Exam : 642-901

Title : Building Scalable Cisco Internetworks

Ver : 10-09-07

QUESTION 1:

Certkiller uses EIGRP as the primary routing protocol in their network. How does EIGRP advertise subnet masks for each destination network?

A. EIGRP advertises a fixed length subnet mask for each destination network.
B. EIGRP advertises only a classful subnet mask for each destination network.
C. EIGRP, like IGRP and RIP, does not advertise a subnet mask for each destination network.
D. EIGRP advertises a prefix length for each destination network.
E. None of the above

Answer: D

Explanation:
Enhanced Interior Gateway Routing Protocol (EIGRP) is a Cisco-proprietary routing protocol based on IGRP. Unlike IGRP, which is a classful routing protocol, EIGRP supports CIDR, allowing network designers to maximize address space by using CIDR and VLSM. Compared to IGRP, EIGRP boasts faster convergence times, improved scalability, and superior handling of routing loops.
The prefix length field signifies the subnet mask to be associated with the network number specified in the destination field. Thus, if an EIGRP router is configured as follows:
1. ip address 172.16.1.1 255.255.255.0
it will advertise 172.16.1.0 with a prefix length of 24.
Likewise, if the router is configured as follows:
1. ip address 172.16.250.1 255.255.255.252
it will advertise 172.16.250.0 with a prefix length of 30.

QUESTION 2:

The Certkiller network uses EIGRP. Identify three characteristics of EIGRP feasible successors? (Select three)

A. Traffic will be load balanced between feasible successors with the same advertised distance.
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. A feasible successor is selected by comparing the advertised distance of a non-successor route to the feasible distance of the best route.

Answer: B, C, E

Explanation:
Once a neighbor relationship has been formed, called an Adjacency, the routers exchange routing update information and each router builds its own topology table. The Updates contain all the routes known by the sender. For each route, the receiving router calculates a distance for that route based on the distance that is conveyed and the cost to that neighbor that advertised the particular route. If the receiving router sees several routes to a particular network with different metrics, then the route with the lowest metric becomes the Feasible Distance (FD) to that network. The Feasible Distance is the metric of a network advertised by the connected neighbor plus the cost of reaching that neighbor. This path with the best metric is entered into the routing table because this is the quickest way to get to that network.
With the other possible routes to a particular network with larger metrics, the receiving router also receives the Reported Distance (RD) to this network via other routers. The Reported Distance being the total metric along a path to a destination network as advertised by an upstream neighbor. The Reported Distance for a particular route is compared with the Feasible Distance that it already has for that route. If the Reported Distance is larger than the Feasible Distance then this route is not entered into the Topology Table as a Feasible Successor. This prevents loops from occurring. If the Reported Distance is smaller than the Feasible Distance, then this path is considered to be a Feasible Successor and is entered into the Topology table. The Successor for a particular route is the neighbor/peer with the lowest metric/distance to that network.
If the receiving router has a Feasible Distance to a particular network and it receives an update from a neighbor with a lower advertised distance (Reported Distance) to that network, then there is a Feasibility Condition. In this instance, the neighbor becomes a Feasible Successor for that route because it is one hop closer to the destination network. There may be a number of Feasible Successors in a meshed network environment, up to 6 of them are entered into the Topology table thereby giving a number of next hop choices for the local router should the neighbor with the lowest metric fail. What you should note here, is that the metric for a neighbor to reach a particular network (i.e. the Reported Distance) must always be less than the metric (Feasible Distance) for the local router to reach that same network. This way routing loops are avoided. This is why routes that have Reported Distances larger than the Feasible Distance are not entered into the Topology table, so that they can never be considered as successors, since the route is likely to loop back through that local router.
Incorrect Answers:
D: The feasible successors are found in the topology table, but not the active routing table.
E: With EIGRP, traffic is load balanced across equal cost links in the routing table, but not between feasible successors.
Reference: http://www.rhyshaden.com/eigrp.htm

QUESTION 3:

Two Certkiller EIGRP routers are attempting to establish themselves as neighbors.
Which EIGRP multicast packet type is responsible for neighbor discovery?

A. Update
B. Query
C. Acknowledgment
D. Reply
E. Hello
F. None of the above

Answer: E

Explanation:
Remember that simple distance vector routers do not establish any relationship with their neighbors. RIP and IGRP routers merely broadcast or multicast updates on configured interfaces. In contrast, EIGRP routers actively establish relationships with their neighbors, much the same way that OSPF routers do.
EIGRP routers establish adjacencies with neighbor routers by using small hello packets. Hellos are sent by default every five seconds. An EIGRP router assumes that as long as it is receiving hello packets from known neighbors, those neighbors (and their routes) remain viable. By forming adjacencies, EIGRP routers do the following:
1. Dynamically learn of new routes that join their network
2. Identify routers that become either unreachable or inoperable
3. Rediscover routers that had previously been unreachable

QUESTION 4:

Certkiller has chosen to use EIGRP for their network routing protocol. Which three statements are true regarding EIGRP? (Select three)

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

Answer: B, C, D

Explanation:
EIGRP routers establish adjacencies with neighbor routers by using small hello packets. Hellos are sent by default every five seconds. An EIGRP router assumes that as long as it is receiving hello packets from known neighbors, those neighbors (and their routes) remain viable. By forming adjacencies, EIGRP routers do the following:
1. Dynamically learn of new routes that join their network
2. Identify routers that become either unreachable or inoperable
3. Rediscover routers that had previously been unreachable

EIGRP routers keep route and topology information readily available in RAM so they can react quickly to changes. Like OSPF, EIGRP keeps this information in several tables, or databases.
* Successor - A successor is a route selected as the primary route to use to reach a destination. Successors are the entries kept in the routing table. Multiple successors for a destination can be retained in the routing table.
* Feasible successor - A feasible successor is a backup route. These routes are selected at the same time the successors are identified, but are kept in the topology table. Multiple feasible successors for a destination can be retained in the topology table.
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 discontiguous 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 5:

You need to configure EIGRP on all routers within the Certkiller network. Which two statements are characteristics of the routing protocol EIGRP? (Select two)

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

Answer: B, D

Explanation:
EIGRP routers establish adjacencies with neighbor routers by using small hello packets. Hellos are sent by default every five seconds. An EIGRP router assumes that as long as it is receiving hello packets from known neighbors, those neighbors (and their routes) remain viable. By forming adjacencies, EIGRP routers do the following:
1. Dynamically learn of new routes that join their network
2. Identify routers that become either unreachable or inoperable

3. Rediscover routers that had previously been unreachable

Both EIGRP and IGRP use the following metric calculation:

metric = [K1 * bandwidth + (K2 * bandwidth)/(256 - load)+(K3 * delay)]* [K5/(reliability + K4)]

The following are the default constant values:

K1 = 1, K2 = 0, K3 = 1, K4 = 0, K5 = 0

When K4 and K5 are 0, the [K5/reliability + K4)] portion of the equation is not factored in to the metric. Therefore, with the default constant values, the metric equation is as follows:

metric = bandwidth + delay

IGRP and EIGRP, which scales the value of 256, use the following equations to determined the values used in the metric calculation:

bandwidth for IGRP = (10000000/bandwidth)
bandwidth for EIGRP = (10000000/bandwidth)*256
delay for IGRP = delay/10
delay for EIGRP = delay/10*256

QUESTION 6:

You have been tasked with setting up summarization in the Certkiller EIGRP network. Which two statements are true about EIGRP manual summarization? (Select two)

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

Answer: A, E

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 discontiguous 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
EIGRP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip summary-address eigrp command, which has the following syntax:

Router(config-if)#ip summary-address eigrp autonomous-system-number
ip-address mask administrative-distance
D 2.1.0.0/16 is a summary, 00:00:22, Null0
Notice that the summary route is sourced from Null0, and not an actual interface.

QUESTION 7:

Router CK1 4 is configured as an EIGRP stub router. What are three key concepts that apply when configuring the EIGRP stub routing feature in a hub and spoke network? (Select three)

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

Answer: A, C, F

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.
More on:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s15/eigrpstb.

QUESTION 8:

You want to reduce the amount of EIGRP traffic across low speed links in the Certkiller network. Which is the most effective technique to contain EIGRP queries?

A. Using a hierarchical addressing scheme
B. Configuring route filters
C. Establishing separate autonomous systems
D. Route summarization
E. None of the above

Answer: D

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 discontiguous 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
EIGRP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip summary-address eigrp command, which has the following syntax:
Router(config-if)#ip summary-address eigrp autonomous-system-number ip-address mask administrative-distance

---

QUESTION 9:

Routers CK1 and CK2 have formed an EIGRP neighbor relationship. In order for two routers to become EIGRP neighbors, which two values must match? (Select two)

A. K values
B. Delay
C. Autonomous system
D. Hello time
E. Hold time
F. Bandwidth

Answer: A, C

Explanation:
Despite being compatible with IGRP, EIGRP uses a different metric calculation and hop-count limitation. EIGRP scales the IGRP metric by a factor of 256. That is because EIGRP uses a

metric that is 32-bits long, and IGRP uses a 24-bit metric. By multiplying or dividing by 256, EIGRP can easily exchange information with IGRP.
EIGRP also imposes a maximum hop limit of 224, which is slightly less than the 255 limit for IGRP. However, this is more than enough to support most of the largest internetworks. To become the neighbor K value should be matched and should belongs to same AS.

Both EIGRP and IGRP use the following metric calculation:

metric = [K1 * bandwidth + (K2 * bandwidth)/(256 - load)+(K3 * delay)]* [K5/(reliability + K4)]

The following are the default constant values:

K1 = 1, K2 = 0, K3 = 1, K4 = 0, K5 = 0

When K4 and K5 are 0, the [K5/reliability + K4)] portion of the equation is not factored in to the metric. Therefore, with the default constant values, the metric equation is as follows:

metric = bandwidth + delay

IGRP and EIGRP, which scales the value of 256, use the following equations to determined the values used in the metric calculation:

bandwidth for IGRP = (10000000/bandwidth)
bandwidth for EIGRP = (10000000/bandwidth)*256
delay for IGRP = delay/10
delay for EIGRP = delay/10*256

QUESTION 10:

Certkiller uses EIGRP as their internal routing protocol. Which three statements are true about EIGRP operation? (Select three)

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

Answer: B, D, E

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 discontiguous 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
EIGRP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip

summary-address eigrp command, which has the following syntax:
Router(config-if)#ip summary-address eigrp autonomous-system-number
ip-address mask administrative-distance

QUESTION 11:

A link failure in the Certkiller network has caused the EIGRP routers to update their routing tables. What action does an EIGRP router take when it cannot find a feasible successor for a network?

A. It examines the topology table for a next best path.
B. It transitions from passive to active state for that network and queries its neighbors.
C. It examines the routing and neighbor tables for the next best path.
D. It transitions from active to passive state for that network and queries its neighbors.
E. None of the above

Answer: B

Explanation:
* Successor - A successor is a route selected as the primary route to use to reach a destination. Successors are the entries kept in the routing table. Multiple successors for a destination can be retained in the routing table.
* Feasible successor - A feasible successor is a backup route. These routes are selected at the same time the successors are identified, but are kept in the topology table. Multiple feasible successors for a destination can be retained in the topology table.

QUESTION 12:

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

A. 1
B. 90
C. 95
D. 0
E. 5
F. 170
G. 255
H. None of the above

Answer: E

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 discontiguous 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
Default Administrative Distance

Source of Route	Default Administrative Distance
Connected Interface	0
Static Route	1
EIGRP Summary	5
External BGP	20
EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EGP	140
External EIGRP	170
Internal BGP	200
Unknown	255

QUESTION 13:

You need to alter the metrics of certain EIGRP routes in the Certkiller network. Which two parameters are used by default to compute the EIGRP metric? (Select two)

A. MTU
B. bandwidth
C. delay
D. reliability
E. load
F. Interface up time

Answer: B, C

Explanation:
Metrics are the mathematics used to select a route. The higher the metric associated with a route, the less desirable it is. For EIGRP, the Bellman-Ford algorithm uses the following equation and creates the overall 24-bit metric assigned to a route:
* metric = [(K1 × bandwidth) + [(K2 × bandwidth) ÷ (256 - load)] + (K3 × delay)] × [K5 ÷ (reliability + K4)]

The elements in this equation are as follows:
* By default, K1 = K3 = 1, K2 = K4 = K5 = 0. Therefore, by default, the metric formula reduces to:
metric = (1 × bandwidth) + (1 × delay)
metric = bandwidth + delay

QUESTION 14:

The Certkiller network is using EIGRP as the network routing protocol. Which of the following statements correctly describe features and characteristics of routing using EIGRP? (Select three)

A. It sends periodic updates every 60 seconds.
B. EIGRP uses DUAL to achieve rapid convergence.
C. Adjacencies exist between master routers (MRs) in each domain.
D. It uses multicast to discover other EIGRP routers on an internetwork.
E. EIGRP provides support for multiple network layer protocols: IPX, AppleTalk, and IP.

Answer: B, D, E

Explanation:
B: Enhanced IGRP uses the Diffusing Update based algorithm (DUAL).
D: EIGRP use hello multicast packets for neighbor discovery/recovery.
E: EIGRP supports IPX, Appletalk, and IP. Separate EIGRP instances can be created for each routed protocol.
Incorrect Answers:
A: EIGRP use triggered updates, not periodic updates.
C: Each EIGRP router maintains a neighbor table that lists adjacent routers. However, there is no concept of master routers in EIGRP.
Reference: "Introduction to Enhanced IGRP (EIGRP)"
http://www.cisco.com/warp/public/459/7.html

QUESTION 15:

While comparing the tables associated with EIGRP and OSPF, you notice some similarities. Which EIGRP table is similar to OSPF's Link State Adjacency table?

A. Neighbor table
B. Routing table
C. Topology table
D. Successor table
E. None of the above

Answer: A

Explanation:
Each EIGRP router maintains a neighbor table that lists adjacent routers. This table is comparable to the neighbors (adjacency) database used by OSPF.
Incorrect Answers:
B: The routing table of EIGRP would compare to the routing table of OSPF
C: Both EIGRP and OSPF contain topology tables, which would compare to each other.
D: EIGRP does not have a successor table. Successors are maintained in the routing and topology tables of EIGRP.
Reference: Building Scalable Cisco Networks (Cisco Press) page 250

QUESTION 16:

DRAG DROP
Match the correct EIGRP term on the left to the slot in the middle that is adjacent to the correct description on the right.

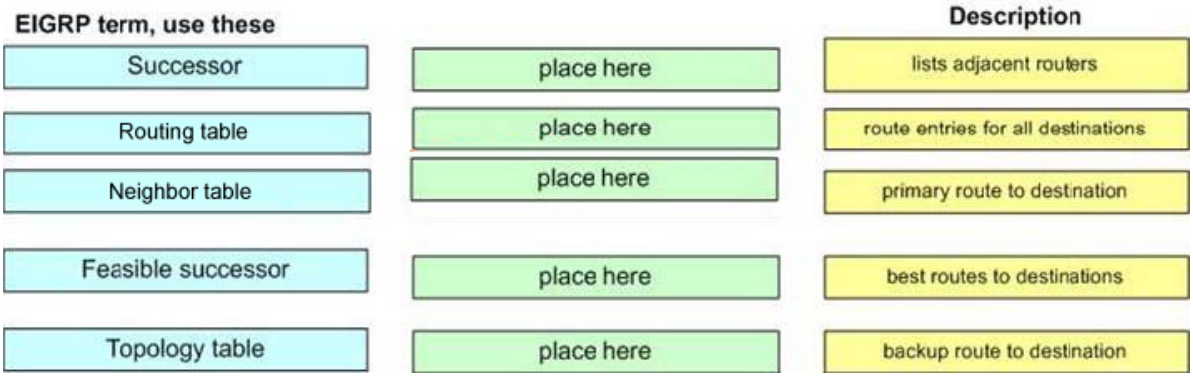

Answer:

Explanation:
Neighbor table - lists adjacent routers
Topology Table - route entries for all destinations.
Successor - primary route used to reach a destination
Routing table - best routes to a destination
Feasible successor - backup route to the destination.

Explanation:
* Neighbor Table - Each EIGRP router maintains a neighbor table that lists adjacent routers. This table is comparable to the neighbors (adjacency) database used by OSPF.
* Topology Table - An EIGRP router maintains a topology table for each network protocol configured: IP, IPX, and AppleTalk. All learned routes to a destination are maintained in the topology table.
* Routing Table - EIGRP choose the best routes to a destination from the topology table and places these routes in the routing table. The router maintains one routing table for each network protocol.
* Successor - This is the primary route used to reach a destination. Successors are kept in the routing table.
* Feasible Successor - This is a neighbor that is downstream with respect to the

destination, but it is not the least-cost path and thus is not used for forwarding data. In other words, this is a backup route to the destination. These routes are selected at the same time as successors, but are kept in the topology table.
Reference: Building Scalable Cisco Networks (Cisco Press) page 250

QUESTION 17:

Different types of EIGRP transmissions are sent between neighboring routers. Which three of the following terms are known as 'reliable packets' in EIGRP? (Select three)

A. Hello
B. ACK
C. Reply
D. Query
E. Update

Answer: C, D, E

Explanation:
Updates are used to convey the reachability of destinations. When a new neighbor is discovered, update packets are sent so the neighbor can build up its topology table. In this case, update packets are unicast. In other cases, such as a link cost change, updates are multicast. Updates are always transmitted reliably.
Queries and replies are sent when destinations go into Active state. Queries are always multicast unless they are sent in response to a received query. In this case, it is unicast back to the successor that originated the query. Replies are always sent in response to queries to indicate to the originator that it does not need to go into Active state because it has feasible successors. Replies are unicast to the originator of the query. Both queries and replies are transmitted reliably.
EIGRP reliable packets are: Update, Query and Reply.
EIGRP unreliable packets are: Hello and Ack.
Incorrect Answers:
A, B: Hellos are multicast for neighbor discovery/recovery. They do not require acknowledgment. A hello with no data is also used as an acknowledgment (ack). Acks are always sent using a unicast address and contain a non-zero acknowledgment number.
Reference: Cisco BSCN version 1.0 study guide, pages 6-18.

QUESTION 18:

Which one of the following statements best describes the way EIGRP advertises subnet mask information to its destination networks?

A. EIGRP advertises a prefix length for each destination network.
B. EIGRP advertises a fixed length subnet mask for each destination network.
C. EIGRP advertises only a classful subnet mask for each destination network.

D. EIGRP, like IGRP and RIP, does not advertise a subnet mask for each destination network.

Answer: A

Explanation:
EIGRP advertises a prefix length to each destination network. It supports VLSM which allows subnet masks to be automatically summarized at the network boundary, and EIGRP can also be configured to summarize on any bit boundary at any interface.
Incorrect Answers:
B: This is incorrect because EIGRP supports Variable Length subnet masks.
C: This is incorrect because EIGRP is considered to be a classless protocol, not classful.
D: EIGRP is more sophisticated then IGRP and RIP and actually does advertise a subnet mask to each destination network.

QUESTION 19:

EIGRP was designed specifically not to use excessive bandwidth for routing updates and functions across WAN links. If EIGRP is configured on an interface, what is the maximum percent of its bandwidth EIGRP would use, by default?

A. 10%
B. 25%
C. 50%
D. 75%
E. 100%

Answer: C

Explanation:
By default, EIGRP will limit itself to using no more than 50% of the available bandwidth. The maximum bandwidth that can be used is a configurable parameter.
Reference: Cisco, Configuration Notes for the Enhanced Implementation of EIGRP.
http://www.cisco.com/warp/public/103/12.html

QUESTION 20:

EIGRP supports Variable Length Subnet Masking (VLSM). Which two features of EGIRP are true regarding EIGRP support for VLSM? (Select two)

A. It advertises a routing mask
B. It is a classful routing protocol
C. It is a classless routing protocol
D. It does not advertise a routing mask.

Answer: A, C

Explanation:
EIGRP includes the routing mask when it advertises routes. Subnet mask information must be included in all routing updates for classless protocols. EIGRP is a classless protocol, meaning that is does not solely rely on the network class in order to find the subnet mask. These two features enable EIGRP to support VLSM.
Incorrect Answers
B: Classful routing protocols do not support VLSM.
D: The subnet mask must be included in the route to support VLSM.

QUESTION 21:

In an effort to decrease the size of the routing tables in the Certkiller network, summarization is being configured on all routers. Where in an EIGRP network is it most appropriate to implement route summarization?

A. At area border routers
B. At autonomous system boundary routers
C. It is done automatically by the DR router in each area
D. Manually at any interface of any EIGRP router within the network
E. In backbone area routers.

Answer: B

Explanation:
Enhanced IGRP performs route summarization at classful network boundaries by default. Automatic route summarization occurs at major network boundaries. ASBRs (autonomous system boundary routers) are used at the major network boundaries.
Incorrect Answers
A, C, E: Area Border Router and DR are used in OSPF, but not in EIGRP. Backbone area routers are also concepts used in OSPF, but not in EIGRP.
D: Manual summarization can be done in any interface at any router within network, but this is not the preferred solution. Summarizing manually at each router can result in black hole routing. However, done properly this is an acceptable practice. If the question calls for more than one answer then the best choice would be B and D.

QUESTION 22:

The Certkiller network is using route summarization to decrease the size of the routing tables. In an EIGRP network, which of the following describes the best method for implementing summarization?

A. At WAN interfaces.
B. Manually at major network boundaries.
C. Dynamically at discontiguous interfaces.
D. Dynamically at major network boundaries.

Answer: D

Explanation:
EIGRP supports arbitrary route summarization. Route summarization takes place automatically at major network boundaries (where network are supernetted).
Note: Route summarization is the consolidation of advertised addresses.
Incorrect Answers
A: Manually configuring summarization at WAN boundaries is not necessarily recommended in all situations. If the routers belong to the same major network boundary then summarization should not be configured.
B: Route summarization is automatic in EIGRP, and it is recommended to keep the defaults.
C: Route summarization should only be applied at contiguous interfaces. Choosing this option would most likely lead to black hole routing, making many networks unreachable.

QUESTION 23:

The Certkiller network is running EIGRP on the T1 links within their NBMA frame relay network. By default what would the hold time be for hello packets across these WAN links?

A. 30 seconds
B. 60 seconds
C. 90 seconds
D. 180 seconds
E. 5 Seconds

Answer: D

Explanation:
EIGRP sends hello packets every 5 seconds on high bandwidth links and every 60 seconds on low bandwidth multipoint links. The hold time is typically three times the hello interval. In this scenario, on slow NBMA media, hold time will be 180 seconds.
Incorrect Answers:
A, B, C: Although the hello timers can be manually configured for these values, they are not the default time.
E: This is the default for high bandwidth links, such as Ethernet networks.
Reference: Enhanced Interior Gateway Routing Protocol
http://www.cisco.com/warp/public/103/eigrp1.html

QUESTION 24:

EIGRP is being used as the routing protocol within the Certkiller network. Which two of the following features allow EIGRP to support classless routing? (Select two)

A. Djikstra's algorithm
B. discontiguous subnets
C. variable length subnet masks
D. periodic update announcements
E. unequal path-cost load balancing

Answer: B, C

Explanation:
B: Discontiguous subnets are supported by EIGRP. This is a classless routing protocol feature.
C: EIGRP support variable length subnet masks (VLSM). This is a classless routing protocol feature.
Incorrect Answers
A: EIGRP use the DUAL algorithm, not the Djikstra's algorithm. The Djikstra algorithm is used by distance vector protocols. EIGRP is considered to be a hybrid routing protocol.
D: EIGRP use periodic hello messages, not update announcements. EIGRP use update packets to convey reachability of destinations. When a new neighbor is discovered, Update packets are sent so the neighbor can build up its. These are not periodic, however.
E: IGRP and EIGRP support unequal cost path load balancing, which is known as variance. However, this is not a classless feature of EIGRP.
Note: Classless routing protocols include the routing mask with the route advertisement. This enables discontiguous subnets and variable length subnet masks.
Reference: Introduction to Enhanced IGRP (EIGRP)
http://www.cisco.com/warp/public/103/1.html
How Does Unequal Cost Path Load Balancing (Variance) Work in IGRP and EIGRP?
http://www.cisco.com/warp/public/103/19.html

QUESTION 25:

EIGRP is being used as the routing protocol on the Certkiller network. While troubleshooting some network connectivity issues, you notice a large number of EIGRP SIA (Stuck In Active) messages. What causes these SIA routes? (Select two)

A. The neighboring router stops receiving ACK packets from this router.
B. The neighboring router starts receiving route updates from this router.
C. The neighboring router is too busy to answer the query (generally caused by high CPU utilization).
D. The neighboring router is having memory problems and cannot allocate the memory to process the query or build the reply packet.

Answer: C, D

Explanation:
SIA routes are due to the fact that reply packets are not received. This could be caused by a router which is unable to send reply packets. The router could have reached the limit of

its capacity, or it could be malfunctioning.
Incorrect Answers
A: Missing replies, not missing ACKs, cause SIA.
B: Routes updates do not cause SIA.
Notes: If a router does not receive a reply to all outstanding queries within 3 minutes, the route goes to the stuck in active (SIA) state. The router then resets the neighbors that fail to reply by going active on all routes known through that neighbor, and it re-advertises all routes to that neighbor.
Reference: Enhanced Interior Gateway Routing Protocol
http://www.cisco.com/warp/public/103/eigrp3.html

QUESTION 26:

EIGRP is being configured on the Certkiller network. After the initial configuration, the EIGRP routers form relationships with its peer routers. Which three of the following statements correctly describe these peer relationships? (Select three)

A. EIGRP will form neighbors if the routers are not adjacent
B. EIGRP will not form neighbors if the metric K-values do not match.
C. EIGRP will not form neighbors if the router AS numbers do not match.
D. EIGRP will form neighbors over primary and secondary interface addressing.
E. EIGRP will form neighbors even though hello and hold timers do not match on the peering interfaces.

Answer: B, C, E

Explanation:
B: To become neighbors the routers metric must be comparable.
C: EIGRP neighbors must have the same AS number.
E: It is possible for two routers to become EIGRP neighbors even though the hello and hold timers don't match.
Incorrect Answers:
A: Neighbor discovery/recovery is the process that routers use to dynamically learn of other routers on their directly attached networks.
D: EIGRP doesn't build peer relationships over secondary addresses.
Reference: Cisco White Paper, Enhanced Interior Gateway Routing Protocol
http://www.cisco.com/warp/public/103/eigrp1.html

QUESTION 27:

The Certkiller network is an EIGRP network consisting of four links (Link A, Link B, Link C, & Link D). The costs associated with each link is A=1, B=3, C=2, D=4. According to the variables below, which link is going to be the feasible successor?

A. Link A
B. Link B

C. Link C
D. Link D

Answer: C

Explanation:
Link A is the best link and it therefore the successor. Link C is the next best link, and it is therefore the feasible successor.
Incorrect Answers
A: Link A is the successor, not the feasible successor.
B: Both Link A and Link C are better than Link B.
D: The feasible successor must have Advertised Destination that is less than the Feasible Distance.
Note: Feasible distance is the best metric along a path to a destination network, including the metric to the neighbor advertising that path. Reported distance, or advertised distance, is the total metric along a path to a destination network as advertised by an upstream neighbor. A feasible successor is a path whose reported distance is less than the feasible distance.
Reference: Enhanced Interior Gateway Routing Protocol
http://www.cisco.com/warp/public/103/eigrp1.html

QUESTION 28:

Part of the routing table of router CK1 is displayed below:
S 62.99.153.0/24 [1/0] via 209.177.64.130
172.209.12.0/32 is subnetted, 1 subnets
D EX 172.209.1
[170/2590720] via 209.179.2.114, 06:47:28, Serial0/0/0.1239
62.113.17.0/24 is variably subnetted, 2 subnets, 2 masks
D EX 99.3.215.0/24
[170/27316] via 209.180.96.45, 09:52:10, FastEthernet11/0/0
[170/27316] via 209.180.96.44, 09:52:10, FastEthernet11/0/0
25.248.17.0/24
[90/1512111] via 209.179.66.25, 10:33:13, Serial0/0/0.1400001
[90/1512111] via 209.179.66.41, 10:33:13, Serial0/0/0.1402001
62.113.1.0/24 is variably subnetted, 12 subnets, 2 masks
D 62.113.1.227/32
[90/2611727] via 209.180.96.45, 10:33:13, FastEthernet1/0/0
[90/2611727] via 209.180.96.44, 10:33:13, FastEthernet1/0/0
S* 0.0.0.0/0 [1/0] via 209.180.96.14
From analyzing the above command output, what is the administrative distance of the external EIGRP routes?

A. 24
B. 32
C. 90

D. 170
E. 27316
F. None of the above

Answer: D

Explanation:
By default an external EIGRP route has a value of 170. By examining the exhibit we see that this default value of the external EIGRP routes (see D-EX in exhibit) indeed is set to 170. The first value within the brackets display the AD, so with a value of [170/27316] the AD is 170 and the metric of the route is 27316.
Incorrect Answers:
A, B: This is the subnet mask used for some of the routes in the table.
C: This is the AD of the internal EIGRP routes, which is the default
E: This is the EIGRP metric of the external EIGRP routes.
Reference: What Is Administrative Distance?
http://www.cisco.com/warp/public/105/admin_distance.html

QUESTION 29:

The Certkiller network is shown below, along with the relevant router configurations:

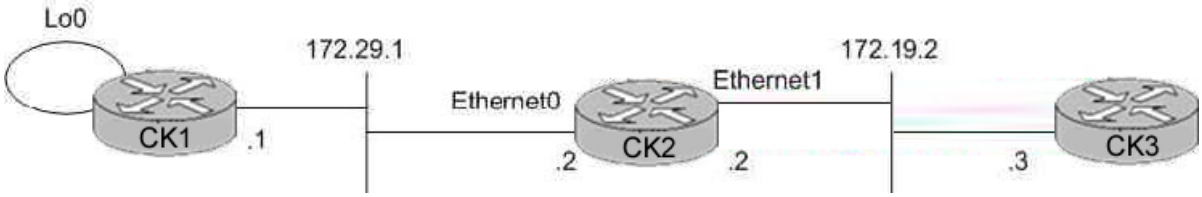


```
CK1 # show run
interface Loopback0
ip address 10.10.10.1 255.255.255.0
!
interface Ethernet0
ip address 172.29.1.1 255.255.255.0
media-type 10BaseT
!!
router eigrp 100
redistribute connected
network 172.29.0.0
auto-summary
no eigrp log-neighgbor-changes
!
ip classless
no ip http server
CK2 # show run
interface Ethernet0
ip address 172.29.1.2 255.255.255.0
```

```
media-type 10BaseT
!
interface Ethernet1
ip address 172.19.2.2 255.255.255.0
media-type 10BaseT
!
router eigrp 100
network 172.19.0.0
network 172.29.0.0
!
ip classless
no ip http server
CK3 # show run
interface Ethernet1/0
ip address 172.19.2.3 255.255.255.0
!
router eigrp 100
network 172.19.0.0
auto-summary
no eigrp log-neighbor-changes
!
ip classless
ip http server
```

With the topology found in the graphic, what will the CK1 loopback 0 be in the CK3 routing table?

A. It will show up in the routing table as D 10.0.0/8.
B. It will show up in the routing table as D EX 10.0.0.0/8.
C. It will show up in the routing table as D 10.0.0./24.
D. It will not show up in CK3 routing table because there is no network command on CK1 .

Answer: B

Explanation:
Because router CK1 is configured with route redistribution, it will redistribute the connected loopback network into EIGRP. Because redistributed routes will show up as external EIGRP routes in the routing table, choice B is correct. Although the loopback interface is using a /24 subnet mask, EIGRP summarizes at network boundaries by default so the network will appear as the class A network of 10.0.0.0/8 in the routing table of the other routers.
Incorrect Answers:
A: The route will be external, since it was redistributed into EIGRP.
C: It will be external because of redistribution, and it will also be summarized since that is the default behavior of EIGRP.

D: Although it was not configured under the EIGRP network command, it would be redistributed because it is a connected route.

QUESTION 30:

EIGRP has been configured as the routing protocol on the Certkiller network.
Which 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 destination network based on bandwidth and delay.

Answer: A, C, D

Explanation:
The default behavior of EIGRP routers is to automatically summarize routes at the network boundary.
EIGRP routers maintain information regarding the successors and the feasible successors to each network destination. This information is useful in the convergence time whenever any of the links fail.
EIGRP neighbors periodically use hellos to establish the relationship. Should any of the neighbors fail, triggered updates are sent to update the network regarding the topology change.
Incorrect Answers:
B: EIGRP networks are inherently flat, as each router is perceived as a peer to every other EIGRP router. IS-IS and OSPF utilize hierarchical network topologies.
E: The Dijkstra algorithm is used by link state routing protocols. EIGRP is considered to be a hybrid routing protocol.

QUESTION 31:

Router Certkiller 1 is the headquarters router in a hub and spoke topology supporting 24 remote offices. Point-to-multipoint Frame Relay EIGRP network is deployed between the headquarters and the remote offices. There is no bandwidth command configured under either the major serial interface or the subinterface on router Certkiller 1. What is the bandwidth of each Frame Relay connection perceived by the EIGRP process?

A. 64 kbps
B. 128 kbps
C. 1.544 Mbps
D. 1.536 Mbps

Answer: A

Explanation:
By default in point-to-multipoint Frame Relay EIGRP network all subinterfaces share the bandwidth and each one receives (throughput of the channel)/(number of subinterfaces) throughput. In this case, the formula is 1544kbps/24 =64kbps.

QUESTION 32:

What administrative distance is given to EIGRP summary routes?

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

Answer: C

Explanation:
Default administrative Distance

Connected interface or static route that identifies the outgoing interface rather than the next hop	0
Static route	1
EIGRP summary route	5
External BGP	20
EIGRP	90
IGRP	100
OSPF	110
RIP	120
External EIGRP	170
Internal BGP	200
An unknown network	255 or infinity

QUESTION 33:

Which of the following statements are true about EIGRP operation? (Choose three)

A. When summarization is configured, the router will add 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

Explanation:
Using the ip summary-address eigrp summarization command will causes the creation of an EIGRP summary default route to the null0 interface with an administrative distance of 5. Caution should be taken when using this as a default route. The low administrative distance of this default route can cause this route to displace default routes learned from other neighbors from the routing table. If the default route learned from the neighbors is displaced by the summary default route, or if the summary route is the only default route present, all traffic destined for the default route will not leave the router, instead, this traffic will be sent to the null0 interface where it is dropped.
By default, Cisco routers will automatically summarize EIGRP routes across major network boundaries. With EIGRP, 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.
Incorrect Answers:
B: When all of the specific routes used within a summary route are deleted from the routing table, the summary route will also be deleted.
D: The metric used in a summary route is the best metric from among the summarized routes, not the maximum (worst).

QUESTION 34:

What action does an EIGRP router take when it cannot find a feasible successor for a network?

A. It examines the routing and neighbor tables for the next best path.
B. It transitions from passive to active state for that network and queries its neighbors.
C. It examines the topology for a next best path.
D. It transitions from active to passive state for the next network and queries its neighbors.

Answer: B

Explanation:
With EIGRP, A topology table entry for a destination can have one of two states. A route is considered in the Passive state when a router is not performing a route recomputation. The route is in Active state when a router is undergoing a route recomputation. If there are always feasible successors, a route never has to go into Active state and avoids a route recomputation.
When there are no feasible successors, a route goes into Active state and a route

recomputation occurs. A route recomputation commences with a router sending a query packet to all neighbors. Neighboring routers can either reply if they have feasible successors for the destination or optionally return a query indicating that they are performing a route recomputation. While in Active state, a router cannot change the next-hop neighbor it is using to forward packets. Once all replies are received for a given query, the destination can transition to Passive state and a new successor can be selected.

QUESTION 35:

The Certkiller EIGRP network is displayed in the following topology diagram:

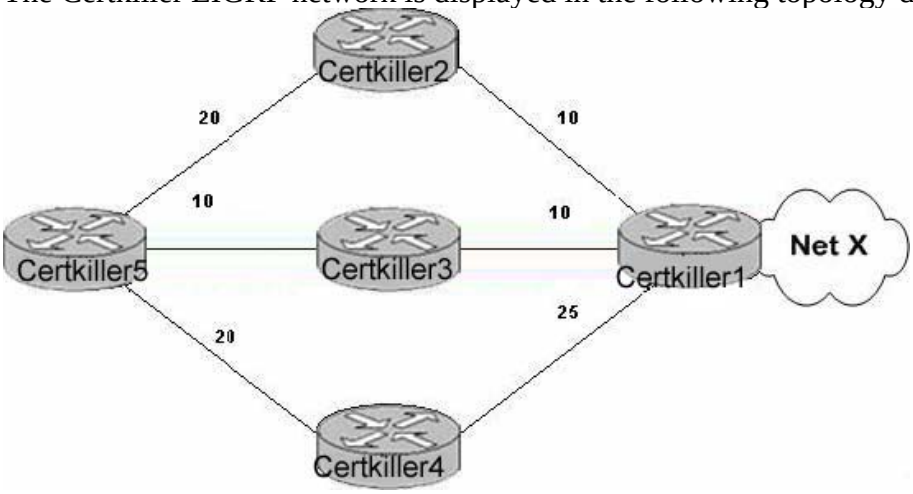

You work as a network technician at Certkiller .com. Study the exhibits carefully. If the command "variance 3" was added to the EIGRP configuration of Certkiller 5, which path or paths would be chosen to route traffic from Certkiller 5 to network X?

A. Certkiller 5- Certkiller 2- Certkiller 1
B. Certkiller 5- Certkiller 2- Certkiller 1 and Certkiller 5- Certkiller 3- Certkiller 1.
C. Certkiller 5- Certkiller 3- Certkiller 1 and Certkiller 5- Certkiller 4- Certkiller 1.
D. Certkiller 5- Certkiller 2- Certkiller 1, Certkiller 5- Certkiller 3- Certkiller 1, and Certkiller 5- Certkiller 4- Certkiller 1.

Answer: B

Explanation:
In this question the variance 3 command is used. Since the lowest metric is 10. 10*3 = 30 meaning we can use all routes with a metric of 30 and under which is answer B(5-2-1 and 5-3-1), not answer C.

QUESTION 36:

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

in 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. TheEIGRP 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 37:

The Certkiller network administrator has issued the "EIGRP stub" command on router CK8 . What is the purpose of the "EIGRP stub" configuration command?

A. to reduce the convergence time by enabling the EIGRP stub router to propagate the EIGRP queries from the EIGRP hub router.
B. to increase scalability by limiting the EIGRP query range.
C. to reduce the size of the routing table by blocking the D EX (External EIGRP) routes into the EIGRP stub router.
D. to reduce the convergence time by enabling the EIGRP stub router to also perform query requests to the EIGRP hub router.
E. None of the above.

Answer: B

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 38:

Router CK1 is an EIGRP router terminating a low speed circuit. When configuring EIGRP to run across a 56 Kbps serial PPP link, what command do you need to put under the serial interface to ensure proper convergence of EIGRP routes?

A. bandwidth 56000
B. bandwidth 56
C. ip bandwidth-percent eigrp 1 56
D. ip bandwidth-percent eigrp 1 56000
E. None of the above

Answer: B

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 39:

Router CK1 is configured as shown below:
router eigrp 100
network 10.0.0.0
eigrp stub
Based on the information shown above, which two types of routes will be advertised? (Select two)

A. Receive-only
B. Stub
C. Static
D. Summary
E. Connected
F. Dynamic

Answer: D, 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 40:

You need to configure EIGRP on a new Certkiller router. Which command should you issue first to configure EIGRP for IP?

A. router eigrp autonomous-system-number
B. ip eigrp autonomous-system-number
C. ip eigrp routing
D. router eigrp process-id
E. None of the above

Answer: A

Explanation:
Despite the complexity of DUAL, configuring EIGRP can be relatively simple. EIGRP configuration commands vary depending on the protocol that is to be routed. An example is IP, IPX, or AppleTalk. This section covers configuration commands for each of these routed protocols, in addition to special controls for IPX SAP.
Perform the following steps to configure EIGRP for IP:
1. Use the following to enable EIGRP and define the autonomous system.
Router(config)#router eigrp autonomous-system-number
The autonomous-system-number is the number that identifies the autonomous system. It is used to indicate all routers that belong within the internetwork. This value must match all routers within the internetwork.
2. Indicate which networks belong to the EIGRP autonomous system on the local router.
Router(config-router)#network network-number

The network-number determines which interfaces of the router are participating in EIGRP and which networks are advertised by the router.

QUESTION 41:

A Certkiller network administrator is troubleshooting an EIGRP configuration across a discontiguous network. What must the administrator do to ensure all of the Certkiller routers have the correct routing information?

