration is to be used
- Plays a key role in site renumbering

QUESTION NO: 287

Which IPv6 address correctly compresses the IPv6 unicast address 2001:0:0:0:0DB8:0:0:417A?

- A. 2001:0DB8:417A
- B. 2001::0DB8::417A
- C. 2001:::0DB8::417A
- D. 2001:0DB8:0:0:417A
- E. 2001::DB8:0:0:417A
- F. 2001:::0DB8:0:0:417A

Answer: E

Explanation:

The point of this question is the about the different form of Ipv6 address.

The IPv6 address is 128 bits long, written as eight 16-bit pieces, separated by colons. Each piece

is represented by four hexadecimal digits. You can compact multiple contiguous fields of zero even further. This is the exception to the rule that at least one digit must be present in every field. You can replace multiple fields of zeros with double colons (::).

Note that :: can replace only one set of contiguous zero fields. Multiple ::s would make the address ambiguous.

QUESTION NO: 288

Refer to the exhibit. What two statements are true? (Choose two.)

```
<output omitted>
!
FastEthernet0/0 is up, line protocol is up
  Link Local Address FE80::205:5FFF:FED3:5808, Interface ID 3
  Area 1, Process ID 1, Instance ID 0, Router ID 172.16.3.3
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec; State BDR, Priority 1
  Designated Router (ID) 172.16.6.6, local address
FE80::205:5FFF:FED3:6408
  Backup Designated router (ID) 172.16.3.3, local address
FE80::205:5FFF:FED3:5808
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:05
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 12, maximum is 12
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 172.16.6.6 (Designated Router)
  Suppress hello for 0 neighbor(s)
```

- A. Interface FastEthernet 0/0 was configured with the ipv6 ospf 1 area 1 command.
- B. OSPF version 2 has been enabled to support IPv6.
- C. The IP address of the backup designated router (BDR) is FE80::205:5FFF:FED3:5808.
- D. The output was generated by the show ip interface command.
- E. The router was configured with the commands: router ospf 1 network 172.16.6.0 0.0.0.255 area 1
- F. This is the designated router (DR) on the FastEthernet 0/0 link.

Answer: A,C

Explanation:

OSPFv3 supports IPv6. The configuration of OSPFv3 is not a subcommand mode of the router ospf command as it is in OSPFv2 configuration. For example, instead of using the network area command to identify networks that are part of the OSPFv3 network, the interfaces are directly configured to specify that IPv6 networks are part of the OSPFv3 network.

The following describes the steps to configure OSPF for IPv6:

Step 1	Complete the OSPF network strategy and planning for your IPv6 network. For example, you must decide whether multiple areas are required.
Step 2	Enable IPv6 unicast routing using the <code>ipv6 unicast-routing</code> command.
Step 3	Enable IPv6 on the interface using the <code>ipv6 ospf area</code> command.
Step 4	(Optional) Configure OSPFv3 interface specific settings, including area, router priority, and OSPFv3 path cost.
Step 5	(Optional) Configure routing specifics from router configuration mode, including router priority, route summarization, and so on.

There are several commonly used OSPFv3 show commands, including the `show ipv6 ospf [process-id] [area-id] interface [interface]` command.

QUESTION NO: 289

Refer to the exhibit. On the basis of the output from the `show ipv6 ospf interface` command, what two statements must be true? (Choose two.)

```
RTA# show ipv6 ospf interface
<Output omitted>
!
Ethernet0/0 is up, line protocol is up
  Link Local Address FE80::A8BB:CCFF:FE00:6E00, Interface ID 2
  Area 0, Process ID 1, Instance ID 0, Router ID 10.10.10.1
  Network Type BROADCAST, Cost:10
  MD5 Authentication SPI 500, secure socket state UP (errors:0)
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 11.11.11.1, local address
FE80::A8BB:CCFF:FE00:6F00
  Backup Designated router (ID) 10.10.10.1, local address
FE80::A8BB:CCFF:FE00:6E00
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 11.11.11.1 (Designated Router)
  Suppress hello for 0 neighbor(s)
```

- A. Interface Ethernet 0/0 was configured with the ipv6 ospf 1 area 1 command.
- B. Interface Ethernet 0/0 has been configured with the ipv6 ospf authentication ipsec spi 500 md5 command.
- C. OSPF version 3 is enabled to support IPv6.
- D. The IP address of the designated router (DR) is FE80::A8BB:CCFF:FE00:6E00.
- E. This is the designated router (DR) on the Ethernet 0/0 link.

Answer: B,C

Reference: http://www.cisco.com/en/US/docs/ios-xml/ios/iproute_ospf/configuration/15-mt/ip6-route-ospfv3-auth-ipsec.pdf (see page 4, step 4)

QUESTION NO: 290

Which statement is true about IPv6?

- A. Only one IPv6 address is assigned per node.
- B. Only one IPv6 address can be assigned to each interface.
- C. Each host can auto configure its address without the aid of a DHCP server.
- D. IPv6 hosts use any cast addresses to assign IP addresses to interfaces.

Answer: C

Reference:

http://www.tcpipguide.com/free/t_DHCPAutoconfigurationAutomaticPrivateIPAddressingA.htm

QUESTION NO: 291

What does the command clear ipv6 ospf process accomplish?

- A. The OSPF adjacencies are cleared and initiated again.
- B. The route table is cleared. Then the OSPF neighbors are reformed.
- C. The shortest path first (SPF) algorithm is performed on the LSA database.
- D. The OSPF database is repopulated. Then the shortest path first (SPF) algorithm is performed.

Answer: D

Explanation: Explanation

The command "clear ipv6 ospf" will clear the present routing table and force the OSPFv3 process to build a new one. This command is often used when something in the network was changed or for debugging purpose.

When the "process" keyword is added, which means "clear ipv6 ospf process", the OSPF

database is cleared and repopulated then the SPF algorithm is performed.

QUESTION NO: 292

In IPv6, the interfaces running OSPF can be configured with multiple address prefixes.

Which statement is true about the IPv6 addresses that can be included into the OSPF process?

- A. Specific addresses cannot be selected for importation into the OSPF process.
- B. Specific addresses can be selected using an ACL.
- C. Specific addresses can be selected using a route map.
- D. Specific addresses can be selected using a prefix list.

Answer: A

Explanation:

http://www.6deploy.eu/workshops/kenya_20080617/IPv6%20Routing_OSPFv3_MP-BGP.pdf
(slide 20, 3rd bullet)

QUESTION NO: 293

Which statement is true about the command `ipv6 ospf 1 area 0`?


- A. It must be issued in router global configuration mode to enable the OSPF process for IPv6.
- B. It must be issued in interface configuration mode to enable the OSPF process for IPv6.
- C. It must be issued before the network command to enable the OSPF process for IPv6.
- D. It must be issued after the network command to enable the OSPF process for IPv6.

Answer: B

Reference: http://www.6deploy.eu/workshops/kenya_20080617/IPv6%20Routing_OSPFv3_MP-BGP.pdf (Slide 29)

QUESTION NO: 294

Refer to the exhibit. Which interoperability technique implemented on the dual-stack routers would allow connectivity between IPv6 sites across automatic created tunnels using the 2002::/16 prefix?



- A. Dual Stack
- B. NAT-PT
- C. 6to4 tunnel
- D. GRE tunnel
- E. ISATAP tunnel

Answer: C

Reference: <http://hcc.dormv6.niu.edu.tw/~hccftp/journal30.pdf> (page 166, see Table 1)

QUESTION NO: 295

Refer to the exhibit. Based on the output from the show command on RT1 which statement is true?

```
RT1# show ipv6 ospf interface
FastEthernet0/0 is up, line protocol is up
Link Local Address FE80::218:B9FF:FE12:2CD1, Interface ID 4
Area 0, Process ID 1, Instance ID 0, Router ID 10.1.1.1
Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State Brother, Priority 1
Designated Router (ID) 10.1.3.1, local address FE80::218:B9FF:FECD:BEF0
Backup Designated router (ID) 10.1.2.1, local address FE80::218:B9FF:FE92:28D8
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:04
Index 1/3/3, flood queue length 0
Next 0x0 (0) 0x0 (0) 0x0 (0)
Last flood scan length is 2, maximum is 2
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 2, Adjacent neighbor count is 2
  Adjacent with neighbor 10.1.3.1 (Designated Router)
  Adjacent with neighbor 10.1.2.1 (Backup Designated Router)
Suppress hello for 0 neighbor(s)
```

- A. OSPFv3 uses global IPv6 addresses to establish neighbor adjacencies.
- B. RT1 has a subnet mask of 64 bits.
- C. RT1 has FastEthernet0/0 set as a DR for network type broadcast.
- D. OSPFv3 uses Link-local addresses to establish neighbor adjacencies.
- E. RT1 does not have a global IPv6 address set on FastEthernet0/0.
- F. OSPFv3 uses IPv4 addresses to establish neighbor adjacencies.