A. The administrator must enable classless routing with the command "ip classless"
B. The administrator must disable automatic summarization with the command "no auto-summary"
C. The administrator must specify a default network with the command "ip default-network"
D. Nothing, EIGRP supports discontiguous networks by default.
E. The administrator must enable manual summarization with the command "ip summary-address"
F. None of the above

Answer: B

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 discontiguous 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 42:

The Certkiller EIGRP network is shown below:

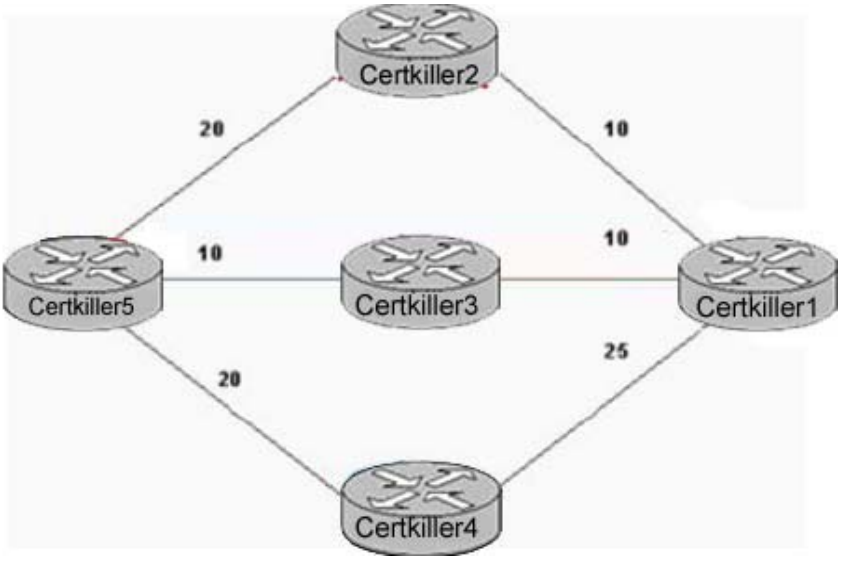

If the command "variance 3" were added to Certkiller 5, which path or paths would be chosen to route traffic to network X, which resides on the Certkiller 1 LAN?

A. Certkiller 5- Certkiller 3- Certkiller 1 and Certkiller 5- Certkiller 4- Certkiller 1
B. Certkiller 5- Certkiller 2- Certkiller 1
C. Certkiller 5- Certkiller 2- Certkiller 1, Certkiller 5- Certkiller 3- Certkiller 1 and Certkiller 5- Certkiller 4- Certkiller 1
D. Certkiller 5- Certkiller 2- Certkiller 1 and Certkiller 5- Certkiller 3- Certkiller 1
E. None of the above

Answer: D

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 43:

You need to configure route summarization on a Certkiller EIGRP router. Which is the correct command format to configure EIGRP summary route?

A. ip summary-route eigrp as-number address mask
B. ip summary-address as-number address mask
C. ip auto-summary as-number address mask
D. ip auto-summary eigrp as-number address mask
E. ip summary-address eigrp as-number address mask
F. None of the above

Answer: E

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 discontiguous 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
EIGRP also enables manual configuration of a prefix to use as a summary address. Manual summary routes are configured on a per-interface basis. The interface that will propagate the route summary must first be selected and then defined with the ip summary-address eigrp command, which has the following syntax:
Router(config-if)#ip summary-address eigrp autonomous-system-number ip-address mask administrative-distance

---

QUESTION 44:

The Certkiller EIGRP network is displayed below:

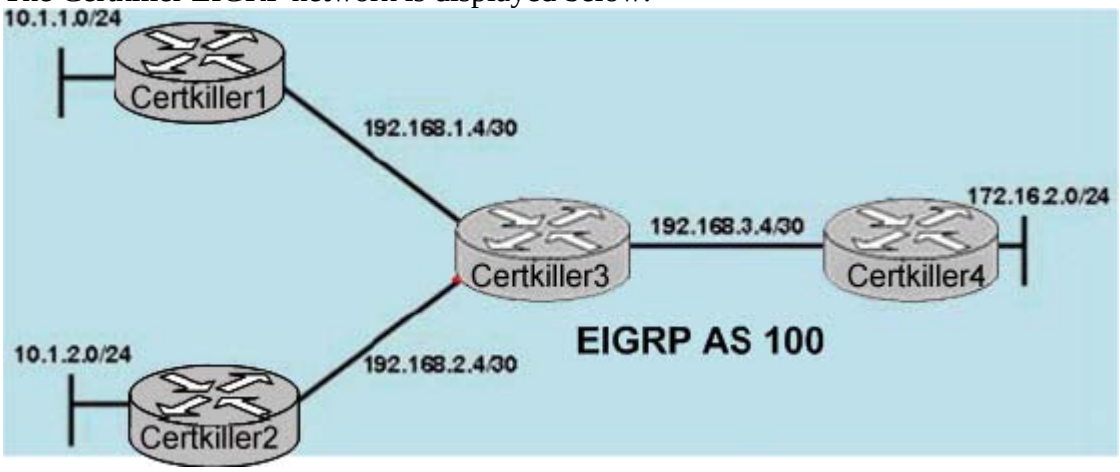

Based on the network shown above, what optional EIGRP configurations will be required in order to achieve full connectivity within AS 100?

A. Use the EIGRP "no auto-summary" command on Certkiller 1 and Certkiller 2.
B. Use the "passive interface" on the Certkiller 1 and Certkiller 2 interface that connects to the 10.1.1.0/24 and 10.1.2.0/24 subnet respectively.
C. Use the EIGRP "no auto-summary" command on Certkiller 3 and Certkiller 4.
D. Use the "passive interface" command between the Certkiller 3 and Certkiller 1 connection and between the Certkiller 3 and Certkiller 2 connection.
E. Use the "variance" command on Certkiller 3.
F. None of the above.

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 discontiguous 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 45:

The Certkiller EIGRP network is shown below:

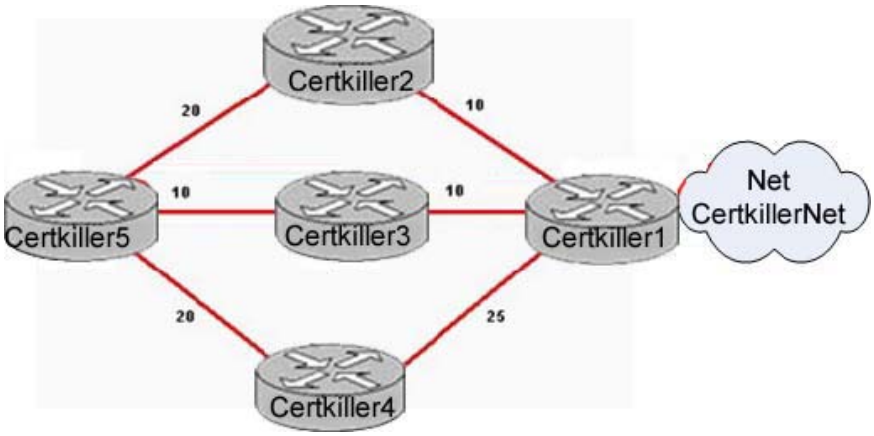

In this network, router Certkiller 5 is configured with the EIGRP "variance 2" command. What path will Router Certkiller 5 take to reach Router Certkiller 1?

A. Both Certkiller 5- Certkiller 2- CertK in1 and Certkiller 5- Certkiller 4- Certkiller 1
B. Only Certkiller 5- Certkiller 2- Certkiller 1
C. Only Certkiller 5- Certkiller 3- Certkiller 1
D. All available paths.
E. Only Certkiller 5- Certkiller 4- Certkiller 1
F. Both Certkiller 5- Certkiller 2- Certkiller 1 and Certkiller 5- Certkiller 3- Certkiller 1
G. None of the above

Answer: F

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 46:

You want to enable load sharing across the Certkiller WAN using EIGRP. Which configuration command is used to enable EIGRP unequal-cost path load balancing?

A. Variance
B. Distance
C. Maximum-paths
D. Default-metric
E. Metric
F. None of the above

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 47:

You are a network administrator on a Cisco router that's running EIGRP. You want to aggregate some of the routes. Which of the following commands below would you enter to configure the summary route you need?

A. ip auto-summary as-number address mask
B. ip summary-address as-number address mask
C. ip auto-summary eigrp as-number address mask
D. ip summary-route eigrp as-number address mask
E. ip summary-address eigrp as-number address mask

Answer: E

Explanation:
The ip summary-address eigrp command is used to configure a summary aggregate address for a specified interface.
Syntax:
ip summary-address eigrp autonomous-system-number address mask
Incorrect Answers:
A, C, D: No such command exists using this syntax.
B: We must specify the appropriate protocol with the eigrp keyword.

QUESTION 48:

Router CK1 has a 256kbps serial interface link to another Certkiller location. The configuration file of this interface is shown below:
interface serial 0/0
bandwidth 56
ip bandwidth-percent eigrp 1 200
From the command output above, how many kbps of bandwidth is allocated for EIGRP traffic?

A. 56 kbps
B. 112 kbps
C. 128 kbps
D. 256 kbps
E. None of the above

Answer: B

Explanation:
The bandwidth-percent command tells EIGRP what percentage of the configured bandwidth it may use. The default is 50 percent. Since the bandwidth command is also used to set the routing protocol metric, it may be set to a particular value for policy reasons. In this case, it is set to 56 kbps even though the actual link is 256 kbps. The bandwidth-percent command can have values greater than 100 if the bandwidth is configured artificially low due to such policy reasons. In this specific case, the bandwidth configured on the interface is 56kbps, so 200 percent of this value is 112 kbps.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/ CK2 07/technologies_tech_note09186a0080094063.shtml

---

QUESTION 49:

The Certkiller EIGRP network is displayed below:

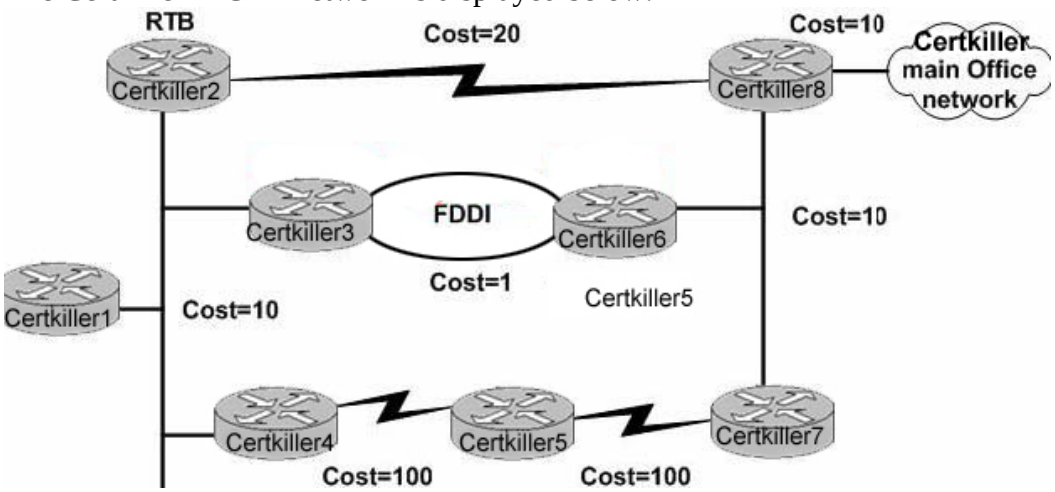

All of the routers in the above network segment are configured with EIGRP. If both

routers Certkiller 2 and Certkiller 3 were to fail, how would Certkiller 1 react?

A. Certkiller 1 will automatically route packets via Certkiller 4 to the Certkiller Main Office Network.
B. Certkiller 1 will place the route via Certkiller 4 into the hold down state.
C. Certkiller 1 will go into the active state for all routers.
D. Certkiller 1 will go into the active state for the route to Certkiller Main Office Network.
E. None of the above

Answer: D

Explanation:
First, Certkiller 1 it checks the topology table for an alternate route then it sends a query to neighboring routers for other routes to the destination. In this case, since both the optimal route and the feasible successor failed, Certkiller 1 must go into the active state first before routing around the failure via Certkiller 4. The re-route will happen only after Certkiller 1 transitions into the active state and learns about the route through Certkiller 4.

QUESTION 50:

You are the network engineer at Certkiller . The Certkiller network has five paths from Router CK1 to a given destination. The local EIGRP metric on Router CK1 for these paths are as follows:
CK1

Path 1: 1500
Path 2: 1500
Path 3: 2000
Path 4: 4000
Path 5: 4000
Variance 3 is configured on Router CK1 . Which paths would be included in Router CK1 's routing table?

A. path 1, 2, and 4
B. path 1, 2, and 3
C. path 1, 2, and 5
D. path 1, 2, 4 and 5
E. path 1, 2, 3, 4, and 5
F. None of the above

Answer: E

Explanation:
If the variance number is higher than the default 1, the EIGRP process multiples the best (lowest) cost or metric value for a path by the number stated as the variance multiplier. All paths to the same destination that have metrics within this new range are now

included in load balancing. Per the question this means: 1500*3=4500. All available options are within 1500-4500.
Reference: Self-Study CCNP BSCI Exam Certification Guide Third Edition P.485, Topic: Load Balancing in EIGRP.

QUESTION 51:

Which one of the following commands would configure EIGRP for IP on your Cisco router?

A. ip eigrp routing
B. router eigrp process-id
C. ip eigrp autonomous-system-number
D. router eigrp autonomous-system-number
E. None of the above

Answer: D

Explanation:
Perform the following steps to configure EIGRP for IP:
Step1 Enable EIGRP and define the autonomous system.
routerCK(config)#router eigrp autonomous-system-number
Step 2 Indicate which networks are part of the EIGRP autonomous system.
routerCK(config-router)#network network-number
Step 3 Define bandwidth of a link for the purposes of sending routing update traffic on the link.
routerCK(config-if)#bandwidth kilobits. This step is optional.

QUESTION 52:

The syntax for configuring variance on an EIGRP router is:
variance number
In the EIGRP command sequence above, what is the function of the number variable?

A. It acts as a multiplier.
B. It defines the limit for how far the metrics can be separated.
C. It indicates how many paths can be used for load balancing.
D. It indicates how many paths can be used for unequal load balancing.
E. None of the above

Answer: A

Explanation:
Every routing protocol supports equal cost path load balancing. IGRP and EIGRP also support unequal cost path load balancing, which is known as variance. The variance

number command instructs the router to include routes with a metric less than or equal to number times the minimum metric route for that destination. Thus number is used as multiplier, even though it defines a limit for the metrics. This limit is used for unequal load balancing.
Incorrect Answers:
B: Number does not directly define a limit.
C, D: Number does not directly indicate the number of paths that can be used.
Reference: http://www.cisco.com/warp/public/103/19.html

QUESTION 53:

The Certkiller EIGRP network is displayed in the diagram below:

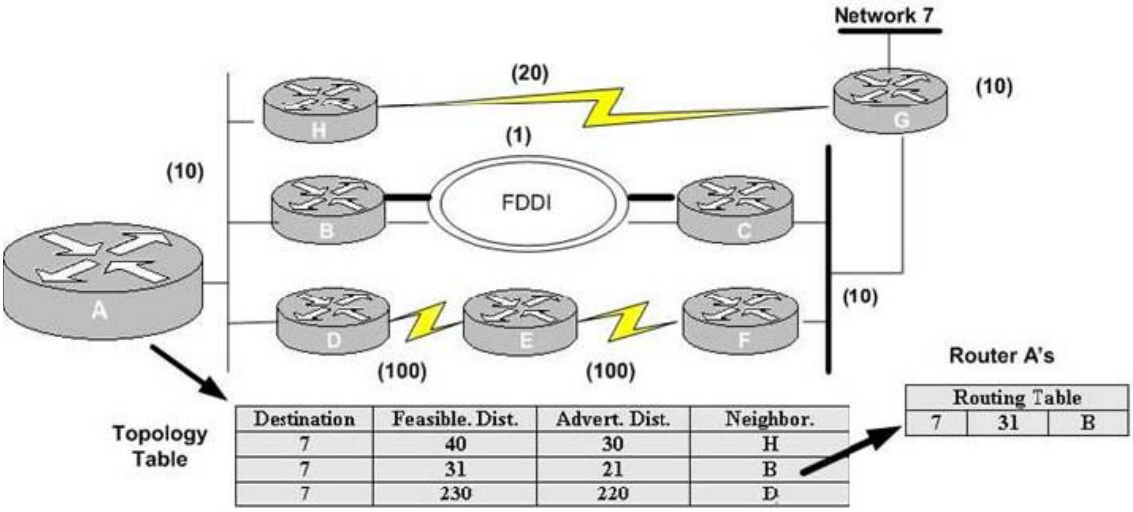

You are the systems administrator of the above EIGRP network and you've elected to shut down router B's FDDI interface. After you execute the shutdown, which router will become the feasible successor to network 7 for router A?

A. B
B. D
C. H
D. There is no feasible successor
E. None of the above

Answer: D

Explanation:
Router H will be the successor, and that route will be placed in the Routing table.
Router A detects the link failure between Router B and network 7. It checks the topology table for a successor. It finds that H is the successor since the advertised distance for H (30) is less than the feasible distance for B (31).
However, there is no next best route - no feasible successor. The candidate route through D has an advertised distance (220) that is higher than the feasible distance of the successor route (40).
Note: Successor is a route selected as the primary route to use to reach a destination.

Successors are the entries kept in the routing table.
A feasible successor is a backup route. These routes are selected at the same time the successors are identified, but they are kept in a topology table,
Incorrect Answers:
A: The FDDI interface of B is down.
B: The candidate route through D has an advertised distance (220) that is higher than the feasible distance of the successor route (40). It cannot be used as a feasible successor.
C: Router H is the successor, not the feasible successor.

QUESTION 54:

You are the network engineer at Certkiller . You want to configure EIGRP to run across a 56 Kbps serial PPP link on the Certkiller network. You also want to ensure the proper convergence of EIGRP routes.
What command should you issue on the serial interface?

A. bandwidth 56
B. bandwidth 56000
C. ip bandwidth-percent eigrp 1 56
D. ip bandwidth-percent eigrp 1 56000
E. None of the above

Answer: A

Explanation:
The bandwidth is a logical construct whose value can have wide-reaching implications on the function of your network. It does not affect the actual speed of the link. In fact, it is practical to configure the bandwidth command only on serial lines, where the speed of the link will vary considerably. This command is useful for routing protocols that use bandwidth as part of the metric. For serial interfaces, this value defaults to a full T1, so it my appear to be a better link than what is actually being used.
Router (config) # interface S0
Router (config-if) # bandwidth speed-of-line
The bandwidth command is always specified in kbps.
EIGRP limits itself to 50 percent of the value specified in the bandwidth command, or if the bandwidth command is not set, the interface defaults. If you need to limit this percentage further, the upper limit that EIGRP uses can be stated as a percentage of the bandwidth command.
The ip bandwidth-percent-eigrp command interacts with the bandwidth command on the interface. You would use this command primarily because in your network, the bandwidth command does not reflect the true speed of the link. The bandwidth command might have been altered to manipulate the routing metric and path selection of a routing protocol, such as IGRP or OSPF. It might be better to use other methods of controlling the routing metric and return the bandwidth to a true value. Otherwise, the ip bandwidth-percent eigrp command is available. It is possible to set a bandwidth percent that is larger than the stated bandwidth. This is with the understanding that, although the

bandwidth might be stated to be 56 kbps, the link is in fact 256 kbps. The following shows the structure of the ip bandwidth-percent eigrp command:
Router (config) #interface S0
Router (config-if) #ip bandwidth-percent eigrp autonomous-system-number percent

---

QUESTION 55:

The Certkiller Network topology is displayed in the diagram below:

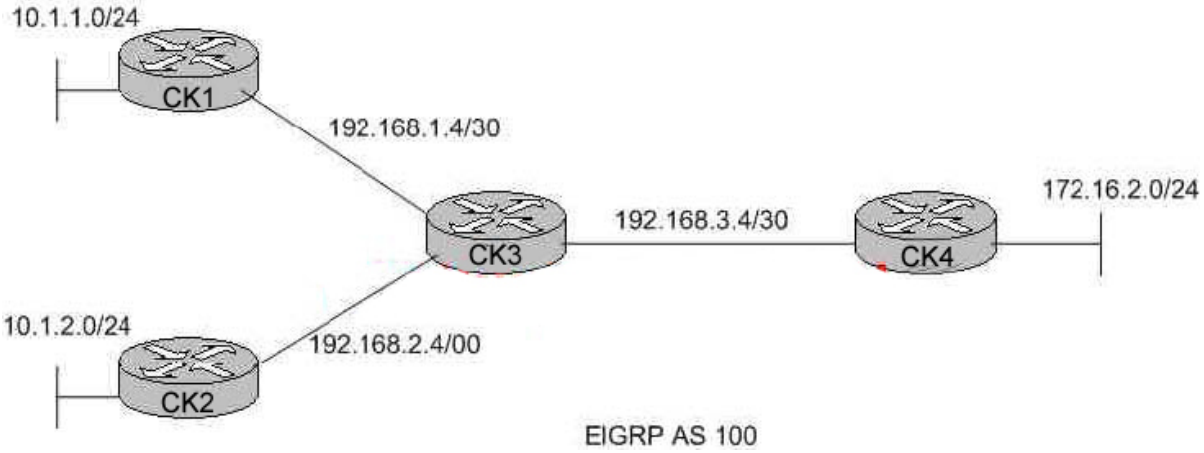

Based on the information above, what optional EIGRP configuration will be required in order to achieve full network connectivity?

A. Use the EIGRP no auto-summary command on CK1 and CK2 .
B. Use the EIGRP no auto-summary command on CK3 and CK4 .
C. Use the passive interface on the CK1 and CK2 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 CK3 and CK1 connection and between the CK3 and CK2 connection.
E. Use the variance command on CK3 .
F. None of the above

Answer: A

Explanation:
A useful feature of EIGRP is automatic route summarization; this summarizes subnets to the classful network boundary. This is enabled by default, you can turn this off per AS by using the following command in router configuration mode:
Router(config-router)#no auto-summary
In the example above, this is needed, otherwise both CK1 and CK2 will advertise the summarized network 10.0.0.0/8 to router CK3 , making their networks unreachable.

QUESTION 56:

You are the network engineer at Certkiller . The Certkiller network is shown in the following graphic:

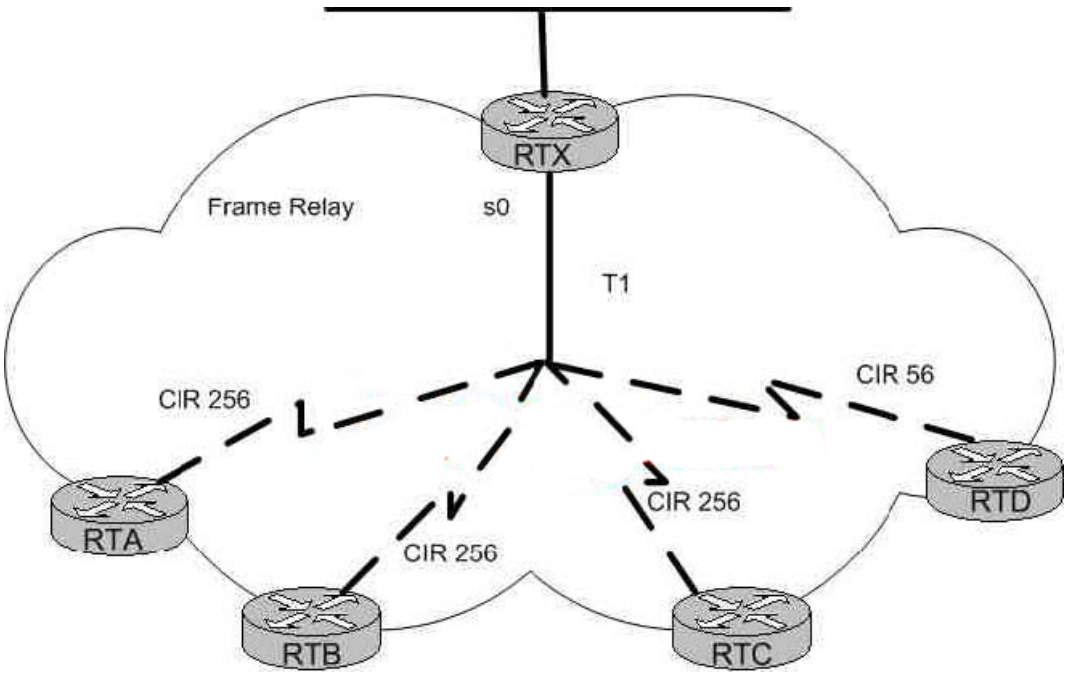

Router RTX is configured with the following commands:
Interface serial 0
Encapsulation frame relay
You want to complete an EIGRP hybrid multipoint configuration on RTX S0 by configuring it with the correct bandwidth parameter.
Which command should you issue?

A. bandwidth 56
B. bandwidth 224
C. bandwidth 256
D. bandwidth 768
E. Non of the above

Answer: B

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 applied to the physical interface.
In the case of the question 4*56=224 - bandwidth 224
Reference: Self-Study CCNP BSCI Exam Certification Guide Third Edition P. 490.

QUESTION 57:

The Certkiller multipoint frame relay network is using EIGRP for the routing protocol. In this network, how is the bandwidth determined for each multipoint neighbor?

A. Bandwidth command per neighbor.

B. The configured CIR per subinterface.
C. The configured CIR divided by the number of neighbors on that interface.
D. Bandwidth of the main interface divided by the number of neighbors on that interface.
E. None of the above

Answer: D

Explanation:
When configuring multipoint interfaces, especially for Frame Relay, remember that all neighbors share the bandwidth equally, regardless of the actual CIR of each individual PVC. For multipoint interfaces, the bandwidth used by the IOS is taken from the configured bandwidth of the main interface, divided by the total number or neighbors on that interface.

QUESTION 58:

Your network consists of a router named R1 that's configured in a hub and spoke topology, supporting 24 remote office via a point-to-multipoint Frame Relay EIGRP network deployment. The bandwidth command was not manually entered on the frame relay main interface or the sub-interfaces. What is the perceived bandwidth (as seen by the EIGRP process) of each Frame Relay connection?

A. 64 kbps
B. 128 kbps
C. 1.544 Mbps
D. 1.536 Mbps
E. None of the above

Answer: A

Explanation:
The default bandwidth for all serial WAN interfaces is 1.544 Mbps. EIGRP by default assumes the bandwidth of a serial interface is a full T1 (1.544 Mbps) if not specified. This includes sub-interfaces. For multipoint frame relay interfaces, the perceived bandwidth used by the IOS is taken from the configured bandwidth, divided by the number of configured neighbors. Since there are 24 remote offices, there are 24 configured frame relay neighbors, so the bandwidth used by the IOS is the full T1 divided by 24, which is 64kbps.

QUESTION 59:

EIGRP has been configured on the WAN links of router CK1 . On one of the serial T1 interfaces, the bandwidth statement was not specified. By default, how will the EIGRP process perceive this T1 link?

A. 256 Kbps
B. 1.544 Mbps
C. It depends as it is set by the PVC
D. It is set by the DLCI
E. None of the above

Answer: B

Explanation:
The default EIGRP link speed is 1.544 Mbps for non-high speed WAN links
Note: The enhanced code uses the "bandwidth" subcommand on interfaces and sub-interfaces in order to determine the rate at which to generate EIGRP packets. This parameter is automatically set on fixed-bandwidth interfaces (such as LANs), but defaults to T1 (1544 Kbps) for all serial media.
Incorrect Answers:
A: The default speed is 1.544 Mbps, not 256 Kbps.
C, D: The default speed is not set by the PVC or the DLCI.

QUESTION 60:

You are in the midst of configuring an NBMA network with EIGRP as the routing protocol. Which of the options below would you use to configure the interface bandwidth for a point-to-point interface?

A. The DLCI assigns the bandwidth for the interface.
B. The sliding window size determines interface bandwidth.
C. You should use the default bandwidth assigned to the interface.
D. You should manually configure bandwidth as the CIR of the PVC.
E. None of the above.

Answer: D

Explanation:
The bandwidth can be configured separately on each subinterface. Since this is NBMA we can assume that Frame Relay is used. For Frame Relay on point-to-point the bandwidth should be set it to the CIR of the PVC.
Note 1: NBMA (Non-broadcast Multi-access) supports many (more than two) routers, but have no broadcast capability. Frame Relay and X.25 are example of NBMA.
Note 2: The CIR (Committed Information Rate) is the committed rate (in bits per second) at which the ingress access interface trunk interfaces, and egress access interface of a Frame Relay network transfer information to the destination Frame Relay end system under normal conditions.
Incorrect Answers:
A: Does not apply.
B: Sliding windows does not apply.
C: The bandwidth can and should be configured separately on each subinterface.

Reference: Configuration Notes for the Enhanced Implementation of EIGRP
http://www.cisco.com/warp/public/103/12.html

QUESTION 61:

Router CK2 is configured for EIGRP as shown below:
router eigrp 100
network 10.0.0.0
eigrp stub
Based on this configuration, which types of routes will be advertised with the EIGRP configuration as shown? (Choose two)

A. Static
B. Receive only
C. Summary
D. Stub
E. Connected
F. Dynamic

Answer: C, E

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

---

QUESTION 62:

The Certkiller EIGRP network is utilizing the EIGRP stub routing feature. Which of the following are key concepts that apply when configuring the EIGRP stub routing feature in a hub and spoke network? (Select 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

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.
This feature should only be used on stub routers. A stub router is defined as a router connected to the network core or distribution layer through which core transit traffic should not flow. A stub router should not have any EIGRP neighbors other than distribution routers. Ignoring this restriction will cause undesirable behavior.
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a0080087026.html

---

QUESTION 63:

The Certkiller WAN is displayed in the following diagram:

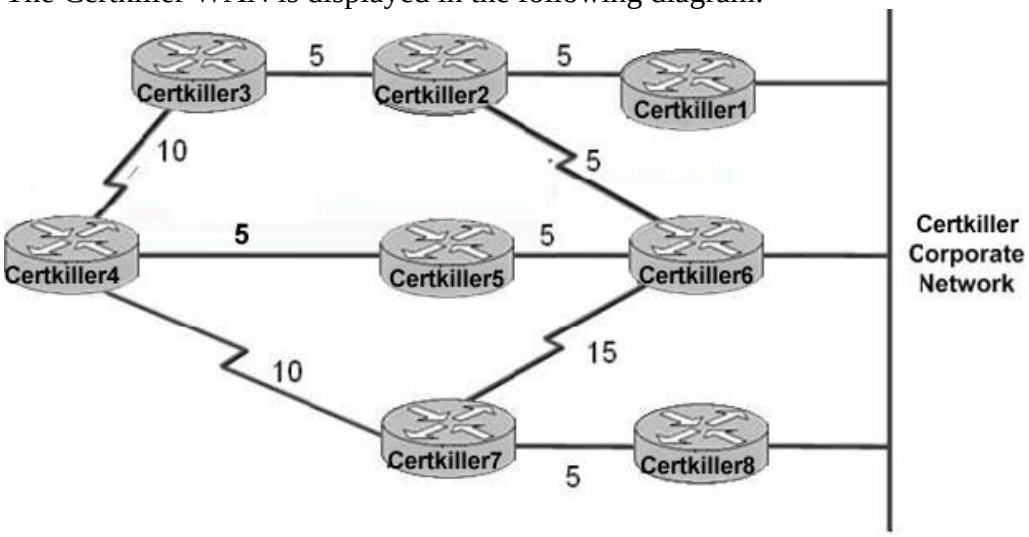

Traffic from Certkiller 4 to the Certkiller Corporate Network is distributed between the links with unequal costs in the EIGRP network by configuring the variance command on Certkiller 4 to 2. How many paths will participate in the load sharing?

A. 1
B. 2
C. 3
D. 4
E. 5
F. 6

Answer: B

Explanation:
EIGRP automatically load balances across links of equal cost. Whether the traffic is sent on a predestination or round-robin basis depends on the internal switching within the router. It is possible to configure EIGRP to load balance across unequal-cost paths using the variance command.
The variance command allows the administrator to identify the metric scope for including additional paths by the use of a multiplier parameter. The command structure follows:
Router(config-router)#variance multiplier
The multiplier argument is the metric value used for load balancing. It can be a value from 1 to 128. The default is 1, which means equal-cost load balancing.
Reference: CCNP Self-Study CCNP BSCI Exam certification guide p.485

QUESTION 64:

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
F. None of the above

Answer: D

Explanation:
EIGRP automatically load balances across links of equal cost. Whether the traffic is sent on a predestination or round-robin basis depends on the internal switching within the router. It is possible to configure EIGRP to load balance across unequal-cost paths using the variance command.
The variance command allows the administrator to identify the metric scope for including additional paths by the use of a multiplier parameter. The command structure follows:
Router(config-router)#variance multiplier
The multiplier argument is the metric value used for load balancing. It can be a value

from 1 to 128. The default is 1, which means equal-cost load balancing.
Reference: CCNP Self-Study CCNP BSCI Exam certification guide p.485

QUESTION 65:

The Certkiller WAN is displayed below:

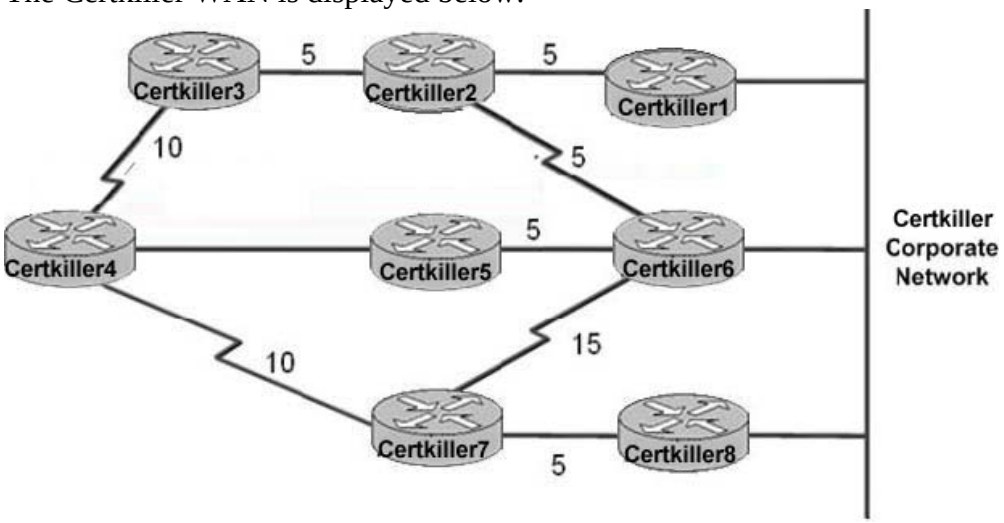

Examine the EIGRP network diagram. All paths should be used for load balancing between Certkiller 4 and the Certkiller Corporate Network. Which value should be assigned to the variance command to accomplish this?

A. 1
B. 2
C. 2.5
D. 3
E. 5
F. 6
G. None of the above

Answer: D

Explanation:
Every routing protocol supports equal cost path load balancing. In addition to that, IGRP and EIGRP also support unequal cost path load balancing. Use the variance command to instruct the router to include routes with a metric less than n times the minimum metric route for that destination, where n is the number specified by the variance command. The variable n can take a value between 1 and 128, with the default being 1, which means equal cost load balancing. Traffic is also distributed among the links with unequal costs, proportionately, with respect to the metric. In this example, the worst route is 2.5 times worse than the best route. Since we can not use 2.5 for the variance (the value must be an integer) we must specify a minimum value of 3.

QUESTION 66:

The Certkiller EIGRP network is displayed in the following diagram:

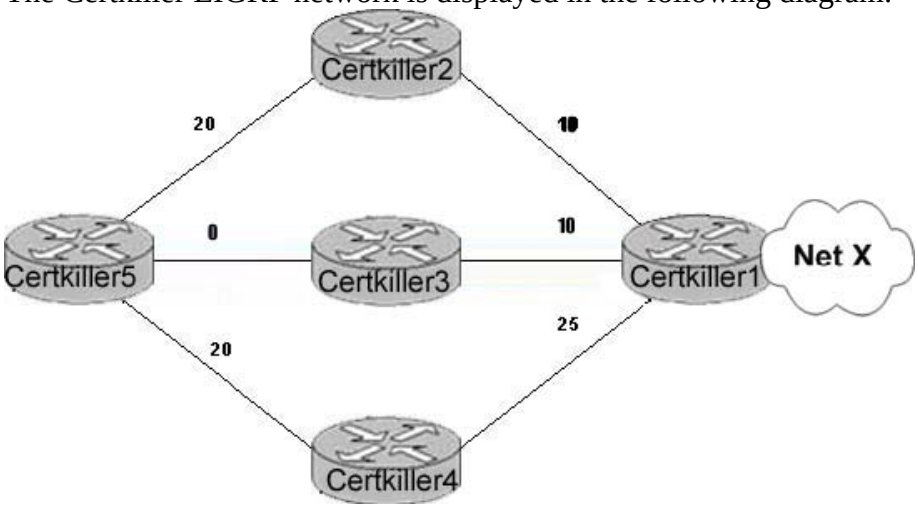

You work as a network technician at Certkiller .com. Study the exhibits carefully. If the following command, variance 2, were added to Certkiller 5, which path or paths would be chosen to route traffic from Certkiller 5 to network X?

A. Only Certkiller 5- Certkiller 4- Certkiller 1
B. Only Certkiller 5- Certkiller 2- Certkiller 1
C. Only Certkiller 5- Certkiller 3- Certkiller 1
D. Both Certkiller 5- Certkiller 2- Certkiller 1 and Certkiller 5- Certkiller 3- Certkiller 1
E. Both Certkiller 5- Certkiller 2- Certkiller 1 and Certkiller 5- Certkiller 4- Certkiller 1
F. All available paths.
G. None of the above

Answer: D

Explanation:
Variance is 2 so 2 *(10+10) = 40

QUESTION 67:

Part of the configuration for a Certkiller router is shown below:
interface serial0
bandwidth 40
ip bandwidth-percent eigrp 1 200
In the configuration in the exhibit, how much bandwidth will be used by the EIGRP updates?

A. 40 kbps
B. 60 kbps
C. 80 kbps
D. 200 kbps
E. 40 Mbps
F. 60 Mbps

G. 80 Mbps
H. 200 Mbps
I. None of the above

Answer: C

Explanation:
By default, a serial interface will assume a bandwidth value of 1544, which is equal to a full T1. This default value can be overridden via the bandwidth configuration statement as shown above. From the file shown in this example, serial 0 is configured as a 40kbps (the bandwidth statement is used in terms of kbps) and EIGRP is configured to use up to twice this amount (200 percent) so up to 80 kbps will be consumed via the EIGRP process.

QUESTION 68:

Five Certkiller routers are connected together as shown in the diagram below:

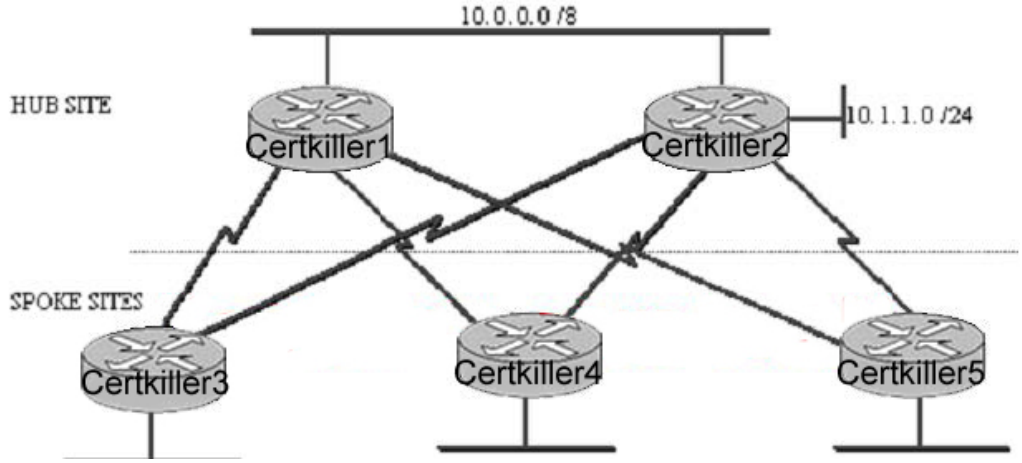

Refer to the exhibit. Certkiller 1 and Certkiller 2 are summarizing the 10.0.0.0/8 network and advertising it to all routers. Which two EIGRP configurations will prevent unnecessary EIGRP updates from being sent and possibly causing 10.0.0.0/8 from being SIA on router Certkiller 1 or Certkiller 2? (Select two)

A. Configure Certkiller 3, Certkiller 4, and Certkiller 5 as EIGRP stub routers.
B. Configure an ACL on the interface of Certkiller 1 and Certkiller 2.
C. Configure Certkiller 1 and Certkiller 2 as EIGRP stub routers.
D. Configure summary routes on Certkiller 1 and Certkiller 2
E. Configure static routing on Certkiller 3, Certkiller 4, and Certkiller 5

Answer: A, D

Explanation:
A dual-homed remote will have two or more distribution (hub) routers. However, the principles of stub routing are the same as they are with a hub and spoke topology.

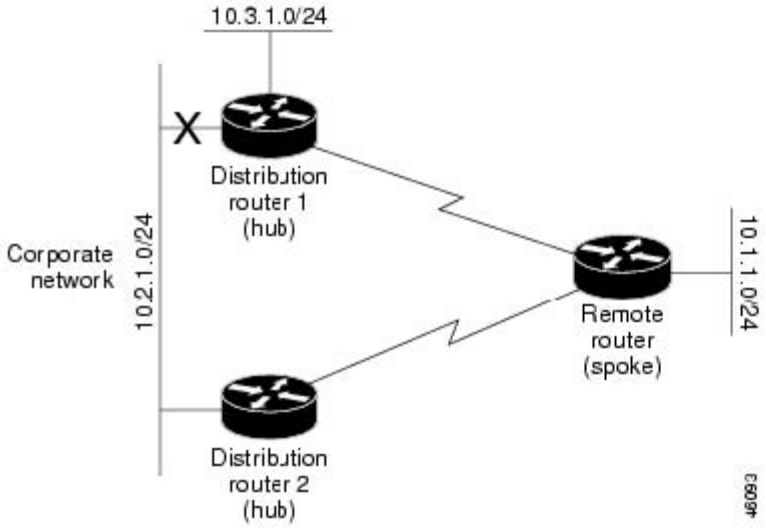

It is not desirable for traffic from distribution router 2 to travel through any remote router in order to reach network 10.3.1.0/24. If the links are sized to handle the load, it would be acceptable to use one of the backup routes. However, most networks of this type have remote routers located at remote offices with relatively slow links. This problem can be prevented if proper summarization is configured on the distribution router and remote router.
It is typically undesirable for traffic from a distribution router to use a remote router as a transit path. A typical connection from a distribution router to a remote router would have much less bandwidth than a connection at the network core. Attempting to use a remote router with a limited bandwidth connection as a transit path would generally produce excessive congestion to the remote router. The EIGRP Stub Routing feature can prevent this problem by preventing the remote router from advertising core routes back to distribution routers. Routes learned by the remote router from distribution router 1 will not be advertised to distribution router 2. Since the remote router will not advertise core routes to distribution router 2, the distribution router will not use the remote router as a transit for traffic destined for the network core.
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a0080087026.html

---

QUESTION 69:

Which command is used to change the administrative distance for EIGRP?

A. Router(config-if)# distance eigrp internal-distance external-distance
B. Router(config-if)# distance weight eigrp internal-distance external-distance
C. Router(config-router)# distance eigrp internal-distance external-distance
D. Router(config-router)# distance weight eigrp internal-distance external-distance

Answer: C

Explanation:
To allow the use of two administrative distances-internal and external-that could be a

better route to a node, use the distance eigrp command in router configuration mode. To reset these values to their defaults, use the no form of this command.
distanceeigrp internal-distance external-distance
no distance eigrp
Syntax Description:

internal-distance	Administrative distance for EIGRP internal routes. Internal routes are those that are learned from another entity within the same autonomous system. The distance can be a value from 1 to 255.
external-distance	Administrative distance for EIGRP external routes. External routes are those for which the best path is learned from a neighbor external to the autonomous system. The distance can be a value from 1 to 255.

Defaults:
internal-distance: 90
external-distance: 170
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_command_reference_chapter09186a00800c

---

QUESTION 70:

EIGRP is being used as the routing protocol within the Certkiller network. There is an issue with EIGRP in this network and you suspect the problem lies with the DUAL process. Which two statements about the EIGRP DUAL process are correct? (Select two)

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

Answer: D, F

Explanation:
The focal point of EIGRP is the Diffusing Update Algorithm (DUAL). This is the route-calculation engine of EIGRP. The full name of this technology is DUAL finite-state machine (FSM). An FSM is an abstract machine, not a mechanical device with moving parts. FSMs define a set of possible states that something can go through, what events cause those states, and what events result from those states. Designers use

FSMs to describe how a device, computer program, or routing algorithm will react to a set of input events. The DUAL FSM contains all the logic used to calculate and compare routes in an EIGRP network.
DUAL tracks all the routes advertised by neighbors in addition to using the composite metric of each route to compare them. DUAL also guarantees that each path is loop-free. Lowest-cost paths are then inserted by DUAL into the routing table.
As noted earlier in the module, EIGRP keeps important route and topology information readily available. This information is stored in a neighbor table and a topology table. These tables supply DUAL with comprehensive route information in case of network disruption. DUAL selects alternate routes quickly by using the information in these tables. If a link goes down, DUAL looks for a feasible successor in its neighbor and topology tables.
A successor is a neighboring router that is currently being used for packet forwarding. A successor also provides the least-cost route to the destination and is not part of a routing loop. Feasible successors provide the next lowest-cost path without introducing routing loops. Feasible successor routes can be used in case the existing route fails. Packets to the destination network are immediately forwarded to the feasible successor. The router is then promoted to the status of successor

---

QUESTION 71:

The following command was issued on router Certkiller 2:

```
Certkiller2#show ip route
     <output commited>
    C    10.1.1.0 is directly connected, Serial0
    D    172.16.1.0 [90/2681856] via 10.1.1.1, Serial0
   D EX 192.168.1.0 [170/2681856] via 10.1.1.1,
     00:00:04,Serial10
```

Given the above output shown above, which statement is true?

A. 192.168.1.0 is a redistributed route into EIGRP.
B. 192.168.1.0 is a summarized route.
C. 192.168.1.0 is a static route.
D. 192.168.1.0 is equal path load balancing with 172.16.1.0.
E. None of the above

Answer: A

Explanation:
To support multiple routing protocols within the same internetwork efficiently, routing information must be shared among the different routing protocols. For example, routes learned from a RIP process may need to be imported into an IGRP process. This process of exchanging routing information between routing protocols is called route redistribution. Such redistribution can be one-way or two-way. One-way routes are where one protocol receives the routes from another. Two-way routes are where both protocols receive routes from each other. Routers that perform redistribution are called boundary

routers because they border two or more autonomous systems or routing domains. This section examines route redistribution in detail, including the use of administrative distance, guidelines for redistribution implementation, and issues with redistribution configuration.
Using multiple routing protocols typically results in increased administrative complexity and overhead. Why would this be done in the first place? Actually, there are several scenarios in which using multiple routing protocols solves more problems than it creates, especially in medium and large sized networks.
The syntax of the redistribute command is as follows:
Router(config-router)#redistribute protocol [process-id] {level-1 | level-1-2 | level-2} [metric metric-value][metric-type type-value] [match {internal | external 1 | external 2}][tag tag-value] [route-map map-tag][weight weight] [subnets]
Example:
routerospf 1
redistribute eigrp 1 metric 66 metric-type 1 subnets
When EIGRP learns the routing information from the different routing protocol it uses D EX symbol to indicate that this routing information has learned from other routing protocol.

QUESTION 72:

A network administrator is troubleshooting an EIGRP connection between router Certkiller 1 with IP address 10.1.2.1 and Certkiller 2 with IP address 10.1.2.2. Given the debug output on Certkiller 1 shown below, which two statements are true?

```
certkiller1#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
```

A. Certkiller 1 received a hello packet with mismatched hello timers.
B. Certkiller 1 will form an adjacency with Certkiller 2.
C. Certkiller 1 received a hello packet with mismatched metric-calculation mechanisms.
D. Certkiller 1 received a hello packet with mismatched autonomous system numbers.
E. Certkiller 1 received a hello packet with mismatched authentication parameters.
F. Certkiller 1 will not form an adjacency with Certkiller 2.

Answer: C, F

Explanation:
Metrics are the mathematics used to select a route. The higher the metric associated with a route, the less desirable it is. For EIGRP, the Bellman-Ford algorithm uses the following equation and creates the overall 24-bit metric assigned to a route:
* metric = [(K1 × bandwidth) + [(K2 × bandwidth) ÷ (256 - load)] + (K3 × delay)] × [K5 ÷ (reliability + K4)]
The elements in this equation are as follows:
* By default, K1 = K3 = 1, K2 = K4 = K5 = 0. Therefore, by default, the metric formula

reduces to:
metric = (1 × bandwidth) + (1 × delay)
metric = bandwidth + delay
K Values should be same to become the EIGRP neighbors.

QUESTION 73:

Study the exhibit below carefully:

```
certkiller1# show ipo 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
```

If the configuration shown below is added to Certkiller 1, which three route entries will EIGRP advertise to neighboring routers? (Select three)
router eigrp 10
network 10.0.0.0
eigrp stub

A. 192.168.20.0/24
B. 10.1.2.0/24
C. 10.1.1.0/24
D. 10.1.3.0/24
E. 10.0.0.0/8

Answer: C, D, 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 74:

Two Certkiller routers are connected together as shown below:

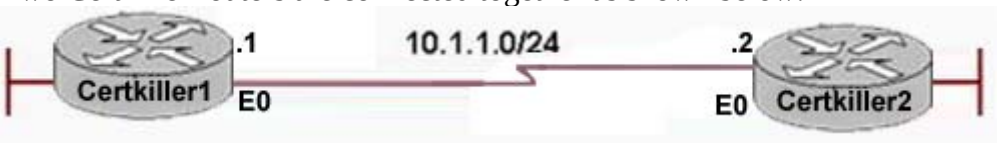

Certkiller 1 configuration exhibit:

```
hostname CertKiller1
!
interface Ethernet0
 ip address 10.1.1.1 255.255.255.0
!
router eigrp 4
 network 10.0.0.0
!
end
```

```
CertKiller1# show ip eigrp
IP-EIGRP neighbors for process
01:20:54: IP-EIGRP: Neighbor 10.1.2.2 not on common subnet for Ethernet0 (10.1.)
01:21:08: IP-EIGRP: Neighbor 10.1.2.2 not on common subnet for Ethernet0 (10.1.)
```

Certkiller 2 configuration exhibit:

```
hostname CertKiller2
!
interface Ethernet0
  ip address 10.1.2.2 255.255.255.0
  ip address 10.1.1.2 255.255.255.0 secondary
!
router eigrp 4
  network 10.0.0.0
!
end
```

```
CertKiller2# show ip eigrp
IP EIGRP neighbor for process 4
H   Address     Interface   Hold   Uptime      SRTT   RTO   Q    Seq    Type
                                   (sec)       (ms)         Cnt  Num
0   10.1.1.1    Et0         12     00:00:35    1      5000  1    0
```

Study the exhibits carefully. EIGRP has been configured on routers Certkiller 1 and Certkiller 2. However, Certkiller 1 does not show Certkiller 2 as a neighbor and does not accept routing updates from Certkiller 2. 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. EIGRP cannot form neighbor relationship and exchange routing updates with a secondary address.
C. EIGRP cannot exchange routing updates with a neighbor's router interface that is configured with two IP addresses.
D. Interface E0 on router Certkiller 1 has not been configured with a secondary IP address of 10.1.2.1/24.
E. None of the above

Answer: B

Explanation:
Remember that simple distance vector routers do not establish any relationship with their neighbors. RIP and IGRP routers merely broadcast or multicast updates on configured interfaces. In contrast, EIGRP routers actively establish relationships with their neighbors, much the same way that OSPF routers do.
EIGRP routers establish adjacencies with neighbor routers by using small hello packets. Hellos are sent by default every five seconds. An EIGRP router assumes that as long as it is receiving hello packets from known neighbors, those neighbors (and their routes) remain viable. By forming adjacencies, EIGRP routers do the following:
1. Dynamically learn of new routes that join their network
2. Identify routers that become either unreachable or inoperable
3. Rediscover routers that had previously been unreachable

QUESTION 75:

While troubleshooting an EIGRP routing problem you notice that one of the Certkiller routers have generated a large number of SIA messages. What are two possible causes for EIGRP Stuck-In-Active routes? (Select 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:
The SIA state means that an EIGRP router has not received a reply to a query from one or more neighbors within the time allotted (approximately 3 minutes). When this happens, EIGRP clears the neighbors that did not send a reply and logs a DUAL-3-SIAerror message for the route that went active
Details: http://www.cisco.com/warp/public/103/18.html#topic1

QUESTION 76:

The Certkiller network is displayed below:

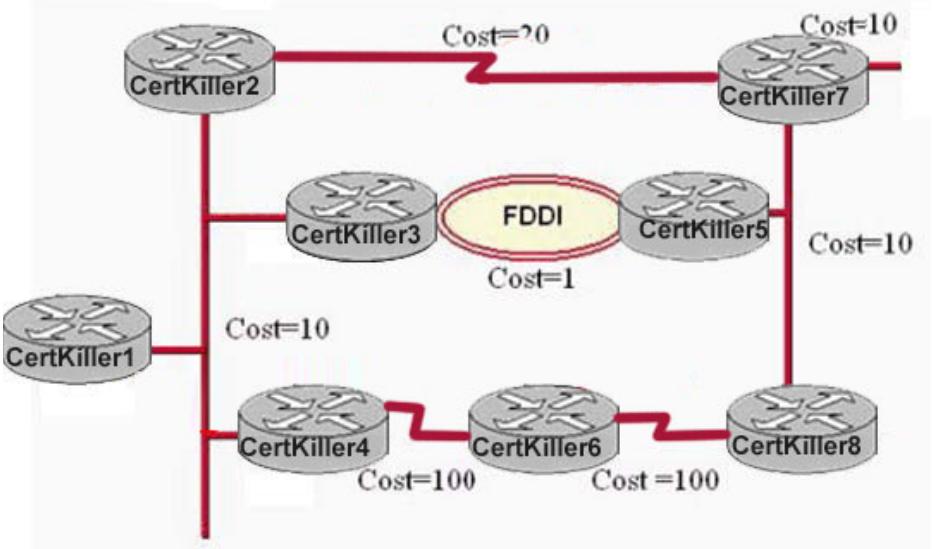

All routers in the diagram are configured with EIGRP. If Certkiller 2 and Certkiller 3 fail, which action will Certkiller 1 take with respect to the HQ network?

A. Certkiller 1 will automatically route packets via Certkiller 4 to the HQ network.
B. Certkiller 1 will go into the active state for the route to HQ network.
C. Certkiller 1 will go into the active state for all routes.
D. Certkiller 1 will place the route via Certkiller 4 into the hold down state.
E. None of the above

Answer: B

QUESTION 77:

EIGRP uses five generic packet types (hello, updates, queries, replies, acknowledgements). If you wished to view the statistics for these packets, which IOS command should you use?

A. debug eigrp packets
B. show ip eigrp traffic
C. show ip eigrp topology
D. show ip eigrp neighbors

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_chapter09186a00800

QUESTION 78:

While troubleshooting a routing problem on the Certkiller EIGRP network you discover that one of the routers is failing to establish adjacencies with its neighbor. What is a likely cause of this problem between neighbors? (Select two)

A. The K-values do not match.
B. The hold times do not match.
C. The hello times do not match.
D. The AS numbers do not match.

Answer: A, D

Explanation:

Peer relationships and adjacencies between routers will not be formed between EIGRP routers if the neighbor resides in a different autonomous system or if the metric-calculation mechanism (K values) is misaligned for that link.
Incorrect Answers:
B, C: It is possible for two routers to become EIGRP neighbors even though the hello and hold timers do not match.

QUESTION 79:

While troubleshooting an EIGRP routing issue, you are seeing a high number of SIA (stuck in active) routes. Which of the following are causes of a route becoming SIA? (Select two)

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

Answer: A, D

Explanation:
The acknowledgement does not reach the destination or they are too delayed. This is normally due to too many routing topology changes, or a router with insufficient memory.
Note: In some circumstances, it takes a very long time for a query to be answered. So long, in fact, that the router that issued the query gives up and clears its connection to the router that isn't answering, effectively restarting the neighbor session. This is known as a stuck in active (SIA) route. 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.
Incorrect Answers:
B: Ack packets don't reply to Query, only Reply do.
C: Does not apply to SIA. This is the normal operation of EIGRP.
Reference:http://www.cisco.com/warp/public/103/eigrp3.html

QUESTION 80:

The Certkiller network is displayed in the diagram below:

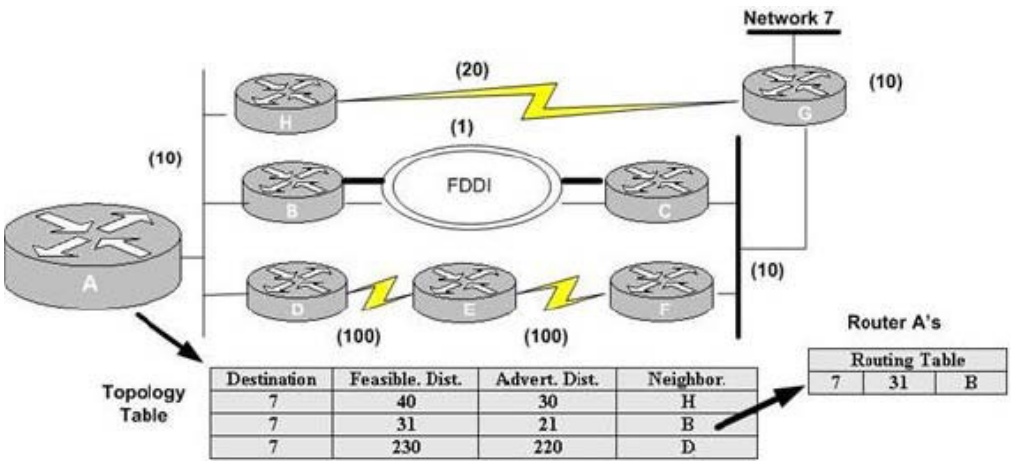

If the FDDI interface in the EIGRP network above were to fail or shut down, which router(s) will become Router A's next-hop to Network 7?

A. B
B. D only
C. H only
D. D and H
E. None of the above

Answer: C

Explanation:
Router H will be the successor, and that route will be placed in the Routing table.
Router A detects the link failure between Router B and network 7. It checks the topology table for a successor. It finds that H is the successor since the advertised distance for H (30) is less than the feasible distance for B (31).
However, there is no next best route - no feasible successor. The candidate route through D has an advertised distance (220) that is higher than the feasible distance of the successor route (40).

QUESTION 81:

The Certkiller network consists of a hub and spoke topology with a main router supporting about 20 regional offices. A point-to-point Frame Relay WAN connects the regional offices to the main office, and EIGRP is deployed as the routing protocol. The committed information rate (CIR) for each of the Frame Relay PVC's is different, and the bandwidth command IS NOT configured on any of the interfaces or subinterfaces. You want to ensure that EIGRP routes everything properly. How should you configure the network?

A. Convert each Frame Relay PVC to point-to-multipoint connection
B. Manually configure the bandwidth of the major interface to the lowest CIR x 24
C. Manually configure the bandwidth of the major interface to the highest CIR x 24
D. Manually configure the bandwidth of each of these PVCs to equal to their respective

CIR.

Answer: D

Explanation:
Although this will require some administrative effort, the only way to ensure that EIGRP properly considers the actual bandwidth to use in the routing decision, each link should be set to the CIR.
Incorrect Answers:
A: This is not a Cisco recommended solution.
B: This will force all PVC to run at a low speed.
C: This would give too high a bandwidth.

QUESTION 82:

The Certkiller EIGRP network is displayed in the diagram below:

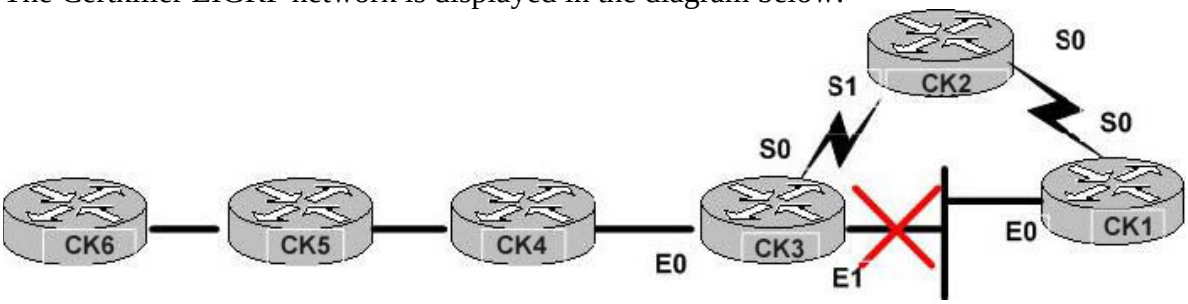

All routers are running EIGRP. Based on this, what will router CK3 do if the link between CK3 and CK1 were to go down? (Select two)

A. It elects a new designated router.
B. It sends a flash update with poison reverse.
C. It checks its topology table for an alternate route.
D. It re-broadcasts its routing table to all other neighboring routers.
E. It sends a query to neighboring routers for other routes to the failed link.

Answer: C, E

Explanation:
The steps of convergence in an EIGRP network are as follows:
1. Router CK3 detects the link failure between CK1 and CK3 . It checks the topology table for a feasible successor, but it doesn't find a qualifying alternate route and enters in an active convergence state. (C)
2. CK3 sends a Query out all interfaces for other routes to the failed link (E). The neighboring routers acknowledge the query.
3. The reply from CK4 indicates no other route to the failed link.
4. CK2 's reply contains a route to the failed link, although it has a higher feasible distance.
5. Router CK3 accepts the new path and metric information, places it in the topology

table, and creates an entry for the routing table.
6. CK3 sends an update about the new route out all interfaces.

QUESTION 83:

The EIGRP topology table for router Certkiller 1 is displayed below:

```
CertKiller1#show ip eigrp topology
IP-EIGRP Toplogy 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, Ethernet0
P 192.168.1.32/28, 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
```

Regarding the command output on Certkiller 1 in the exhibit, which statements are true? (Select three)

A. Certkiller 1 is in AS 200
B. Certkiller 1 will load balance between three paths to reach the 192.168.1.48/28 prefix, because all three paths have the same AD of 40512000.
C. The best path for Certkiller 1 to reach the 192.168.1.48/28 prefix is via 192.168.1.66.
D. 40512000 is the advertised metric 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 Certkiller 1 is in the query process for those routes.

Answer: A, C, D

Explanation:
The Certkiller 1 router resides in AS 200, as displayed by the "IP EIGRP topology for process 200" output.
Regarding the numbers specified in the parenthesis, the first number is the EIGRP metric that represents the cost to the destination. The second number is the EIGRP metric that this peer advertised.
Based on this, the best path to the 192.168.1.48/28 destination is via 192.168.1.66, because the metric is less than the alternatives.

QUESTION 84:

The following show command was issued on router CK1 :

```
CK1#show ip eigrp topology
IP-EIGRP Topology Table for AS(24)/ID(1.1.1.1)

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

P 10.2.1.0/24, 1 successors, FD is 28160
       via Connected, FastEthernet0/0
P 10.1.1.0/24, 1 successors, FD is 20512000
       via Connected Serial0/0
P 172.20.0.0/16,   successors, FD is 40992000
       via 10.1.1.2 (40992000/40480000), Serial0/0
P 172.16.0.0/16, 1 successors, FD is 23072000
       via 10.1.1.2 (23072000/2560256), Serial0/0
P 172.17.0.0/16, 1 successors, FD is 25632000
       via 10.1.1.2 (25632000/5120256), Serial0/0
P 172.18.0.0/16, 1 successors, FD is 32032000
       via 10.1.1.2 (32032000/11520256), Serial0/0
A 172.19.0.0/16, 1 successors, FD is Inaccessible, Q
   2 replies, active never, query-origin: Successor Origin
       via 10.1.1.3 (391248640/390736640), r, Serial0/0
     Remaining replies:
       via 10.1.1.2, r, Serial0/0
CK1#
```

Based on the exhibited command output, which two statements are true? (Choose two)

A. The EIGRP network is stable.
B. The router at 10.1.1.3 has not replied to the CK1 query packet.
C. The route to 172.19.0.0/16 is undergoing recomputation.
D. The route to 172.19.0.0/16 is stuck-in-active.
E. CK1 has sent a query packet to 10.1.1.2.

Answer: B, C

Explanation:
The output shown above is a snapshot from an EIGRP router re-computing the route to d 172.19.0.0/16. As can be seen above, a query was originated by CK1 to the neighbor at 10.1.1.3, but the query has not yet been replied to by this router.

QUESTION 85:

The Certkiller OSPF network has been segmented into numerous different areas. Which three choices below are advantages to creating multiple areas in OSPF? (Select three)

A. Fewer hello packets
B. Fewer adjacencies needed

C. Less frequent SPF calculations
D. Smaller routing tables
E. Reduced Link State Update (LSU) overhead

Answer: C, D, E

Explanation:
Three issues can overwhelm an OSPF router in a heavily populated OSPF network:
1. high demand for router processing and memory resources
2. large routing tables
3. large topology tables
In a very large internetwork, changes are inevitable. OSPF routers are likely to run SPF calculations frequently, which deprive the router of precious CPU cycles and memory resources.
Not only is the routing table frequently recalculated in a large OSPF network, but it also risks being overstuffed with multiple paths and hundreds of routes. Full routing tables make routers less efficient. Finally, the link-state database, which must contain a complete topology of the network, will also threaten to consume resources and slow down the router.
1. Reduced frequency of SPF calculations - Because detailed route information is kept within each area, it is not necessary to flood all link-state changes to all other areas. Therefore, only those routers affected by a change need to run the SPF calculation.
2. Smaller routing tables - When using multiple areas, detailed route entries for specific networks within an area are kept inside the area. Rather than advertise these explicit routes outside the area, the routes can be summarized into one or more summary routes. Advertising these summaries reduces the amount of LSAs propagated between areas but allows all networks to remain reachable.
3. Reduced link-state update (LSU) overhead - LSUs can contain a variety of LSA types, including link-state information and summary information. Rather than send an LSU about each network to every area, advertise a single route or a few summarized routes between areas to reduce the overhead associated with LSUs that cross multiple areas.

QUESTION 86:

The Certkiller network administrator has enabled OSPF across an NBMA network and has issued the command "ip ospf network nonbroadcast" on router CK1 . Given those facts, which two statements are true? (Select two)

A. All routers must be configured in a fully meshed topology with all other routers.
B. DR and BDR elections will not occur.
C. DR and BDR elections will occur.
D. The neighbor command is required to build adjacencies.
E. Interfaces will automatically detect and build adjacencies with neighbor routers.

Answer: C, 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_addressip_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 87:

The Certkiller network has just completed migrating to the OSPF routing protocol. Which two statements are true of the OSPF link-state routing protocol? (Select two)

A. Using the DUAL algorithm, each OSPF router independently calculates its best paths to all destinations in the network.
B. OSPF sends summaries of individual link-state entries every 30 minutes to ensure LSDB synchronization.
C. 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.
D. Using the Bellman-Ford algorithm, each OSPF router independently calculates its best paths to all destinations in the network.
E. OSPF sends updates every 10 seconds.
F. OSPF sends triggered updates when a network change occurs.

Answer: B, F

Explanation:
Open Shortest Path First (OSPF)is an open standards routing protocol. It is important to recognize that Cisco's implementation of OSPF is a standards-based version. This means that Cisco based its version of OSPF on the open standards. While doing so, Cisco has also added features to its version of OSPF that may not be found in other implementations of OSPF. This becomes important when interoperability is needed.
OSPF has become one of the most widely used routing protocols in existence today because of the ability to implement it across multi-vendor platforms. OSPF utilizes Dijkstra's Shortest Path First (SPF) algorithm, which allows for faster network convergence. The popularity of OSPF is continuing to grow with the advent of Multi-Protocol Label Switching (MPLS). Currently, the only routing protocols MPLS traffic engineering has extensions for are OSPF and IS-IS.
John Moy heads up the working group of OSPF. Two RFCs define OSPF: Version 1 is defined by RFC 1131, and Version 2 is defined by RFC 2328. Version 2 is the only version to make it to an operational status. However, many vendors modify OSPF. OSPF is known as a link-state routing protocol. The Dijkstra algorithm is used to calculate the shortest path through the network. Within OSPF, links become synonymous with interfaces.
Some of the advantages of OSPF are as follows:
* Support of hierarchical network design through the use of areas.
* The use of link-state databases reduces the chances of routing loops.
* Full support of classless routing behavior.
* Decreased size in routing tables through the use of manual route summarization. Automatic route summarization is not supported by OSPF.
* Routing updates are sent only when the information is needed, decreasing the use of network bandwidth for routing updates.
* Utilization of multicast packets decreases the impact on routers not running OSPF and end stations.
* Support of authentication, which allows the user to implement more secure networks.
OSPF is a robust protocol, and due to its robustness, you must learn many terms in order to understand the operation of OSPF. The next section covers the terminology necessary to enable you to understand the many operations and procedures performed by the OSPF process.

QUESTION 88:

Which of the three areas below are valid OSPF area types? (Select three)

A. Stub
B. Active
C. Remote
D. Backbone
E. Ordinary or standard

Answer: A, D, E

Explanation:
OSPF Area Types:
The characteristics that are assigned to an area control the type of route information that it can receive. For example, the size of routing tables may need to be minimized in an OSPF area. In this case configure the routers to operate in an area that does not accept external routing information, Type 5 LSAs.
The following are several area types that are possible:
1. Standard area - A standard area can accept link updates and route summaries.
2. Backbone area (transit area) - When interconnecting multiple areas, the backbone area is the central entity to which all other areas connect. The backbone area is always Area 0. All other areas must connect to this area to exchange route information. The OSPF backbone has all the properties of a standard OSPF area.
3. Stub area -A stub area is an area that does not accept information about routes external to the autonomous system, the OSPF internetwork, such as routes from non-OSPF sources. If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0.
4. Totally stubby area - A totally stubby area is an area that does not accept external autonomous system (AS) routes and summary routes from other areas internal to the autonomous system. Instead, if the router needs to send a packet to a network external to the area, it sends it using a 0.0.0.0/0 default route. Totally stubby areas are a Cisco proprietary feature.
5. Not-so-stubby area (NSSA) - An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs.

QUESTION 89:

OSPF networks create Link State Advertisements (LSAs). What kind of LSA is conceived by an Area Border Router (ABR) containing route summarization of area subnets?

A. Router link, type 1.
B. Router link, type 8.
C. Network summary link, type 3.
D. AS external summary link, type 4.

Answer: C

Explanation:
The network summary link entry is a LSA type 3 that is originated by ABRs. It describes the links between the ABR and the internal routers of a local area. These entries are flooded throughout the backbone area to the other ABRs.
Incorrect Answers
A: Router link, LSA type 1 is generated by each router for each area it belongs to. It describes the states of the router's link to the area.

B: Router LSA is type 1.
D. External link LSA is type 5.

QUESTION 90:

Router CK1 has an interface in one OSPF area and another interface in a different OSPF area. What kind of router is CK1 ?

A. ABR
B. ASBR
C. Internal router
D. Backbone router
E. None of the above

Answer: A

Explanation:
An ABR (Area Border Router) shares an interface with at least one other OSPF area.
Incorrect Answers:
B: An ASBR (Autonomous System Border Router) have at least one interface in a non-OSPF network.
C: An internal router is only connected to routers in the internal area.
D: A backbone router has at least one interface in the backbone, also known as Area 0.

QUESTION 91:

When an OSPF ASBR announces external networks, what kind of LSA type does it use?

A. Type 1 LSA
B. Type 2 LSA
C. Type 3 LSA
D. Type 4 LSA
E. Type 5 LSA

Answer: E

Explanation:
In OSPF, Type 5 LSA is used to announce external networks.
The OSPF Router LSA types are described here:
Router LSA (type 1) - information about the router and its directly connected links; flooded only w/in the area
Network LSA (type 2) - information about a LAN and the routers connected to it; originated by the
DR; flooded only w/in the area
Summary LSA (type 3) - describe networks reachable outside the area; originated by

the area border routers (ABR)
ASBR Summary LSA (type 4) - reachability information to the Autonomous System Border
Routers outside the area; originated by the ABRs
External LSA (type 5) - information about destinations outside the OSPF domain; originated by the ASBRs; flooded throughout the whole network
Other LSA type exist for multicast information (type 6) and other extensions, including NSSA area (type 7).

QUESTION 92:

The Certkiller OSPF network is displayed in the diagram below:

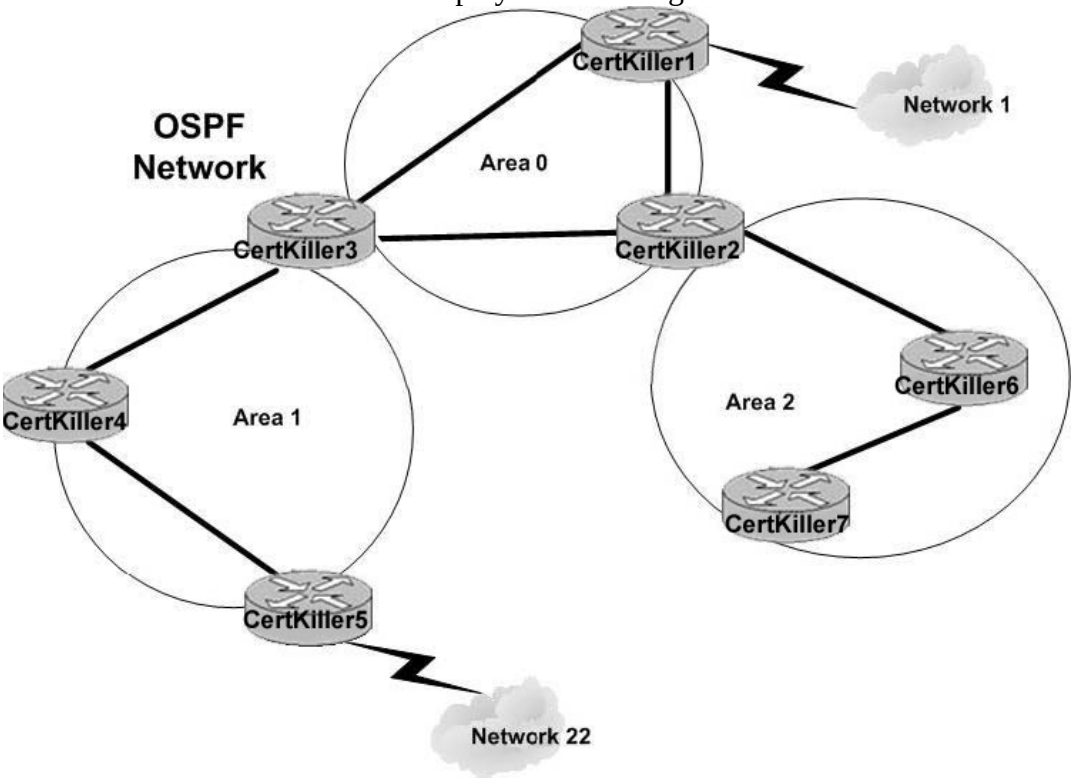

The status of external Network 22 has just changed. The router Certkiller 5 uses a type 7 LSA to update area 1 of this change. In OSPF terminology, what kind of area is area 1 classified as?

A. Backbone area
B. Transit area
C. Stubby area
D. Totally stubby area
E. Not-so-stubby area

Answer: E

Explanation:

NSSA use type 7 LSAs. The different OSPF area types are described below.
1. Stub Areas: These areas do not accept routes belonging to external autonomous systems (AS); however, these areas have inter-area and intra-area routes. In order to reach the outside networks, the routers in the stub area use a default route which is injected into the area by the Area Border Router (ABR).
2. Normal Areas: These areas can either be standard areas or transit (backbone) areas. Standard areas are defined as areas that can accept intra-area, inter-area and external routes.
3. Backbone area is the central area to which all other areas in OSPF connect.
4. Totally Stub Areas: These areas do not allow routes other than intra-area and the default routes to be propagated within the area. The ABR injects a default route into the area and all the routers belonging to this area use the default route to send any traffic outside the area.
5. NSSA: This type of area allows the flexibility of importing a few external routes into the area while still trying to retain the stub characteristic. Assume that one of the routers in the stub area is connected to an external AS running a different routing protocol, it now becomes the ASBR, and hence the area can no more be called a stub area. However, if the area is configured as a NSSA, then the ASBR generates a NSSA external link-state advertisement (LSA) (Type-7) which can be flooded throughout the NSSA area. These Type-7 LSAs are converted into Type-5 LSAs at the NSSA ABR and flooded throughout the OSPF domain
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/ CK4 80/technologies_tech_note09186a0080094a74.shtml

---

QUESTION 93:

What kind of information can you find in a type 3 OSPF network summary link LSA?

A. Summary of routes in the AS.
B. Summary of link states in an OSPF area.
C. Summary of IP subnets in an OSPF area.
D. Summary of metric coast from ABR to ASBR.
E. None of the above.

Answer: C

Explanation:
A type 3 area LSA is Inter-area prefix LSA for ABR It advertises internal networks to routers in other areas (interarea routes). Type 3 LSAs may represent a single network or a set of networks summarized into one advertisement. Only ABRs generate summary LSAs. In OSPF for IPv6, addresses for these LSAs are expressed as prefix, prefix length instead of address, mask. The default route is expressed as a prefix with length 0.
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide_chapter09186a0080145c56.h

QUESTION 94:

DRAG DROP
Drag the OSPF area type icons on the left, and drop them into the centre target area that adjoins the corresponding description on the right.

Answer:

Explanation:
1. Stub Areas: These areas do not accept routes belonging to external autonomous systems (AS); however, these areas have inter-area and intra-area routes. In order to reach the outside networks, the routers in the stub area use a default route which is injected into the area by the Area Border Router (ABR).
2. Normal Areas: These areas can either be standard areas or transit (backbone) areas. Standard areas are defined as areas that can accept intra-area, inter-area and external routes.
3. The Backbone area is the central area to which all other areas in OSPF connect.
4. Totally Stub Areas: These areas do not allow routes other than intra-area and the default routes to be propagated within the area. The ABR injects a default route into the area and all the routers belonging to this area use the default route to send any traffic outside the area.
5. NSSA: This type of area allows the flexibility of importing a few external routes into the area while still trying to retain the stub characteristic. Assume that one of the routers in the stub area is connected to an external AS running a different routing protocol, it now becomes the ASBR, and hence the area can no more be called a stub area. However, if the area is configured as a NSSA, then the ASBR generates a NSSA external link-state advertisement (LSA) (Type-7) which can be flooded throughout the NSSA area. These

Type-7 LSAs are converted into Type-5 LSAs at the NSSA ABR and flooded throughout the OSPF domain
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/ CK4 80/technologies_tech_note09186a0080094a74.shtml

---

QUESTION 95:

In a multi-area OSPF network inter-area summarization produces type's 3 & 4 Link State Advertisements (LSAs). At which router are these LSA types produced?

A. At the ABR.
B. At the ASBR.
C. At the backbone DR.
D. Each internal router calculates and sends network updates.
E. None of the above.

Answer: A

Explanation: Network summary link entries, LSA type 3 or LSA type 4 are originated by ABRs. They describe the links between the ABR and the internal routers of a local area. These entries are flooded throughout the backbone area to the other ABRs.
Note: Type-3 describes routes to networks within the local area and are sent to the backbone area. Type-4 describes reachability to ASBRs. These link entries are not flooded through totally stubby areas.
Incorrect Answers:
B: Autonomous system external link entries, LSA type 5, are originated by ASBRs. They describe routes to destinations external to the autonomous system.
C: Backbone DRs does not generate these LSAs.
D: Internal routers produce router link entries, LSA type 1. They describe the states of the router's link to the area.

QUESTION 96:

What does Cisco recommend for those who set up networks with multiple OSPF areas? (Select two)

A. There should not be more than three areas per route.
B. Area 0 must be larger than any subsequent OSPF area.
C. A router can be a DR or BDR for more than one LAN.
D. You should not run more than one instance of the OSPF process on an ABR.

Answer: C, D

Explanation:
C: A router can be a DR or BDR for more than one LAN. Since the loopback address is

normally used as the router ID, this ID is used for all LAN segments.
D: Running multiple OSPF processes on the same router is not recommended because it creates multiple database instances that add extra overhead.
Incorrect Answers:
A: There is no such limit. The hierarchical structure could be larger.
B: There is no requirement that Area 0 must be the largest OSPF area.

QUESTION 97:

When designing and implementing OSPF in multiple areas; what should be taken in consideration? (Select two)

A. Each area must connect to area 0.
B. Each area must have a unique AS number.
C. Remote areas must be configured as stub or NSSA areas.
D. Traffic between two areas must travel across the backbone area.

Answer: A, D

Explanation:
In OSPF, all areas must be connected to a backbone area, which is area 0. It is also required that traffic from one area to another non-backbone area must traverse area 0, since all areas must connect to area 0.
Incorrect Answers
B: OSPF routers within the same Autonomous System should be configured with the same AS number. The AS number is also called the OSPF process ID.
C: Remote areas are not required to be either stub or not so stubby areas.

QUESTION 98:

In OSPF, a Designated Router (DR) floods LSAs to all the routers on the segment that it has a neighbor relationships with. Over what kind of link does this happen and what kind of LSA type is sent?

A. Router link, type 1.
B. Network link, type 2.
C. External link, type 5.
D. Network summary link, type 3.
E. None of the above

Answer: B

Explanation:
A network Link LSA type 2 is generated by DRs in multi-access networks. They describe the set of routers attached to a particular network. Flooding occurs within the area that contains the network only.