Answer: D

Explanation:

OSPFv3 is an updated version of OSPF designed to accommodate IPv6 natively. OSPFv3 uses the multicast address FF02::5 and FF02::6, but like EIGRP it uses its link-local address as the source address in advertisements.

QUESTION NO: 296

To enable BGP tunneling over an IPv4 backbone, the IPv4 address 192.168.30.1 is converted into a valid IPv6 address. Which three IPv6 addresses are acceptable formats for the IPv4 address? (Choose three.)

- A. 192.168.30.1:0:0:0:0:0:0:0
- B. 0:0:0:0:0:0:192.168.30.1
- C. ::192.168.30.1
- D. C0A8:1E01::
- E. 192.168.30.1::
- F. ::C0A8:1E01

Answer: B,C,F

Reference: <http://tools.ietf.org/html/rfc4291> (See topic: 2.5.5.1)

QUESTION NO: 297

What is IPv6 router solicitation?

- A. a request made by a node to join a specified multicast group
- B. a request made by a node for its IP address
- C. a request made by a node for the IP address of the DHCP server
- D. a request made by a node for the IP address of the local router

Answer: D

Reference: <http://computernetworkingnotes.com/ipv6-features-concepts-and-configurations/ipv6-neighbor-discovery.html> (See router solicitation)

QUESTION NO: 298

Your Company trainee asks you, in the context of IPv6 and OSPF, what best describes a type 9 LSA? What should you tell her?

- A. Link LSA
- B. Interarea prefix LSA for ABRs
- C. Router LSA
- D. Switch LSA
- E. Intra-area prefix LSA
- F. None of the above

Answer: E

Reference: <http://packetpushers.net/a-look-at-the-new-lsa-types-in-ospfv3-with-vyatta-and-cisco/> (See LSA Type 9)

QUESTION NO: 299

Company plans on migrating their network from IPv4 to IPv6 in the near future. Which three techniques can be used to transition from IPv4 to IPv6? (Select three)

- A. Dual stack
- B. NAT
- C. Flow label
- D. Mobile IP
- E. 6to4 tunneling
- F. Anycast
- G. MBGP

Answer: A,B,E

Reference: <http://www.inethk.asia/docs/christian-ipv6-isp-perspective-ISOC.pdf> (slide 9)

QUESTION NO: 300

Which command must be globally enabled on a Cisco router to support IPv6?

- A. ip routing ipv6
- B. ipv6 unicast-routing
- C. ipv6 routing
- D. ip classless
- E. ipv6 cef

Answer: B

Reference: <http://computernetworkingnotes.com/ipv6-features-concepts-and-configurations/routing-with-ipv6.html> (See the first para)

QUESTION NO: 301

What number is a valid representation for the

200F:0000:AB00:0000:0000:0000:0000/56 IPv6 prefix?

- A. 200F:0:0:AB/56
- B. 200F:0:AB00::/56
- C. 200F::AB00/56
- D. 200F:AB/56

Answer: B

Explanation:

The 0s are truncated.

QUESTION NO: 302

The company network is in the process of migrating the IP address scheme to use IPv6. Which of the following address types are associated with IPv6? (Select three)

- A. Unicast
- B. Private
- C. Broadcast
- D. Public
- E. Multicast

F. Anycast

Answer: A,E,F

Reference:

http://www.cisco.com/en/US/solutions/collateral/ns340/ns414/ns742/ns824/sbaBN_IPv6addrG.pdf
(see page 3, IPv6 address types)

QUESTION NO: 303

Company has migrated to IPv6 in their network. Which three IPv6 notations represent the same address? (Select three)

- A. 2031::130F::9C0:876A:130B
- B. 2031:0000:130F:0000:0000:09C0:876A:130B
- C. 2031:0:130F:::9C0:876A:130B
- D. 2031::130F:0::9C0:876A:130B
- E. 2031:0:130F:0:0:09C0:876A:130B
- F. 2031:0:130F::9C0:876A:130B

Answer: B,E,F

Explanation:

With IP version 6, octets containing all zero's can be simply represented as :, while consecutive zero fields can be represented as ::. ANSWER choices C and E are simply the shorthand version of the fully written IPv6 address shown in choice.

QUESTION NO: 304

Which statement is true concerning 6to4 tunneling?

- A. IPv4 traffic is encapsulated with an IPv6 header.
- B. The edge routers can use any locally configured IPv6 address.
- C. Hosts and routers inside a 6to4 site will need a special code.
- D. An edge router must use IPv6 address of 2002::/16 in its prefix.

Answer: D

Explanation: Explanation

6to4 tunnels use IPv6 addresses that concatenate 2002::/16 with the 32-bit IPv4 address of the edge router, creating a 48-bit prefix.

QUESTION NO: 305

Which two statements are true about using IPv4 and IPv6 simultaneously on a network segment? (Choose two.)

- A. Hosts can be configured to receive both IPv4 and IPv6 addresses via DHCP.
- B. Host configuration options for IPv4 can be either statically assigned or assigned via DHCP. Host configuration options for IPv6 can be statically assigned only.
- C. IPv6 allows a host to create its own IPv6 address that will allow it to communicate to other devices on a network configured via DHCP. IPv4 does not provide a similar capability for hosts.
- D. IPv4 and IPv6 addresses can be simultaneously assigned to a host but not to a router interface.
- E. IPv6 provides for more host IP addresses but IPv4 provides for more network addresses.

Answer: A,C

Explanation:

Like DHCP in IPv4, IPv6 hosts can also be configured to acquire connectivity parameters from DHCPv6 servers. IPv4 clients use DHCP broadcasts to locate DHCP servers, and since broadcasts are extinct in IPv6, clients use specialized multicasts to locate DHCPv6 servers. These multicasts use the reserved address FF02::1:2. One notable difference between DHCP and DHCPv6 is that while DHCP can inform clients which node to use as the default gateway, DHCPv6 does not do this.

QUESTION NO: 306

Which statement describes the difference between a manually configured IPv6 in IPv4 tunnel versus an automatic 6to4 tunnel?

- A. A manually configured IPv6 in IPv4 tunnel allows multiple IPv4 destinations.
- B. An automatic 6to4 tunnel allows multiple IPv4 destinations.
- C. A manually configured IPv6 in IPv4 tunnel does not require dual-stack (IPv4 and IPv6) routers at the tunnel endpoints.
- D. An automatic 6to4 tunnel does not require dual-stack (IPv4 and IPv6) routers at the tunnel endpoints.

Answer: B

Explanation:

An automatic 6to4 tunnel allows isolated IPv6 domains to be connected over an IPv4 network to remote IPv6 networks. The key difference between automatic 6to4 tunnels and manually configured tunnels is that the tunnel is not point-to-point; it is point-to-multipoint -> it allows

multiple IPv4 destinations .

Manually 6to4 is point-to-point -> only allows one IPv4 destination.

Configuring 6to4 (manually and automatic) requires dual-stack routers (which supports both IPv4 & IPv6) at the tunnel endpoints because they are border routers between IPv4 & IPv6 networks.

(Reference: http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-tunnel_ps6441_TSD_Products_Configuration_Guide_Chapter.html#wp1055515)

QUESTION NO: 307


In a comparison of an IPv4 header with an IPv6 header, which three statements are true? (Choose three.)

- A. An IPv4 header includes a checksum. However, an IPv6 header does not include one.
- B. A router has to recompute the checksum of an IPv6 packet when decrementing the TTL.
- C. An IPv6 header is half the size of an IPv4 header.
- D. An IPv6 header has twice as many octets as an IPv4 header.
- E. An IPv6 header is simpler and more efficient than an IPv4 header.
- F. The 128-bit IPv6 address makes the IPv6 header more complicated than an IPv4 header.

Answer: A,D,E

Explanation: Explanation

The image below shows the differences between an IPv4 header and an IPv6 header:



(Reference and a good resource, too:

http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_9-3/ipv6_internals.html)

QUESTION NO: 308 DRAG DROP

Click and drag the correct techniques for transitioning networks or devices from IPv4 to IPv6 from the left to the target zone on the right.

Source	Target
NAT-PT	IPv4 to IPv6 Transition Methods
6to4 tunnels	
GRE tunnels	
route tagging	
IPsec tunnels	
ISATAP tunnels	

Answer:

Click and drag the correct techniques for transitioning networks or devices from IPv4 to IPv6 from the left to the target zone on the right.

Source	Target
NAT-PT	IPv4 to IPv6 Transition Methods
6to4 tunnels	
GRE tunnels	
route tagging	
IPsec tunnels	
ISATAP tunnels	

Explanation:

IPv4 to IPv6 Transition Methods

NAT-PT
6to4 tunnels
GRE tunnels
ISATAP tunnels

QUESTION NO: 309

Under which circumstance will a branch ISR router contain interface vlan configurations?

- A. performing inter-VLAN routing
- B. performing 802.1Q trunking
- C. performing ISL trunking
- D. Ethernet Switch Module installed
- E. ADSL WIC installed
- F. running Call Manager Express

Answer: D

Explanation:

In smaller offices, a single ISR may be used for a both remote connectivity and inter-VLAN routing. In that case, know that an Ethernet Switch Module would be required for the ISR router

QUESTION NO: 310

Refer to the exhibit.

```
R1#show ip route
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area. * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 212.50.185.126 to network 0.0.0.0

D    212.50.167.0/24 [90/160000] via 212.50.185.82, 00:05:55, Ethernet1/0
    212.50.166.0/24 is variably subnetted, 4 subnets, 2 masks
D    212.50.166.0/24 is a summary, 00:05:55, Null0
C    212.50.166.1/32 is directly connected, Loopback1
C    212.50.166.2/32 is directly connected, Loopback2
C    212.50.166.20/32 is directly connected, Loopback20
    212.50.185.0/27 is subnetted, 3 subnets
C    212.50.185.64 is directly connected, Ethernet1/0
C    212.50.185.96 is directly connected, Ethernet0/0
C    212.50.185.32 is directly connected, Ethernet2/0
D*EX 0.0.0.0/0 [170/2174976] via 212.50.185.126, 00:05:55, Ethernet0/0
    [170/2174976] via 212.50.185.125, 00:05:55, Ethernet0/0
```

How would you confirm on R1 that load balancing is actually occurring on the default-network (0.0.0.0)?

- A. Use ping and the show ip route command to confirm the timers for each default network resets to 0.
- B. Load balancing does not occur over default networks; the second route will only be used for failover.
- C. Use an extended ping along with repeated show ip route commands to confirm the gateway of last resort address toggles back and forth.
- D. Use the traceroute command to an address that is not explicitly in the routing table.

Answer: D

Explanation: Explanation

The most simple method to test load balancing is to use the "traceroute" command. If load balancing is working correctly, we will see different paths to reach the destination each time we use that command.

Unknown address will be routed via the default-network 0.0.0.0 so we must use an address that is not explicitly in the routing table.

QUESTION NO: 311

Refer to the exhibit.

```
R1#show route-map divert
route-map divert, permit, sequence 1
  Match clauses:
    ip address (access-lists): 101
  Set clauses:
    ip next-hop 212.50.185.126
    ip next-hop recursive 192.0.0.1
    ip default next-hop 212.50.185.125
  Policy routing matches: 0 packets, 0 bytes
```

Based upon the configuration, you need to understand why the policy routing match counts are not increasing. Which would be the first logical step to take?

- A. Confirm if there are other problematic route-map statements that precede divert.
- B. Check the access list for log hits.
- C. Check the routing table for 212.50.185.126.
- D. Remove any two of the set clauses. (Multiple set clause entries will cause PBR to use the routing table.)

Answer: B

Explanation: Explanation

First we should check the access-list log, if the hit count does not increase then no packets are

matched the access-list -> the policy based routing match counts will not increase.

QUESTION NO: 312

An IPv6 overlay tunnel is required to communicate with isolated IPv6 networks across an IPv4 infrastructure. There are currently five IPv6 overlay tunnel types. Which three IPv6 overlay tunnel statements are true? (Choose three.)

- A.** Overlay tunnels can only be configured between border routers capable of supporting IPv4 and IPv6.
- B.** Overlay tunnels can be configured between border routers or between a border router and a host capable of supporting IPv4 and IPv6.
- C.** Cisco IOS supports manual, generic, routing encapsulation (GRE), IPv6-compatible, 4to6, and multiprotocol Label Switching (MPLS) Overlay tunneling mechanism.
- D.** Cisco IOS supports manual, generic routing encapsulation (GRE), IPv4-compatible, 6to4, and Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) overlay tunneling mechanisms.
- E.** A manual overlay tunnel supports point-to-multipoint tunnels capable of carrying IPv6 and Connectionless Network Service (CLNS) packets.
- F.** Overlay tunneling encapsulates IPv6 packets in IPv4 packets for delivery across an IPv4 infrastructure.

Answer: B,D,F

Explanation:

B: Overlay tunnels can be configured between border routers or between a border router and a host capable of supporting IPv4 and IPv6.

D. Cisco IOS supports manual, generic routing encapsulation (GRE), IPv4-compatible, 6to4, and Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) overlay tunneling mechanisms.

F: Overlay tunneling encapsulates IPv6 packets in IPv4 packets for delivery across an IPv4 infrastructure.

QUESTION NO: 313

Refer to the exhibit.


```
access-list 101 permit ip host 172.16.0.0 0.0.255.255 any
!
route-map divert permit 10
 match ip address 101
 set ip next-hop 212.50.185.126
 set ip next-hop recursive 192.0.0.1
 set ip default next-hop 212.50.185.125
!
interface GigabitEthernet0/1
 ip address 172.16.10.1 255.255.255.0
 ip policy route-map divert
```

Which command would verify if PBR reacts to packets sourced from 172.16.0.0/16?

- A. show ip route
- B. show policy-map
- C. show access-lists
- D. show route-map

Answer: D

Explanation: Explanation

The “show route-map “route-map name” displays the policy routing match counts so we can learn if PBR reacts to packets sourced from 172.16.0.0/16 or not.

```
R1#show route-map divert
route-map divert, permit, sequence 1
Match clauses:
 ip address (access-lists): 101
Set clauses:
 ip next-hop 212.50.185.126
 ip next-hop recursive 192.0.0.1
 ip default next-hop 212.50.185.125
Policy routing matches: 0 packets, 0 bytes
```

QUESTION NO: 314

Which of the following NSAP addresses is a private, locally administered address?


- A. 39.0f01.0002.0000.0c00.1111.00
- B. 48.0f01.0002.0000.0c00.1111.00
- C. 49.0004.30ac.0000.3090.c7df.00
- D. 52.0f01.0002.0000.0c00.1111.00

Answer: C

Reference: <https://sites.google.com/site/amitsciscozone/home/is-is/nsap-addresses>

QUESTION NO: 315

A policy needs to be implemented on Router B so that any traffic sourced from 172.16.10.0/24 will be forwarded to Router C.



Which configuration on Router B will achieve the desired effect?

A. access-list 1 permit 172.16.10.0 0.0.0.255

!

interface s0

ip policy route-map policy

!

route-map policy permit 10

match ip address 1

set ip next-hop 172.16.12.3

B. access-list 1 permit 172.16.10.0 0.0.0.255

!

interface e0

ip policy route-map policy

!

route-map policy permit 10

match ip address 1

set ip next-hop 172.16.12.2

C. access-list 1 permit 172.16.10.0 0.0.0.255

!

```
interface e0
ip policy route-map policy
!
route-map policy permit 10
match ip address 1
set ip next-hop 172.16.14.4
D. access-list 1 deny 172.16.10.0 0.0.0.255
!
interface s0
ip policy route-map policy
!
route-map policy permit 10
match ip address 1
set ip next-hop 172.16.12.2
```


Answer: A

Explanation: Explanation

The “next-hop” IP address should be the E1 interface of router C (172.16.12.3).

QUESTION NO: 316

Given the network diagram.



Which address would successfully summarize only the networks seen?

- A. 192.168.0.0/24
- B. 192.168.8.0/20
- C. 192.168.8.0/21
- D. 192.168.12.0/20
- E. 192.168.16.0/21
- F. These networks cannot be summarized.

Answer: C

Explanation:

let's suppose it is a /20 then we would have addresses from 192.168.0.1 to 192.168.15.255
Now let's suppose it is a /21 then we would have addresses from 192.168.8.1 to 192.168.15.255
So both summaries encompass the networks we want to summarize but the second one is the most restrictive one as it only encompasses the networks we were asked to summarize and not others so it is the correct summary.

In fact just count the number of subnets which is 8 and find the exponent of 2 which is 8, that gives you 3 and to find out the summary mask just do $/24 - 3 = /21$.

if you make it in binary you'll find out the answer too:

the interesting octet is 3rd one so let's convert in binary.