Incorrect Answers:
A: Router link, LSA type 1 is generated by each router for each area it belongs to. It describes the states of the router's link to the area.
C: External link, LSA type 5 is originated by the ASBR. It describes routes to destinations external to the autonomous system. It is flooded throughout an OSPF autonomous system except for stub and totally stubby areas.
D: Network summary link entry, LSA type 3 is originated by ABRs. It describes the links between the ABR and the internal routers of a local area. These entries are flooded throughout the backbone area to the other ABRs.

QUESTION 99:

In an OSPF network, neighbor relationships allow the topology of the network to scale. What are two reasons as to why this can help a network topology scale? (Select two)

A. Neighbor adjacencies control distribution of routing protocol updates.
B. Routing table information does not flood the network until holddown timers have expired.
C. The hello protocol is a more efficient means of sending routing updates than table exchanges used in RIPv1.
D. Topological database is maintained with incremental updates, with full exchange occurring only every 30 minutes.

Answer: A, C

Explanation:
When a network topology changes and becomes more sophisticated, fewer routers are burdened with the responsibility of having to relearn the entire network. The hello protocol is also efficient because rather then focusing on knowing what all the routers know, the routers only focus on who is still around, therefore available for routing.

QUESTION 100:

In an OSPF network; which IP address multicasts all OSPF Designated Routers and Backup Designated Routers?

A. 224.0.0.5
B. 224.0.0.6
C. 224.0.0.9
D. 224.0.0.11
E. 224.0.0.2

Answer: B

Explanation:
224.0.06 is the IP address used by all OSPF DRs and BDRs.

Incorrect Answers:
A: This is the multicast address used by all OSPF routers.
C: This address is used by RIP version 2 routers.
D: This multicast address is used by mobile agents.
E: This multicast IP address is used by all routers on the subnet.
Reference: Building Scalable Cisco Networks (Cisco Press) page 114

QUESTION 101:

What kind of OSPF Link State Advertisement (LSA) is found in a type 3 network summary link?

A. Summary of routes in the AS.
B. Summary of link states in an OSPF area.
C. Summary of IP subnets in an OSPF area.
D. Summary of metric costs from ABR to ASBR.

Answer: C

Explanation:
Type 3 LSAs are INTER-AREA advertisements and they are not used within the local area.
The various LSA types used in OSPF are described below:
LS Type Link State ID _____ 1 The originating router's Router ID. 2 The IP interface address of the network's Designated Router. 3 The destination network's IP address. 4 The Router ID of the described AS boundary router. 5 The destination network's IP address.

QUESTION 102:

You are the administrator of the large Certkiller OSPF and are considering dividing it to multiple areas. What are two advantages of configuring multiple OSPF areas? (Select two)

A. It eliminates the need for a DR or BDR in the OSPF network.
B. It eliminates security concerns by segregating portions of the network.
C. Type 1 and 2 LSAs are confined to a single area, reducing routing overhead.
D. Area members have smaller topological databases than if the network was one large area, requiring less CPU to derive routes.

Answer: C, D

Explanation:
C: Link State update (LSU) traffic is reduced. Rather than send an LSU about each network within an area, you can advertise a single or fewer summarized routes between areas to reduce the overhead associated with linkstate updates when they are crossing

areas.
Note: All LSA types, except the AS-external-LSAs (LS type = 5), are flooded throughout a single area only.
D: Topological database size would decrease and fewer CPU cycles would be required to calculate routes.
Incorrect Answers:
A: A DR and BDR for each area would still be required.
B: Security concerns would not be eliminated. There are no security advantages in using more than one area in an OSPF network.

QUESTION 103:

OSPF has been configured on router CK1 , which lies on a multi point network. On this multi-access network, what type of router floods type 2 Link State Advertisements (LSAs) into an OSPF area?

A. DR
B. ABR
C. BDR
D. ASBR
E. None of the above

Answer: A

Explanation:
Type 2 LSAs are generated by Designated Routers (DRs) in multi-access networks. They describe the set of routers attached to a particular network and are flooded within the area that contains the network only.
Note: Instead of routers exchanging link-state information with every other router on the segment, each router sends the link-state information to the DR and Backup Designated Router (BDR). The DR sends each router's link-state information to all other routers in the network. This flooding process significantly reduces the router-related traffic on a segment.
Incorrect Answers:
B: Area Border Routers generate Type 3 and Type 4 LSAs.
C: BDRs are just used for backup, should the DR fail.
D: Autonomous System Boundary Routers (ASBRs) generate Type 5 LSAs.

QUESTION 104:

You want router one of the OSPF areas on your network to be able to accept type 7 external routers, but not type 5 external routes. What kind of area should be configured to accomplish this?

A. A not-so stubby area
B. A stubby area

C. A totally stubby area
D. A backbone area
E. An on-demand area

Answer: A

Explanation:
The NSSA external LSA -Identified as Type 7, these LSAs are created by the ASBR residing in a not so stubby area (NSSA). This LSA is similar to an autonomous system external LSA, except that this LSA is contained within the NSSA area and is not propagated into other areas, but it is converted into a Type 5 LSA by the ABR.
An NSSA is an area that is seen as a stub area but can receive external routes, which it will not propagate into the backbone area and thus the rest of the OSPF domain. Another LSA, Type 7, is created specifically for the NSSA. This LSA can be originated and communicated throughout the area, but it will not be propagated into other areas, including Area 0. If the information is to be propagated throughout the AS, it is translated into an LSA Type 5 at the NSSA ABR.

QUESTION 105:

DRAG DROP
Match the OSPF area type on the left to the type of OSPF routes (as would be seen in the routing table) it will support on the right hand side:

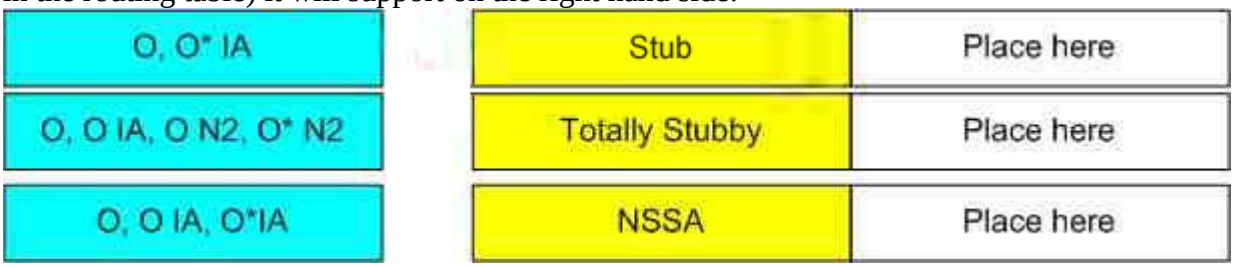

Answer:

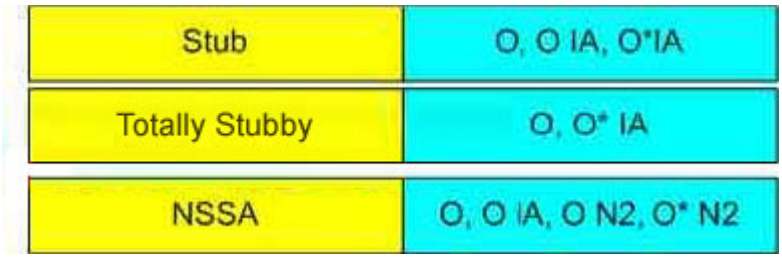

QUESTION 106:

Which of the following OSPF routes are supported by the Totally Stubby type?
(Select two)

A. O
B. O IA
C. O N2

D. O* N2
E. O* IA

Answer: A, E

Explanation:
The only routes that are seen in a totally stubby area are intra-area routes (O) and the default Intra area route (O*IA).
Example routing table of a totally stubby router:
CK1 #show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate
default
Gateway of last resort is not set
203.250.15.0 255.255.255.252 is subnetted, 1 subnets
C 203.250.15.0 is directly connected, Serial0
131.108.0.0 255.255.255.240 is subnetted, 1 subnets
O 131.108.79.208 [110/74] via 203.250.15.1, 00:31:27, Serial0
O*IA 0.0.0.0 0.0.0.0 [110/74] via 203.250.15.1, 00:00:00, Serial0
Note that the only routes that show up are the intra-area routes (O) and the default-route 0.0.0.0. The external and inter-area routes have been blocked. The cost of the default route is now 74 (64 for a T1 line + 10 advertised by RTC). No configuration is needed on RTE in this case. The area is already stub, and the no-summary command does not affect the Hello packet at all as the stub command does.

QUESTION 107:

Which type of OSPF network requires manual configuration of OSPF neighbors?

A. broadcast multi-access
B. nonbroadcast multi-access
C. point-to-point
D. point-to-multipoint

Answer: B

Explanation:
Non-broadcast multi-access networks, as found in frame-relay and ATM, require the neighbors to be explicitly specified. This is the case where the PVCs are configured on the physical interface using frame map commands. When sub-interfaces are used, the neighbor commands do not need to be configured.
Note: On the physical interface of an NBMA network, the command "ip ospf

network-type point-to-multipoint" command can be used instead of specifying the neighbors individually.

QUESTION 108:

Which of the following OSPF routes are supported by the Stub Area type? (Select three)

A. O
B. O IA
C. O N2
D. O* N2
E. O* IA

Answer: A, B, E

Explanation:
When an area is configured as a stub area, it will learn about all routes, except for the external routes.
Example routing table of a stub router:
CK2 #show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate
default
Gateway of last resort is 203.250.15.1 to network 0.0.0.0
203.250.15.0 255.255.255.252 is subnetted, 1 subnets
C 203.250.15.0 is directly connected, Serial0
O IA 203.250.14.0 [110/74] via 203.250.15.1, 00:26:58, Serial0
128.213.0.0 255.255.255.252 is subnetted, 1 subnets
O IA 128.213.63.0 [110/84] via 203.250.15.1, 00:26:59, Serial0
131.108.0.0 255.255.255.240 is subnetted, 1 subnets
O 131.108.79.208 [110/74] via 203.250.15.1, 00:26:59, Serial0
O*IA 0.0.0.0 0.0.0.0 [110/65] via 203.250.15.1, 00:26:59, Serial0
Note that all the routes show up except the external routes which were replaced by a default route of 0.0.0.0.

QUESTION 109:

Which of the following OSPF routes are supported by the Not-So-Stubby-Area (NSSA) type? (Select three)

A. O N1
B. O IA
C. O N2

D. O* N2
E. O* IA

Answer: B, C, D

Explanation:
The various route types used by OSPF are as follows:
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
An OSPF NSSA will receive inter-area, external type 2, and default external type 2 routes.

QUESTION 110:

Routers CK1 , CK2 , and CK3 are part of OSPF area 100 as shown below:

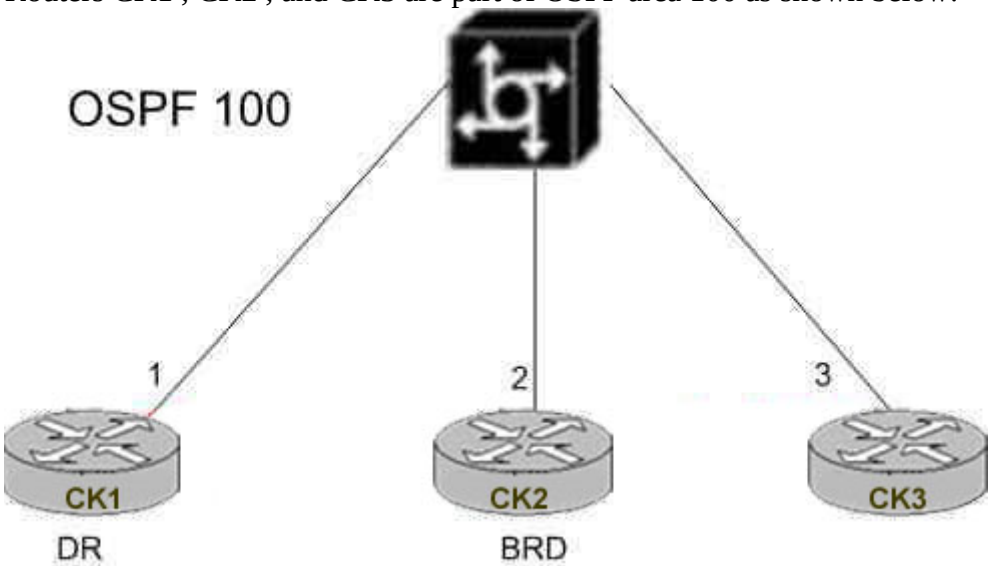

In the Certkiller network shown above, router CK1 was elected the DR and CK2 was elected the BDR. Assume that CK1 fails, and that CK2 takes the place of the DR while CK3 becomes the new BDR. What will happen when CK1 comes back online?

A. CK1 will take the place of DR immediately upon establishing its adjacencies.
B. CK1 will take the place of DR only if CK2 fails.
C. CK1 will take the place of DR only if both CK2 and CK3 fail.
D. A new election will take place establishing an all new DR and BDR based on configured priority levels an 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 111:

Which of the following are features are characteristics of an OSPF ABR? (Choose three)

A. maintains one link-state database for all areas
B. routes traffic to and from other areas
C. maintains a separate link-state database for each area
D. originates type 3 link-state advertisements which can be summarized
E. has at least one physical interface in area 0.
F. has at least one interface in an external internetwork.

Answer: B, C, D

Explanation:
ABR forwards summary LSAs. It forwards both Type 3 LSAs and Type 4 LSAs. Type 3 LSAs are forwarded to the other ABRs, and Type 4 LSAs are forwarded to the ASBRs. ABR also forwards Type 3 LSAs from other areas into its own area. If the ABR has multiple links in the same area, it also forwards Type 1 and Type 2 LSAs in its capacity as an internal router.

QUESTION 112:

The Certkiller OSPF network is shown in the diagram below:

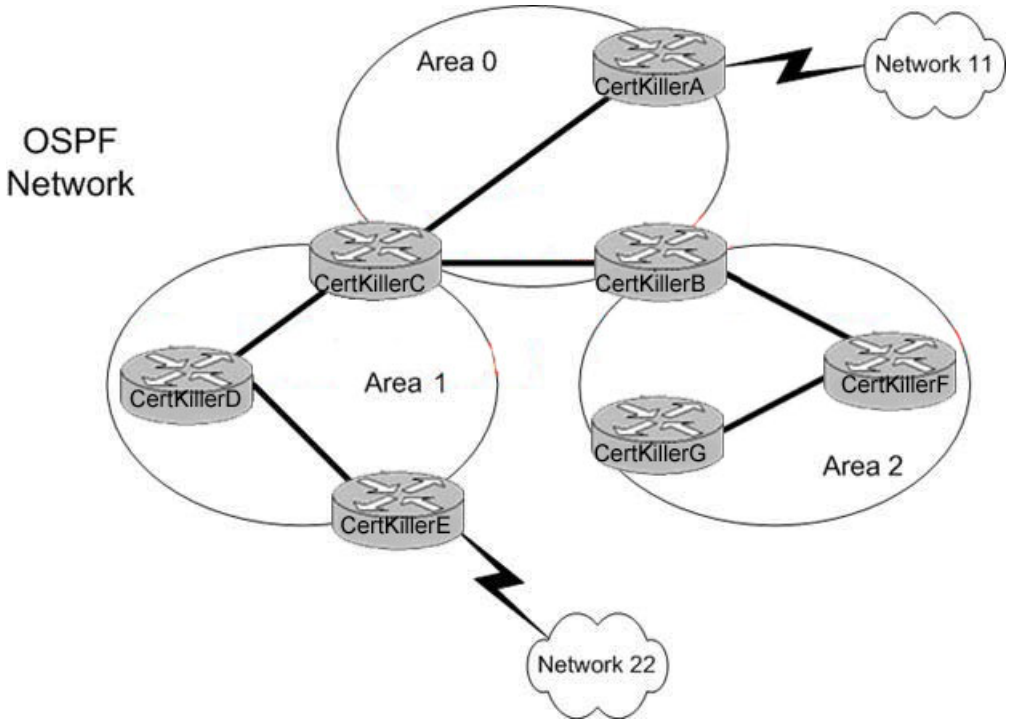

Which type of Link State Announcement will Certkiller E use to announce changes in network 22 to Area 1?

A. Type 1 LSA
B. Type 2 LSA
C. Type 3 LSA
D. Type 4 LSA
E. Type 5 LSA

Answer: E

Explanation:
The external link LSA is originated by AS boundary routers and is flooded throughout the AS. Each external advertisement describes a router to a destination in another autonomous system. Default routes for the AS can also be described by AS external advertisements. This is identified as a Type 5 LSA.

QUESTION 113:

What is the OSPF cost for the 100 Mbps Ethernet media type?

A. 1
B. 2
C. 6
D. 10
E. 100

Answer: A

Explanation:
By default, OSPF uses the formula of 100,000,000 (100M) divided by bandwidth, or 100,000,000/BW. For a 100M fast Ethernet connection, the formula is 100M/100M = 1.

QUESTION 114:

Which of the following OSPF routes are supported by the Not-So-Stubby-Area (NSSA) type? (Select three) (801)

A. O
B. O N2
C. O* N2
D. O E1
E. O E2

Answer: A, B, C

Explanation:
The various route types used by OSPF are as follows:
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
An OSPF NSSA will receive inter-area, external type 2, OSPF routes.

QUESTION 115:

Which three are advantages to creating multiple areas in OSPF? (Select 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:
OSPF employs the concept of an area, which is a grouping of contiguous OSPF networks and hosts. OSPF areas are logical subdivisions of OSPF autonomous systems whose internal topology is hidden to routers outside the area. Areas allow an additional level of hierarchy different from that provided by IP network classes, and they can be used to

aggregate routing information and mask the details of a network. These features make OSPF particularly scalable to large networks. When using areas, OSPF routers need to maintain information only within their own area so less frequent SPF calculations are required, smaller routing tables occur with area summarizations, and the link state update overhead is reduced.

QUESTION 116:

Which of the following items must match in an OSPF hello packet in order for two routers to become OSPF neighbors? (Choose four)

A. Neighbors
B. Hello/dead intervals
C. Area ID
D. Address of DR/BDR
E. Stub area flag
F. Authentication password

Answer: B, C, E, F

Explanation:
The following hello parameters must match in order for two OSPF routers to become adjacent neighbors:
1. OSPF area number
2. OSPF area type, such as stub or NSSA flag
3. Subnet and subnet mask
4. OSPF HELLO and Dead timer values
5. Authentication values/passwords

QUESTION 117:

SIMULATION
Refer to the following network topology exhibit:

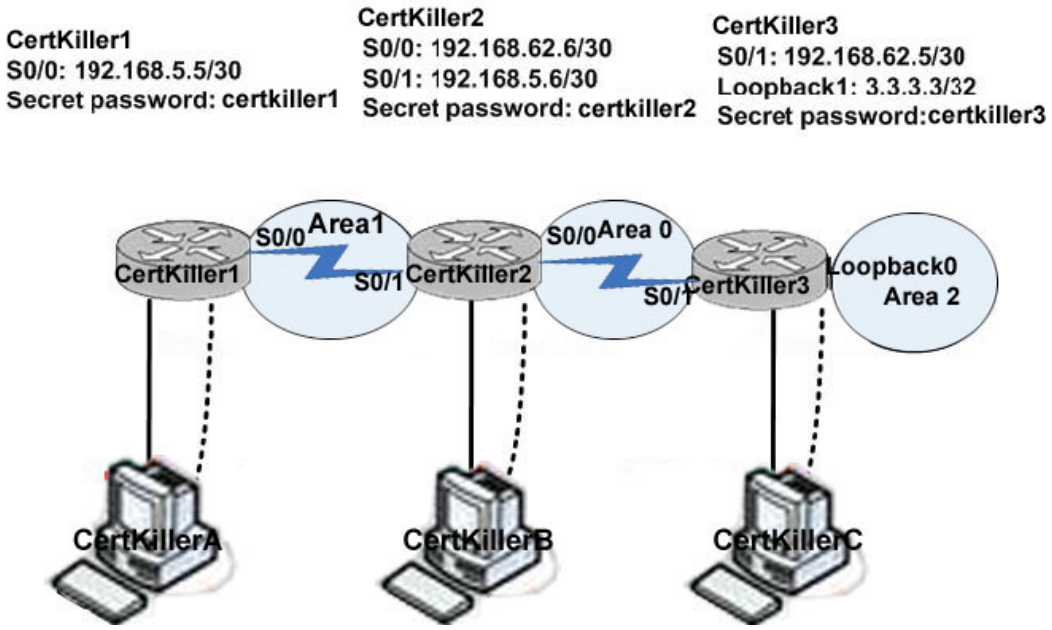

You work as a network technician at Certkiller .com. Certkiller .com is a large international company with offices in USA, Europe, and Japan. You work at the Tokyo office. You are administering three routers Certkiller 1, Certkiller 2, and Certkiller 3. Certkiller 2 and Certkiller 3 have been configured for OSPF.
The current situation is such as Area 0 and Area 2 have been configured correctly as follows:
* the S0/1 interface of Certkiller 3 is in Area 0.
* The S0/0 interface of Certkiller 2 is in Area 0.
* The loopback address of Certkiller 3 is in Area2.
Your boss at the Certkiller .com Tokyo office, Miss Certkiller, has asked you to Area 1 as follows:
* the S0/0 interface of Certkiller 1 should be in Area 1
* the S0/1 interface of Certkiller 2 should be in Area 1
* no other interfaces should be in Area 1
* no external routes or inter-area, with the exception of the default route, should be receive by Area 1.
As a final advice Certkiller tells you to make sure that you use the appropriate OSPF mask.

Answer:

Explanation:
Certkiller 1# conf t
Certkiller 1(config)# router ospf 1
Certkiller 1(config-router)# network 192.168.5.4 0.0.0.3 area 1
Certkiller 1(config-router)# area 1 stub
Certkiller 1(config-router)# end
Certkiller 1# copy run start
Certkiller 2# conf t
Certkiller 2(config)# router ospf 1

Certkiller 2(config-router)# network 192.168.5.4 0.0.0.3 area 1
Certkiller 2(config-router)# area 1 stub no-summary
Certkiller 2(config-router)# end
Certkiller 2# copy run start
1. Stub area -A stub area is an area that does not accept information about routes external to the autonomous system, the OSPF internetwork, such as routes from non-OSPF sources. If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0.
2. Totally stubby area - A totally stubby area is an area that does not accept external autonomous system (AS) routes and summary routes from other areas internal to the autonomous system. Instead, if the router needs to send a packet to a network external to the area, it sends it using a 0.0.0.0/0 default route. Totally stubby areas are a Cisco proprietary feature.
To configure an area as a stub or totally stubby area, use the following syntax on all router interfaces that are configured to belong to that area:
Router(config-router)#area area-id stub
The no-summary keyword creates a totally stubby area. The area stub command is configured on each router in the stub location, which is essential for the routers to become neighbors and exchange routing information. When this command is configured, the stub routers exchange Hello packets with the E bit set to 0. The E bit is in the Options field of the Hello packet. It indicates that the area is a stub area. The state of this bit must be agreed upon otherwise the routers will not become neighbors.

QUESTION 118:

Router CK1 is an OSPF router connected to area 1. 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 metric-type 2
E. router ospf 1
area 1 nssa default-information originate
F. None of the above

Answer: B

Explanation:
Not-so-stubby area (NSSA) - An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7

LSA routes into Type 5 LSAs
To configure an OSPF area as a NSSA, configure all OSPF router interfaces that belong to the area using the following command syntax:
Router(config-router)#area area-id nssa [no-summary]
Typically, use the optional keyword no-summary when configuring NSSA on an ABR. This prevents Type 3 and Type 4 summary routes from flooding the NSSA area and minimizes the routing tables within the area. In effect, the no-summary keyword makes the NSSA totally stubby.

QUESTION 119:

The Certkiller multi-protocol WAN is shown below:

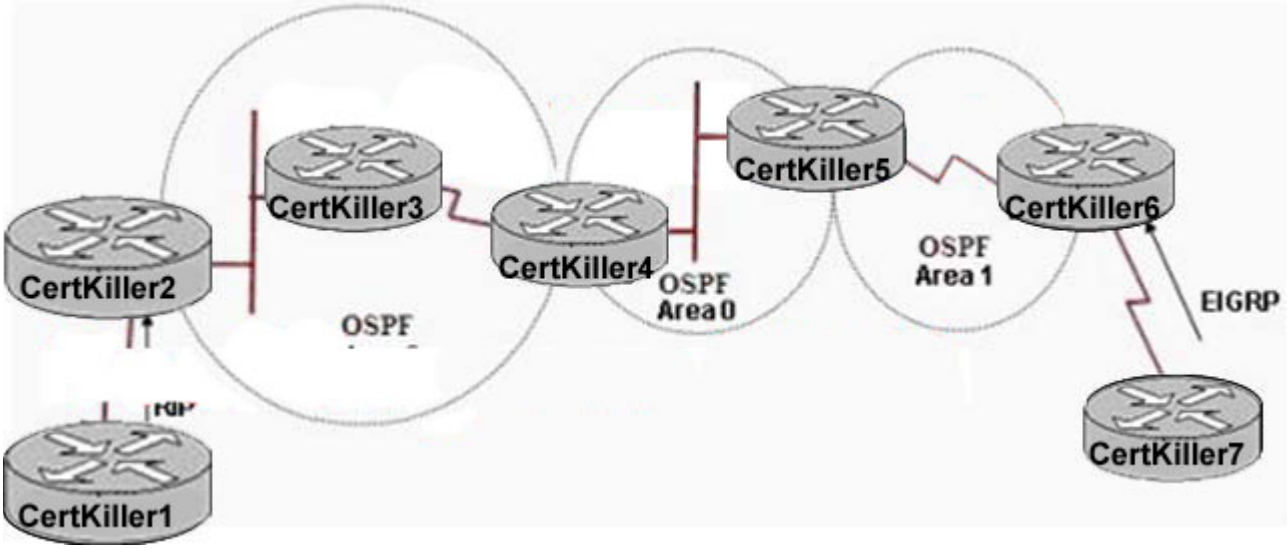

In this network, routers Certkiller 2, Certkiller 3, Certkiller 4, and Certkiller 5 have OSPF enabled. What should be configured on the routers in area 1 to ensure that all default summary routes and redistributed EIGRP routes will be forwarded from Certkiller 6 to area 1, and only a default route for all other OSPF routes will be forwarded from Certkiller 5 to area 1.

A. Certkiller 5(config-router)# area 1 nssa no-summary
Certkiller 6(config-router)# area 1 nssa
B. Certkiller 5(config-router)# area 1 stub no-summary
Certkiller 6(config-router)# area 1 stub
C. Certkiller 5(config-router)# area 1 nssa
Certkiller 6(config-router)# area 1 nssa
D. Certkiller 5(config-router)# area 1 stub
Certkiller 6(config-router)# area 1 stub

Answer: A

Explanation:
Not-so-stubby area (NSSA)
- An NSSA is an area that is similar to a stub area but allows for importing external

routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs
To configure an OSPF area as a NSSA, configure all OSPF router interfaces that belong to the area using the following command syntax:
Router(config-router)#area area-id nssa [no-summary]
Typically, use the optional keyword no-summary when configuring NSSA on an ABR. This prevents Type 3 and Type 4 summary routes from flooding the NSSA area and minimizes the routing tables within the area. In effect, the no-summary keyword makes the NSSA totally stubby.

QUESTION 120:

Assuming that you are configuring an ABR in an OSPF area, which IOS command would you execute if your goal was to summarize the networks advertised out of the area?

A. summary-address address mask
B. area area-id range address mask
C. auto-summary address mask area area-id
D. network network-number wildcard mask area area-id
E. None of the above

Answer: B

Explanation:
The area area-id range address mask command consolidates IA (intra-area) routes on an ABR. This command instructs the ABR to summarize routes for a specific area before injecting them into a different area.
Incorrect Answers:
A: The summary-address address mask command consolidates external routes (inter-area) on an ASBR
C: Auto-summarization is not useful here.
D: The network command cannot be used for this purpose. This command is simply used to add a network to the OSPF routing process.

QUESTION 121:

Which two of the following characteristics are defined by the network command? (Select two)

A. The OSPF area ID
B. The OSPF router ID
C. The OSPF process ID
D. Which interface belongs to which OSPF area

Answer: A, D

Explanation:
The network command designates the OSPF area for an interface with the specified IP address.
Syntax: network address wildcard-mask area area-id

QUESTION 122:

Router CK1 is configured for OSPF as shown below:
router ospf 76
network 172.22.23.0 0.0.0.0 area 1
network 172.18.0.0 0.0.255.255 area 0
area 0 range 172.18.0.0 255.255.0.0
area 1 range 172.22.23.0 255.255.255.0
Which of the following statements are true regarding the above configuration?
(Select three)

A. The OSPF router ID is 76.
B. This is an area border router.
C. The designated router priority is 76.
D. This router connects area 1 to the backbone area.
E. Any router interfaces with an address of 172.18.x.x are in area 0.

Answer: B, D, E

Explanation:
An interface may belong to only one area. If a router has multiple interfaces and if any of those interfaces belong to different areas, the router is considered as an area border router. The networks that follow the network command are connected to each other and Area O is always the backbone area. Finally, the command "network 172.18.0.0 0.0.255.255 area 0" identifies that all interfaces with IP address of 172.18.0.0 area within area 0.
Incorrect Answers:
A, C: The command router ospf 76 identifies the process ID as 76 and not the router ID nor the router priority.
Catherine Paquet and Daine Teare, "Building Scalable Cisco Networks" (Cisco Press 2001), pp 178-181.

QUESTION 123:

When configuring a multi-area OSPF network to summarize routes, what additional command is required by ASBR's that is not needed by ABR's?

A. area range command
B. ospf summarize command
C. aggregate-route command
D. summary-address command

E. None of the above

Answer: D

Explanation:
Using this command for OSPF causes an OSPF autonomous system boundary router (ASBR) to advertise one external route as an aggregate for all redistributed routes that are covered by the address. For OSPF, this command summarizes only routes from other routing protocols that are being redistributed into OSPF.
Incorrect Answers:
A: The area range command is used only with area border routers (ABRs). 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.
B, C: There are no such commands for OSPF.
Reference: OSPF Commands

QUESTION 124:

A Certkiller OSPF router is configured in the following manner:
router ospf 200
network 203.42.67.0 0.0.0.255 area 7
network 203.42.68.0 0.0.0.255 area 0
area 7 stub no-summary
area 7 default-cost 30
Which of the following statements are true regarding this configuration? (Select two)

A. Area 7 is a totally stubby area.
B. If the backbone becomes discontiguous, traffic can be routed through area 7.
C. Redistribution of other routing protocols takes place at the area designated router.
D. Area 7 non-ABR routers contain only intra-area routing information and a default route.

Answer: A, D

Explanation:
An extension to stub areas is what is called "totally stubby areas". Cisco indicates this by adding a "no-summary" keyword to the stub area configuration. A totally stubby area is one that blocks external routes and summary routes (inter-area routes) from going into the area.
Note: The area stub command is used to define an area as a stub area.
Syntax: area area-id stub [no-summary]
The no-summary optional parameter prevents an ABR from sending summary link advertisements into the stub area.
Reference: OSPF Design Guide
http://www.cisco.com/warp/public/104/3.html

QUESTION 125:

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 above

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.
Incorrect Answers:
A: For a virtual link to work both ends need to be configured.
B, C: These commands are not specific to an ABR only.

QUESTION 126:

Router CK1 is an OSPF ABR that connects area 3 to the network. You wish to configure area 3 to summarize the IP address range 172.16.20.192 - 172.16.20.223. Which two of the following IOS commands would you execute to accomplish this? (Select two)

A. network 172.16.20.192 0.0.0.31 area 3
B. area 3 range 172.16.20.192 172.16.20.223
C. area 3 range 172.16.20.192 255.255.255.224
D. network 172.16.20.192 255.255.255.224 area 3

Answer: A, C

Explanation:
A: When configuring multiple OSPF areas, make sure to associate the correct network addresses with the desired area ID. Syntax: network address wildcard-mask area area-id.
We must use a wildcard mask (0.0.0.31) and not a network mask (255.255.255.224)
C: We must instruct the ABR to summarize routes for a specific area before injecting them into a different area. Syntax: area area-id range address mask
Incorrect Answers:
B: This is the wrong syntax. We should use a network mask to specify the address range.
D: We should use a wildcard mask, not a network mask with the network command.

QUESTION 127:

Router CK1 is an ABR that is part of the Certkiller OSPF network. CK1 has been configured with the "area range" command. What does this command provide?

A. It provides the range of areas in the OSPF internetwork.
B. It provides the range of areas in which this ABR participates.
C. It provides a summary of networks to be advertised outside the area.
D. It provides a summary of networks outside the area and inside the AS.

Answer: C

Explanation:
The area range command creates/deletes an area address range entry and optionally specifies whether to advertise the addresses. It is used for summarizing routes into an area.
Syntax: [no] area area-id range address mask [do-not-advertise]
Description: Address ranges are used to aggregate address ranges from within the area into one single advertisement sent into the adjacent areas, or to prevent the advertisement of networks.
An OSPF area is defined as a list of address ranges. Each address range consists of an address and mask pair. Area range entries are used to aggregate network information before the advertisements are flooded into the backbone. A single summary LSA is originated for each range.

QUESTION 128:

The Certkiller OSPF network is depicted below:

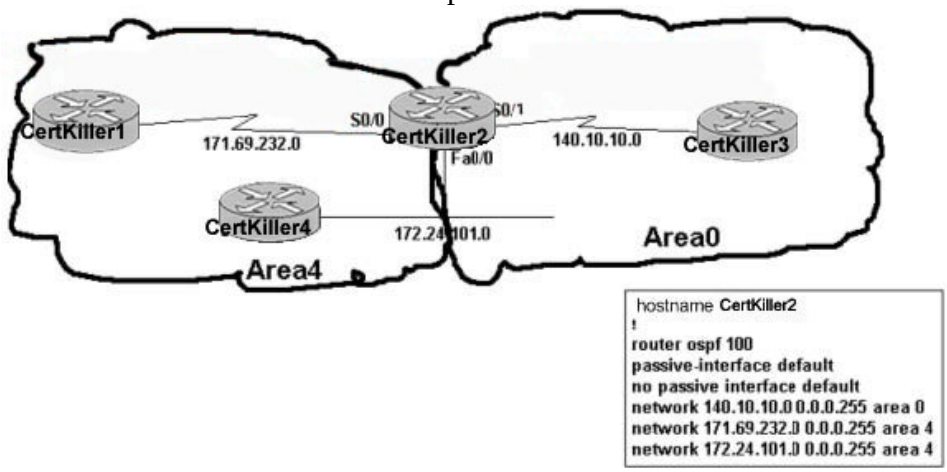

Which two statements concerning Certkiller 2 are true about the above configuration? (Choose two.)

A. The interfaces S0/1 and Fa0/0 will not process any routing update that the router receives.

B. The interfaces S0/1 and Fa0/0 are configured as passive interfaces.
C. The S0/0 interface will not process any routing update that it receives.
D. The S0/0 interface will not sending out routing updates.
E. Certkiller 3 will hear about the 171.69.232.0 and 172.24.101.0 networks through OSPF update.
F. Certkiller 1 and Certkiller 4 will hear about the 140.10.0.0 network through OSPF update.

Answer: E, F

Explanation:
This is a bit of a trick question, since the statement "no passive interface default" command will effectively cancel out the previous command, which will result in normal OSPF operation. Due to this, all networks will be advertised normally via the OSFP process.

QUESTION 129:

The Certkiller OSPF network is shown below:

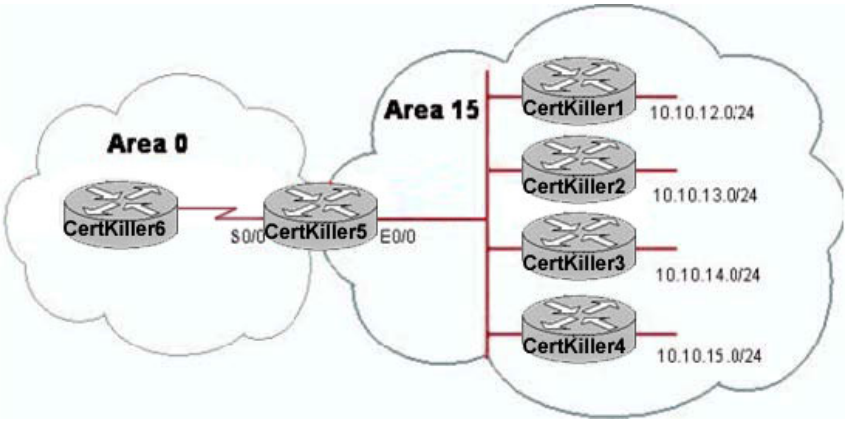

Based on the info shown above, which command will summarize area 15?

A. on the ABR E0/0 interface: area 15 range 10.10.8.0 255.255.248.0
B. on the ABR S0/0 interface: area 15 range 10.10.8.0 255.255.252.0
C. on the ABR in the OSPF config-router mode: area 15 range 10.10.12.0 255.255.252.0
D. on all routers in area 15 in the OSPF config-router mode: area 15 range 10.10.12.0 255.255.248.0
E. on the ABR S0/0 interface: summary-address 10.10.8.0 255.255.248.0
F. on the ABR in the OSPF config-router mode: summary-address 10.10.12.0 255.255.252.0

Answer: C

Explanation:
To consolidate and summarize routes at an area boundary, use the area range command in router configuration mode. The area range command is used only with Area Border

Routers (ABRs). 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 C specifies the correct syntax and route summarization network mask.

QUESTION 130:

Certkiller Network Topology:

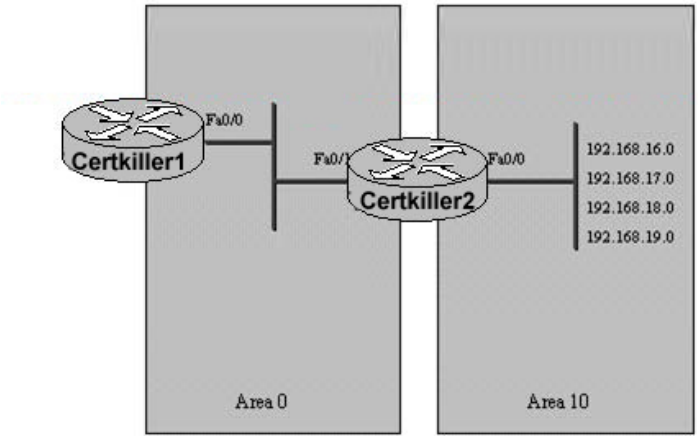

You work as a network technician at Certkiller .com. Study the exhibit above carefully. Which command will Certkiller 2 use to summarize routes for the 192.168.16.0/22 supernet before injecting them into Area 0?

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 Certkiller 2 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 131:

OSPF is being configured over the Certkiller frame relay network as displayed in the diagram below:

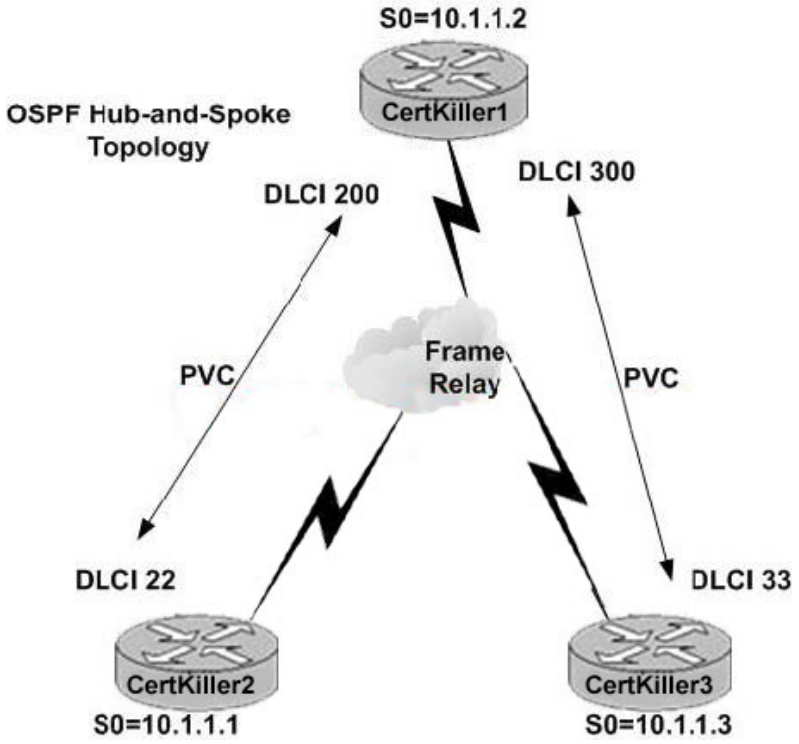

Which three of the following commands would you enter on the Certkiller 1 serial 0 interface, if you wanted to implement OSPF on this network? (Select three)

A. ip ospf network point-to-point
B. ip ospf network point-to-multipoint
C. frame-relay map ip 10.1.1.1 200
D. frame-relay map ip 10.1.1.3 300
E. frame-relay map ip 10.1.1.1 200 broadcast
F. frame-relay map ip 10.1.1.3 300 broadcast

Answer: B, E, F

Explanation:
The ip ospf network command, typed under the interface configuration mode, is used to specify the OSPF network configuration and sets the network mode to point-to-multipoint
DLCI- Data-link connection identifier (DLCI) number.
Broadcast - Forwards broadcasts to the specified IP address. This keyword is needed on the frame relay mappings in order to transmit the OSPF information across the frame relay network.
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a0080087b42.html

---

QUESTION 132:

Router CK1 is configured as shown below:

```
interface serial0
ip address 110.1.1.1 255.255.255.0
encapsulation frame-relay
ip ospf network point-to-multipoint
router ospf 51
network 110.1.1.0 0.0.0255 area 0
```

Based on the information above, which two of the following statements are true?
(Select two)

A. DB/BDR elections do not take place.
B. The router is restricted to a hub and spoke topology.
C. The area 0 NBMA cloud is configured as more than one subnet.
D. OSPF neighbor statements are not necessary.
E. The OSPF hello timers need to be adjusted manually.