8 **0000** 1000

9 0000 1001

10 0000 1010

11 0000 1011

12 0000 1100

13 0000 1101


14 0000 1110

15 0000 1111

I've bolded the bits in common there are 5 so $16+5=21$ which is mask and address is 192.168.8.0

QUESTION NO: 317

Refer to the Exhibit.



What is the correct configuration to enable router P4 to exchange RIP routing updates with router P1 but not with router P3?

- A.** P4(Config)# interface fa0/0
P4(Config-if)# neighbor 192.168.10.3
P4(Config-if)# passive-interface fa0/0
- B.** P4(Config)# router rip
P4(Config-router)# neighbor 192.168.10.3
P4(Config-router)# passive-interface fa0/0
- C.** P4(Config)# interface fa0/0
P4(Config-if)# neighbor 192.168.10.3
P4(Config-if)# passive interface 192.168.10.34
- D.** P4(Config)# router rip
P4(Config-router)# neighbor 192.168.10.34 no broadcast
P4(Config-router)# passive-interface fa0/0

Answer: B

Explanation:

When you configure router P1 to be the neighbor of P4 with a passive interface, the RIP routing updates will be exchanged with the neighbor only.

QUESTION NO: 318

Which two statements are true about 6to4 tunnels? (Choose two.)

- A.** In a 6to4 tunnel, the first two bytes of the IPv6 address will be 0x2002 and the next four bytes will be the hexadecimal equivalent of the IPv4 address.
- B.** In a 6to4 tunnel, the first two bytes of the IPv6 address will be locally derived and the next two bytes will be the hexadecimal equivalent of the IPv4 address.
- C.** In a 6to4 tunnel, the IPv4 address 192.168.99.1 would be converted to the 2002:c0a8:6301::/48 IPv6 address.
- D.** In a 6to4 tunnel, the IPv4 address 192.168.99.1 would be converted to the 2002:c0a8:6301::/16 IPv6 address.
- E.** In a 6to4 tunnel, the IPv4 address 192.168.99.1 would be converted to the 2002:1315:4463:1::/64 IPv6 address.

Answer: A,C

Explanation:

In a 6to4 tunnel, the first two bytes of the IPv6 address will be 0x2002 and the next four bytes will be the hexadecimal equivalent of the IPv4 address. The IPv4 address 192.168.99.1 would be converted to the 2002:c0a8:6301::/48 IPv6 address.

QUESTION NO: 319

Which two statements about 6to4 tunneling are accurate? (Choose two.)

- A.** Prepending a reserved IPv6 code to the hexadecimal representation of 192.168.0.1 facilitates 6to4 tunneling.
- B.** Each 6to4 site receives a /48 prefix in a 6to4 tunnel.
- C.** 2002::/48 is the address range specifically assigned to 6to4.
- D.** Prepending 0x2002 with the IPv4 address creates an IPv6 address that is used in 6to4 tunneling.
- E.** 6to4 is a manual tunnel method.

Answer: B,D

Reference: <http://msdn.microsoft.com/en-us/library/aa916400.aspx>

QUESTION NO: 320

What are three reasons to control routing updates via route filtering? (Choose three).

- A. to hide certain networks from the rest of the organization
- B. for easier implementation
- C. to control network overhead on the wire
- D. for simple security
- E. to prevent adjacencies from forming


Answer: A,C,D

Explanation:

Route filtering hide certain networks from the rest of the organization and it also controls network overhead. Not only this, it also provides security to the routing updates.

QUESTION NO: 321

Refer to the exhibit.



Which interoperability technique implemented on the router would allow Host-1 to communicate with Host-2?

- A. Dual Stack
- B. NAT-PT
- C. 6to4 tunnel
- D. GRE tunnel
- E. ISATAP tunnel


Answer: B

Reference:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps8802/ps6969/ps1835/prod_white_paper09186a008011ff51.html

QUESTION NO: 322

Refer to the exhibit.



Which statement is true?

- A. RTA will redistribute the RIP routers into the NSSA as type 7 LSAs. RTB will translate the type 7 LSAs into type LSAs and flood them throughout the OSPF backbone.
- B. RTA will redistribute the RIP routers into the NSSA as type 7 LSAs. RTB will flood the type 7 LSAs throughout the backbone.
- C. RTA will redistribute the RIP routers into the NSSA as type 5 LSAs. RTB will flood the type 5 LSAs throughout the backbone.
- D. RTA will redistribute the RIP routers into the NSSA as type 5 LSAs. RTB will translate the type of 5 LSAs unto type 7 LSAs and flood them throughout the OSPF backbone.
- E. RTA will not redistribute the RIP routers into the NSSA.

Answer: A

Reference:

<https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=7&cad=rja&ved=0CFoQFjAG&url=http%3A%2F%2Fdocs.dsl.net.pk%2FCisco%2FOSPF%2Fcis185-BSCI-lecture5-MultiArea-OSPF->

Part2.ppt&ei=8GWsUcX8EpG5hAeu34HAAg&usg=AFQjCNGj_IgwFjjBrjXr5hyCuTvY2icPag&sig2=wOzZ7YkgFK9aVKP8VB_b_Q&bvm=bv.47244034,d.ZG4 (See slide 9, 10 and 11)

QUESTION NO: 323

To configure 6to4 tunneling on a dual-stack edge router. Which three of the following are valid components in 6to4 Tunneling configuration? (Choose Three)

- A. IPv4 Tunnel IP address
- B. Tunnel mode (6to4)
- C. Tunnel Keepalives
- D. IPv4 Tunnel Destination
- E. IPv4 Tunnel Source
- F. 6to4 IPv6 address (within 2002::/16)

Answer: B,E,F

Reference: <http://packetlife.net/blog/2010/mar/15/6to4-ipv6-tunneling/> (see 6to4 configuration, all three bullets)

QUESTION NO: 324

Which two among the following are used to indicate external type of route in routing table? (Choose two.)

- A. D EX
- B. IA
- C. O E2
- D. R E2
- E. i L2

Answer: A,C

Reference: <http://www.xpresslearn.com/cisco/routing/route-table-explanation>

QUESTION NO: 325

Which three statements about configuring OSPF in a IPv6 network are true? (Choose three.)

- A. OSPF version 2 will support IPv6.
- B. OSPF version 3 will support IPv6.
- C. Multiple instances of OSPF for IPv6 can be run on a link.
- D. Networks must be explicitly configured using the network command in router OSPF configuration mode.

E. IPv4 addresses cannot be used as the router ID in OSPF for IPv6.

F. The interface command `ipv6 ospf <process-id> area <area-id>` is all that is required to enable OSPF for IPv6 on an interface.

Answer: B,C,F

Reference: <http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-ospf.html>

QUESTION NO: 326

Refer to the exhibit.

```
<output omitted>

ipv6 unicast routing

interface fastethernet 0/0
 ip address 192.168.200.1 255.255.255.0
 ipv6 address 3ffe:b00:c18:1::3/127
```

Which two statements are true about the router configuration? (Choose two.)


- A. This configuration allows applications on the same segment to communicate via IPv4 or IPv6.
- B. This configuration is referred to as a dual-stack 6to4 tunnel.
- C. This configuration is referred to as a dual stack.
- D. This configuration will attempt to route packets using IPv4 first, and if that fails, then IPv6.

Answer: A,C

Explanation: This router demonstrates an example of an IPv6 Dual Stack configuration. Dual stack (Figure 1 below) runs both IPv4 and IPv6 protocol stacks on a router in parallel, making it similar to the multiprotocol network environments of the past, which often ran Internetwork Packet Exchange (IPX), AppleTalk, IP, and other protocols concurrently. The technique of deploying IPv6 using dual-stack backbones allows IPv4 and IPv6 applications to coexist in a dual IP layer routing backbone. The IPv4 communication uses the IPv4 protocol stack, and the IPv6 communication uses the IPv6 stack.

As a transition strategy, dual stack is ideal for campus networks with a mixture of IPv4 and IPv6 applications.

Figure 1: Dual-Stack Example

**Reference:**

http://www.cisco.com/en/US/products/ps6553/products_white_paper0900aecd805c54d0.shtml

QUESTION NO: 327

When implementing a 6to4 tunnel, which IPv6 address is the correct translation of the IPv4 address 192.168.99.1?

- A. c0a8:6301:2002::/48
- B. 2002:c0a8:6301::/48
- C. 2002:c0a8:6301::/8
- D. 2002::/16

Answer: B


Explanation:

16 bits for the most significant 6to4 reserved bits (2002::/16) plus 32 bits source ipv4 address (translated in HEX format) = 48 bits,

Topic 5, Implement an IPv4 or IPv6 based redistribution solution, given a network design and a set of requirements

QUESTION NO: 328

Refer to the exhibit. Will redistributed RIP routes from OSPF Area 2 be allowed in Area 1?



- A. Because Area 1 is an NSSA, redistributed RIP routes will not be allowed.
- B. Redistributed RIP routes will be allowed in Area 1 because they will be changed into type 5 LSAs in Area 0 and passed on into Area 1.
- C. Because NSSA will discard type 7 LSAs, redistributed RIP routes will not be allowed in Area 1.
- D. Redistributed RIP routes will be allowed in Area 1 because they will be changed into type 7 LSAs in Area 0 and passed on into Area 1.
- E. RIP routes will be allowed in Area 1 only if they are first redistributed into EIGRP.


Answer: A

Explanation:

The following are several area types that are possible on OSPF:

QUESTION NO: 329

Study this exhibit below carefully.



What is the effect of the distribute-list command in the R1 configuration?

- A. R1 will permit only the 10.0.0.0/24 route in the R2 RIP updates
- B. R1 will not filter any routes because there is no exact prefix match
- C. R1 will filter the 10.1.0.0/24 and the 172.24.1.0/24 routes from the R2 RIP updates
- D. R1 will filter only the 172.24.1.0/24 route from the R2 RIP updates

Answer: C

Explanation:

The command "distribute-list 10 in Serial0 will create an incoming distribute list for interface serial 0 and refers to access list 10. So it will permit routing updates from 10.0.x.x network while other entries (in this case the 10.1.0.0/24 and 172.24.1.0/24 networks) will be filtered out from the routing update received on interface S0.

QUESTION NO: 330

Which three route filtering statements are true? (Choose three)

- A. After the router rip and passive-interface s0/0 commands have been issued, the s0/0 interface will not send any RIP updates, but will receive routing updates on that interface.
- B. After the router eigrp 10 and passive-interface s0/0 commands have been issued, the s0/0 interface will not send any EIGRP updates, but will receive routing updates on that interface
- C. After the router ospf 10 and passive-interface s0/0 commands have been issued , the s0/0

interface will not send any OSPF updates, but will receive routing updates on that interface

D. When you use the passive-interface command with RIPv2, multicasts are sent out the specified interface

E. When you use the passive-interface command with EIGRP, hello messages are not sent out the specified interface

F. When you use the passive-interface command with OSPF, hello messages are not sent out the specified interface

Answer: A,E,F

Explanation:

Passive-interface command is used in all routing protocols to disable sending updates out from a specific interface. However the command behavior varies from one protocol to another"

- In RIP, this command will not allow sending multicast updates via a specific interface but will allow listening to incoming updates from other RIP speaking neighbors. This means that the router will still be able to receive updates on that passive interface and use them in its routing table.

In EIGRP and OSPF the passive-interface command stops sending outgoing hello packets, hence the router can not form any neighbor relationship via the passive interface. This behavior stops both outgoing and incoming routing updates.

QUESTION NO: 331

Router RTA is configured as follows:

```
RTA (config)#router rip
```


```
RTA(config-router)#network 10.0.0.0
```

```
RTA(config-router)#distribute-list 44 in interface BRIO
```

```
RTA(config-router)#exit
```

```
RTA(config)#access-list 44 deny 172.16.1.0 0.0.0.255
```

```
RTA(config)#access-list 44 permit any
```

What are the effects of this RIP configuration on router RTA? (Choose two)

- A. no routing updates will be sent from router RTA on interface BRI0 to router RTX
- B. router RTA will not advertise the 10.0.0.0 network to router RTX
- C. the route to network 172.16.1.0 will not be entered into the routing table on router RTA
- D. user traffic from the 172.16.1.0 network is denied by access-list 44
- E. the routing table on router RTA will be updated with the route to router RTW

Answer: C,E

Explanation:

Distribute list are used to filter routing updates and they are based on access lists. In this case, an access list of 44 was created to deny the route from network 172.16.1.0/24 so this route will not be entered into the routing table of RTA. But the route from RTW can be entered because it is not filtered by the access list A and B are not correct because the distribute list is applied to the inbound direction of interface BRI0 so outgoing routing updated will not be filtered.

Distribute list just filters routing updates so user traffic from network 172.16.1.0 will not be denied.

QUESTION NO: 332

Refer to the exhibit.

```
<output omitted>
!
router ospf 10
 redistribute rip route-map rip-in
!
<output omitted>
!
route-map rip-in permit 10
 match ip address 10 20
 set metric 100
 set metric-type type-1
route-map rip-in deny 20
 match ip address 30

route-map rip-in permit 30
 set metric 200
 set metric-type type-2
!
access-list 10 permit 10.0.10.0 0.0.0.255
access-list 20 permit 192.168.1.0 0.0.0.255
access-list 30 permit 10.0.0.0 0.255.255.255
```

Which two statements are correct regarding the routes to be redistributed into OSPF?

(Choose two.)

- A. The network 192.168.1.0 will be allowed and assigned a metric of 100.
- B. The network 192.168.1.0 will be allowed and assigned a metric of 200.
- C. All networks except 10.0.0.0/8 will be allowed and assigned a metric of 200.
- D. The network 172.16.0.0/16 will be allowed and assigned a metric of 200.
- E. The network 10.0.10.0/24 will be allowed and assigned a metric of 200.

Answer: A,D

Explanation:

QUESTION NO: 333

Into which two types of areas would an area border router (ABR) inject a default route?

(Choose two.)

- A. stub
- B. the autonomous system of an exterior gateway protocol (EGP)

- C. NSSA
- D. totally stubby
- E. the autonomous system of a different interior gateway protocol (IGP)
- F. area 0

Answer: A,D

Explanation: Explanation

Both stub area & totally stubby area allow an ABR to inject a default route. The main difference between these 2 types of areas is:

+ Stub area replaces LSA Type 5 (External LSA – created by an ASBR to advertise network from another autonomous system) with a default route + Totally stubby area replaces both LSA Type 5 and LSA Type 3 (Summary LSA – created by an ABR to advertise network from other areas, but still within the AS, sometimes called interarea routes) with a default route.

Below summarizes the LSA Types allowed and not allowed in area types:

Area Type	Type 1 & 2 (within area)	Type 3 (from other areas)	Type 4	Type 5	Type 7
Standard & backbone	Yes	Yes	Yes	Yes	No
Stub	Yes	Yes	No	No	No
Totally stubby	Yes	No	No	No	No
NSSA	Yes	Yes	No	No	Yes
Totally stubby NSSA	Yes	No	No	No	Yes

QUESTION NO: 334

What two situations could require the use of multiple routing protocols? (Choose two)

- A. when using UNIX host-based routers
- B. when smaller broadcast domains are desired
- C. because having multiple routing protocols confuses hackers
- D. when migrating from an older Interior Gateway Protocol (IGP) to a new IGP
- E. when all equipment is manufactured by Cisco
- F. when there are multiple paths to destination networks

Answer: A,D

Explanation:

Simple routing protocols work well for simple networks, but networks grow and become more complex. While running a single routing protocol throughout your entire IP internetwork is desirable, multiprotocol routing is common for a number of reasons, including company mergers, multiple departments managed by multiple network administrators, multivendor environments, or simply because the original routing protocol is no longer the best choice. Often, the multiple protocols are redistributed into each other during a migration period from one protocol to the other.

QUESTION NO: 335

R3#show run | include defaultip

default-network 140.140.0.0

ip default-network 130.130.0.0

R3#show ip route | begin Gateway

Gateway of last resort is 0.0.0.0 to network 130.130.0.0

116.0.0.0/8 is variably subnetted, 5 subnets, 3 masks

C 116.16.37.0/30 is directly connected, Serial1/0.2

C 116.16.32.0/30 is directly connected, Serial2/0.2

C 116.16.34.0/28 is directly connected, Serial1/0.1

C 116.16.35.0/28 is directly connected, Serial2/0.1

S 116.0.0.0/8 [1/0] via 116.16.34.0

* 140.140.0.0/32 is subnetted, 3 subnets

O 140.140.1.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1

O 140.140.3.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1

O 140.140.2.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1

* 130.130.0.0/16 is variably subnetted, 4 subnets, 2 masks

D* 130.130.0.0/16 is a summary, 00:30:04, Null0

C 130.130.1.0/24 is directly connected, Ethernet0/0

C 130.130.2.0/24 is directly connected, Ethernet0/1

C 130.130.3.0/24 is directly connected, Ethernet1/0

D 150.150.0.0/16 [90/679936] via 116.16.35.5, 00:02:58, Serial2/0.1

Refer to the exhibit. Why is the 140.140.0.0 network not used as the gateway of last resort even though it is configured first?