Answer: A, D

Explanation:
In an OSPF Point-to-Multipoint environment, DB/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 133:

You have been assigned the task of connecting two office networks together via a frame relay network, and running OSPF across this network. What kind of configuration structure would you use to accomplish this?

A. Point-to-point over sub-interfaces.
B. Point-to-multipoint star configuration.
C. Point-to-multipoint using a single subnet.
D. Point-to-multipoint nonbroadcast using a single subnet.

E. None of the above.

Answer: A

Explanation:
Using point to point subinterfaces will eliminate the issues surrounding split horizons over NBMA networks. Configuring the network in this way is the Cisco recommended solution.
Reference: RFC1586, Guidelines for Running OSPF over Frame Relay Network.

QUESTION 134:

A Certkiller router is configured as shown below:
interface serial 0
ip address 164.67.36.1 255.255.255.224
encapsulation frame-relay
ip ospf network non-broadcast
!
router ospf 1
network 164.67.36.0 0.0.0.31 area 0
neighbor 164.67.36.2
neighbor 164.67.36.3
Based on this configuration, which of the following statements are true? (Select two)

A. There can be no DR or BDR in this configuration.
B. This is a point-to-point configuration over Frame Relay.
C. The network mode is nonbroadcast multiaccess (NBMA).
D. The DR and BDR need a static list of neighbors due to non-broadcast.

Answer: C, D

Explanation:
The configuration printout shows an OSPF configuration across a frame relay network that is not configured for subinterfaces. With this setup you are required to specify neighbors so that a DR and BDR can be elected. In addition, all locations are part of the same IP subnet so the OSPF neighbors need to be explicitly assigned at the host location of an NBMA network.
Note: Had subinterfaces been used, the network would be seen as a series of logical point to point links, so the OSPF neighbors would not need to be specified.

QUESTION 135:

Refer to the Certkiller network exhibits shown below.
Network topology exhibit:

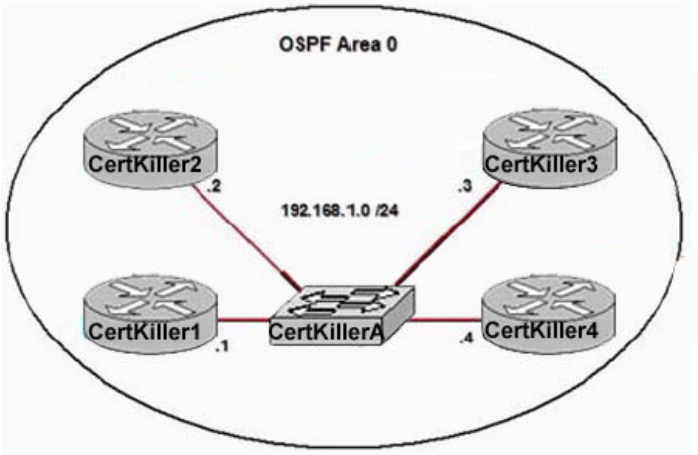

Configuration exhibit:

```
CertKiller4#show ip ospf neighbor

Neighbor ID     Sri     State           Dial Time   Address         Interface
192.168.1.2       1     output          00:00:37    192.168.1.2     Ethernet0
192.168.1.1       1     omitted         00:00:38    192.168.1.1     Ethernet0
192.168.1.3       1                     00:00:32    192.168.1.3     Ethernet0
```

Study the exhibits shown above carefully. All routers have simultaneously been reloaded, and the DR election has concluded as expected. Based on this information, which state is Certkiller 3 in?

A. 2WAY/BDR
B. 2WAY/DR
C. FULL/DR
D. 2WAY/DROTHER
E. FULL/DROTHER
F. FULL/BDR
G. None of the above

Answer: F

Explanation:
Each OSPF interface (multi-access networks only) possesses a configurable router priority. The Cisco default is 1. If you don't want a router interface to participate in the DR/BDR election, set the router priority to 0 using the ip ospf priority command in interface configuration mode.
Let's go through the steps that occur in any given router when the DR and BDR are elected for a specific multi-access network of which it is a member:
1. The local router creates a list of eligible neighbors. The eligible neighbors are those with which the local router has entered at least a 2Way state. The local router adds itself to this list and to all lists that are formed from this list in subsequent steps, for which it qualifies.
2. From this list, all routers whose participating interfaces have a router priority of 0 are removed. These routers will be among the DROthers on this network.
3. The local router makes note of the current DR and BDR values for later comparison.
4. A list of all routers not claiming to be the DR (their own Hello packets do not list them

as the DR) is compiled from the list resulting from step 2.
5. The local router will select the BDR from the list in step 4, based on the following criteria in order:
* If one or more of the routers in the list have declared themselves the BDR, then the one of these with the highest priority is selected to be the BDR.
* If all router priorities are equal, the router with the highest router ID becomes the BDR.
* If no router in the resulting list from step 4 has declared itself the BDR, then the router with the highest router priority is selected to be the BDR.
* If all routers have the same router priority, then the router with the highest router ID is selected to be the BDR.
6. A list of all routers claiming to be the DR (their own Hello packets list them as the DR) is compiled from the list resulting from step 2.
7. The local router will select the DR from the list in step 6, based on the following criteria in order:
* The router with the highest router priority is selected to be the DR.
* If all router priorities are equal, the router with the highest router ID is selected to be the DR.
* If the resulting list from step 6 is empty, meaning that no router has declared itself the DR, then the BDR that was selected in step 5 becomes the DR.
8. If the local router's DR status has been altered by the preceding steps, either causing it to become, or cease to be, the DR, based on a comparison to the results of step 3, then it will repeat steps 4 through 7. This serves to make sure that the local router does not declare itself both the DR and the BDR, because both may be declared by default. If selected to be the DR, the local router will definitely not make the BDR list the next time around in step 4. Conversely, if dethroned as the DR, the local router can become eligible to be the BDR the next time around, by making the list in step 4.

QUESTION 136:

The following output was seen on router Certkiller 1:

```
CertKiller1#show ip route ospf

o IA  6.0.0.0/0  [110/65] via 5.0.0.2, 00:00:18, Serial2/1/0
O*N2 0.0.0.0/0 [110/1] via 5.0.0.2, 00:00:18, Serial2/1/0
```

Based on the output shown above, which of the following statements is true?

A. A default route is configured on the local router.
B. OSPF router 5.0.0.2 is an ABR.
C. The default route is learned from an OSPF neighbor.
D. Network 6.0.0.0/8 was learned from an OSPF neighbor within the area.
E. None of the above.

Answer: D

Explanation:

The various types routes used by OSPF:

O	OSPF networks from within the same area as the router. These are networks learned from router and network LSAs.
O IA	OSPF interarea. This is sent out by the ABRs and is created from the summary link LSA (type 3 and type 4). These routes will not be seen on a router within a totally stubby area because it will not receive LSAs external to the area.
O N1	OSPF NSSA external type 1
O N2	OSPF NSSA external type 2
O E1	OSPF external type 1. These routes are generated by the ASBR and show routes that are external to the autonomous system. The cost of this external route is the summarization of the external cost, plus the cost of the path to the ASBR. These routes will not be seen in a stub or totally stubby area.
O E2	OSPF external type 2. These routes do not take into account the cost of the path to the ASBR. They consider only the external cost.

QUESTION 137:

The Certkiller OSPF network is shown in the following exhibit:

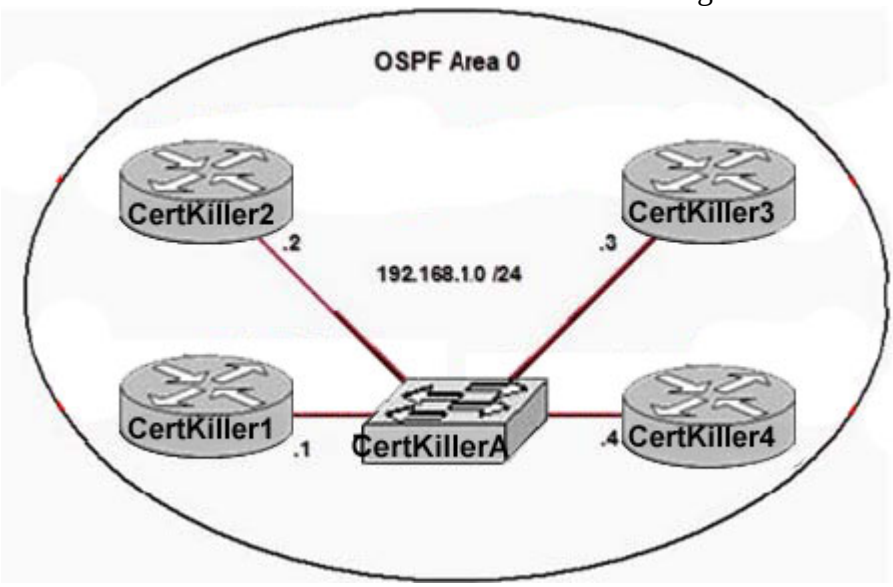

Certkiller 4 configuration exhibit:

```
CertKiller4#show ip ospf neighbor

Neighbor ID     Sri     State          Dial Time    Address        Interface
192.168.1.2       1     output         00:00:37     192.168.1.2    Ethernet0
192.168.1.1       1     omitted        00:00:38     192.168.1.1    Ethernet0
192.168.1.3       1                    00:00:32     192.168.1.3    Ethernet0
```

Study the exhibits shown above carefully. All routers have simultaneously been reloaded and the DR election has concluded as expected. What state is router Certkiller 2 in?

A. 2WAY/BDR
B. FULL/BDR
C. 2WAY/DROTHER

D. 2WAY/DR
E. FULL/DROTHER
F. FULL/DR
G. None of the above

Answer: E

Explanation:
Each OSPF interface (multi-access networks only) possesses a configurable router priority. The Cisco default is 1. If you don't want a router interface to participate in the DR/BDR election, set the router priority to 0 using the ip ospf priority command in interface configuration mode.
Let's go through the steps that occur in any given router when the DR and BDR are elected for a specific multi-access network of which it is a member:
9. The local router creates a list of eligible neighbors. The eligible neighbors are those with which the local router has entered at least a 2Way state. The local router adds itself to this list and to all lists that are formed from this list in subsequent steps, for which it qualifies.
10. From this list, all routers whose participating interfaces have a router priority of 0 are removed. These routers will be among the DROthers on this network.
11. The local router makes note of the current DR and BDR values for later comparison.
12. A list of all routers not claiming to be the DR (their own Hello packets do not list them as the DR) is compiled from the list resulting from step 2.
13. The local router will select the BDR from the list in step 4, based on the following criteria in order:
* If one or more of the routers in the list have declared themselves the BDR, then the one of these with the highest priority is selected to be the BDR.
* If all router priorities are equal, the router with the highest router ID becomes the BDR.
* If no router in the resulting list from step 4 has declared itself the BDR, then the router with the highest router priority is selected to be the BDR.
* If all routers have the same router priority, then the router with the highest router ID is selected to be the BDR.
14. A list of all routers claiming to be the DR (their own Hello packets list them as the DR) is compiled from the list resulting from step 2.
15. The local router will select the DR from the list in step 6, based on the following criteria in order:
* The router with the highest router priority is selected to be the DR.
* If all router priorities are equal, the router with the highest router ID is selected to be the DR.
* If the resulting list from step 6 is empty, meaning that no router has declared itself the DR, then the BDR that was selected in step 5 becomes the DR.
16. If the local router's DR status has been altered by the preceding steps, either causing it to become, or cease to be, the DR, based on a comparison to the results of step 3, then it will repeat steps 4 through 7. This serves to make sure that the local router does not declare itself both the DR and the BDR, because both may be declared by default. If selected to be the DR, the local router will definitely not make the BDR list the next time

around in step 4. Conversely, if dethroned as the DR, the local router can become eligible to be the BDR the next time around, by making the list in step 4.

QUESTION 138:

Three Certkiller routers and connected via a frame relay network as shown below:

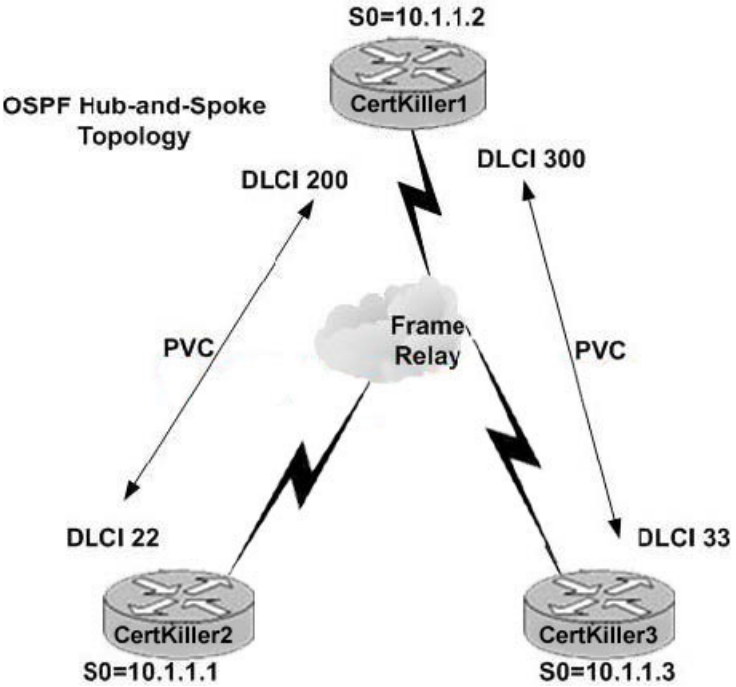

Examine the s0 interface on router Certkiller 1. What is the OSPF network type by default?

A. Nonbroadcast
B. Broadcast
C. Point-to-multipoint
D. Point-to-point
E. Point-to-multipoint nonbroadcast

Answer: A

Explanation:
Based on the information above, the physical serial 0 interface is used, with no subinterfaces.
When OSPF is configured on a physical interface (for example, interface S0) of a nonbroadcast multiaccess technology such as Frame Relay, the default network type of NON_BROADCAST is assigned. When OSPF is configured on point-to-point subinterfaces, the default interface type of POINT_TO_POINT is assigned. When OSPF is configured on multipoint subinterfaces, the default interface type of NON_BROADCAST is assigned.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/technologies_configuration_example09186a0080094054.shtml

QUESTION 139:

Router CK1 has been configured for OSPF on all of the interfaces. A logical loopback interface has not been created in CK1 . By default, what will determine the OSPF router ID of router CK1 ?

A. It defaults to 255.
B. The name set by the hostname command.
C. The highest IP address configured in the router.
D. The priority number of the router set by the priority command.
E. The lowest IP address configured in the router.

Answer: C

Explanation:
The show ip ospf interface command verifies that interfaces have been configured in the intended areas. If no loopback address is specified, the interface with the highest address is the taken router ID. It also gives the timer intervals, including the hello interval, and shows the neighbor adjacencies.
Note: the router ID used is not pre-emptive, meaning that if a loopback address is configured or another interface is configured with a higher IP address, the router ID will not change until it is restarted.
Reference: Building Scalable Cisco Networks (Cisco Press) page 134

QUESTION 140:

While troubleshooting some connectivity issues, you issue the "show ip ospf database" in order to examine the link state database. Which three of the statements below are true regarding the OSPF link state database? (Select three)

A. Each router has an identical link state database.
B. External routes are imported into a separate link state database.
C. Synchronization of link state databases is maintained via flooding of LSAs.
D. Information in the link state database is used to build a routing table by calculating a shortest-path tree.
E. By default, link state databases are refreshed every 10 minutes in the absence of topology changes.

Answer: A, C, D

Explanation:
The Link state database is a collection of link state advertisement for all routers and networks. Each router in the OSPF network maintains an identical database. LSA flooding occurs whenever there is a change in the OSPF topology, ensuring that the databases are synchronized. OSPF also uses the SPF algorithm to build the database

tables.
Incorrect Answers:
B. Only one link state database is maintained, and it is used for all OSPF routes.
E. The default refresh time is 30 minutes.
Reference: Building Scalable Cisco Networks (Cisco Press) page 178.

QUESTION 141:

When designing OSPF networks, it is not recommended to have an excessive number of routers in any single area. Which two of the following statements describe the consequences of having too many routers in an area? (Select two)

A. Excessive LSA traffic.
B. Frequent routing table recalculation.
C. Frequent adjacencies table recalculation.
D. Unacceptable number of reachability errors.

Answer: A, B

Explanation:
In OSPF, if too many routers are in the same area they spend too much time communicating with each other, because anytime a topology change in any single route occurs, LSAs are flooded to all other routers within the area. In OSPF the amount of overhead used increases exponentially with the number of routers in an area. Too many routers can cause an overwhelming amount of LSA and OSPF database calculations.
Incorrect Answers:
C: The number of routers in an area has little or no impact with the number of neighbor adjacencies that must remain established for any given router.
D. An excessive number of routers in any single area will not alone increase the number of errors within the network.

QUESTION 142:

Router CK1 has OSPF configured on its fast Ethernet interface. What is the OSPF cost associated with this 100Mbps Ethernet connection?

A. 1
B. 6
C. 10
D. 100
E. None of the above

Answer: A

Explanation:
The OSPF cost of an interface is inversely proportional to the bandwidth of that interface.

A higher bandwidth indicates a lower cost. There is a higher cost and time delay involved in crossing a 56k serial line than crossing a 10M ethernet line. The formula used to calculate the cost is of any single OSPF network is:
Cost= 100,000,000/bandwith in bps.
Since a 100M Ethernet link is 100,000,000 bps we have:
10000 0000 / 10000 0000 = 1
Reference:http://www.cisco.com/warp/public/104/2.html

QUESTION 143:

The Certkiller OSPF network is displayed below:

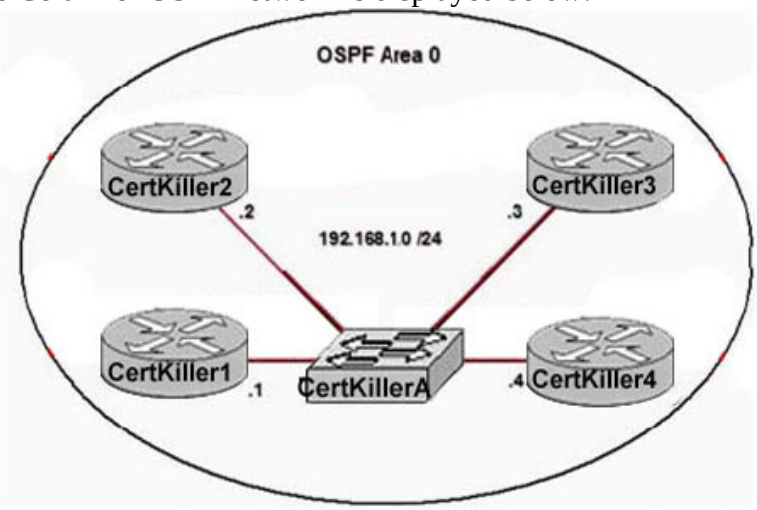

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.2	1	output omitted	00:00:34	192.168.1.2	Ethernet0
192.168.1.4	1		00:00:32	192.168.1.4	Ethernet0
192.168.1.3	1		00:00:39	192.168.1.3	Ethernet0

Refer to the exhibit. All routers have converged and neighbor relationships have been established. Which state is Certkiller 4 in?

A. 2WAY/DROTHER
B. 2WAY/BDR
C. 2WAY/DR
D. FULL/DROTHER
E. FULL/BDR
F. FULL/DR

Answer: F

Explanation:
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. Based on the output above, the dead timer shows that the neighbor relationships have already been established, so the state will be Full for all neighbors, and since the OSPF priorities are all set to 1, the router with the highest router ID will be the DR.

QUESTION 144:

When designing OSPF networks, maintaining stability in an area is important.
Which of the following describes a reason for this?

A. Instability causes more LSAs to be sent, requiring more CPU to recalculate routes.
B. Convergence cannot happen until holddown timers expire, so routing loops can occur.
C. Flooding the area topological database instances consumes excessive bandwidth.
D. Summary link LSAs cannot be sent until all routers in the OSPF area have the same topological database.
E. None of the above

Answer: A

Explanation:
Instability of routes (links) would force sending of LSAs (Link State Advertisements), and CPU time would be required to recalculate the routes.
Incorrect Answers:
B: OSPF does not use hold down timers. Holddown timers are used by distance vector protocols, such as RIP.
C: The topological database is not distributed, only link changes.
D: Summary-link LSAs are not sent to all routers.
Summary-link LSAs originate from area border routers, and flood throughout the LSA's associated area. Each summary-LSA describes a route to a destination outside the area but within the AS.
Reference:http://www.faqs.org/rfcs/rfc2328.html

QUESTION 145:

Router CK1 is an ABR in the Certkiller OSPF network. What does an ABR connect to?

A. Multiple OSPF areas
B. OSPF and RIP networks
C. Multiple designated routers
D. Multiple OSPF autonomous systems
E. Multiple redistributed networks.

Answer: A

Explanation:
An ABR (Area Border Router) shares an interface with at least one other OSPF area. A router that contains network/area statements for two or more different areas is an ABR.

QUESTION 146:

You are a systems administrator of a large multi-area OSPF network, and you've just created a new area for an upcoming remote network. Ordinarily OSPF areas are be connected to the backbone, area 0. However, circumstances dictate you to connect it to the existing area 2 at this time.
Which conditions have to be met in order to make this configuration work? (Select three)

A. There must be a virtual link.
B. Area 2 must be a stub area.
C. Area 2 cannot be a stub area.
D. Area 2 must attach directly to area 0.
E. Network summary link LSAs must be disabled.

Answer: A, C, D

Explanation:
A: There must be a virtual link from the new area to the backbone, area 0. The virtual link provides the disconnected area a logical path to the backbone.
C: The area through which you configure the virtual link, known as a transit area (here area 2), must have full routing information. It cannot be a stub area.
D: The transit area, area 2, must attach directly to area 0.
Reference: OSPF Virtual Link
http://www.cisco.com/warp/public/104/ospfdb7.html
Incorrect Answers:
B: The transit area, area 2, cannot be a stub area.
E: This is not a requirement.

QUESTION 147:

A new location is being added to the Certkiller network and a new OSPF area is being created for it. What kind of OSPF router would you use to connect this new OSPF area to the backbone area?

A. ABR
B. Stub router
C. Internal router
D. Backbone router
E. None of the above

Answer: A

Explanation:
ABRs are used to connect non-backbone areas to the backbone area (area 0).
Incorrect Answers:

B: Stub routers do not apply. An area is stub, not a router and not stub routers are typically a type of ABR, but not every ABR is configured as a stub.
C: Internal routers are only uses within an area.
D: Backbone routers sit on the perimeter of the backbone area. They have at least one interface connected to area 0. However, backbone routers do not necessarily connect to other areas.

QUESTION 148:

The Certkiller OSPF network is displayed in the diagram below:

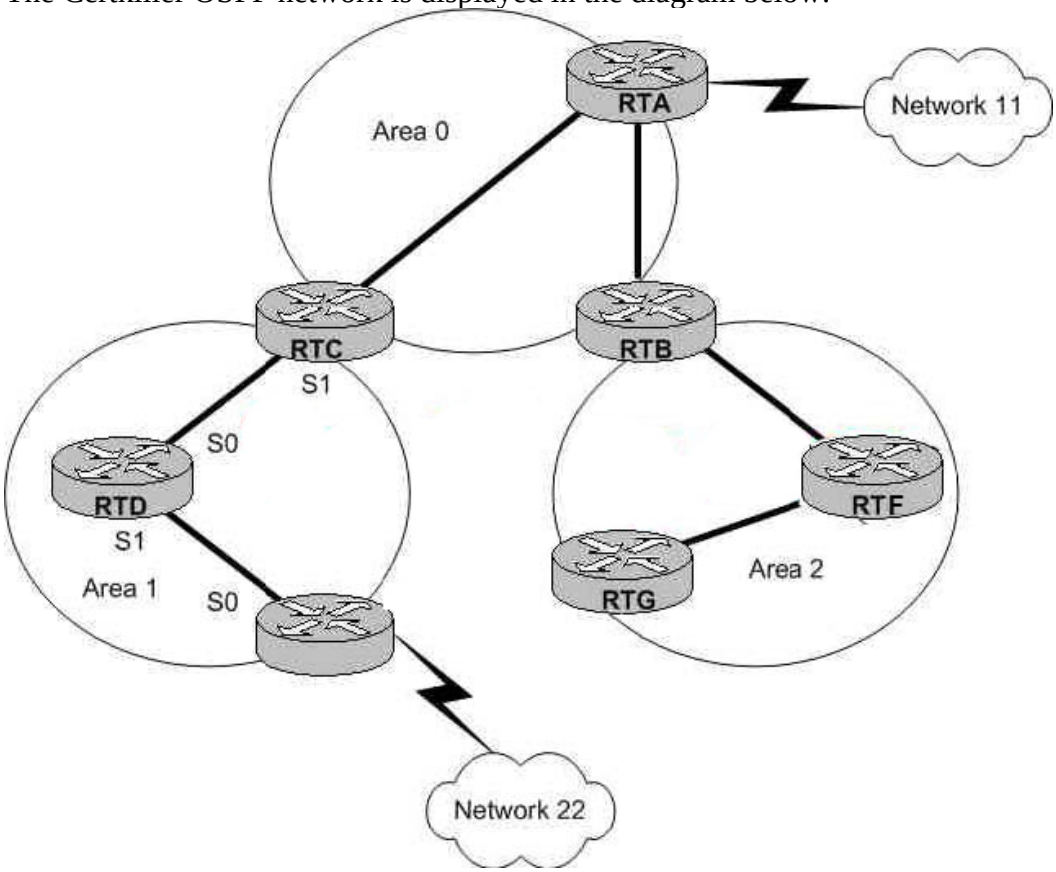

Based on the information above, what kind of router is RTD in OSPF terminology?

A. Designated router
B. Internal router
C. Backbone router
D. Area border router
E. Autonomous system boundary router
F. None of the above.

Answer: B

Explanation:
An area is a set of networks and hosts within an AS that have been administratively

grouped together. We recommend that you configure an area as a collection of contiguous IP subnetted networks. Routers that are wholly contained within an area are called internal routers. All interfaces on internal routers are directly connected to networks within the area.
Incorrect Answers:
A: Designated routers are defined on multi-access networks. In a network consisting of point to point circuits, the DR/BDR election process does not take place.
C: Backbone routers are routers that are contained within area 0.
D: Routers that belong to more than one area are called area border routers.
E: Routers that exchange routing information with routers in other ASs are called AS boundary routers.

QUESTION 149:

Which command should you use to verify what networks are being routed by a given OSPF process?

A. show ip ospf
B. show ip route
C. show ip protocol
D. show ip ospf database
E. None of the above

Answer: C

Explanation:
The information displayed by the show ip protocols command is useful in debugging routing operations. Information in the Routing Information Sources field of the show ip protocols output can help you identify a router suspected of delivering bad routing information. For OSPF routers, this command will display the routed networks.
Incorrect Answers:
A: To display general information about Open Shortest Path First (OSPF) routing processes, use the show ip ospf command in EXEC mode. This command will display the areas assigned and other useful information, but not the networks being routed.
Example:
CK1 # show ip ospf
Routing Process "ospf 201" with ID 192.42.110.200
Supports only single TOS(TOS0) route
It is an area border and autonomous system boundary router
Redistributing External Routes from,
igrp 200 with metric mapped to 2, includes subnets in redistribution
rip with metric mapped to 2
igrp 2 with metric mapped to 100
igrp 32 with metric mapped to 1
Number of areas in this router is 3
Area 192.42.110.0

Number of interfaces in this area is 1
Area has simple password authentication
SPF algorithm executed 6 times
B: This will display the active routing table, but not the networks that are being routed.
D: The OSPF database does not display the networks being routed.

QUESTION 150:

You have a multi-area OSPF network and you're concerned because one of the sites is having connectivity problem to resources in a different area. Which IOS privileged mode command would you enter to confirm that your network: A) has a path to its ABR, B) has a path to its ASBR, and C) the SPF calculation is functional?

A. show ip protocols
B. show running-config
C. show ip ospf neighbor
D. show ip ospf border-routers

Answer: D

Explanation:
The show ip ospf border-routers command displays the internal OSPF routing table entries to an area border router (ABR) and autonomous system boundary router (ASBR). The SPF No in the output is the internal number of SPF calculation that installs this route.
Example:
Router Certkiller # show ip ospf border-routers
OSPF Process 109 internal Routing Table
Destination Next Hop Cost Type Rte Type Area SPF No
160.89.97.53 144.144.1.53 10 ABR INTRA 0.0.0.3 3
160.89.103.51 160.89.96.51 10 ABR INTRA 0.0.0.3 3
160.89.103.52 160.89.96.51 20 ASBR INTER 0.0.0.3 3
160.89.103.52 144.144.1.53 22 ASBR INTER 0.0.0.3 3
Incorrect Answers:
A: The show ip protocols command only displays routing protocol parameters and current timer values.
B: The show running-config command displays the currently used configuration mode. The required information will not be displayed.
C: The show ip ospf neighbor command displays OSPF-neighbor information on a per-interface basis. It does not include ABR, ASBR or SPF information.

QUESTION 151:

You have logged onto the console of a router running OSPF, and for your analysis log you need information on: the filters, default metric, maximum paths, and the number of areas configured. Which of the IOS commands below should you use?

A. show ip protocol
B. show ip route
C. show ip ospf interface
D. show ip ospf
E. Show ip routing

Answer: A

Explanation:
The show ip protocols command, displays parameters about timers, filters, metrics, network, and other information for the entire router.
Reference: Building Scalable Cisco Networks (Cisco Press) page 133

QUESTION 152:

Which IOS command tells you how many times the OSPF shortest path first algorithm was executed?

A. show ip protocol
B. show ip ospf interface
C. show ip ospf
D. show ip ospf database
E. None of the above

Answer: C

Explanation:
The following table describes the output of the "show ip ospf" command and their meanings:

Table52 show ip ospf Field Descriptions

Field	Description
Routing process "ospf 201" with ID 10.0.0.1	Process ID and OSPF router ID.
Supports...	Number of types of service supported (Type 0 only).
SPF schedule delay	Delay time of SPF calculations.
Minimum LSA interval	Minimum interval between link-state advertisements.
LSA group pacing timer	Configured LSA group pacing timer (in seconds).
Interface flood pacing timer	Configured LSA flood pacing timer (in

	milliseconds).
Retransmission pacing timer	Configured LSA retransmission pacing timer (in milliseconds).
Number of...	Number and type of link-state advertisements that have been received.
Number of external LSA	Number of external link-state advertisements.
Number of opaque AS LSA	Number of opaque link-state advertisements.
Number of DCbitless external and opaque AS LSA	Number of demand circuit external and opaque link-state advertisements.
Number of DoNotAge external and opaque AS LSA	Number of do not age external and opaque link-state advertisements.
Number of areas in this router is	Number of areas configured for the router.
External flood list length	External flood list length.

Reference:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/iprrp_r/ip2_s3g.htm#wp1036469

QUESTION 153:

An OSPF link can be in multiple states at any given moment (ie. Exstart, exchange, full). Which two IOS commands let you view the state of the link? (Select two)

A. show ip ospf
B. show ip protocols
C. show ip ospf neighbor
D. show ip ospf interface

Answer: C, D

Explanation:
The link state exstart is an OSPF link state (see note below). We need retrieve OSPF link state information.
C: The output of the show ip ospf neighbor command is used To display OSPF-neighbor information on a per-interface basis. It includes link state information.
D: The show ip ospf interface command is used to display OSPF-related interface information for a particular interface. This includes the link state of the specified interface.
Note: exstart state: After two OSPF neighboring routers establish bi-directional

communication and complete DR/BDR election (on multi-access networks), the routers transition to the exstart state.
Incorrect Answers:
A: The show ip ospf command is used to display general information about OSPF routing processes. However, it does not include any link state information.
B: The command "show ip protocols" displays the parameters and current state of the active routing protocol process. It does not show any link state information.

QUESTION 154:

While performing a routine inspection of your OSPF network you come to the realization that the Designated Router (DR) may be overloaded, and you're considering changing the DR. What are three ways you could manipulate the election of the DR? (Select three)

A. Use of the priority command.
B. Use of the router-id command.
C. Assignment of the loopback address.
D. Assigning a lower IP address
E. By adding additional memory to the desired router.

Answer: A, B, C

Explanation:
A: The OSPF router priority is only used on multi-access networks such as LAN's. This establishes whether the router is eligible to become the Designated Router (DR) for the LAN. A priority of zero means that the router is not eligible to become DR.
B: The router ID is used to break ties during the DR and BDR election processes if the priority values are equal.
C: The loopback address is used in DR election. The highest IP address on the router is the router ID. If a loopback address is configured, then it is the router ID.
Incorrect Answers:
D: The highest IP address is used as the router ID, not the lowest.
E: This would not make it a more likely DR candidate, did the DR election process does not take into consideration the amount of memory.

QUESTION 155:

You are the network engineer at Certkiller . The Certkiller network topology is shown in the following graphic:

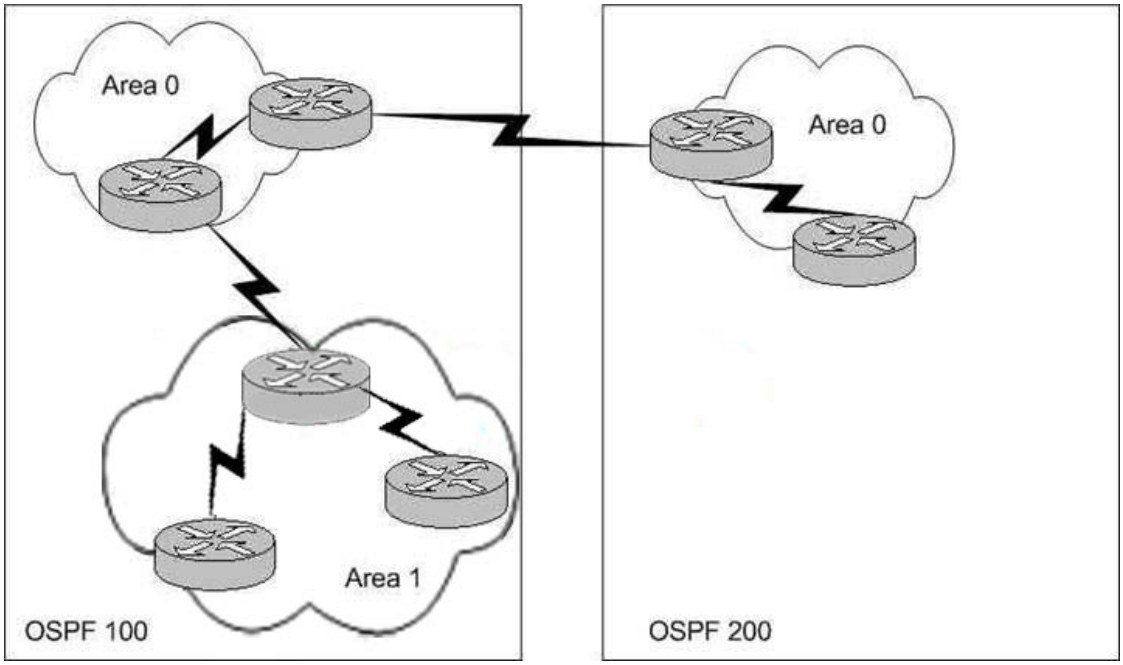

OSFP area 1 has been configured to receive only the intra-area routes and a default route from the area 1 ABR. During routine maintenance, you issue the "show ip route" command on area 1 internal routers and notice that O and O IA routes appear, but no O E2 routes appear.
What is the probable cause of this problem?

A. At the ABR for area 1, the area 1 stub command is missing the no-summary option.
B. All the routers in area 1 are missing the area 1 stub no summary option.
C. At the internal routers in area 1, the area 1 stub command is missing.
D. At the ABR for area 1, the area 1 stub command is missing.
E. The virtual link that transits area 1 is configured incorrectly.
F. Area 1 should be configured as a not-so-stubby-area with the area 1 nssa command.

Answer: A

Explanation:
An autonomous system boundary router (ASBR) advertises external destinations throughout the OSPF autonomous system. In many cases, external link states make up a large percentage of the link states in the databases of every router. A stub area is an area in which you don't allow advertisements of external routes, thus reducing the size of the database even more. Instead, a default summary route (0.0.0.0) is inserted into the stub area in order to reach these external routes. If you have no external routes in your network, then you have no need to define stub areas.

QUESTION 156:

The complex Certkiller OSPF network topology is displayed in the following diagram:

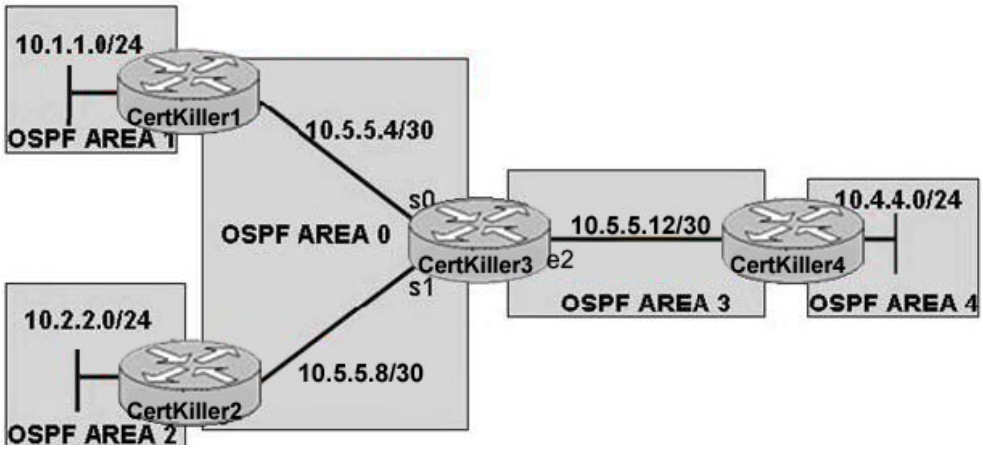

The routing table of Certkiller 3 contains all the proper OSPF routes except the 10.4.4.0/24 OSPF route. Which show command on Certkiller 3 will be most useful to determine the cause of this problem and why?

A. Perform show ip ospf virtual-link to examine if the virtual link between Certkiller 3 and Certkiller 4 is configured and in the up state.
B. Perform show run to examine if area 3 is configured as a stubby area.
C. Perform show run to examine if area 3 is configured as a nssa area.
D. Perform show ip ospf neighbor to examine which router is elected as the DR on the e2 interface.
E. Perform show ip ospf neighbor to examine the OSPF database entries.

Answer: A

Explanation:
All OSPF areas need to be directly connected to the backbone (area 0) in order to ensure proper OSPF operation. If an area is not able to be directly connected with area 0, a virtual link needs to be configured in order to logically connect that area to the backbone network. In this case, a virtual link needs to be configured between routers Certkiller 3 and Certkiller 4. This will logically extend the OSPF area 3 into area 4, connecting it into the backbone. In this example, if the route to 10.4.4.0/24 is missing, a misconfigured or broken virtual link is the most likely cause of the problem.

QUESTION 157:

Which command would display OSPF parameters such as filters, default metric, maximum paths, and number of areas configured on a router?

A. show ip protocol
B. show ip route
C. show ip ospf interface
D. show ip ospf
E. show ip interface
F. None of the above

Answer: A

Explanation:
The "show ip protocol" command displays values about routing timers and network information associated with the entire router.
This includes, the AS number associated with the routing process, number of areas configured on the router, the metric, and the maximum paths.

QUESTION 158:

What can cause an OSPF virtual link not to be established?