- A. The last default-network statement will always be preferred.
- B. A route to the 140.140.0.0 network does not exist in the routing table.
- C. Default-network selection will always prefer the statement with the lowest IP address.
- D. A router will load balance across multiple default-networks; repeatedly issuing the show ip route command would show the gateway of last resort changing between the two networks.

Answer: B

Explanation:

As you can see in the exhibit, 140.140.0.0 doesn't appear in the routing table.

QUESTION NO: 336

How is network layer addressing accomplished in the OSI protocol suite?

- A. Internet Protocol address
- B. Media Access Control address
- C. Packet Layer Protocol address
- D. Network Service Access Point address
- E. Authority and Format Identifier address

Answer: D

Explanation:

OSI network-layer addressing is implemented by using two types of hierarchical addresses: network service access-point addresses and network-entity titles.

A network service-access point (NSAP) is a conceptual point on the boundary between the network and the transport layers. The NSAP is the location at which OSI network services are provided to the transport layer. Each transport-layer entity is assigned a single NSAP, which is individually addressed in an OSI internetwork using NSAP addresses.

QUESTION NO: 337

Refer to the exhibit.

```
<output omitted>
!
router rip
  distribute-list 2 out ethernet 0
  distribute-list 1 out
!
access-list 1 permit 10.0.0.0 0.255.255.255
access-list 2 permit 10.0.1.0 0.0.0.255
!
<output omitted>
```

On the basis of the partial configuration, which two statements are correct? (Choose two.)

- A. Only routes matching 10.0.1.0/24 will be advertised out Ethernet 0.
- B. Only routes 10.0.1.0/24 will be sent out all interfaces.
- C. Only routes 10.0.1.0/24 will be allowed in the routing table.
- D. Only routes matching 10.0.0.0/8 will be advertised out Ethernet 0.
- E. Only routes matching 10.0.0.0/8 will be advertised out interfaces other than Ethernet 0.
- F. All routes will be advertised out interfaces other than Ethernet 0.

Answer: A,E

Explanation: Explanation

In this case, the following algorithm is used when multiple distribute-lists are used:

1. First check which interface is being sent out. If it is Ethernet 0, distribute-list 2 is applied first. If the network is denied then no further checking is done for this network. But if distribute-list 2 permits that network then distribute-list 1 is also checked. If both distribute-lists allow that network then it will be sent out.
2. If the interface is not Ethernet 0 then only distribute-list 1 is applied.

Now let's take some examples. + If the advertised network is 10.0.1.0/24, it will be sent out all interfaces, including Ethernet 0. + If the advertised network is 10.0.2.0/24, it will be sent out all interfaces, excepting Ethernet 0. + If the advertised network is 11.0.0.0/8, it will be dropped.

Note: It is possible to define one interface-specific distribute-list per interface and one protocol-specific distribute-list for each process/autonomous-system.

(For more information, please read:

http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080208748.shtml)

QUESTION NO: 338

Which routing protocol will continue to receive and process routing updates from neighbors after

the passive interface router configuration command is entered?

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

Answer: B

Reference: <http://www.routeralley.com/ra/docs/rip.pdf>

QUESTION NO: 339

Which three statements are true when configuring redistribution for OSPF? (Choose three)

- A. The default metric is 10.
- B. The default metric is 20.
- C. The default metric type is 2.
- D. The default metric type is 1.
- E. Subnets do not redistribute by default.
- F. Subnets redistribute by default.

Answer: B,C,E

Reference:

http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094e9e.shtml (see redistributing routes into OSPF)

QUESTION NO: 340

What is the benefit of deploying IPv6 in a campus network using dual stack mode?

- A. Dual Stack Mode takes advantage of IPv6 over IPv4 tunnel within a network.
- B. IPv4 and IPv6 run alongside one another and have no dependency on each other to function
- C. IPv4 and IPv6 share network resources.
- D. IPv6 can depend on existing IPv4 routing, QoS, security, and multicast policies.

Answer: B

Explanation:

Deploying IPv6 in the campus using the dual-stack model offers several advantages over the hybrid and service block models. The primary advantage of a dual stack model is that it does not require tunneling within the campus network. The dual stack model runs the two protocols as ships in the night, meaning that IPv4 and IPv6 run alongside one another and have no dependency on each other to function except that they share network resources. Both have independent routing.

QUESTION NO: 341

To configure 6to4 on a dual-stack edge router. Which three of the following are valid in 6to4 Tunneling configuration? (Choose three)

- A. IPv4 Tunnel IP address
- B. Tunnel mode (6to4)
- C. Tunnel Keepalives
- D. IPv4 Tunnel Destination
- E. IPv4 Tunnel Source.
- F. 6to4 IPv6 address (within 2002 /16)

Answer: B,E,F

Reference: <http://packetlife.net/blog/2010/mar/15/6to4-ipv6-tunneling/> (see 6to4 configuration, all three bullets)

QUESTION NO: 342

A network administrator is troubleshooting a redistribution of OSPF routes into EIGRP.

```
router eigrp 1
```

```
network 10.0.0.0
```

```
!
```

```
router ospf 1
```

```
network 172.10.0.0 0.0.255.255 area 0
```

```
redistribute eigrp 1
```

Given the exhibited commands, which statement is true?

- A. Redistributed routes will have an external type of 1 and a metric of 1.
- B. Redistributed routes will have an external type of 2 and a metric of 20.
- C. Redistributed routes will maintain their original OSPF routing metric.
- D. Redistributed routes will have a default metric of 0 and will be treated as reachable and advertised.
- E. Redistributed routes will have a default metric of 0 but will be treated as unreachable and not advertised.

Answer: B

Explanation:

By default, all routes redistributed into OSPF will be tagged as external type 2 (E2) with a metric of 20, except for BGP routes (with a metric of 1).

Note: The cost of a type 2 route is always the external cost, irrespective of the interior cost to reach that route. A type 1 cost is the addition of the external cost and the internal cost used to reach that route.

QUESTION NO: 343

Which three steps are most helpful in verifying proper route redistribution? (Choose three.)

- A. On the routers not performing the route redistribution, use the show ip route command to see if the redistributed routes show up.
- B. On the ASBR router performing the route redistribution, use the show ip protocol command to verify the redistribution configurations.
- C. On the ASBR router performing the route redistribution, use the show ip route command to verify that the proper routes from each routing protocol are there.
- D. On the routers not performing the route redistribution, use the show ip protocols command to verify the routing information sources.
- E. On the routers not performing the route redistribution, use the debug ip routing command to verify the routing updates from the ASBR.

Answer: A,B,C

Explanation:

In order to verify proper route redistribution, use the "show ip route" command on all routers within the network, as well as the ASBR, to verify that the routes are properly being advertised to all routers. In addition, issuing the "show ip protocol" can be used on the router performing the redistribution to verify that routes are being redistributed into each other.

QUESTION NO: 344

A router is configured for redistribution to advertise EIGRP routes into OSPF on a boundary router. Given the configuration:

```
router ospf 1
```

```
redistribute eigrp 1 metric 25 subnets
```

What is the function of the 25 parameter in the redistribute command?

- A. It specifies the seed cost to be applied to the redistributed routes.
- B. It specifies the administrative distance on the redistributed routes.
- C. It specifies the metric limit of 25 subnets in each OSPF route advertisement.
- D. It specifies a new process-id to inject the EIGRP routes into OSPF.


Answer: A

Reference:

<http://www.hh.se/download/18.1a63903b12e0cb97563800024/1297326173185/Manipulating+Routing+Updates.pdf> (slide 25, configuring redistribution)

QUESTION NO: 345

Refer to the exhibit.



R1 and R2 belong to the RIP routing domain that includes the networks 10.20.0.0/16 and 10.21.0.0/16. R3 and R4 are performing two-way route redistribution between OSPF and RIP. A

network administrator has discovered that R2 is receiving OSPF routes for the networks 10.20.0.0/16 and 10.21.0.0/16 and a routing loop has occurred.

Which action will correct this problem?

- A. Apply an inbound ACL to the R2 serial interface.
- B. Change the RIP administrative distance on R3 to 110.
- C. Configure distribute-lists on R3 and R4.
- D. Set the OSPF default metric to 20.
- E. Change the OSPF administrative distance on R3 to 110.

Answer: C


Explanation:

Distribute List is Like an access-list, use to deny or permit the routing update to pass through a router/interface. Distribute List allow you apply an access list to a routing updates.

It can be apply on in or out bond of an interface under a routing process. e.g in fig. R1 want to send a routing update to it neighbor, this update will go through from interface S0/0, router will check, is there some Distribute List apply to this interface. If there is a Distribute List which would contain the allow route to pass through this interface.

QUESTION NO: 346

Observe the exhibit.



If the command variance 3 were added to RTE, which path or paths would be chosen to route traffic to network X?


- A. E-B-A
- B. E-B-A and E-C-A
- C. E-C-A and E-D-A
- D. E-B-A, E-C-A and E-D-A

Answer: B

Explanation: Advertised distance of RTD is greater than FD of RTE-RTC-RTA, so the route through D will not be used.

QUESTION NO: 347

Which command should be added to RTB under router bgp 100 to allow only the external OSPF routes to be redistributed to RTC?



- A. redistribute ospf 1
- B. redistribute ospf 1 match external 1
- C. redistribute ospf 1 match external 2
- D. redistribute ospf 1 match external 1 external 2


Answer: D

Explanation:

Use the external keyword along with the redistribute command under router bgp to redistribute OSPF external routes into BGP. With the external keyword, you have three choices: 1. redistribute both external type-1 and type-2 (Default) 2. redistribute type-1 3. redistribute type-2 Enter the commands in the configuration mode as described here: RTB(config-router)# router bgp 100 RTB(config-router)# redistribute ospf 1 match external.

QUESTION NO: 348

Router E is configured with the EIGRP variance 2 command.



What path will Router E take to reach Router A?

- A. only E-D-A
- B. only E-B-A
- C. only E-C-A
- D. both E-B-A and E-C-A
- E. both E-B-A and E-D-A
- F. all available paths.


Answer: D

Explanation:

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

QUESTION NO: 349

A network administrator recently redistributed RIP routes into an OSPF domain. However, the administrator wants to configure the network so that instead of 32 external type-5 LSAs flooding into the OSPF network, there is only one.



What must the administrator do to accomplish this?

- A. Configure summarization on R1 with area 1 range 172.16.32.0 255.255.224.0
- B. Configure summarization on R1 with summary-address 172.16.32.0 255.255.224.0
- C. Configure area 1 as a stub area with area 1 stub
- D. Configure area 1 as a NSSA area with area 1 stub nssa

Answer: B

Explanation:


In many cases, the router doesn't even need specific routes to each and every subnet (for example, 172.16.1.0/24). It would be just as happy if it knew how to get to the major network (for example, 172.16.0.0/16) and let another router take it from there. In our telephone network example, the local telephone switch should only need to know to route a phone call to the switch for the called area code. Similarly, a router's ability to take a group of subnetworks and summarize them as one network (in other words, one advertisement) is called route summarization. Besides reducing the number of routing entries that a router must keep track of, route summarization can also help protect an external router from making multiple changes to its routing table due to instability within a particular subnet. For example, let's say that we were working on a router that connected to 172.16.2.0/24. As we were working on the router, we rebooted it several times. If we were not summarizing our routes, an external router would see each time 172.16.2.0/24 went away and came back. Each time, it would have to modify its own routing table. However, if our external router were receiving only a summary route (i.e., 172.16.0.0/16), then it wouldn't have to be concerned with our work on one particular subnet. This is especially a problem for EIGRP, which can create stuck in active (SIA) routes that can lead to a network melt-down.

Summarization Example We have the following networks that we want to advertise as a single

summary route: * 172.16.100.0/24 * 172.16.101.0/24 * 172.16.102.0/24 * 172.16.103.0/24 * 172.16.104.0/24 * 172.16.105.0/24 * 172.16.106.0/24

QUESTION NO: 350

Refer to the exhibit.



Routers R2, R3, R4, and R5 have OSPF enabled. What should be configured on the routers in area 1 to ensure that all default summary routes and redistributed EIGRP routes will be forwarded from R6 to area 1, and only a default route for all other OSPF routes will be forwarded from R5 to area 1.

- A. R5(config-router)# area 1 stub
R6(config-router)# area 1 stub
- B. R5(config-router)# area 1 stub no-summary
R6(config-router)# area 1 stub
- C. R5(config-router)# area 1 nssa
R6(config-router)# area 1 nssa
- D. R5(config-router)# area 1 nssa no-summary
R6(config-router)# area 1 nssa

Answer: D


Explanation:

External RIP routes are being routed in OSPF area 1 where they are injected as type 7 so we use (area 1 NSSA) command on the ASBR(R2) and (Area 1 NSSA no-summary) command on R3 and R4.

You can verify issuing the command "show ip ospf database" and you will see the type 7 lsa's on ASBR(R2) and LSA's Type 5 and 7 on both the ABR routers(R3 ,R4)

QUESTION NO: 351

Refer to the exhibit.



If R1 is configured for 6to4 tunneling, what will the prefix of its IPv6 network be?

- A. 1723:1100:1::/48
- B. FFFF:AC1F:6401::/16
- C. AC1F:6401::/32
- D. 2002:AC1F:6401::/48
- E. 3FFE:AC1F:6401::/32

Answer: D

Reference: <http://networklessons.com/ipv6/how-to-configure-ipv6-automatic-6to4-tunneling/>

Topic 6, Implement Layer 3 Path Control Solution**QUESTION NO: 352**

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

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

Answer: A,C

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

QUESTION NO: 353

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

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

Answer: C

Explanation:

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

QUESTION NO: 354

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

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

Answer: B

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

QUESTION NO: 355

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

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

Answer: C

Explanation:

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

QUESTION NO: 356

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


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

Answer: B

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

QUESTION NO: 357

Refer to the exhibit.

**Router B1 Configuration**

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

Router B1 Configuration (con't)

```
Interface f0/0
 ip address 192.168.1.1 255.255.255.0
interface Tunnel0
 ip address 172.16.1.1 255.255.255.0
 crypto map tunnel
 tunnel source F0/0
 tunnel destination 172.16.1.2
 tunnel path-mtu-discovery
 ip ospf mtu-ignore
router ospf 200
 network 10.10.1.1 0.0.0.224 area 1
 network 172.16.10.1 0.0.0.240 area 1
 network 192.168.1.0 0.0.0.255 area 0
```

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

- A. The crypto isakmp configuration is not correct.
- B. The crypto map configuration is not correct.
- C. The network 172.16.1.0 is not included in the OSPF process.
- D. The interface tunnel configuration is not correct.

Answer: A

Explanation:

The probably cause of lack connectivity is the wrong key value used

QUESTION NO: 358

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

- A. business-grade firewall
- B. premium support via SMART net
- C. site-to-site VPN for remote offices

- D. Cisco IOS software-based
- E. email security
- F. XML support

Answer: A,C,E

Reference:

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

QUESTION NO: 359

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

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

Answer: D

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

QUESTION NO: 360

What two features are benefits of using GRE tunnels with IPsec over using an IPsec tunnel alone in building-to-building site-to-site VPNs? (Choose two.)

- A. allows dynamic routing securely over the tunnel
- B. IKE keepalives are unidirectional and sent every ten seconds
- C. reduces IPsec headers overhead since tunnel mode is used
- D. supports non-IP traffic over the tunnel
- E. uses Virtual Tunnel Interface (VTI) to simplify the IPsec VPN configuration

Answer: A,D

Reference:

http://ptgmedia.pearsoncmg.com/images/9781587201509/samplechapter/158720150X_CH14.pdf

(page 332)

QUESTION NO: 361

Which statement is true about an IPsec/GRE tunnel?

- A.** The GRE tunnel source and destination addresses are specified within the IPsec transform set.
- B.** An IPsec/GRE tunnel must use IPsec tunnel mode.
- C.** GRE encapsulation occurs before the IPsec encryption process.
- D.** Crypto map ACL is not needed to match which traffic will be protected.

Answer: C

Reference: <http://www.cisco.com/en/US/ts/fn/620/fn62394.html>

QUESTION NO: 362

Which of the following is a GRE Tunnel characteristic?

- A.** GRE impose more CPU overhead than IPsec on VPN gateways
- B.** GRE tunnels can run through IPsec tunnels.
- C.** GRE Tunnel doesn't have support for IPv6
- D.** GRE consists of two sub-protocols: Encapsulated Security Payload (ESP) and Authentication Header (AH).

Answer: B

Explanation:

If you run an IPsec tunnel through a GRE tunnel then we call it as "IPsec over GRE"

QUESTION NO: 363

What are the four main steps in configuring a GRE tunnel over IPsec on Cisco routers?