A. There are DR/BDR election errors on the OSPF virtual link.
B. The ABRs are using the wrong OSPF router ID in the virtual link configuration.
C. The virtual link transit area is not Area 0.
D. The virtual link is set up over an NBMA network.
E. There is a mismatch of the OSPF network type on the virtual link configuration.
F. None of the above

Answer: B

Explanation:
Every OSPF area must be contiguous to the Backbone Area 0. Non-backbone areas can be thought of "leaves" to the OSPF tree where the backbone is the trunk.
If it is not physically possible to have an area connected to the backbone, you must configure a Virtual Link in the two routers which are providing the tunnel to the backbone. These two routers do not need to be physically connected, but they must share a common area called the "transit area".
Each virtual link section is identified by the Router ID (NOT the interface IP address) of the virtual link neighbor (NOT the Router ID of THIS router). The Router ID of the neighbor is the largest IP interface address associated with that router, and the IP address used for the virtual link must specify the router ID of the peer. You can request the router ID of the virtual neighbor by issuing the command "show ospf rtrid".
Reference:
http://www.cisco.com/en/US/products/hw/routers/ps4081/products_configuration_example09186a0080094260.s

---

QUESTION 159:

While verifying the OSPF operation on your network, you issue the following command:
show ip ospf interface
What information parameters will you get after you execute this command? (Select two)

A. The router ID
B. The summary link counts
C. The neighbor adjacencies
D. The link-state update interval

Answer: A, C

Explanation:
A: The router ID is displayed.
C: Neighbor adjacencies are included in the output of this command.
D: OSPF use the hello interval as update interval. The hello interval is displayed by this command.

QUESTION 160:

Which IOS command would you use to find out which networks are routed by a particular OSPF process?

A. show ospf
B. show ip route
C. show ip protocols
D. show ip ospf database
E. None of the above

Answer: C

Explanation:
The show ip protocols command display current routing protocols. It displays the parameters and current state of the active routing protocol process. The output includes a list of the networks routing for individual ospf processes.
Sample output:
Rt Certkiller # show ip protocols
Routing Protocol is "ospf 200"
Sending updates every 0 seconds
Invalid after 0 seconds, hold down 0, flushed after 0
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Redistributing: ospf 200
Routing for Networks:
172.6.31.5/32
Routing Information Sources:
Gateway Distance Last Update
Distance: (default is 110)
Incorrect Answers:
A: The show ospf command displays summary information regarding the global OSPF configuration.

B: The show ip route command displays the IP routing table.
D: The show ip ospf database command displays the contents of the topological database maintained by the router. The command also shows the router ID and the OSPF process ID. However, the output does not include the networks routing for individual ospf processes.

QUESTION 161:

The routing table of router CK is shown below:
CK#show ip routeCodes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default<Some output deleted> Gateway of last resort is 30.64.0.2 to network 0.0.0.0 30.0.0.0/8 is variably subnetted, 9 subnets, 2 masksO IA 30.2.0.0/16 [110/74] via 30.64.0.2, 00:09:13, Ethernet0C 30.1.3.0/24 is directly connected, Serial0O IA 30.3.0.0/16 [110/148] via 30.64.0.2, 00:05:22, Ethernet0C 30.1.2.0/24 is directly connected, Serial1
Based on the output above, which of the following statements are true regarding the routing table? (Select two)

A. The area is a stub area.
B. The area is totally stubby.
C. Network 30 is using VLSM
D. The routing table is for an ABR

Answer: A, C

Explanation:
Incorrect Answers:
A: This is a stub area.
C: The exhibits states that the 30.0.0.0/8 subnet is variably subnetted and the networks used are /16 and /24, so VLSM is used
Incorrect Answers:
B: The area is stub, not totally stubby.
D: This is not an ABR (Area border router) since there are no external routes, only internal and directly connected.

QUESTION 162:

While troubleshooting an OSPF routing issue you issue the "show ip ospf interface" command. Which of the following parameters are provided from this? (Select two)

A. Router ID
B. Summary link counts
C. Neighbor adjacencies
D. Link-state update interval

Answer: A, C

Explanation:
The show ip ospf interface command verifies that interfaces have been configured in the intended areas. If no loopback address is specified, the interface with the highest address is the taken router ID. It also gives the timer intervals, including the hello interval, and shows the neighbor adjacencies.
Reference: Building Scalable Cisco Networks (Cisco Press) page 134.

QUESTION 163:

Which command should you use to obtain information on OSPF link state advertisements? In particular: which LSA's have been sent, which LSA's have been received, and the time when the last LSA was received?

A. show ip ospf database
B. show ip ospf neighbors
C. show ip ospf protocols
D. show ip ospf interfaces
E. None of the above

Answer: A

Explanation:
The show ip ospf database command is used to display lists of information related to the OSPF topological, the link state, database for a specific router.
Sample output:
R_ Certkiller #show ip ospf database
OSPF Router with ID (192.168.0.12) (Process ID 1)
Router Link States (Area 0)
Link ID ADV Router Age Seq# Checksum Link count
192.168.0.10 192.1680.10 817 0x80000003 0xFF56 1
192.168.0.11 192.1680.11 817 0x80000003 0xFD55 1
192.168.0.12 192.168.0.12 816 0x80000003 0xFB54 1
192.168.0.13 192.168.0.13 816 0x80000003 0xF953 1
192.168.0.14 192.168.0.14 817 0x80000003 0xD990 1
Net Link States (Area 0)
Link ID ADV Router Age Seq# Checksum
192.168.0.14 192.168.0.14 812 0x80000002 0x4AC8
Incorrect Answers:
B: The show ip ospf neighbor is used to display OSPF-neighbor information on a per-interface basis. It does show the required information though.
Sample output:
Neighbor ID Pri State Dead Time Address Interface
192.168.0.13 1 2WAY/DROTHER 00:00:31 192.168.0.13 Ethernet0

192.168.0.14 1 FULL/BDR 00:00:38 192.168.0.14 Ethernet0
C: There is no such command.
D: The show ip ospf interface command is used to display OSPF-related interface information. It displays the circuit name and state, IP address, network mask, broadcast address, redundancy, Internet Control Message Protocol (ICMP) settings, and RIP settings. However, it does not display LSAs.
Sample output:
RouterCK# show ip interfaces
IP Interface Summary:
Circuit Name: VLAN2 State: active
IP Address: 172.16.1.200 Network Mask: 255.255.0.0
Broadcast Address: 172.16.255.255 Redundancy: disabled
ICMP Redirect: enabled ICMP Unreachable: enabled
RIP: enabled

QUESTION 164:

The following output was seen on router Certkiller 1 and Certkiller 3:

```
CertKiller1#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


CertKiller3#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
```

The Certkiller 1 and Certkiller 3 routers are OSPF neighbors over the Ethernet 0/0 connection. Based on the show ip ospf neighbor output from the Certkiller 1 and Certkiller 3 routers, which statement is true?

A. Certkiller 1 is the DR because it has a higher OSPF router priority.
B. Certkiller 1 is the DR because it has a lower OSPF router ID.
C. Certkiller 3 is the DR because it has a higher OSPF router priority.
D. Certkiller 3 is the DR because it has a lower OSPF router ID.
E. Both Certkiller 1 and Certkiller 3 are using the default OSPF router priority.

Answer: A

Explanation:
The output shown above displays information about the neighbors, so from Certkiller 3 we see that the priority of Certkiller 1 is 2, and that Certkiller 1 is the DR. From Certkiller 1 we can see that its neighbor (Certkiller 3) is the BDR and has an OSPF priority of 1.

QUESTION 165:

Which three types of OSPF route entries can be found in the routing table of an internal OSPF router within an OSPF not so stubby area? (Select three)

A. O
B. O*IA
C. O*OA
D. O E1
E. O E2
F. O N1

Answer: A, B, F

Explanation:
The various route types used by OSPF are as follows:
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
An OSPF NSSA will receive inter-area, external type 2, OSPF routes.

QUESTION 166:

Three Certkiller routers are connected together as shown in the diagram below:

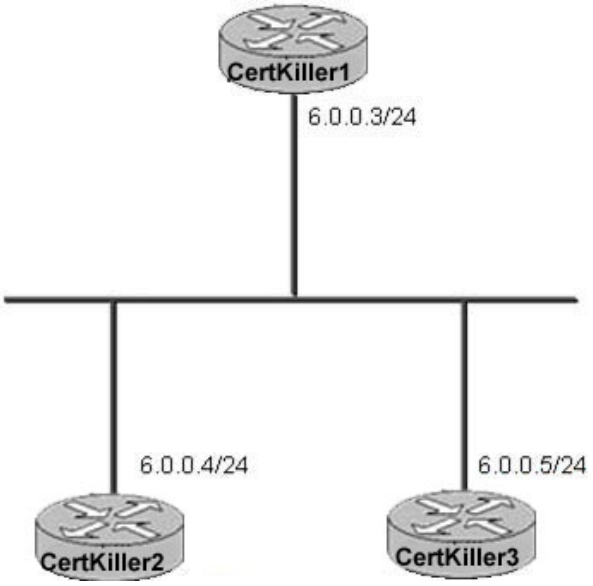

Exhibit

```
CertKiller1#show ip route
...
C    10.1.3.0/24 is directly connected, Serial2
O    10.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1
C    10.1.1.0/24 is directly connected, Serial1
S    192.168.20.0/24 is directly connected, Ethernet0
```

You work as a network technician at Certkiller .com. Study the exhibits carefully. If the following command, "default information-originate", is added to the OSPF configuration on Certkiller 1, what will the neighboring routers see in their route table?

A. S* 0.0.0.0/0 [1/0] via 6.0.0.3, 00:28:00
B. O E1 0.0.0.0 [1/0] via 6.0.0.3, 00:28:00
C. O E2 0.0.0.0 [1/0] via 6.0.0.3, 00:28:00
D. No default routes will be propagated to neighboring routers.

Answer: D

Explanation:
Default routes injected into a normal area can be originated by any OSPF router. The OSPF router does not, by default, generate a default route into the OSPF domain. In order for OSPF to generate a default route, you must use the default-information originate command. However, this will only advertise a default route to the OSPF neighbors if the router itself also contains a default route. Based on the routing table shown in this example, a default route does not already exist.
There are two ways to advertise a default route into a normal 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 167:

You need to choose the best routing protocol for use in the Certkiller network. Which three statements are correct about the differences in IS-IS and OSPF? (Select three)

A. IS-IS is more CPU-intensive than is OSPF.
B. New additions to the protocol are easily implemented in OSPF but not with IS-IS.
C. IS-IS LSP contains TLV fields and OSPF LSU contains the LSAs.
D. For greater fine tuning there are more IS-IS timers.
E. OSPF has more area types than IS-IS.

Answer: C, D, E

Explanation:
The configuration of OSPF is based on a central backbone, Area 0, with all other areas being physically attached to Area 0. Because of this, certain design constraints will inevitably exist. A good, consistent IP addressing structure is necessary when this type of hierarchical model is used. It is used to summarize addresses into the backbone, and reduce the amount of information that is carried in the backbone and advertised across

the network. In comparison, IS-IS also has a hierarchy with Level 1 and Level 2 routers. With IS-IS, the area borders lie on the links. However, significantly fewer link-state PDUs (LSPs) are used. Therefore, many more routers, up to 1000, can reside in a single area. This capability makes IS-IS more scalable than OSPF. IS-IS allows a more flexible approach to extending the backbone. Adding Level 2 routers can extend the backbone. This process is less complex than with OSPF.
With regard to CPU use and the processing of routing updates, IS-IS is more efficient. Not only are there fewer LSPs to process, as compared to OSPF LSAs, but also the mechanism by which IS-IS installs and withdraws prefixes is less intensive.
Both OSPF and IS-IS are link-state protocols and therefore provide fast convergence. The convergence time depends on a number of factors, such as timers, number of nodes, and types of routers. Based on the default timers, IS-IS will detect a failure quicker than OSPF and therefore should converge more rapidly. If there are many neighbors and adjacencies to consider, the convergence time depends on the processing power of the router. IS-IS is typically less CPU intensive than OSPF.
The timers in IS-IS allow more tuning than OSPF. There are more timers to adjust, and therefore finer granularity can be achieved. By tuning the timers, convergence time can be significantly decreased. However, this speed may be at the expense of stability, so a compromise may have to be made. A network engineer should understand the implications of adjusting these timers.

QUESTION 168:

You need to design an addressing scheme in the Certkiller IS-IS network. How is network layer addressing accomplished in the OSI protocol suite?

A. Packet Layer Protocol address
B. Internet Protocol address
C. Network Service Access Point address
D. Authority and Format Identifier address
E. Media Access Control address
F. None of the above

Answer: C

Explanation:
The 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. The NSAP address identifies any system in an OSI network. Various NSAP formats are used for various systems. Different protocols may use different representations of the NSAP.
The NSAP address is the network layer address for CLNS packets. As with DECnet Phase V, one NSAP address is used for each device, not for each interface. LSPs, Hello PDUs, and other routing PDUs are OSI formatted PDUs. Therefore, every IS-IS router requires an OSI address. IS-IS uses the OSI address in the LSPs to identify the router, build the topology table, and build the underlying IS-IS routing tree. NSAP addresses

contain the OSI address of the device and provide a link to upper-layer processes. The NSAP address can be thought of as equivalent to the combination of an IP address and upper-layer protocol identifier in an IP header.

QUESTION 169:

IS-IS uses the concept of NSAP addresses. Which of the following NSAP addresses is a private, locally administered address?

A. 49.0004.30ac.0000.3090.c7df.00
B. 39.0f01.0002.0000.0c00.1111.00
C. 52.0f01.0002.0000.0c00.1111.00
D. 48.0f01.0002.0000.0c00.1111.00
E. None of the above

Answer: A

Explanation:
An NSAP address consists of three parts. They are the area address, the system ID, and the NSAP selector byte. The total length is between eight and 20 bytes. The area address is a variable length field composed of high order octets, excluding the system ID and the selector byte. The system ID is the ES or IS identifier in an area, similar to the OSPF router ID. The system ID has a fixed length of six bytes as engineered in the Cisco IOS. The N-selector byte is a service identifier. The role of N-selector byte is analogous to that of a port or socket in TCP/IP.
The NET is used by routers to identify themselves in the LSPs and forms the basis for OSI route calculation. Addresses starting with value 49, AFI = 49, are considered as private addresses. Private addresses are analogous to those specified by RFC 1918 for IP addresses. These addresses are routed by IS-IS. However, this group of addresses should not be advertised to other CLNS networks.

QUESTION 170:

When comparing OSPF and IS-IS, which three characteristics apply to IS-IS but not to OSPF? (Select three)

A. IS-IS encapsulates PDUs directly into a data-link frame
B. IS-IS uses a default IOS metric of 10 on each interface
C. IS-IS runs PRC (Partial Route Calculations) to calculate IP reachability information
D. IS-IS uses stubby areas to improve network scalability
E. IS-IS uses an on-demand circuit to reduce the hello and LSA flooding across switched WAN links, such as ISDN
F. IS-IS uses a DIS and a backup DIS to present the psuedo-node on the LAN

Answer: A, B, C

Explanation:
The configuration of OSPF is based on a central backbone, Area 0, with all other areas being physically attached to Area 0. Because of this, certain design constraints will inevitably exist. A good, consistent IP addressing structure is necessary when this type of hierarchical model is used. It is used to summarize addresses into the backbone, and reduce the amount of information that is carried in the backbone and advertised across the network. In comparison, IS-IS also has a hierarchy with Level 1 and Level 2 routers. With IS-IS, the area borders lie on the links. However, significantly fewer link-state PDUs (LSPs) are used. Therefore, many more routers, up to 1000, can reside in a single area. This capability makes IS-IS more scalable than OSPF. IS-IS allows a more flexible approach to extending the backbone. Adding Level 2 routers can extend the backbone. This process is less complex than with OSPF.
With regard to CPU use and the processing of routing updates, IS-IS is more efficient. Not only are there fewer LSPs to process, as compared to OSPF LSAs, but also the mechanism by which IS-IS installs and withdraws prefixes is less intensive.
Both OSPF and IS-IS are link-state protocols and therefore provide fast convergence. The convergence time depends on a number of factors, such as timers, number of nodes, and types of routers. Based on the default timers, IS-IS will detect a failure quicker than OSPF and therefore should converge more rapidly. If there are many neighbors and adjacencies to consider, the convergence time depends on the processing power of the router. IS-IS is typically less CPU intensive than OSPF

QUESTION 171:

You need to determine the best routing protocol to be used within the Certkiller network. What are two routing protocols defined by the OSI protocol suite that operate at the network layer? (Select two)

A. Routing Information Protocol
B. Intermediate System-to-Intermediate System
C. End System-to-End System
D. End System-to-Intermediate System
E. Interior Gateway Routing Protocol

Answer: B, D

Explanation:
1. Intermediate System to Intermediate System:
The OSI routing protocol that learns the location of the networks within the autonomous system so that data can be forwarded to the remote hosts. With IS-IS routing, OSI distinguishes between Level 1, Level 2, and Level 3 routing to simplify router design and operation. Level 1 ISs communicate with other Level 1 ISs in the same area. Level 2 ISs route between Level 1 areas and form an intradomain routing backbone. Level 3 routing is done between separate domains. Hierarchical routing simplifies backbone design, because Level 1 ISs only need to know how to get to the nearest Level 2 IS.
2. End System-to-Intermediate System

The protocol by which the OSI ES and the IS communicate to dynamically learn Layer 2 adjacencies. ES-IS is an analogous to Address Resolution Protocol (ARP) in IP. Although not technically a routing protocol, ES-IS is included here because it is commonly used with routing protocols to provide end-to-end data movement through an internetwork. Routing between end systems and intermediate systems is sometimes referred to as Level 0 routing.

QUESTION 172:

IS-IS is often considered an alternative to OSPF in many networks. Which two statements identify similarities between IS-IS and OSPF? (Select two)

A. Support for address summarization between areas
B. Support for multiple areas per router
C. Support for classless routing
D. Support for designated intermediate systems and backup designated intermediate systems
E. Support for both DIS adjacencies and neighbor adjacencies

Answer: A, C

Explanation:
Similarities between OSPF and IS-IS
* Both of them are link-state routing protocols.
* To maintain their link-state databases, both routing protocols use the Dijkstra SPF algorithm.
* Both, through the use of areas, support a hierarchical network topology.
* They both use Hello packets to form adjacencies with their neighbors.
* For broadcast multi-access networks, they both elect a designated router (DR).
* They both support VLSM and the summarization of areas.
* Both allow the use of authentication to ensure a more secure network.
1. Both allow multiple instances per device, OSPF with the process ID and IS-IS with a tag after the router isis command.

QUESTION 173:

Certkiller is using IS-IS in their network domain. In IS-IS routing, which level is used to route between different areas within the same domain?

A. Level 2
B. Level 0
C. Level 1
D. Level 3
E. None of the above

Answer: A

Explanation:
An IS-IS network is termed a domain, analogous to an autonomous system (AS) in TCP/IP. The following shows IS-IS within the domain as a two-level hierarchy:
1. Level 1 (L1) ISs, closely equivalent to OSPF internal non-backbone routers, are responsible for routing to ESs inside an area. L1 ISs enable communication between ESs in an area.
2. Level 2 (L2) ISs, closely equivalent to backbone routers in OSPF, route between areas only.
3. Level 1 and Level 2 (L1L2) Intermediate ISs, closely equivalent to area border routers (ABRs) in OSPF, route between areas and the backbone. They participate in the L1 intra-area routing and the L2 inter-area routing.
L2 routers are also referred to as area routers because they interconnect the L1 areas. L2 routers store a separate LSPD, which contains only the inter-area topology information.

QUESTION 174:

Router CK1 is configured as a level 1 router in the Certkiller IS-IS network. IS-IS Level 1 routers use LSPs for which purpose?

A. To build a topology database for areas outside the AS
B. To establish adjacency with ES's
C. To build a topology database for the local and remote areas
D. To build a topology database for remote areas
E. To establish adjacency with L1 IS's
F. To build a topology database for the local area only
G. None of the above

Answer: F

Explanation:
An IS-IS network is termed a domain, analogous to an autonomous system (AS) in TCP/IP. The following shows IS-IS within the domain as a two-level hierarchy:
1. Level 1 (L1) ISs, closely equivalent to OSPF internal non-backbone routers, are responsible for routing to ESs inside an area. L1 ISs enable communication between ESs in an area.
L1 routers are also referred to as station routers because they enable stations to communicate with each other and the rest of the network. A contiguous group of L1 routers defines an area. The L1 routers maintain the L1 link-state PDU database (LSPD), which defines the picture of the area itself and the exit points to neighboring areas.

QUESTION 175:

Certkiller uses the IS-IS routing protocol within their network. In IS-IS routing, when would the use of Level 1 IS routing be required?

A. When routing between end systems and intermediate systems
B. When routing between Level 1 areas in different domains
C. When routing between intermediate systems in different areas
D. When routing between Level 1 areas in the same domain
E. When routing between intermediate systems in the same area
F. None of the above

Answer: E

Explanation:
An IS-IS network is termed a domain, analogous to an autonomous system (AS) in TCP/IP. The following shows IS-IS within the domain as a two-level hierarchy:
1. Level 1 (L1) ISs, closely equivalent to OSPF internal non-backbone routers, are responsible for routing to ESs inside an area. L1 ISs enable communication between ESs in an area.
L1 routers are also referred to as station routers because they enable stations to communicate with each other and the rest of the network. A contiguous group of L1 routers defines an area. The L1 routers maintain the L1 link-state PDU database (LSPD), which defines the picture of the area itself and the exit points to neighboring areas.

QUESTION 176:

The Certkiller IS-IS network is displayed below:

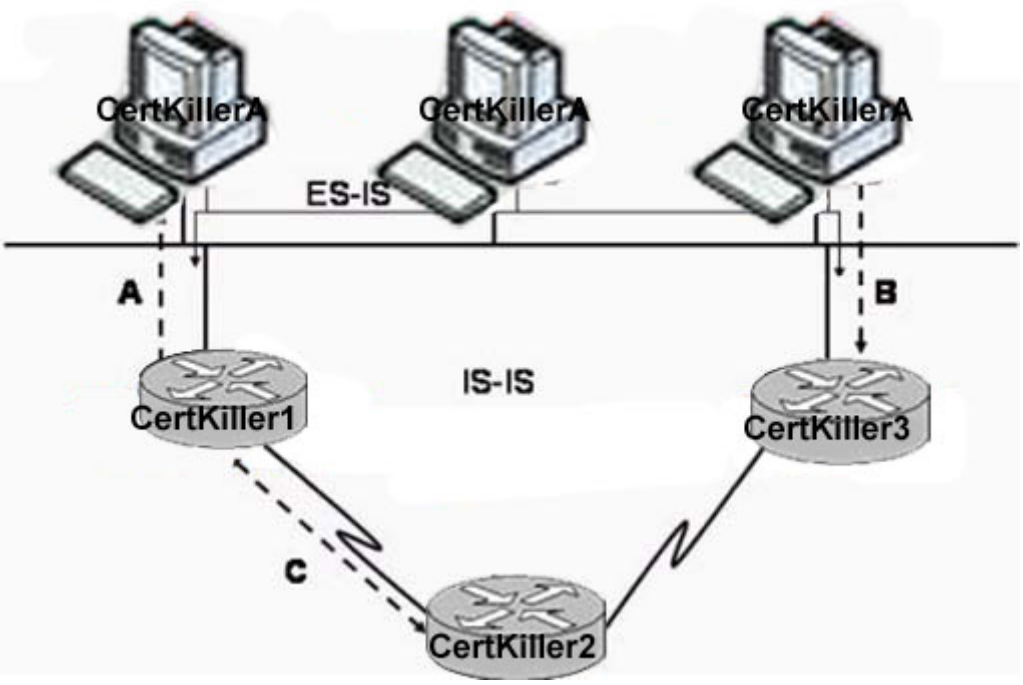

Given the above diagram, which statement correctly matches the letter to the type of IS-IS hello PDU?

A. A=ESH; B=ISH; C=IIH
B. A=ISH; B=IIH; C=ESH

C. A=IIH; B=ISH; C=ESH
D. A=IIH; B=ESH; C=ISH
E. A=ISH; B=ESH; C=IIH
F. None of the above

Answer: E

Explanation:
There are four categories of IS-IS PDUs:
1. Hello PDU (ESH, ISH, IS-IS Hello [IIH]) - Used to establish and maintain adjacencies. ESHs are sent from ESs to ISs. ISHs are sent from ISs to ESs. IIHs are sent between ISs. Note that ESH and ISH PDUs are ES-IS PDUs, not IS-IS PDUs.
2. LSP - Used by IS-IS to distribute link-state information. There are independent pseudonode and non-pseudonode LSPs for both Level 1 and Level 2.
3. Complete Sequence Number PDU (CSNP) - Used to distribute a complete link-state database on the router. CSNPs are used to inform other routers of LSPs that may be outdated or missing from their own database. This ensures that all routers have the same information and are synchronized. The packets are similar to an OSPF database description packet.
4. Partial Sequence Number PDU (PSNP) - Used to acknowledge and request link-state information.

QUESTION 177:

Router CK1 is a Designated Intermediate System on a Certkiller network. What is periodically sent by a DIS on a LAN to ensure that all adjacent neighbors' IS-IS link-state databases are synchronized?

A. Hello
B. Partial SNP (PSNP)
C. Complete SNP (CSNP)
D. Link-state summary
E. Database description packet (DDP)
F. Database query
G. None of the above

Answer: C

Explanation:

Explanation:
There are four categories of IS-IS PDUs:
1. Hello PDU (ESH, ISH, IS-IS Hello [IIH]) - Used to establish and maintain adjacencies. ESHs are sent from ESs to ISs. ISHs are sent from ISs to ESs. IIHs are sent between ISs. Note that ESH and ISH PDUs are ES-IS PDUs, not IS-IS PDUs.
2. LSP - Used by IS-IS to distribute link-state information. There are independent

pseudonode and non-pseudonode LSPs for both Level 1 and Level 2.
3. Complete Sequence Number PDU (CSNP) - Used to distribute a complete link-state database on the router. CSNPs are used to inform other routers of LSPs that may be outdated or missing from their own database. This ensures that all routers have the same information and are synchronized. The packets are similar to an OSPF database description packet.
4. Partial Sequence Number PDU (PSNP) - Used to acknowledge and request link-state information.

QUESTION 178:

You need to determine the best routing protocol for your network. Which two statements are true about IS-IS routing, when compared to OSPF? (Select two)

A. IS-IS and OSPF are both Open Standard, link-state routing protocols which support VLSM.
B. OSPF is more scalable than IS-IS because of its ability to identify normal, stub, and NSSA areas.
C. Based on the default timers, OSPF detects a failure faster than IS-IS does.
D. IS-IS is more efficient than OSPF in the use of CPU resources.
E. OSPF default timers permit more tuning than IS-IS does.

Answer: A, D

Explanation:
IS-IS is the dynamic link-state routing protocol for the OSI protocol stack. As such, IS-IS distributes routing information for routing CLNP data for the ISO CLNS environment. When IS-IS is used strictly for the ISO CLNS environment, it is referred to as ISO IS-IS.
Differences between IS-IS and OSPF
Although IS-IS and OSPF share many common features, they do have quite a few differences:
* Whereas OSPF routers can be part of multiple areas, an IS-IS router belongs to only one area per routing process.
*
* In OSPF, the boundaries of areas are set in the router. The boundaries of areas are on the network connections between routers for IS-IS, reiterating that each router is in only one area per routing process.
* IS-IS utilizes CLNS protocol data units (PDUs) to send information between routers instead of using IP packets, like OSPF does.
* IS-IS allows for the preempting of DRs, where OSPF does not.
* OSPF DROthers do not form adjacencies with other DROthers on broadcast multi-access networks, while in the same environment, all IS-IS intermediate systems form adjacencies with one another.
* The backbone of an IS-IS network is designated by the type of routers in it instead of being designated by an area number (0, in the case of OSPF).

QUESTION 179:

Router CK1 is an IS-IS device that needs to route within the same area. Which two provide intra-area routing services? (Select two)

A. L2 ES
B. L2 IS
C. L1 IS
D. L1 ES
E. L1/L2 IS

Answer: C, E

Explanation:
An IS-IS network is termed a domain, analogous to an autonomous system (AS) in TCP/IP. The following shows IS-IS within the domain as a two-level hierarchy:
1. Level 1 (L1) ISs, closely equivalent to OSPF internal non-backbone routers, are responsible for routing to ESs inside an area. L1 ISs enable communication between ESs in an area. L1 routers are also referred to as station routers because they enable stations to communicate with each other and the rest of the network. A contiguous group of L1 routers defines an area. The L1 routers maintain the L1 link-state PDU database (LSPD), which defines the picture of the area itself and the exit points to neighboring areas.
2. Level 2 (L2) ISs, closely equivalent to backbone routers in OSPF, route between areas only.
3. Level 1 and Level 2 (L1L2) Intermediate ISs, closely equivalent to area border routers (ABRs) in OSPF, route between areas and the backbone. They participate in the L1 intra-area routing and the L2 inter-area routing.
L1L2 routers support an L1 function to communicate with the other L1 routers in their area and maintain the L1 LSP information in an L1 LSPD. They inform other L1 routers that they are an exit point for the area. They also support an L2 function to communicate with the rest of the backbone and maintain an L2 topology database separately from their L1 LSPD. IS-IS does not share the concept of an Area 0 with OSPF. An IS-IS domain appears as a set of distinct areas interconnected by a chain of L2 routers, weaving their way through and between the Level 1 areas.

QUESTION 180:

Router CK1 is an IS-IS device that needs to route traffic between different areas. Which two provide inter-area routing services? (Select two)

A. L1/L2 IS
B. L2 IS
C. L1 IS
D. L2 ES
E. L1 ES
F. None of the above

Answer: A, B

Explanation:
* Level 1 and Level 2 (L1L2) Intermediate ISs, closely equivalent to area border routers (ABRs) in OSPF, route between areas and the backbone. They participate in the L1 intra-area routing and the L2 inter-area routing. L1L2 routers support an L1 function to communicate with the other L1 routers in their area and maintain the L1 LSP information in an L1 LSPD. They inform other L1 routers that they are an exit point for the area. They also support an L2 function to communicate with the rest of the backbone and maintain an L2 topology database separately from their L1 LSPD. IS-IS does not share the concept of an Area 0 with OSPF. An IS-IS domain appears as a set of distinct areas interconnected by a chain of L2 routers, weaving their way through and between the Level 1 areas.
* Level 2 (L2) ISs, closely equivalent to backbone routers in OSPF, route between areas only. L2 routers are also referred to as area routers because they interconnect the L1 areas. L2 routers store a separate LSPD, which contains only the inter-area topology information.

QUESTION 181:

A network administrator has decided to use IS-IS instead of OSPF in the Certkiller network. Which three are benefits of IS-IS over OSPF? (Select three)

A. IS-IS requires fewer neighbor relationships in a broadcast multiaccess network
B. IS-IS supports more routers in an area
C. IS-IS supports route tags
D. IS-IS does not require Hello packets to establish neighbor relationships
E. IS-IS produces fewer link state advertisements for a given network
F. IS-IS supports network layer protocols other than IP

Answer: B, E, F

Explanation:
IS-IS is the dynamic link-state routing protocol for the OSI protocol stack. As such, IS-IS distributes routing information for routing CLNP data for the ISO CLNS environment. When IS-IS is used strictly for the ISO CLNS environment, it is referred to as ISO IS-IS.
Differences between IS-IS and OSPF.
Although IS-IS and OSPF share many common features, they do have quite a few differences:
* Whereas OSPF routers can be part of multiple areas, an IS-IS router belongs to only one area per routing process.
* In OSPF, the boundaries of areas are set in the router. The boundaries of areas are on the network connections between routers for IS-IS, reiterating that each router is in only one area per routing process.
* IS-IS utilizes CLNS protocol data units (PDUs) to send information between routers

instead of using IP packets, like OSPF does.
* IS-IS allows for the preempting of DRs, where OSPF does not.
* OSPF DROthers do not form adjacencies with other DROthers on broadcast multi-access networks, while in the same environment, all IS-IS intermediate systems form adjacencies with one another.
* The backbone of an IS-IS network is designated by the type of routers in it instead of being designated by an area number (0, in the case of OSPF).

QUESTION 182:

One of the Certkiller IS-IS routers is using the NSAP 39.0100.0102.0001.0c00.1211.00. Given this NSAP, which portion is interpreted by IS-IS as the area?

A. 0001
B. 39
C. 0001.0c00
D. 0001.0c00.1211
E. 39.0100.0102
F. 39.0100
G. None of the above

Answer: E

Explanation:
A NET is a NSAP where the last byte is always zero. On a router running IS-IS, a NET can be 8 to 20 B in length. The last byte is always the n-selector and must be zero. The n-selector indicates which transport entity the packet is sent to. An n-selector of zero indicates no transport entity, and means that the packet is for the routing software of the system. The six bytes directly preceding the n-selector are the system ID. The system ID length is a fixed size and cannot be changed. The system ID must be unique throughout each area (Level 1) and throughout the backbone (Level 2). All bytes preceding the system ID are the area ID. The area ID field is the portion of the NSAP preceding the system ID.
The following example shows a router with area ID 47.0004.004d.0001 and system ID 0001.0c11.1110:
47.0004.004d.0001.0001.0c11.1110.00
In our example, the six bytes preceding the N selector is 0001.0c00.1211 which is the system ID. The portion preceding this, 39.0100.0102, is the area ID.

QUESTION 183:

Routers CK1 and CK2 have established an IS-IS neighbor relationship. What does IS-IS use to establish and maintain neighbor relationships between Intermediate Systems?

A. IIH
B. LSP
C. CLNP
D. CLNS
E. ISH
F. None of the above

Answer: A

Explanation:
IS-IS utilizes Hello PDUs to discover its neighbors and to form adjacencies with them. After the formation of adjacencies, Hello PDUs are sent out every 10 seconds by default to maintain the adjacencies. The Hello PDU contains information about the router, the router's capabilities, and certain parameters about the interface sending the Hello PDU. Once the two routers agree on their respective capabilities and the parameters set forth, the routers will form an adjacency.
Hello PDU (ESH, ISH, IS-IS Hello [IIH])
- Used to establish and maintain adjacencies. ESHs are sent from ESs to ISs. ISHs are sent from ISs to ESs. IIHs are sent between ISs. Note that ESH and ISH PDUs are ES-IS PDUs, not IS-IS PDUs.

QUESTION 184:

You need to decide on the best routing protocol for the Certkiller network and have narrowed your choices to IS-IS and OSPF. Which two are characteristics of the IS-IS protocol but not OSPF? (Select two)

A. IS-IS provides for network scalability by allowing the network to be separated into areas
B. IS-IS utilizes SPF algorithm
C. IS-IS supports demand circuit routing
D. IS-IS forms adjacencies with all neighbors
E. IS-IS supports three layers of hierarchical routing
F. IS-IS provides routing support for multiple network layer protocols

Answer: D, F

Explanation:
IS-IS is the dynamic link-state routing protocol for the OSI protocol stack. As such, IS-IS distributes routing information for routing CLNP data for the ISO CLNS environment. When IS-IS is used strictly for the ISO CLNS environment, it is referred to as ISO IS-IS.
Differences between IS-IS and OSPF.
Although IS-IS and OSPF share many common features, they do have quite a few differences:
* Whereas OSPF routers can be part of multiple areas, an IS-IS router belongs to only one area per routing process.

* In OSPF, the boundaries of areas are set in the router. The boundaries of areas are on the network connections between routers for IS-IS, reiterating that each router is in only one area per routing process.
* IS-IS utilizes CLNS protocol data units (PDUs) to send information between routers instead of using IP packets, like OSPF does.
* IS-IS allows for the preempting of DRs, where OSPF does not.
* OSPF DROthers do not form adjacencies with other DROthers on broadcast multi-access networks, while in the same environment, all IS-IS intermediate systems form adjacencies with one another.
* The backbone of an IS-IS network is designated by the type of routers in it instead of being designated by an area number (0, in the case of OSPF).

QUESTION 185:

Router CK1 is a Level-2 router in the Certkiller IS-IS network. A Level-2 IS can establish an adjacency with which routers? (Select three)

A. Any Level-2 IS in any area
B. Any Level-1/Level-2 IS in the same area
C. Any Level-1 IS in any area
D. Any Level-1 IS in the same area
E. Any Level-1/Level-2 IS in any area

Answer: A, B, E

Explanation:
Level 2 (L2) ISs, closely equivalent to backbone routers in OSPF, route between areas only. L2 routers are also referred to as area routers because they interconnect the L1 areas. L2 routers store a separate LSPD, which contains only the inter-area topology information.

QUESTION 186:

The Certkiller IS-IS network utilized locally administered addresses. Which IS-IS NET represents a locally administered private address?

A. 39.0040.0010.0c99.1112.00
B. 47.0010.0000.0a11.3564.00
C. 45.0004.0000.0d35.4554.00
D. 49.0001.0000.0c12.3456.00
E. None of the above

Answer: D

Explanation:
An NSAP address consists of three parts. They are the area address, the system ID, and

the NSAP selector byte. The total length is between eight and 20 bytes. The area address is a variable length field composed of high order octets, excluding the system ID and the selector byte. The system ID is the ES or IS identifier in an area, similar to the OSPF router ID. The system ID has a fixed length of six bytes as engineered in the Cisco IOS. The N-selector byte is a service identifier. The role of N-selector byte is analogous to that of a port or socket in TCP/IP.
The NET is used by routers to identify themselves in the LSPs and forms the basis for OSI route calculation. Addresses starting with value 49, AFI = 49, are considered as private addresses. Private addresses are analogous to those specified by RFC 1918 for IP addresses. These addresses are routed by IS-IS. However, this group of addresses should not be advertised to other CLNS networks.

QUESTION 187:

Router CK1 is the DIS in the Certkiller IS-IS network. In an IS-IS environment, what happens when the designated IS router crashes?

A. A new election process occurs immediately, establishing a new DIS until a router with a higher priority or MAC address establishes an adjacency.
B. The elected backup designated router takes the place of the DIS temporarily without the necessity of a new election, until the original DIS comes back online.
C. The elected backup designated router takes the place of the DIS indefinitely without the necessity of a new election.
D. A new election process occurs immediately, establishing a new DIS that will remain in place indefinitely.
E. None of the above

Answer: A

Explanation:
The idea behind the Designated Intermediate System (DIS) is similar to the one behind the Designated Router (DR) in OSPF. The DIS creates and acts on behalf of a pseudonode, a virtual node. All the routers on the LAN, including the DIS, form an adjacency with the pseudonode, or PSN. Instead of flooding and database synchronization taking place over n x (n-1) adjacencies, the pseudonode enables the reduction in flooding and database synchronization to occur only over the adjacencies formed with the pseudonode. On a LAN, one of the routers will be elected the DIS based on interface priority. The default priority is 64. The configurable range is 0 to 127. If all interface priorities are the same, the router with the highest SNPA is selected. MAC addresses are the SNPAs on LANs. On Frame Relay networks, the local DLCI is the SNP
A. If the SNPA is a DLCI and is the same at both sides of a link, the router with the higher system ID in the NSAP address will become the DIS.
The DIS election is preemptive, unlike DR election with OSPF. If a new router boots on the LAN with a higher interface priority, it becomes the DIS, purges the old pseudonode LSP, and a new set of LSPs will be flooded. The DIS Hello interval, at 3.3 seconds, is

three times faster than the interval for other routers on the LAN. This allows for quick detection of DIS failure and immediate replacement on the LAN. Remember that there is no concept of backup DIS in IS-IS.

QUESTION 188:

Certkiller has implemented IS-IS as their internal routing protocol. Which three statements below regarding IS-IS are true? (Select three)

A. L1 routers have no knowledge about routes outside their area.
B. To route packets to another area, L1 routers must forward the packets to the L2 router of the destination area.
C. To route packets to another area, L1 routers must forward the packets to an L1/L2 router within their area.
D. L2 routers form adjacencies with L1 and L1/L2 neighbors.
E. L1/L2 routers maintain a separate Level 1 link-state database and a Level 2 link-state database; they do not advertise L2 routes to L1 routers.

Answer: A, C, E

Explanation:
L1 routers A Level 1 (L1) router is a router in a non-backbone area. L1 routers know only about intra-area routes. All they know about inter-area routes is a default route to the L1/L2 router for the area. All routers within a Level 1 area contain the same link-state database. These routers receive link-state PDUs (LSPs) only from within the area. They will not receive LSPs from other areas. L1 routers will not receive any information from an L2 router. The L1 router would be the equivalent of an internal router for OSPF.
L2 routers Level 2 (L2) routers are the backbone routers. They handle all of the inter-area traffic. An L2 router can belong only to the backbone area. L2 routers will send LSPs to all other L2 routers and to all L1/L2 routers, regardless of the area the L1/L2 router belongs to. The L2 router can be compared to a backbone router for OSPF.
L1/L2 routersLevel 1/Level 2 (L1/L2) routers are similar in function to an OSPF ABR. L1/L2 routers will send LSPs to both L1 and L2 routers. The LSPs that the L1/L2 router sends to L1 routers help it to maintain its Level 1 link-state database. The LSPs that the L1/L2 router sends to L2 routers help it to maintain its Level 2 link-state database. The L1/L2 router contains two link-state databases, and information stored in the Level 2 link-state database is not shared with any L1 routers.