(Choose Four)

- A.** Configure a physical interface or create a loopback interface to use as the tunnel endpoint.

- B.** Create the GRE tunnel interfaces.
- C.** Add the tunnel interfaces to the routing process so that it exchanges routing updates across that interface.
- D.** Add the tunnel subnet to the routing process so that it exchanges routing updates across that interface.
- E.** Add all subnets to the crypto access-list, so that IPsec encrypts the GRE tunnel traffic.
- F.** Add GRE traffic to the crypto access-list, so that IPsec encrypts the GRE tunnel traffic.

Answer: A,B,D,F

Reference:

http://www.cisco.com/public/technotes/smba/en/us/internet/Sec_GRE_tunnel_ovr_IPSEC.pdf

QUESTION NO: 364

A network administrator uses GRE over IPsec to connect two branches together via VPN tunnel. Which one of the following is the reason for using GRE over IPsec?


- A.** GRE over IPsec provides better QoS mechanism and is faster than other WAN technologies.
- B.** GRE over IPsec decreases the overhead of the header.
- C.** GRE supports use of routing protocol, while IPsec supports encryption.
- D.** GRE supports encryption, while IPsec supports use of routing protocol.

Answer: C

Reference: <http://www.dslreports.com/faq/8228> (4th para)

QUESTION NO: 365

When policy-based routing (PBR) is being configured,



Which three criteria can the set command specify? (Choose three.)

- A. all interfaces through which the packets can be routed
- B. all interfaces in the path toward the destination
- C. adjacent next hop router in the path toward the destination
- D. all routers in the path toward the destination
- E. all networks in the path toward the destination
- F. type of service and precedence in the IP packets

Answer: A,C,F

Explanation:

Configuring Policy-Based Routing (PBR):

You can configure PBR by following these steps. Some of the steps may be omitted depending on your application for PBR. For this example, note the set condition options listed in step 3 (answers are bolded).

Step1

Define and configure the route map needed for the policy. This is accomplished with the route-map command, as discussed previously.

Step2

Define and configure the match statements the route map will use. The most common match statements used are the following:

match ip address [access-list number]

The match ip address is used to call a standard, extended, or expanded-range ACL.

match length [min_packet_length_0-2147483647] [max_packet_length_0-2147483647]

The match length is used to match the Layer 3 packet length, in bytes, with all associated headers and trailers included. You must enter the minimum and maximum packet length. Use the match length command to policy route traffic based on packet size. You can deploy this to route traffic with large or small packet sizes to specific areas of the network.

Step3

Configure and define the new routing policy with set commands. Multiple set commands may be used; if multiple commands are used, they are executed in the following order:

set ip {precedence [value_0-7 | name] | tos [value_0-8 | name]}

set ip next-hop ip_address

set interface interface_name

set ip default next-hop ip_address

set default interface interface_name

Set ip precedence {[1-7]}[routine|critical|flash|flash-override|immediate|internet|network|priority]}

Reference: <http://www.ciscopress.com/articles/article.asp?p=102092&rl=1>

QUESTION NO: 366

Router R1, a branch router, connects to the Internet using DSL. Some traffic flows through a GRE and IPsec tunnel, over the DSL connection, destined for an Enterprise network. Which of the following answers best describes the router's logic that tells the router, for a given packet, to apply GRE encapsulation to the packet?

- A.** When the packet received on the LAN interface is permitted by the ACL listed on the tunnel gre *acl* command under the incoming interface
- B.** When routing the packet, matching a route whose outgoing interface is the GRE tunnel interface

- C. When routing the packet, matching a route whose outgoing interface is the IPsec tunnel interface
- D. When permitted by an ACL that was referenced in the associated crypto map

Answer: B

Explanation:

As for the correct answer, the process of routing a packet out a GRE tunnel interface triggers the GRE encapsulation action. As for the incorrect answers: There is no `tunnel gre acl` command. There is no IPsec tunnel interface. Finally, one answer refers to logic that would describe a router's logic when determining whether to encapsulate a packet into an IPsec tunnel.

QUESTION NO: 367

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

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

Answer: C


Explanation:

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

New Questions

QUESTION NO: 368

Refer to the exhibit.

**Router B1 Configuration**

```

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

```

Router B1 Configuration (con't)

```

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

```

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

What did the TAC engineer configure incorrectly?


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

Answer: B

Explanation:

QUESTION NO: 369

Refer to the exhibit.

**Router B1 Configuration**

```
crypto ipsec transform-set 10 esp-sha-hmac esp-3des
crypto map tunnel 1 ipsec-isakmp
 set transform-set 10
 set peer 192.168.2.1
 match address 102
crypto isakmp policy 1
 authentication pre-share
 group 2
crypto isakmp key ***** address 172.16.1.2
access-list 102 permit gre host 192.168.1.1 host 192.168.2.1
access-list 102 permit esp host 192.168.1.1 host 192.168.2.1
access-list 102 permit udp host 192.168.1.1 eq isakmp host 192.168.2.1
```

Router B1 Configuration (con't)

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

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

What did the TAC engineer configure incorrectly?

- A. The crypto isakmp configuration is not correct.
- B. The crypto map configuration is not correct.
- C. The interface tunnel configuration is not correct.
- D. The network configuration is not correct; network 172.16.1.0 is missing.

Answer: A

Explanation:

QUESTION NO: 370

Refer to the exhibit.

```

<output omitted>
!
ipv6 unicast-routing
!
interface Ethernet0
 ip address 192.168.99.1 255.255.255.0
!
interface Tunnel2002
 no ip address
 no ip redirects
 ipv6 address 2002:COA8:6301::1/128
 tunnel source ethernet0
 tunnel mode ipv6ip 6to4
!
ipv6 route 2002::/16 Tunnel2002
ipv6 route ::/0 2002:COA8:2101::1
!
<output omitted>

```

Given the partial configuration in the exhibit, which IPv6 statement is true?


- A. The configuration is an example of an encrypted IPv6 VPN tunnel.
- B. The configuration is an example of a one to one IPv6 tunnel.
- C. The configuration is an example of a 6to4 tunnel.
- D. The configuration is an example of a 4to6 tunnel.

Answer: C

Explanation:

QUESTION NO: 371

Refer to the exhibit.



In the network diagram, Area 1 is defined as a stub area. Because redistribution is not allowed in the stub area, EIGRP routes cannot be propagated into the OSPF domain. How does defining area 1 as a not-so-stubby area (NSSA) make it possible to inject EIGRP routes into the OSPF NSSA domain?

- A. by creating type 5 LSAs
- B. by creating type 7 LSAs
- C. by creating a link between the EIGRP domain and the RIP domain, and redistributing EIGRP into RIP
- D. by manually changing the routing metric of EIGRP so that it matches the routing metric of OSPF

Answer: B

Explanation:

QUESTION NO: 372

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


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

Answer: A,B,D

Explanation:

QUESTION NO: 373

Refer to the exhibit.



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

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

Answer: B

Explanation:

QUESTION NO: 374 DRAG DROP

Drag each OSPF state to the correct definition.

init	No information has been received, but Hello packets can still be sent to the neighbor.
loading	A Hello packet is received, but the ID of the receiving router was not included in the Hello packet.
exstart	Each router see its own Router ID in the neighbor field of the Hello packet; there is a DR / BDR election.
full	The routers and their DR and BDR establish a master-slave relationship.
2-Way	Routers exchange DBD packets that describe the contents of the entire link-state database.
down	Based on the information provided by the DBDs, routers send link-state request packets.
exchange	All the router and network LSAs are exchanged and the router databases are synchronized.

Answer:

Drag each OSPF state to the correct definition.

init	down
loading	init
exstart	2-Way
full	exstart
2-Way	exchange
down	loading
exchange	full

Explanation:

down
init
2-Way
exstart
exchange
loading
full

QUESTION NO: 375 DRAG DROP

Drag each OSPF router type to the approximate description on the left. Not all types are used.	
internal routers	have all interfaces in one area and maintain identical LSDBs
external routers	have interfaces attached to multiple areas, maintain separate LSDBs for each area
backbone routers	have at least one interface connected to area 0
ABRs	have at least one interface attached to an external internetwork such as EIGRP
ASBRs	
peer routers	

Answer:

Drag each OSPF router type to the approximate description on the left. Not all types are used.

internal routers

external routers

backbone routers

ABRs

ASBRs

peer routers

internal routers

ABRs

backbone routers

ASBRs

Explanation:

internal routers

ABRs

backbone routers

ASBRs