QUESTION 189:

There are several varieties of IS-IS routers. Two types of IS-IS provide intra-area routing services, meaning that they are used to route traffic within the same area. Which of the following two match this description? (Select two)

A. L1 IS
B. L1 ES

C. L2 IS
D. L2 ES
E. L1/L2 IS

Answer: A, E

Explanation:
L1 IS and L1/L2 IS routers provide intra-area routing services. A Level 1 router knows the topology only of its own area and has Level 1 or Level1/Level2 neighbors in this area. It has a Level 1 link-state database with all the information for intra-area routing. It uses the closest Level 2-capable router in its own area to send packets out of the area, a scenario that may result in suboptimal routing.
Incorrect Answers:
B, D: End systems are OSI terms used for end hosts, such as PCs or servers. They do not provide for any type of routing functionality.
C: L2 routers make up the backbone of the IS-IS network, and they are used to provide for inter-area routing.

QUESTION 190:

In OSPF, the backbone area is area 0. How could you describe the backbone area of an IS-IS domain?

A. A contiguous collection of L1 routers.
B. A contiguous collection of L1 and L1/L2 routers.
C. A contiguous collection of L1 and L2 routers.
D. A contiguous collection of L2 and L1/L2 routers.
E. None of the above.

Answer: D

Explanation:
IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different area. L2 routers are used for inter-area routing, and a L1 router can only route traffic to other areas via an L2 or and L1/L2 IS-IS router.

QUESTION 191:

Which one of the following statements correctly describes a characteristic of an IS-IS backbone?

A. A router in all areas.
B. All routers in one area.
C. A chain of L2 and L1/L2 IS-IS routers.
D. An unbroken string of L1 and LS IS-IS routers.

E. A central area to which all other areas are attached.
F. None of the above

Answer: C

Explanation:
IS-IS does not have a backbone area like the OSPF area 0. The IS-IS backbone is a contiguous collection of Level 2-capable routers, each of which can be in a different areaL1 routers are used for internal routing within the area and so only L2 and L1/L2 routers form the backbone in IS-IS.

QUESTION 192:

The routing protocol IS-IS uses areas. What must an administrator connect each IS-IS area to?

A. Area 0
B. Area 1
C. Level-1 backbone
D. Level-2 backbone
E. External IS-IS areas

Answer: D

Explanation:
Small IS-IS networks are built as a single area that includes all the routers in the network. As the network grows larger, it is usually reorganized into a backbone area made up of the connected set of all Level 2 routers from all areas, which are in turn connected to local areas. Routers that are solely used as L1 routers must route all traffic that is destined outside of the area to a L1/L2 or a L2 router. The contiguous set of L2 routers make up the OSPF backbone, similar to area 0 in OSPF networks.
Incorrect Answers:
A, B: Area 0 or Area 1 has no special significance in IS-IS. In IS-IS, areas are broken down into levels and are not numbered like OSPF areas.
C: Level 2, not Level 1.
E: This is not a requirement.

QUESTION 193:

Which packet type is used on the Certkiller IS-IS network to acknowledge LSPs on point-to-point links and to request missing pieces of information in the IS-IS link-state database?

A. Complete SNP (CSNP)
B. Partial SNP (PSNP)
C. Hello

D. Database query
E. Database description packet (DDP)
F. None of the above

Answer: B

Explanation:
The different IS-IS packet types are described below:
CSNP-Complete Sequence Number PDUs are used to tell other routers about all the LSPs in one router database; similar to an OSPF database descriptor packet
PSNP-Partial Sequence Number PDUs are used to request an LSP and acknowledge receipt of an LSP
PDUs (PSNP's) every 3 seconds and complete sequence number PDUs (CSNPs) every 10 seconds that describe the full database of the DIS. Another router can ask the DIS for a missing LSP using a partial sequence number packet (PSNP), or in turn give the DIS a new LSP.
Reference: CCNP Self-Study CCNP BSCI Exam Certification Guide p.361

QUESTION 194:

In IS-IS routing, which level is used to route between different areas within the same domain?

A. Level 0
B. Level 1
C. Level 2
D. Level 3

Answer: C

Explanation:
Small IS-IS networks are built as a single area that includes all the routers in the network. As the network grows larger, it is usually reorganized into a backbone area made up of the connected set of all Level2 routers from all areas, which is in turn connected to local areas. Within a local area, routers know how to reach all system IDs. Between areas, routers know how to reach the backbone, and the backbone routers know how to reach other areas.
Routers establish Level1 adjacencies to perform routing within a local area (intra-area routing). Routers establish Level2 adjacencies to perform routing between different areas (inter-area routing).

QUESTION 195:

In your network you have two Level-1/ Level-2 IS-IS routers named CK1 and CK2 . The two routers are connected together via a WAN link, and they both belong to the same area. What kind of an adjacency can you expect them to establish?

A. Level-1 only
B. Level-2 only
C. Both Level-1 and Level-2
D. None unless statically configured
E. WAN links do not support IS-IS adjacencies

Answer: C

Explanation:
IS-IS routers that are configured as L1/L2, which is the default, will run two separate SPF algorithms, one for each level. Neighbors on point-to-point networks always become adjacent unless they do not see themselves in their neighbors' hello PDU and match on certain parameters. On broadcast networks and nonbroadcast multiaccess (NBMA) networks, the DIS (Designated Intermediate System) will become adjacent with its neighbors.
Two routers will become neighbors if the following parameters are agreed upon:
1. Level 1-The two routers sharing a common network segment must have their interfaces configured to be in the same area if they are to have a Level 1 adjacency.
2. Level 2-The two routers sharing a common network segment must be configured as Level 2 if they are in different areas and want to become neighbors.
Within an Area Level 1 routing is used. Routing between areas is referred to as Level 2 routing.
A Level 1/Level 2 router may have neighbors in any area. It has two link-state databases: a Level 1 link-state database for intra-area routing and a Level 2 link-state database for inter-area routing. A Level 1/Level 2 router runs two SPFs and may require more memory and processing as a result.

QUESTION 196:

You have a network segment with an IS-IS level 1-2 router, and you're thinking of upgrading your network. What is the maximum number of Level-2 routing processes that this router can belong to?

A. 1
B. 2
C. 30
D. 40
E. limited only by the router's resources
F. Depends on the total number of IDB blocks.

Answer: A

Explanation:
The Cisco IOS software can handle simultaneous operation of up to 30 dynamic IP routing processes. The combination of routing processes on a router or access server

consists of the following protocols (with the limits noted):
1. Up to 30 IGRP routing processes
2. Up to 30 OSPF routing processes
3. One RIP routing process
4. One IS-IS process
5. One BGP routing process
Up to 30 EGP routing processes

QUESTION 197:

Which of the following types of routers can a Level-2 IS establish adjacencies with?
(Select three)

A. Any Level-1 IS in any area
B. Any Level-2 IS in any area
C. Any Level-1 IS in the same area
D. Any Level-1/Level-2 IS in any area
E. Any Level-1/Level-2 IS in the same area

Answer: B, D, E

Explanation:
The various adjacency types that an IS-IS router can be are shown below:

level-1	A Level 1 adjacency may be established if there is at least one area address in common between this system and its neighbors. Level 2 adjacencies will never be established over this interface.
level-1-2	A Level 1 and Level 2 adjacency is established if the neighbor is also configured as level-1-2 and there is at least one area in common. If there is no area in common, a Level 2 adjacency is established. This is the default.
level-2-only	Level 2 adjacencies are established if the other routers are L2 or L1L2 routers and their interfaces are configured for L1L2 or L2. Level 1 adjacencies will never be established over this interface.

Reference:http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/cs/csprtn1/csisis.htm

---

QUESTION 198:

One of the characteristics of the routing protocol IS-IS is its
ability to use level-1 routing. Where in an IS-IS network does this happen?

A. Between domains.
B. Between areas.
C. Between intermediate systems in the same area.
D. Between end systems and intermediate systems in the same area.

Answer: C

Explanation:
A two-level hierarchy is used to support large routing domains. A large domain may be administratively divided into areas. Each system resides in exactly one area.1 Routing within an area is referred to as Level 1 routing. Routing between areas is referred to as Level 2 routing. A Level 2 Intermediate System (IS) keeps track of the paths to destination areas. A Level 1 IS keeps track of the routing within its own area. For a packet destined for another area, a Level 1 IS sends the packet to the nearest Level 2 IS in its own area, regardless of what the destination area is. Then the packet travels via Level 2 routing to the destination area, where it may travel via Level 1 routing to the destination. It should be noted that selecting an exit from an area based on Level 1 routing to the closest Level 2 IS might result in suboptimal routing.2
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/ CK3 81/technologies_white_paper09186a00800a3e6f.shtml

---

QUESTION 199:

Your newly appointed Certkiller trainee wants to know why Level-3 area routing is not a supported feature of integrated IS-IS on Cisco routers. What will your reply be?

A. The system ID on a Cisco router is limited to 6 bytes.
B. The NET on a Cisco router is restricted to a maximum of 8 bytes.
C. The lack of a Domain portion of the NSAP only accommodates for 2 levels of routing hierarchy.
D. Cisco routers cannot route CLNS data that use the ISO/IEC 10589 standard of NSAP addressing.
E. Since the NSAP service identifier (N-SEL) must always be set to 00, no other service types are available.

Answer: C

Explanation:
Integrated IS-IS is a version of the OSI IS-IS routing protocol that uses a single routing algorithm to support more network layer protocols than just CLNP. Integrated IS-IS sometimes is called Dual IS-IS, named after a version designed for IP and CLNP networks. Only one IS-IS process is allowed whether you run it in integrated mode, ISO CLNS only or IP only.

QUESTION 200:

Which of the following three statements correctly describe the characteristics of IS-IS? (Select three)

A. L1 routers have no knowledge about routes outside their area.
B. L1/L2 routers maintain a separate Level 1 link-state database and a Level 2 link-state database; they do not advertise L2 routes to L1 routers.
C. To route packets to another area, L1 routers must forward the packets to the L2 router of the destination area.
D. To route packets to another area, L1 routers must forward the packets to an L1/L2 router within their area.
E. L2 routers form adjacencies with L1 and L1/L2 neighbors.

Answer: A, B, D

Explanation:
Intermediate System-to-Intermediate System (IS-IS) is an OSI link-state hierarchical routing protocol that floods the network with link-state information to build a complete, consistent picture of network topology. To simplify router design and operation, IS-IS distinguishes between Level 1 and Level 2 ISs. Level 1 ISs communicate with other Level 1 ISs in the same area. Level 2 ISs route between Level 1 areas and form an intradomain routing backbone. Hierarchical routing simplifies backbone design because Level 1 ISs need to know only how to get to the nearest Level 2 IS. The backbone routing protocol also can change without impacting the intra-area routing protocol.
Reference:http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/osi_rout.htm

---

QUESTION 201:

On router CK1 , IS-IS is configured on all Ethernet interfaces. In IS-IS, what network type identifies a LAN interface?

A. broadcast
B. point-to-point
C. pseudo-node
D. non-broadcast
E. point-to-multipoint

Answer: A

Explanation:
The types of networks that IS-IS defines include Point-to-point networks and Broadcast networks. Broadcast networks are used on all LAN interfaces.
Reference: http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm
Incorrect Answers:
B: Point-to-point networks, such as serial lines, connect a single pair of routers.

C: A Designated Intermediate System (DIS) creates a pseudonode (a virtual node), and all the routers on a LAN, including the DIS, form an adjacency with the pseudonode instead of forming n*(n-1) order adjacencies with each other in a full mesh. DISs are not used by default however.
D: Non-broadcast is not used by IS-IS.
E: Point-to-multipoint is not used by IS-IS.

QUESTION 202:

The Certkiller network is utilizing IS-IS for the routing protocol. By default, what is the metric used by the Certkiller routers on IS-IS interfaces?

A. Delay
B. Error
C. Default
D. Expense
E. Bandwidth
F. Hops
G. Cost

Answer: G

Explanation:
Explanation:
The original IS-IS specification defines four different types of metrics. All routers support the default metric of cost. Delay, expense, and error are optional metrics. The delay metric measures transit delay, the expense metric measures the monetary cost of link utilization, and the error metric measures the residual error probability associated with a link.
The Cisco implementation uses cost only. If the optional metrics were implemented, there would be a link-state database for each metric and SPF would be run for each link-state database.

QUESTION 203:

In the IS-IS routing protocol, what metric is supported exclusively by the Cisco IOS?

A. Cost
B. Delay
C. Expense
D. Error
E. None of the above

Answer: A

Explanation:
The original IS-IS specification defines four different types of metrics. Cost, being the default metric, is supported by all routers. Delay, expense, and error are optional metrics. The delay metric measures transit delay, the expense metric measures the monetary cost of link utilization, and the error metric measures the residual error probability associated with a link.
The Cisco implementation uses cost only. If the optional metrics were implemented, there would be a link-state database for each metric and SPF would be run for each link-state database.
IS-IS uses a single required default metric with a maximum path value of 1024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links are calculated by summing link values. Maximum metric values were set at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation will be reasonably efficient. IS-IS also defines three optional metrics (costs): delay, expense, and error.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/ CK3 81/technologies_white_paper09186a00800a3e6f.shtml

---

QUESTION 204:

The Certkiller network is displayed in the diagram below:

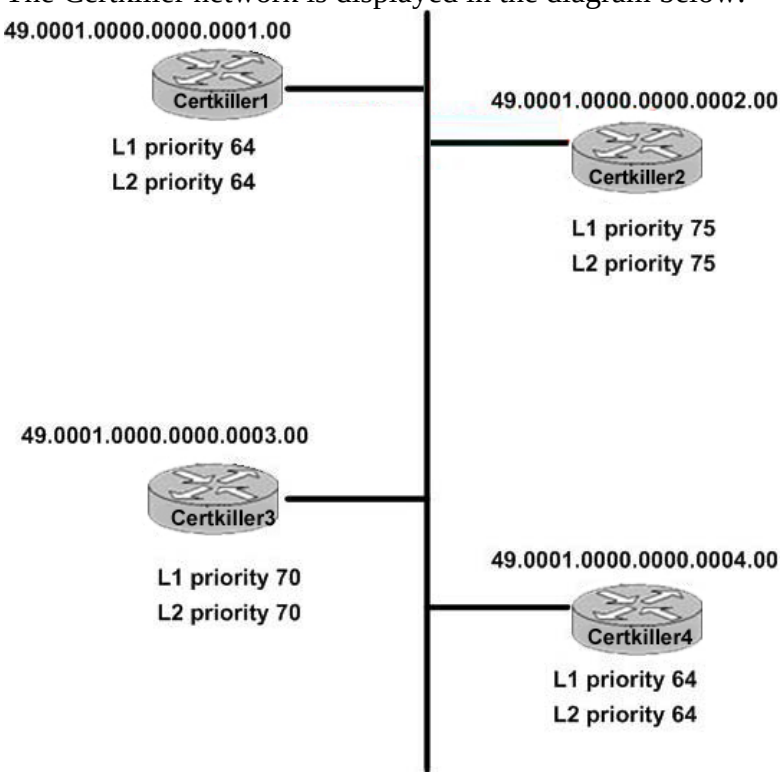

In the network segment of the above exhibit, which router is going to become the level-2 DIS (Designated Intermediate System)?

A. Certkiller 1

B. Certkiller 2
C. Certkiller 3
D. Certkiller 4
E. None

Answer: B

Explanation:
Since the Level 2 priority of Certkiller 2 is higher, it will become the DIS. The DIS election process is described below:
On a LAN, one of the routers elects itself the DIS, based on interface priority (the default is 64). If all interface priorities are the same, the router with the highest subnetwork point of attachment (SNPA) is selected. The SNPA is the MAC address on a LAN, and the local data link connection identifier (DLCI) on a Frame Relay network. If the SNPA is a DLCI and is the same at both sides of a link, the router with the higher system ID becomes the DIS. Every IS-IS router interface is assigned both a L1 priority and a L2 priority in the range of 0 to 127.
The DIS election is preemptive (unlike OSPF). If a new router boots on the LAN with a higher interface priority, the new router becomes the DIS. It purges the old pseudonode LSP and floods a new set of LSPs.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/technologies_tech_note09186a0080094b42.shtml#subfirstthree

---

QUESTION 205:

On a point-to-point IS-IS network connection, the PSNP (Partial Sequence Number Packet) is responsible for which of the following?

A. Acknowledging LSPs.
B. Replacing IIH packets.
C. Establishing adjacencies.
D. Sending link-state changes.
E. None of the above

Answer: A

Explanation:
Partial sequence number PDUs (PSNPs) are used to request an LSP (or LSPs) and acknowledge receipt of an LSP (or LSPs).
Reference: Cisco, Introduction to Intermediate, System-to-Intermediate System Protocol
http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm

---

QUESTION 206:

Which of the following metrics can be used by IS-IS when determining the best path to a destination? (Choose four)

A. Default
B. Load
C. Delay
D. Expense
E. Reliability
F. Error
G. Uptime

Answer: A, C, D, F

Explanation:
IS-IS uses a single required default metric with a maximum path value of 1024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links are calculated by summing link values. Maximum metric values were set at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation will be reasonably efficient. IS-IS also defines three optional metrics (costs): delay, expense, and error. The delay cost metric reflects the amount of delay on the link. The expense cost metric reflects the communications cost associated with using the link.
The error cost metric reflects the error rate of the link. IS-IS maintains a mapping of these four metrics to the quality of service (QoS) option in the CLNP packet header. IS-IS uses these mappings to compute routes through the network.

QUESTION 207:

The Certkiller IS-IS network is displayed below:

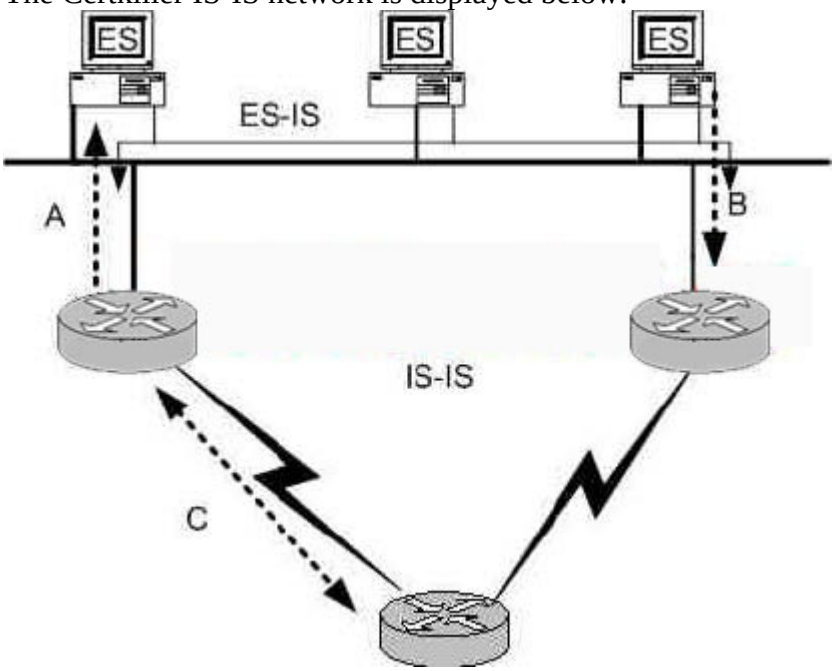

Given the above diagram, which statement correctly matches the letter to the type

of IS-IS hello PDU?

A. A=ISH; B=IIH; C:ESH
B. A=IIH; B=ESH; C=ISH
C. A=ISH; B=ESH; C:IIH
D. A=ESH; B=ISH; C=IIH
E. A=IIH; B=ISH; C=ESH

Answer: C

Explanation:
OSI CLNS is a network layer service similar to bare IP service. A CLNS entity communicates over Connectionless Network Protocol (CLNP) with its peer CLNS entity.
In the OSI architecture there are "systems": Routers are ISs, and hosts are End Systems (ESs).
ESs themselves have no routing information; they discover ISs (routers) by listening to Intermediate System Hellos (ISHs) and sending traffic to any random router. ESs send End System Hellos (ESHs); they do not choose a designated router to handle all traffic, and optimal routing is accomplished via redirects.
ISs discover ESs by listening to ESHs, and ISs send ISHs to ESs.
Intermediate System-to-Intermediate System Hello (IIH)-Used by routers to detect neighbors and form adjacencies. In addition to the IIH, which is an IS-IS protocol data unit (PDU), there is an ISH and an ESH, which are End System-to-Intermediate System (ES-IS) PDUs.

QUESTION 208:

DRAG DROP
As a Certkiller .com network technician you are required to drag the appropriate IS-IS routing item to its correct definition.

As a CertKiller.com network technician you are required to drag the appropriate IS-IS routing them to its correct definition.

Between ES's and IS's on the same subnet OSI routing begins at this level.	Drop answer here
Rouling between separate domains. Similar to BGP.	Drop answer here
Between IS's on the same area. Also called area routing.	Drop answer here
Called inter-area routing.	Drop answer here

Select from these

Level 1/2 routing | Level 2 routing | Level 3 routing | Level 0 routing | Level 1 routing

Answer:

As a CertKiller.com network technician you are required to drag the appropriate IS-IS routing them to its correct definition.

Between ES's and IS's on the same subnet OSI routing begins at this level.	Level 0 routing
Rouling between separate domains. Similar to BGP.	Level 3 routing
Between IS's on the same area. Also called area routing.	Level 1 routing
Called inter-area routing.	Level 2 routing

Select from these

Level 1/2 routing

Explanation:
Level 1 routing Level 1 routing is routing between intermediate systems within the same area. Basically, Level 1 routing is intra-area routing, and it occurs between all routers contained within the same area.
Level 2 routing Level 2 routing occurs between intermediate systems in different areas. All Level 2 routing will cross the backbone at some point. You can think of Level 2 routing as inter-area routing, and it occurs between routers in different IS-IS areas.
Level 3 routing Level 3 routing is routing between different routing domains. This type of routing will occur when traffic needs to leave the IS-IS routing domain to reach another routing domain. Level 3 routing is also known as internetwork routing.

QUESTION 209:

What does IS-IS use to establish and maintain neighbor relationships between ISs?

A. IIH
B. LSP
C. CLNS
D. CLNP
E. ISH

Answer: A

Explanation:
IIH (IS-IS Hello) - These packets create and maintain neighbor relationships and adjacencies. There are three types of Integrated IS-IS Hello packet. The type of packet is defined in the fixed header under the Type field and allows the packet to be handed off to the appropriate process.
Reference: CCNP BSCI Exam Certification Guide p.370

QUESTION 210:

In an IS-IS environment, what happens when the designated IS router crashes?

A. The elected backup designated router takes the place of the DIS indefinitely without the necessity of a new election.
B. The elected backup designated router takes the place of the DIS indefinitely without the necessity of a new election, until the original DIS comes back online.
C. A new election process occurs immediately, establishing a new DIS until a router with a higher priority or MAC address establishes an adjacency.
D. A new election process occurs immediately, establishing a new DIS that will remain in place indefinitely.

Answer: C

Explanation:
There is no backup designated router in IS-IS. Therefore, if the DIS meets an untimely death, a new DIS would be elected, based on priority or highest MAC address. If another router comes online with a higher priority, it will dislodge the existing DIS and rule in its place. This behavior is different from that of OSPF. Once a new DIS is elected, the link-state databases are purged and new LSPs are flooded.
Reference: CCNP Self-Study CCNP BSCI Exam certification guide p.806

QUESTION 211:

IS-IS Level 1 routers use LSPs for which purpose?

A. to build topology database for the local area only
B. to build topology database for the local and remote areas

C. to build topology database for remote areas
D. to build topology database for areas outside the AS
E. to establish with L1 IS's
F. to establish adjacency with ES's

Answer: A

Explanation:
IS-IS Level 1 routers shares LSAs with L1 routers in the same area.

QUESTION 212:

In IS-IS routing, when would the use of Level 1 IS routing be required?

A. When routing between end systems and intermediate systems.
B. When routing between Level 1 areas in different domains.
C. When routing between Level 1 areas in the same domain.
D. When routing between intermediate systems in the same area.
E. When routing between intermediate systems in different areas.

Answer: D

Explanation:
Layer 1 links connect the routers in the same are, so layer 1 routing takes place between intermediate systems in the same area. Layer 2 routing takes is between intermediate systems in different areas.

QUESTION 213:

Router CK1 is configured as a Level-1 IS-IS router. Which of the following routers can router CK1 establish an adjacency with? (Select two)

A. Any Level-1 IS in any area.
B. Any Level-2 IS in any area.
C. Any Level-1 IS in the same area.
D. Any Level-2 IS in the same area.
E. Any Level-1/Level-2 IS in the same area.

Answer: C, E

Explanation:
A Level-1 IS router can establish adjacencies with other routers Level-1 and Level-1/Level-2 IS routers within the same area.
Incorrect Answers:
A: Level-1 adjacencies can only be established within the same area.
B, D: Level-2 adjacencies require Level-2 IS routers.

QUESTION 214:

Router CK1 is an IS-IS Level-1/Level2 router. Which of the following type of router will CK1 NOT be able to establish an adjacency with?

A. Any Level-1 IS in any area.
B. Any Level-2 IS in any area.
C. Any Level-1 IS in the same area.
D. Any Level-1/Level-2 IS in any area.
E. Any Level-1/Level-2 IS in the same area

Answer: A

Explanation:
Level-1 routers can only establish adjacencies with other Level-1 routers within the same area. This is because L1 routers in different areas must connect through a L2 capable router, so two L1 routers must be in the same area in order to become adjacent.
Incorrect Answers:
B, C, D, E: IS-IS has a two-level hierarchy. Contiguous Level 2-capable routers form the backbone. Both Level 2 and Level 1 routers live in areas. Routers can be Level 1 (L1), Level 2 (L2), or both (L1/L2). Within Cisco IOS(r) Software, the default configuration is both Level 1 and Level 2 at the same time which allows an IS-IS network to run with minimal configuration in a plug-and-play fashion. L1/L2 routers act as 2 separate routers in terms of adjacency building, so it will become adjacent with other L2 routers, as well as other L1 routers that are in the same area.

QUESTION 215:

Study the exhibit shown below carefully. IS-IS is configured on all routers. What two NET addresses could be used to support an IS-IS adjacency between routers Certkiller 1 and Certkiller 2?

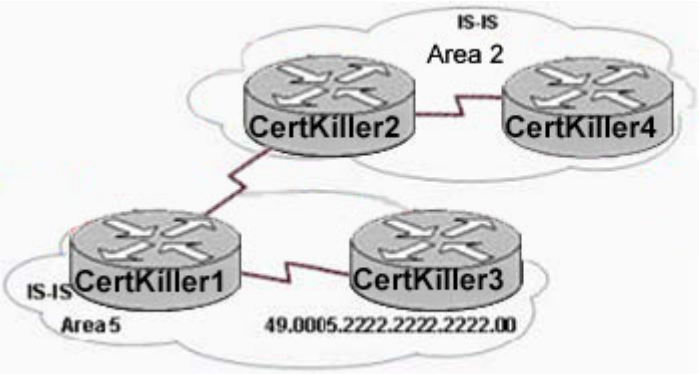

A. Certkiller 1 - 49.0005.1111.1111.1111.00
Certkiller 2 - 47.0002.2222.2222.2222.00
B. Certkiller 1 - 49.0005.1111.1111.1111.00
Certkiller 2 - 49.0002.2222.2222.2222.00

C. Certkiller 1 - 49.0005.1111.1111.1111.00
Certkiller 2 - 47.0005.2222.2222.2222.00
D. Certkiller 1 - 49.0002.1111.1111.1111.00
Certkiller 2 - 49.0002.2222.2222.2222.00

Answer: B

Explanation:
Network Entity Titles
Area ID The area ID is typically a one- or two-octet field that precedes the system ID. The area ID is used to signify the area the router belongs to. The area ID can span up to two octets if need be. For the BSCI, we will concentrate on a one-octet area ID. For the NSAP addressing scheme used with IS-IS, the area ID is everything to the left of the system ID.
System ID
The system ID is used to identify the router. It is similar to the router ID in OSPF. The system ID can be up to eight octets in length. Cisco supports only a six-octet field for the system ID. Whether or not the vendor you are using supports more than six octets, the number of octets used must be the same throughout the IS-IS routing domain. Normally, the system ID will be set to one of the MAC addresses on the router. This is not a rule, and the system ID can be represented in a number of ways. The system ID that is used for a router must be unique throughout the IS-IS routing domain. There must be only one system ID set per router, regardless of the number of routing processes. The system ID must be common across all routing processes on the same router.
SEL The NSAP selector byte (SEL) is a one-octet field that represents the service being offered at the network level of the device. For our implementation of IS-IS, the SEL will always be set to 00. 00 represents the router. A good comparison to the SEL is the TCP or UDP port number being associated with an IP address, forming a socket that identifies a unique application instance on the internetwork. Basically, it tells you what service is being referenced for the particular address.

QUESTION 216:

SIMULATION
Refer to the following network topology exhibit:

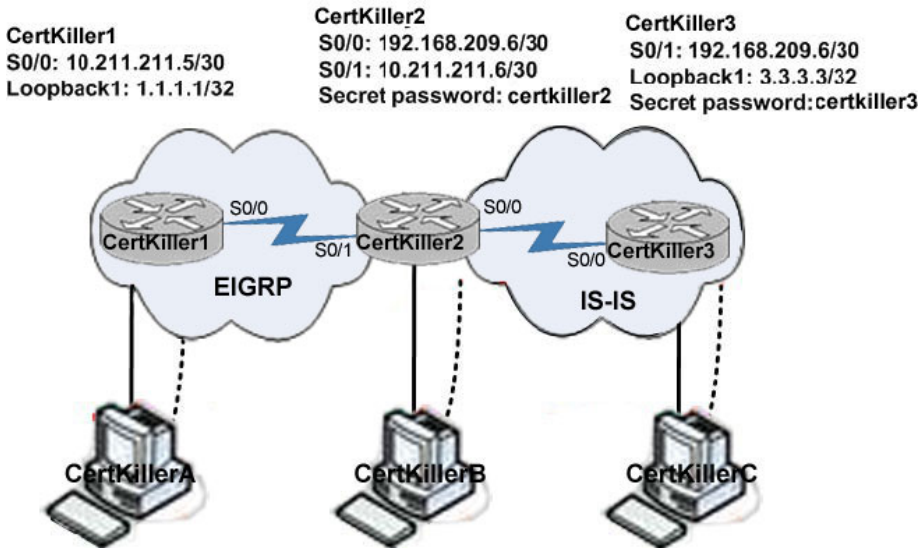

You work as a network engineer at Certkiller .com. Certkiller .com is a large international company with offices on all continents. Certkiller .com has recently purchased the Snafu Inc. company. Certkiller .com and Snafu Inc. have not been using the same routing protocols. Certkiller .com has been using EIGRP, while Snafu Inc. has been using IS-IS. Your boss at Certkiller .com, Miss Certkiller, has asked you to implement a resolution for the problem of sharing data among the two companies. As a first step you set up an intermediate router named Certkiller 2 between the two corporate networks. Your idea is to use Certkiller 2 to stage redistribution of information between the Certkiller .com and Snafu Inc. networks. You are now required to configure the redistribution according to the following requirements:

1. Seed metric for EIGRP must have the following characteristics:
Bandwidth=512 Kbps
2. Delay=100
3. Reliability=255
4. Load=1
5. MTU=1500

Seed metric for IS-IS must be set to 50.
Only redistribute Level-1 IS-IS routes into EIGRP.
Redistribute EIGPR routes into IS-IS as Level-1 routes.

Answer:

Explanation:
Certkiller 2# conf t
Certkiller 2(config)# router isis
Certkiller 2(config-router)# redistribute eigrp 100 level-1 metric 50
Certkiller 2(config)# router eigrp 100
Certkiller 2(config-router)# redistribute isis level-1 metric 512 10 255 1 1500
Certkiller 2(config-router)# redistribute connected
Certkiller 2# copy run start
Redistribution is the process of allowing routing information known in one routing

protocol to be shared with another routing protocol. It should be noted here that routing protocols are the only items that can be redistributed into. You can redistribute such items as connected interfaces, static routes, and default routes into a routing protocol. There are two types of redistribution available: one-way redistribution and mutual redistribution. One-way redistribution occurs when routing information contained in a routing protocol is shared with another routing protocol, but the other routing protocol doesn't share its routing information with the sharing protocol. Mutual redistribution is the process of two routing protocols sharing their routing information with each other. IGRP and EIGRPIGRP and EIGRP are Cisco proprietary routing protocols. Both use a composite metric made up of the following:
* Bandwidth
* Delay
* Reliability
* Load
IGRP is a classful routing protocol, whereas EIGRP is a classless routing protocol. Remember: If IGRP and EIGRP are running on the same router and both of them are using the same autonomous system number, they will automatically redistribute with each other. If they do not have the same autonomous system number, you will need to manually redistribute between them. These are the only IGPs we will discuss in this study guide that behave in this manner.
R1#conf tEnter configuration commands, one per line. End with CNTL/Z.R1(config)# router eigrp 100R1(config-router)#redistribute isis level-2 metric 100 100 200 1 1500 R1(config-router)#^ZR1#This configuration will redistribute only the IS-IS Level 2 routes into EIGRP AS 100. If we wanted all IS-IS routes to redistribute into EIGRP AS 100, we would need to use the following configuration:
R1#conf tEnter configuration commands, one per line. End with CNTL/Z.R1(config)# router eigrp 100 R1(config-router)#redistribute isis level-1-2 metric 100 100 200 1 1500R1(config-router)#^ZR1#That's really all there is to basic IS-IS redistribution. So far, we have described basic redistribution into all of the IGPs covered in this study guide. We're now going to look at redistributing connected interfaces, static routes, and default routes into the different routing protocols. We will then conclude the redistribution section with a look at VLSM to FLSM redistribution and incorporating filters into redistribution.

QUESTION 217:

You have been tasked with setting up IS-IS on a new Certkiller router. What are the basic configuration steps to enable IS-IS?

A. Configure the network net-id command(s) under router isis and enable IS-IS on each interface with the "ip router isis" command.
B. Configure the net system-id and the network net-id commands under router isis and enable IS-IS on each interface with the "ip router isis" command.
C. Configure the net system-id and the network net-id commands under router isis.
D. Configure the network net-id command(s) and the is-type level-1-2 command under router isis.

E. Configure the net system-id command under router isis and enable IS-IS on each interface with the "ip router isis" command.
F. None of the above.

Answer: E

Explanation:
To enable Integrated IS-IS on a router for IP routing is easy. There are many more commands used to tune the IS-IS processes. However, only the following three commands are required to start Integrated IS-IS:
1. Enable IS-IS as an IP routing protocol, using the command router isis, and assign a tag if there are multiple IS-IS processes. If the tag is omitted, a tag of zero (0) is assumed.
2. Identify the router for IS-IS by assigning a NET to the router with the net command.
3. Enable IS-IS on the interfaces participating in IS-IS, using the command ip router isis. This is slightly different to most other IP routing protocols where the participating interfaces are specified by network statements. There is no network statement under the IS-IS process. If there are multiple IS-IS processes, interfaces must state which process they belong to by specifying the appropriate tag.

QUESTION 218:

You need to troubleshoot some issues in the Certkiller IS-IS network. Which command lists the system IDs of all known IS-IS routers?

A. show clns neighbors detail
B. show isis topology
C. show clns neighbors
D. show is-is neighbors detail
E. show isis database
F. None of the above

Answer: B

Explanation:
The command show isis database gives you information about all the LSP information stored in the router's IS-IS link-state database. Here is a sample output from Houston's

IS-IS link-state database:

```
r13#show clns neighbors

System Id    Interface    SNPA            State   Holdtime   Type Protocol
r1           Fa0/0        0004.c136.66e0  Up      9          L1L2 IS-IS
r14          Se0/1        *HDLC*          Up      27         L1L2 IS-IS

r13#show isis database

IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num   LSP Checksum   LSP Holdtime     ATT/P/OL
r1.00-00              0x0000000D    0x0BE2         535              0/0/0
r1.01-00              0x00000009    0xC13E         324              0/0/0
r13.00-00           * 0x0000000D    0x801D         1005             0/0/0
r14.00-00             0x0000000D    0x72E1         613              0/0/0
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num   LSP Checksum   LSP Holdtime     ATT/P/OL
r1.00-00              0x0000000F    0x6A0F         778              0/0/0
r1.01-00              0x0000000B    0x4D39         995              0/0/0
r13.00-00           * 0x0000000F    0x24B0         543              0/0/0
r14.00-00             0x0000000F    0x3083         651              0/0/0
```

QUESTION 219:

The following output was seen on a Certkiller ISIS router:

```
CertKiller1#show ip route isis
     10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
 ! L2   10.200.200.14/32 [115/30] via 10.1.0.2, Serial1/0
 i L1   10.200.200.13/32 [115/20] via 10.1.1.3, Ethernet0/0
 i L1   10.1.3.0/24 [115/20] via 10.1.1.3, Ethernet0/0
 ! L2   10.1.2.0/23 [115/20] via 10.1.0.2, Serial1/0
 ! sd   10.1.2.0/23 [115/10] via 0.0.0.0, Null0
```

Based on the "show ip route isis" output on Certkiller 1 shown above, which statement is true?

A. The i su 10.1.0.0/23 route is a summary route.
B. The i su 10.1.0.0/23 route is an IS-IS external route.
C. The Certkiller 1 IS-IS router is an ASBR.
D. The Certkiller 1 IS-IS router is performing route aggregation and is suppressing the more specific 10.1.0.0/23 prefix.
E. The i su 10.1.0.0/23 route is a suppressed route.
F. The Certkiller 1 IS-IS router is an ABR that belongs to multiple IS-IS areas.

Answer: A

Explanation:
Example:
RouterA#show ip routeCodes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGPD - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter areaN1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGPi - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate defaultU - per-user static route, o - ODRGateway of last resort is not seti L1 192.168.50.0/24 [115/20] via 192.168.40.1, Serial1.1 i L2 192.168.1.0/24 [115/40] via 192.168.30.2, Serial0.1 i L2 192.168.10.0/24 [115/30] via 192.168.30.2, Serial0.1 i L2 192.168.20.0/24 [115/20] via 192.168.30.2, Serial0.1 C 192.168.40.0/24 is directly

connected, Serial1.1C 192.168.30.0/24 is directly connected, Serial0.1As can be seen from this routing table, there is one Level 1 route and three Level 2 routes that were learned through IS-IS. You know that a route is a Level 1 IS-IS route by the code L1, and you recognize a Level 2 IS-IS route by the code L2. The code I at the beginning of the line means that the route is an IS-IS route. I su means summary Address.

QUESTION 220:

The following output was seen on a Certkiller ISIS router:

```
CertKiller3# show isis route

IS-IS Level-1 Router  Table - version 4
System Id        Next-Icp        Interface    SNPA            Metric  State
CertKiller1      CertKiller1     Et0/0        aabb.cc00.3300  10      Up     L2-IS
CertKiller3      --
Default route out 1f area - (via 1 12-attached IS)
System Id        Next-Icp        Interface    SNPA            Metric  State
                 CertKiller1     Et0/0        aabb.cc00.3300  10      Up
```

In reference to the Certkiller 3 "show isis route" output shown above, which statement is true?

A. Certkiller 3 has been configured with a nondefault IS-IS metric.
B. Certkiller 1 has been configured with a nondefault IS-IS metric.
C. Certkiller 1 is a level-1 only IS-IS router.
D. Certkiller 3 is a level-2 only IS-IS router.
E. Certkiller 3 routing table should contain i L2 entries.
F. Certkiller 1 is the exit point out of the area for P1 Certkiller 3.
G. None of the above

Answer: F

Explanation:
The show isis route command shows the L1 routes to IS-IS neighbors. R1 has visibility of the other L1 routers in its area. The L1L2 routers appear in the L1 routing table by virtue of their L1 connection. There is a note at the end of their entry to show that they also act as L2. The closest L1L2 router also appears as the default route out of area. Again, the next-hop IS, its SNPA, is the interface over which that next hop is reached. The cumulative metric to that destination is shown for all IS routes. The neighbors show that their state is 'up' and the Hello process has established an adjacency.

QUESTION 221:

The Certkiller administrator has decided to configure Integrated IS-IS on the network. What must the administrator do to enable the interfaces to distribute IP information using IS-IS?

A. All configured IP networks will automatically be enabled when IS-IS is configured with the "router isis" global configuration command.
B. Each interface must be enabled to support IS-IS with the "ip router isis" interface

configuration command.
C. Each network must be identified with the "net interface" configuration command.
D. The networks configured on each interface must be associated to the IS-IS routing protocol using the "network" router configuration command.
E. None of the above.

Answer: B

Explanation:
To enable Integrated IS-IS on a router for IP routing is easy. There are many more commands used to tune the IS-IS processes. However, only the following three commands are required to start Integrated IS-IS:
1. Enable IS-IS as an IP routing protocol, using the command router isis, and assign a tag if there are multiple IS-IS processes. If the tag is omitted, a tag of zero (0) is assumed.
2. Identify the router for IS-IS by assigning a NET to the router with the net command.
3. Enable IS-IS on the interfaces participating in IS-IS, using the command ip router isis
. This is slightly different to most other IP routing protocols where the participating interfaces are specified by network statements. There is no network statement under the IS-IS process. If there are multiple IS-IS processes, interfaces must state which process they belong to by specifying the appropriate tag.

QUESTION 222:

Which IOS command would you enter to if you wanted to see the IS-IS level-2 routing database table?

A. show isis route
B. show clns route
C. show isis database
D. show clns neighbors
E. None of the above

Answer: C

Explanation:
The show isis database (detail) command displays the contents of the IS-IS database. By default, this command will display all of the L1 and L2 database information.

QUESTION 223:

Which IOS command displays the Level-1 routing table in Integrated IS-IS?

A. show isis route
B. show clns route
C. show isis database

D. show clns neighbors
E. None of the above

Answer: A

Explanation:
The "show isisroute" command displays the IS-IS Level 1 forwarding table for IS-IS-learned routes.

QUESTION 224:

The Certkiller network is shown below, along with the relevant information on router CK1 :

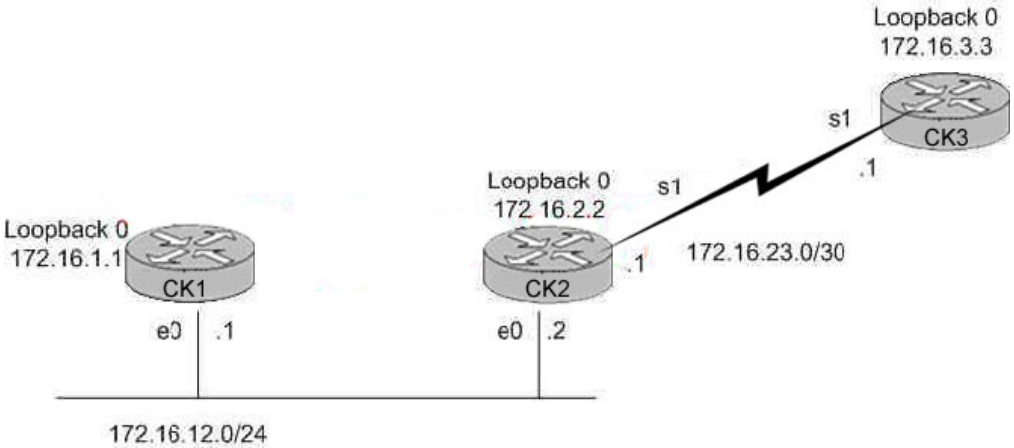

CK1 #show clns interface ethernet 0
Ethernet0 is up, line protocol is up
Checksums enabled, MTU 1497, Encapsulation
SAP
Routing Protocol: ISIS
Circuit Type: level-1-2
Interface number 0x0, local circuit ID 0x1
Level-1 Metric: 10, Priority: 64, Circuit ID: CK2 .02
Number of active level-1 adjacencies: 1
Level-2 Metric: 10, Priority: 64, Circuit ID: CK2 .01
Number of active level-2 adjacencies: 1
Next ISIS LAN Level-1 Hello in 5 seconds
Next ISIS LAN Level-2 Hello in 2 seconds
Based on the information above, which statement is correct about IS-IS?

A. Router CK1 is the DIS for both L1 and L1
B. Router CK1 is the backup DIS for both L1 and L2.
C. Router CK2 is the DIS for both L1 and L2.
D. Router CK1 is the DIS for L1 and a different router (CK2) is the DIS for L2.
E. Router CK1 is the DIS for L2 and a different router (CK2) is the DIS for L1.

Answer: C

Explanation:
It is the DIS that generates the pseudonode Link State Packet (LSP) and is denoted with a non-zero LSP-ID - R2.01

QUESTION 225:

In order to configure ISIS to route IP traffic, the following command was entered into router CK1 :
ip router isis
In which IOS configuration mode was this command typed in?

A. Line configuration mode
B. Router configuration mode
C. Global configuration mode
D. Interface configuration mode
E. Root configuration mode

Answer: D

Explanation:
To configure an IS-IS routing process for IP on an interface, use the ip router isis interface configuration command.
Note: To enable IS-IS, perform the following tasks starting in global configuration mode:
Step 1: router isis
Enable IS-IS routing and specify an IS-IS process for IP, which places you in router configuration mode.
Step 2: net network-entity-title
Configure NETs for the routing process; you can specify a name for a NET as well as an address.
Step 3: interfacetype number
Enter interface configuration mode.
Step 4: ip router isis [tag]
Specify the interfaces that should be actively routing IS-IS.
Incorrect Answers:
A, B; C: The ip router isis cannot be used in either line, router or Global configuration mode.
E: This mode does not exist.

QUESTION 226:

The Certkiller IS-IS network has been configured with mesh groups. Which of the following describe an advantage of utilizing mesh groups?

A. Mesh groups optimize LSP flooding.

B. Mesh groups keep the routers more secure.
C. Mesh groups help the routers form adjacencies.
D. Mesh group speed the flow of data across WAN links.
E. None of the above.

Answer: A

Explanation: The mesh group feature is a mechanism to reduce flooding of LSPs in nonbroadcast multi-access (NBMA) networks with highly meshed, point to point topologies.
Reference: RFC 2973, IS-IS Mesh Groups

QUESTION 227:

When configuring IS-IS, why is it important to manually configure the link costs for each of the associated links within the network?

A. Because by default, all links have a cost of 10 regardless of the bandwidth.
B. It is not important because the default provides for optimal routing.
C. Because there is no default link cost. The link cost must be configured on each interface.
D. Because by default, all LAN interfaces have a cost of 10 and all WAN interfaces have a cost of 50.

Answer: A

Explanation:
The original IS-IS specification defines four different types of metrics: cost, delay, expense, and error. The Cisco implementation uses cost only. All links use the metric of 10 by default.
The original IS-IS specification defines four different types of metrics. Cost, being the default metric, is supported by all routers. Delay, expense, and error are optional metrics. The delay metric measures transit delay, the expense metric measures the monetary cost of link utilization, and the error metric measures the residual error probability associated with a link.
The Cisco implementation uses cost only. If the optional metrics were implemented, there would be a link-state database for each metric and SPF would be run for each link-state database.
IS-IS uses a single required default metric with a maximum path value of 1024. The metric is arbitrary and typically is assigned by a network administrator. Any single link can have a maximum value of 64, and path links are calculated by summing link values. Maximum metric values were set at these levels to provide the granularity to support various link types while at the same time ensuring that the shortest-path algorithm used for route computation will be reasonably efficient. IS-IS also defines three optional metrics (costs): delay, expense, and error.
Reference: Introduction to Intermediate System-to-Intermediate System Protocol,

http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm
Incorrect Answers:
B: While some routing protocols calculate the link metric automatically based on bandwidth (OSPF) or bandwidth/delay (Enhanced Interior Gateway Routing Protocol [EIGRP]), there is no automatic calculation for IS-IS.
C: The default link cost is 10.
D: There is no such distinction between LAN and WAN interfaces.

QUESTION 228:

Which of the following best describe the addresses that are used at the network layer of the OSI model?

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
F. None of the above

Answer: D

Explanation:
The term "network address" is used to refer to the Network Service Access Point (NSAP) at which the OSI Network Service is made available to a Network Service user by the Network Service provider.
Incorrect Answers:
A: IP is a protocol that is based on the OSI model.
B: MAC addresses operate at a lower level in the OSI model.
C, E: Do not apply.
Reference: RFC 941, Addendum to the Network Service Definition Covering Network Layer Addressing

---

QUESTION 229:

The OSI IS-IS NSAP address 47.040C.0061.040C.0056.0D12.00 is applied to the Ethernet interface of router CK1 . What is the area ID of the address?

A. 00
B. 47
C. 47.040C
D. 47.040C.0061
E. 040C.0056.0D12
F. None of the above

Answer: D

Explanation:
An NSAP address (figure 7) has two major parts: the initial domain part (IDP) and the domain specific part (DSP) (Figure 7). The IDP consists of a 1-byte authority and format identifier (AFI) and a variable-length initial domain identifier (IDI), and the DSP is a string of digits identifying a particular transport implementation of a specified AFI authority. Everything to the left of the system ID can be thought of as the area address of a network node.
Figure 7The NSAP address

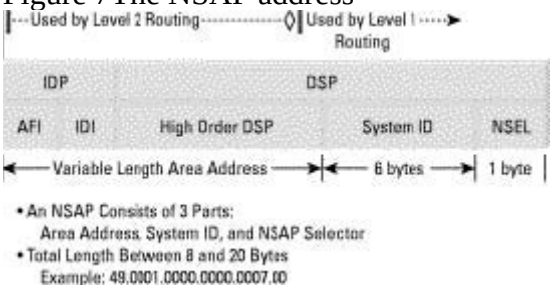

Incorrect Answers:
A: 00 is the SEL.
B, C: 47 and 47.040C is just a part of the Domain
E: 040C.0056.0D12 is the SystemID.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/ CK3 81/technologies_white_paper09186a00800a3e6f.shtml

---

QUESTION 230:

Which of the answer choices below represents a valid private, locally administered NSAP 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

Explanation:
The private NSAP addresses have AFI beginning with49, as these are locally administered addresses..
AESA Network Service Access Point (NSAP) ATM Addresses
There are 3 types of private ATM addresses:
1. NSAP encoding format for E.164 addresses - The authority and format identifier (AFI) is 45. These addresses are used in establishing ISDN calls by public networks, and they are normally used in public telephony.
2. Data Country Code (DCC) AESA
- The AFI is 39. These addresses are to be used in public networks. For example, the

initial domain identifier (IDI) value 0x84.0f identifies the United States.
3. International Code Designator (ICD) AESA - The AFI is 47. These addresses are used in private organizations, and the ICD field indicates the code set or organization. Cisco uses by default ICD addresses.
Incorrect Answers:
A: Addresses beginning with 39 is administered by the ISO for ISO Data Country Code.
B, D: 48 and 52 are unassigned AFI addresses.
Reference:
http://www.cisco.com/en/US/tech/ CK3 9/ CK4 9/technologies_tech_note09186a00800c9761.shtml

---

QUESTION 231:

You are a guest lecturer at the Certkiller Academy and are preparing a lesson on the Cisco IS-IS. Which of the following points can you include in your discussion of the Cisco IS-IS NSAP address System IDs? (Select three)

A. System IDs can vary in size within a domain.
B. The System ID identifies a node in an IS-IS network.
C. The System ID must be unique within a Level-1 area.
D. The System ID must be unique within a Level-2 area.
E. The System ID must be the MAC address of the router.

Answer: B, C, D

Explanation:
B: Each system ID within an area must be unique. It is used to identify a IS-IS node.
C: All Level 1 routers and hosts in an area must have an NSAP with the same area address.
D: Level 2 routers advertise their own area addresses (NSAP) to the other Level 2 routers in the backbone.
Incorrect Answers:
A: All ISs and ESs in a routing domain must have system IDs of the same length. Furthermore, Cisco implements a fixed length of 6 bytes for the system ID.
E: There are several techniques for creating unique system IDs
* Start numbering 1, 2, 3, 4, and so on.
* Use Media Access Control (MAC) addresses.
* Convert and use the loopback IP address: 192.168.11.1 --> 192.168.011.001--> 1921.6801.1001.
Reference: Introduction to Intermediate System-to-Intermediate System Protocol
http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/insys_wp.htm

---

QUESTION 232:

When comparing and contrasting the NET and NSAP used in IS-IS routing, which of the following statements is true?

A. A NET is an NSAP address with the N-selector byte set to 00.
B. Network Entity Titles do not have to start and stop on byte boundaries.
C. The System ID field of the NSAP address does not uniquely identify a node.
D. To identify a Domain, an NSAP address must be used because a NET can only identify an Area.
E. A NET is a special version of an NSAP address restricted to 8 bytes for the Area Address, System ID and the N-Selector byte.
F. None of the above.

Answer: A

Explanation:
NETs and NSAPs are ISO addresses. The differences between the NET and NSAP addresses are subtle. The NET address is the address of the host, where the value in the NSEL field is set to 0x00. Therefore, there is no upper-layer protocol identified within the host. With no application identified within the end host, the packet can be routed to the destination, but it cannot be handed off to a process after it has been delivered. However, routers do not have upper-layer protocols to identify because they are transitory ISs. Therefore, the NSAP of the router is referred to as a NET because the NSEL field is set to 00. The NSAP is the full ISO address. It not only defines the area and destination host within the area, but also specifies where to send the incoming packet after it has reached the host. The NSEL field at the end of the ISO address specifies the upper-layer protocol and is similar to the Protocol field of the IP header.
Network Entity Title (NET) In IS-IS, this is the ISO address of the system, but not to the process destination within the system. The NET describes both the area and system ID of a system in the IS-IS network but excludes the NSEL, which is set to 0x00. If the NSEL identifies the process within the system, the ISO address is called the "NSAP address."
Network Service Access point (NSAP) In IS-IS, this describes a service at the network layer to which the packet is to be directed. The NSAP is the NET address with the NSEL field set to a positive value, a value other than 0x00.
Network entity title (NET)
The NET describes both the area and system ID of a system in the IS-IS network but excludes the NSEL, which defines the NSAP address of the system.
Network service access point (NSAP)
Describes a service at the network layer to which the packet is to be directed. The NSAP is the NET address with the SEL field set to a value other than 0x00.

An ISO Address

IDP		DSP		
AFI (1 octet)	IDI	High Order DSP	System ID (1-8 octets)	NSEL (1 octet)
AREA			ID	SEL

QUESTION 233:

What are the basic configuration steps needed to enable IS-IS?

A. Configure the net system-id command under router isis and enable IS-IS on each interface with the ip router isis command.
B. Configure the network net-id command(s) under router isis and enable IS-IS on each interface with the ip router isis command.
C. Configure the network net-id command(s) and the is-type level-1-2 command under router isis.
D. Configure the net system-id and the network net-id commands under router isis.
E. Configure the net system.-id and the network net-id commands under router isis and enable IS-IS on each interface with the ip router isis command.

Answer: A

Explanation:
In order to enable IS-IS for IP on a Cisco router and have it exchange routing information with other IS-IS enabled routers, you must perform the following two tasks:
1. Enable the IS-IS process and assign area
2. Enable IS-IS for IP routing on an interface
The sample configuration below configures an IS-IS router with the following parameters:
1. Area 49.0001
2. Level 1 (L1) and Level 2 (L2) routers (this is the default unless otherwise specified)
3. No optional parameters
4. Running IS-IS for IP only
5. Loopback interfaces (loopbacks are advertised by IS-IS, not IS-IS enabled)

```
Router CK1
!
interface Loopback0
ip address 172.16.1.1 255.255.255.255
!--- Creates loopback interface and assigns
!--- IP address to interface Loopback0.
!
interface Ethernet0
ip address 172.16.12.1 255.255.255.0
ip router isis
!--- Assigns IP address to interface Ethernet0
!--- and enables IS-IS for IP on the interface.
!
router isis
passive-interface Loopback0
net 49.0001.1720.1600.1001.00
!
!--- Enables the IS-IS process on the router,
!--- makes loopback interface passive
!--- (does not send IS-IS packets on interface),
```

!--- and assigns area and system ID to router.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/technologies_configuration_example09186a0080093f38.shtml

---

QUESTION 234:

An administrator has decided to configure Integrated IS-IS on the network.
What must the administrator do to enable the interfaces to distribute IP
information using IS-IS?

A. The networks configured on each interface must be associated to the IS-IS routing
protocol using the network router configuration command.
B. Each interface must be enabled to support IS-IS with the ip router isis interface
configuration command.
C. All configured IP networks will automatically be enabled when IS-IS is configured
with the router isis global configuration command.
D. Each network must be identified with the net interface configuration command.
E. None of the above

Answer: B

Explanation:
The preparation for configuring any routing protocol requires a thorough understanding
of the network topology and a coherent addressing scheme. When you have these, the
basic configuration of the Integrated IS-IS is as follows:
1) Enable the routing process Integrated IS-IS with the router isis command.
2) Configure the Network Entity Title (NET) address, thus assigning the area with the net
network-address router subcommand.
3) Enable Integrated IS-IS for IP on the relevant interfaces with the IP router isis
interface subcommand.
Reference: Cisco press 642-801 Study Guide p.397

QUESTION 235:

You're troubleshooting a router that's running Integrated IS-IS and you suspect
something abnormal with the neighboring routers. Which IOS command would you
enter to display the adjacencies?

A. show isis route
B. show clns route
C. show isis database
D. show clns neighbors
E. None of the above

Answer: D

Explanation:
The show clns neighbors command displays ES and IS neighbors. The output includes adjacency information.
Example:
router CertK # show clns neighbors
System Id SNPA Interface State Holdtime Type Protocol
0000.0000.0007 aa00.0400.6408 Ethernet0 Init 277 IS ES-IS
0000.0C00.0C35 0000.0c00.0c36 Ethernet1 Up 91 L1 IS-IS
0800.2B16.24EA aa00.0400.2d05 Ethernet0 Up 29 L1L2 IS-IS
0800.2B14.060E aa00.0400.9205 Ethernet0 Up 1698 ES ES-IS
0000.0C00.3E51 *HDLC* Serial1 Up 28 L2 IS-IS
0000.0C00.62E6 0000.0c00.62e7 Ethernet1 Up 22 L1 IS-IS
0A00.0400.2D05 aa00.0400.2d05 Ethernet0 Init 24 IS ES-IS
Incorrect Answers
A: The show isis routes command is used to display the IS-IS Level 1 forwarding table for IS-IS learned routes.
B: The show clns route command is used to display all of the destinations to which this router knows how to route packets.
C: The show isis database command is used to display the Intermediate System-to-Intermediate System (IS-IS) link state database.

QUESTION 236:

The IS-IS routing table for Certkiller 1 is displayed in the diagram below:

```
CertKiller1#show ip route isis
     10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
i L2    10.200.200.14/32 [115/30] via 10.1.0.2, Serial1/0
i L1    10.200.200.13/32 [115/20] via 10.1.1.3, Ethernet0/0
i L1    10.1.3.0/24 [115/20] via 10.1.1.3, Ethernet0/0
i L2    10.1.2.0/23 [115/20] via 10.1.0.2, Serial1/0
i su    10.1.0.0/23 [115/10] via 0.0.0.0, Null0
```

Regarding the command output on Certkiller 1 in the exhibit, which statement is true?

A. The 10.1.0.0/23 route is an IS-IS external route
B. The Certkiller 1 IS-IS router is an ASBR.
C. The 10.1.0.0/23 route is a suppressed route.
D. The 10.1.0.0/23 route is a summary route.
E. The Certkiller 1 IS-IS router is an ASBR that belongs to multiple IS-IS areas.
F. The Certkiller 1 IS-IS router is performing route aggregation and is suppressing the more specific 10.1.0.0/23 prefix.

Answer: D

Explanation:

The "su" in the routing table denotes an IS-IS summary null route. This route is automatically created within the router when IS-IS route summarization is configured.

QUESTION 237:

The "show isis route" command was issued on router Certkiller 3 as shown below:

```
Certkiller 3#show isis route

IS-IS Level-1 Routing Table   Version
System Id       Next-Hop        Interface   SNPA            Metric  State
Certkiller3     Certkiller1     Et0/0       aabb.cc00.3300  10      Up     L2-IS
Certkiller1     --

Default route out of area - (via 1 L2-attached IS)
System Id       Next-Hop        Interface   SNPA            Metric  State
                Certkiller1     Et0/0       aabb.cc00.3300  10      Up
```

Based on the information above, which statement is true?

A. Certkiller 1 is the exit point out of the area for Certkiller 3
B. Certkiller 1 is a level-1 only IS-IS router.
C. Certkiller 1 has been configured with a non-default IS-IS metric.
D. Certkiller 3 is a level-2 only IS-IS router.
E. Certkiller 3 has been configured with a non-default IS-IS metric.
F. Certkiller 3 routing table should contain L2 entries.
G. None of the above

Answer: A

Explanation:
In the IS-IS routing table of router Certkiller 3, the default router is displayed as Certkiller 1, which is a Level 2 router. Level 2 routers are used for routing inter-area traffic, so Certkiller 3 will use this router as the exit point for all traffic destined to another area.
Incorrect Answers:
B: Certkiller 1 is a Level 2 router.
C, E: The default IS-IS metric for all links is 10, which is the value shown here.
D, F: Certkiller 3 is a level 1 router.

QUESTION 238:

The routing table of Certkiller 1 is displayed below:

```
CertKiller1# show ip route

Gateway of last resorf is not set

140.140.0.0 is subnetted (mask is 255.255.255.0), 3 subnets

C       140.140.64.0 255.255.255.0 is possibly down, routing via 0.0.0.0, Ethernet 0
i  L2   140.140.67.0 [115/20] via 140.140.64.240, 0:00:12, Ethernet 0
i  L2   140.140.66.0 [115/20] via 140.140.64.240, 0:00:12, Ethernet 0
```

Based on this information, how was network 140.140.66.0 discovered?

A. It is an IGRP route that has been redistributed in from an external protocol.
B. It is an IS-IS route via inter-area routing.
C. It is an IS-IS route via intra-area routing.
D. It is an IBGP route.
E. It is a static route.

Answer: B

Explanation:
Level-2 routers are used for routing traffic between areas (inter-area). Based on the routing table of Certkiller 1, the 140.140.66.0/24 network is known via a Level 2 IS-IS router.

QUESTION 239:

Exhibit:
Certkiller A# show ip route isis
10.0.0.0/8 is variably subnetted, 7 subnets, 3 masks
I L2 10.200.200.14/32 [115/30] via 10.1.0.2, Serial 1/0
I L1 10.200.200.13/32 [115/20] via 10.1.1.3, Ethernet 0/0
I L1 10.1.3.0/24 [115/20] via 10.1.1.3, Ethernet 0/0
I L2 10.1.2.0/23 [115/20] via 10.1.0.2, Serial 1/0
I su 10.1.0.0/23 [115/10] via 0.0.0.0, Null0
Based on the show ip route isisoutput on Router Certkiller A, which statement is true?

A. The "I su" 10.1.0.0/23 route is an IS-IS external route.
B. The Certkiller A IS-IS router is an ASBR.
C. The "I su" 10.1.0.0/23 route is a suppressed route.
D. The "I su" 10.1.0.0/23 route is a summary route.
E. The R1 IS-IS router is an ABR that belongs to multiple IS-IS areas.
F. The R1 IS-IS route is performing route aggregation and is suppressing the more 10.1.0.0/23 prefix.

Answer: D

Explanation:
When viewing the "show IP route ISIS" output, the "su" entry represents a summarized route. This is the direct result of the "summary-address" router configuration command. When creating the summary route. The Cisco IOS automatically creates the summary route and point it to interface Null 0.

QUESTION 240:

The Certkiller network is displayed below:

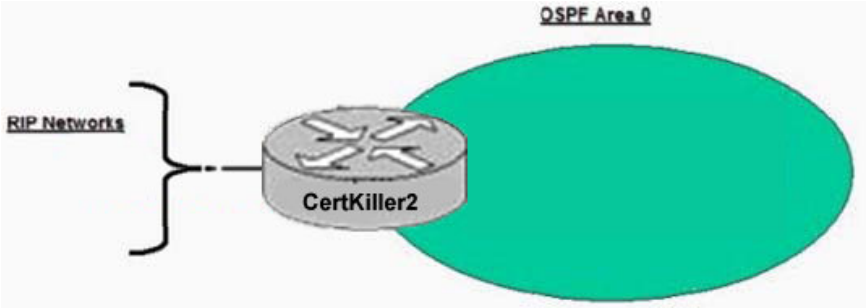

Router Certkiller 2 was configured as shown below:

```
CertKiller2(config)# route-map REDIST-MAP permit 10
CertKiller2(config-route-map)# match ip address 101
CertKiller2(config-route-map)# set metric 100
CertKiller2(config-route-map)# set metric-type type-1
CertKiller2(config-route-map)# exit

CertKiller2(config)# route-map REDIST-MAP deny 20
CertKiller2(config-route-map)# match ip address 110
CertKiller2(config-route-map)# exit

CertKiller2(config)# route-map REDIST-MAP permit 30
CertKiller2(config-route-map)# set metric 5000
CertKiller2(config-route-map)# set metric-type type-2
CertKiller2(config-route-map)# exit

CertKiller2(config)# router ospf 10
CertKiller2(config-router)# redistribute rip subnets route-map REDIST-MAP
```

Based on the exhibits above, which statement is true.

A. All routes matching access list 110 will be forwarded.
B. There is a set entry missing in the route-map REDIST-MAP deny 20 statement.
C. The set metric-type entry is not a valid route-map command.
D. There is a match entry missing in the route-map REDIST-MAP permit 30 statement.
E. All routes not matching access list 101 or 110 will be flagged as an OSPF external type 2 with a metric of 5000.
F. None of the above.

Answer: E

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.

Route maps operate similar to access lists, by examining one line at a time and when a match is found, action is taken. Route maps are different from numbered access lists because they can be modified without changing the entire list. Each route map statement is given a number. If a sequence number is not specified, the first route map condition will automatically be numbered as ten (10). The second condition will automatically be numbered as 20, and so on. The optional sequence number can be used to indicate the position that a new route map is to have in the list of route maps already configured with the same name.
After entering the route-map command, enter set and match commands in the route-map configuration mode. Each route-map command has a list of match and set commands associated with it. The match commands specify the match criteria. They are the conditions that should be tested to determine whether to take action. The set commands specify the set actions. They are the actions to be performed if the match criteria are met.

QUESTION 241:

The Certkiller WAN is displayed below:

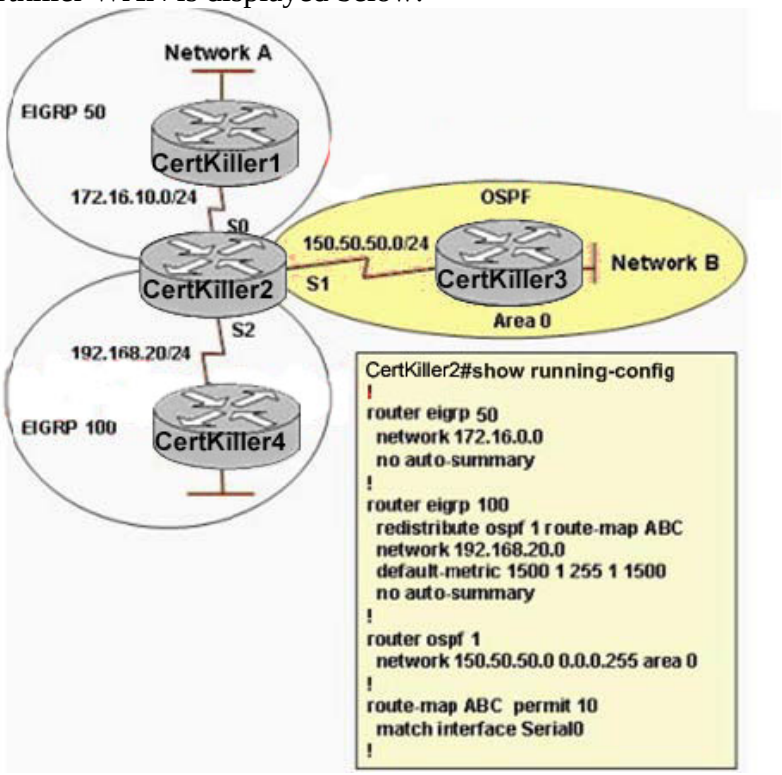

Based on the configuration shown above and the default behavior of the routing protocols, which network will be present in the routing table of Certkiller 4?

A. Only networks belonging to EIGRP 100 will be present in the routing table of Certkiller 4.
B. Network B, which is redistributed using the redistribute statement under EIGRP 100.
C. Network A, which is automatically redistributed between two different EIGRP autonomous systems.
D. Network A, which is automatically redistributed between two different EIGRP

autonomous systems and Network B which is redistributed using the redistribute statement.

Answer: A

Explanation:
To support multiple routing protocols within the same internetwork efficiently, routing information must be shared among the different routing protocols. For example, routes learned from a RIP process may need to be imported into an IGRP process. This process of exchanging routing information between routing protocols is called route redistribution. Such redistribution can be one-way or two-way. One-way routes are where one protocol receives the routes from another. Two-way routes are where both protocols receive routes from each other. Routers that perform redistribution are called boundary routers because they border two or more autonomous systems or routing domains. This section examines route redistribution in detail, including the use of administrative distance, guidelines for redistribution implementation, and issues with redistribution configuration. In exhibit eigrp is redistributed into ospf but ospf is not redistributed into eigrp so Certkiller 4 has the routing table only belongs to eigrp 100.

QUESTION 242:

The Certkiller 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.

Based on the network shown above, what must the administrator do to accomplish this?

A. Configure summarization on Certkiller 1 with area 1 range 172.16.32.0 255.255.224.0
B. Configure summarization on Certkiller 1 with summary-address 172.16.32.0 255.255.224.0
C. Configure area 1 as a NSSA area with area 1 stub nssa
D. Configure area 1 as a stub area with area 1 stub
E. None of the above

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

To determine what the summary route would be for these networks, we can follow a simple two-step process:

1. Write out each of the numbers in binary, as shown in the table.

IP Network Address	Binary Equivalent
172.16.100.0	**10101100.0001000.01100**100.0
172.16.101.0	**10101100.0001000.01100**101.0
172.16.102.0	**10101100.0001000.01100**110.0
172.16.103.0	**10101100.0001000.01100**111.0
172.16.104.0	**10101100.0001000.01101**000.0
172.16.105.0	**10101100.0001000.01101**001.0
172.16.106.0	**10101100.0001000.01101**010.0

2. Examine the table to determine the maximum number of bits (starting from the left) that all of the addresses have in common. (Where they are lined up, we boldfaced them to make them easier for you to see.) The number of common bits is the prefix length for the summarized address (/20).

In this example, we can see from the table that all of the addresses have the first 20 bits in common. The decimal equivalent of these first 20 bits is 172.16.96.0. So, we can write our new summarized address as 172.16.96.0/20. If we were to later add a network 172.16.98.0, it would need to be behind the router summarizing this address space. If we didn't, it could cause problems.

QUESTION 243:

A Certkiller router is configured as shown below:

```
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
```

Based on the output shown above, which three statements accurately describe the result of applying the exhibited route map? (Select three)

A. The map redistributes into EIGRP all routes that match the pfx prefix list and the five metric values-40000, 1000, 255, 1, and 1500.
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.
E. The map prohibits the redistribution of all external OSPF routes with tag 6 set.
F. All routes that do no match clauses 10 and 20 of the route map are redistributed with their tags set to 8.

Answer: C, D, F

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.
Route maps operate similar to access lists, by examining one line at a time and when a match is found, action is taken. Route maps are different from numbered access lists because they can be modified without changing the entire list. Each route map statement is given a number. If a sequence number is not specified, the first route map condition will automatically be numbered as ten (10). The second condition will automatically be numbered as 20, and so on. The optional sequence number can be used to indicate the position that a new route map is to have in the list of route maps already configured with

the same name.
After entering the route-map command, enter set and match commands in the route-map configuration mode. Each route-map command has a list of match and set commands associated with it. The match commands specify the match criteria. They are the conditions that should be tested to determine whether to take action. The set commands specify the set actions. They are the actions to be performed if the match criteria are met.

QUESTION 244:

Study the exhibit below carefully:

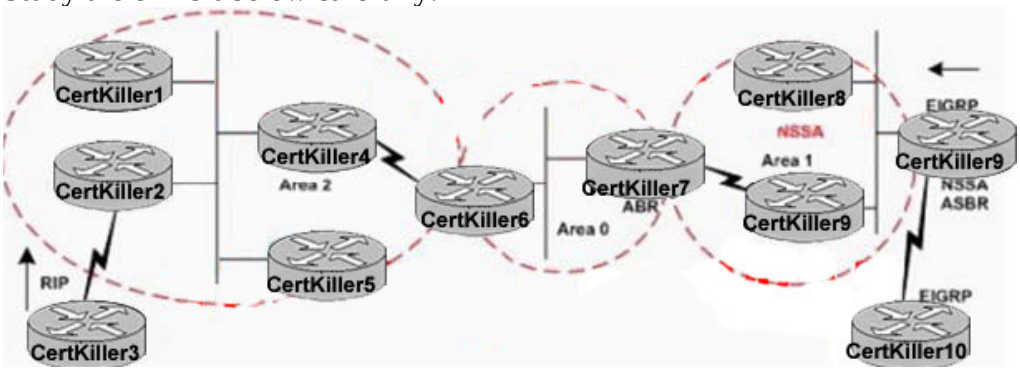

Based on the network shown above, will redistributed RIP routes from OSPF Area 2 be allowed in Area 1?

A. RIP routes will be allowed in Area 1 only if they are first redistributed into EIGRP.
B. Because Area 1 is an NSSA, redistributed RIP routes will not be allowed.
C. 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.
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. Because NSSA will discard type 7 LSAs, redistributed RIP routes will not be allowed in Area 1.

Answer: B

Explanation:
The following are several area types that are possible on OSPF:
1. Standard area - A standard area can accept link updates and route summaries.
2. Backbone area (transit area) - When interconnecting multiple areas, the backbone area is the central entity to which all other areas connect. The backbone area is always Area 0. All other areas must connect to this area to exchange route information. The OSPF backbone has all the properties of a standard OSPF area.
3. Stub area -A stub area is an area that does not accept information about routes external to the autonomous system, the OSPF internetwork, such as routes from non-OSPF sources. If routers need to reach networks outside the autonomous system, they use a default route. A default route is noted as 0.0.0.0/0.
4. Totally stubby area - A totally stubby area is an area that does not accept external autonomous system (AS) routes and summary routes from other areas internal to the

autonomous system. Instead, if the router needs to send a packet to a network external to the area, it sends it using a 0.0.0.0/0 default route. Totally stubby areas are a Cisco proprietary feature.
5. Not-so-stubby area (NSSA) - An NSSA is an area that is similar to a stub area but allows for importing external routes as Type 7 LSAs and translation of specific Type 7 LSA routes into Type 5 LSAs.

QUESTION 245:

The Certkiller network is shown in the following topology exhibit:

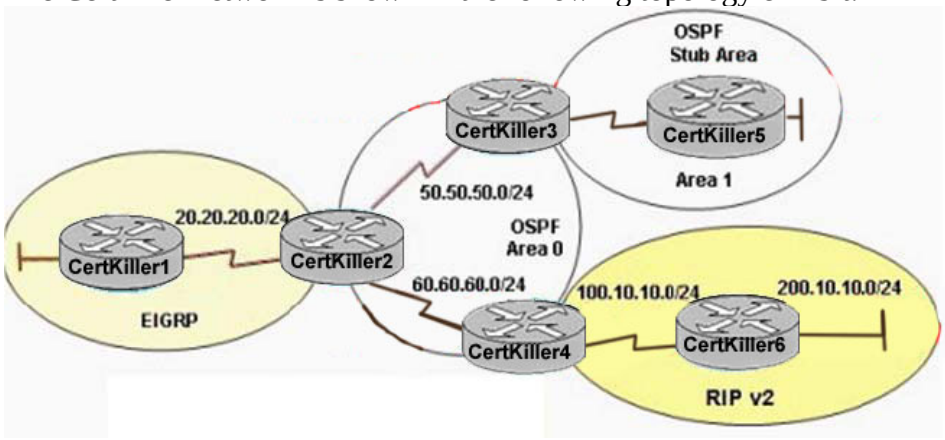

Certkiller 2 configuration exhibit:

```
CertKiller2#show running-.config
!
router eigrp 100
 redistribute ospf 1 route-map RED
 network 20.0.0.0
 default-metric 1500 1 255 1 1500
 no auto-summary
 !
 router ospf 1
 network 50.50.50.0 0.0.0.255 area 0
 network 60.60.60.0 0.0.0.255 area 0
!
access-list 10 deny 100.10.10.0 0.0.0.255
access-list 10 deny 200.10.10.0 0.0.0.255
access-list 10 permit any
 !
 route-map ABC  permit 10
  match ip address 10
```

Refer to the network exhibit topology exhibit and the partial configuration exhibit of router Certkiller 2. On router Certkiller 4, all RIP routes are redistributed into the OSPF domain. A second redistribution is configured on router Certkiller 2 using a route map. Based on the configuration on Certkiller 2, which EIGRP external routes will be present in the routing table of Certkiller 1?

A. There will be no EIGRP external routes in the routing table of Certkiller 1.
B. The routes originating from the RIP routing domain.
C. Only routes originating in the OSPF routing domain.
D. All routes originating from RIP and OSPF routing domains.
E. None of the above.

Answer: C

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.
Route maps operate similar to access lists, by examining one line at a time and when a match is found, action is taken. Route maps are different from numbered access lists because they can be modified without changing the entire list. Each route map statement is given a number. If a sequence number is not specified, the first route map condition will automatically be numbered as ten (10). The second condition will automatically be numbered as 20, and so on. The optional sequence number can be used to indicate the position that a new route map is to have in the list of route maps already configured with the same name.
In Exhibit Access-list is created to deny from 100.10.0.0 and 200.10.10.0 (RIP Domain) and that is called by route-map RED.
While redistributing OSPF routes into EIGRP RED rout-map is called so it denies to publish the RIP domain network into the EIGRP.

QUESTION 246:

The Certkiller network administrator is troubleshooting a redistribution problem with RIP routes into OSPF. Given the exhibited configuration commands shown below, which statement is true?

```
router rip
   network 10.0.0.0
!
reuter ospf 5
   network 172.10.0.0  0.0.255.255 area 0
   redistribute rip
```

A. Redistributed routes will maintain their original RIP routing metric.
B. Redistributed routes will have a default metric of 0 and will be treated as unreachable and not advertised.
C. Redistributed routes will be tagged as external type 2 (E2) with a metric of 20.
D. Redistributed routes will be tagged as external type 1 (E1) with a metric of 30.
E. Redistributed routes will have a default metric of 0 but will not be treated as reachable and will be advertised.
F. None of the above

Answer: C

Explanation:
To support multiple routing protocols within the same internetwork efficiently, routing information must be shared among the different routing protocols. For example, routes learned from a RIP process may need to be imported into an IGRP process. This process of exchanging routing information between routing protocols is called route redistribution. Such redistribution can be one-way or two-way. One-way routes are where one protocol receives the routes from another. Two-way routes are where both protocols receive routes from each other. Routers that perform redistribution are called boundary routers because they border two or more autonomous systems or routing domains. This section examines route redistribution in detail, including the use of administrative distance, guidelines for redistribution implementation, and issues with redistribution configuration.
The syntax of the redistribute command is as follows:
Router(config-router)#redistribute protocol [process-id] {level-1 | level-1-2 | level-2} [metric metric-value][metric-type type-value] [match {internal | external 1 | external 2}][tag tag-value] [route-map map-tag][weight weight] [subnets]
OSPF defines two external metric types: Type 1 and Type 2. By default, routes are redistributed into OSPF as Type 2. From the perspective of the router computing the cost to a route external to the OSPF autonomous system, a Type 1 external route (marked with E1 in the IP routing table) has a metric that is the sum of the internal OSPF cost, all the way back to the ASBR, and the external redistributed cost. From the viewpoint of the same router, a Type 2 external route (marked with E2 in the IP routing table) has a metric equal only to the redistributed cost that the ASBR originally advertised, with no premium for pathways back internal to the ASBR

QUESTION 247:

The Certkiller network administrator is troubleshooting a redistribution of OSPF routes into EIGRP. Given the exhibited commands shown below, which statement is true?

```
router ospf 5
 network 10.0.0.0  0.0.255.255 area 0

router eigrp 100
   network 172.10.0.0
   redistribute ospf 5
```

A. Redistributed routes will have an external type of 2 and a metric of 20.
B. Redistributed routes will have a default metric of 0 but will be treated as unreachable and not advertised.
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 an external type of 1 and a metric of 1.
F. None of the above.

Answer: B

Explanation:
IGRP and EIGRPIGRP and EIGRP are Cisco proprietary routing protocols. Both use a composite metric made up of the following:
* Bandwidth
* Delay
* Reliability
* Load
IGRP is a classful routing protocol, whereas EIGRP is a classless routing protocol.
Remember: If IGRP and EIGRP are running on the same router and both of them are using the same autonomous system number, they will automatically redistribute with each other. If they do not have the same autonomous system number, you will need to manually redistribute between them. These are the only IGPs we will discuss in this study guide that behave in this manner.
Like IGRP, EIGRP requires you to specify a metric except when the redistribution occurs with another instance of EIGRP or IGRP. All instances of EIGRP and IGRP use the same composite style of metrics, meaning that the metric will carry across the redistribution. Redistributing any other routing protocol, except EIGRP or IGRP, into EIGRP will require you to set the metric. If a metric is not specified, the metric value defaults to 0. Because EIGRP uses a composite metric, it will not understand a metric of 0 and the routes will not be used. Be smart and always supply a metric for routes being redistributed into EIGRP.
Example:
RouterA(config)#router ripRouterA(config-router)#redistribute igrp 100 metric 1RouterA(config-router)#^Z
Now we will configure the same type of redistribution, but this time we will use the default-metric command:
RouterA#conf tEnter configuration commands, one per line. End with CNTL/Z.RouterA(config)#router ripRouterA(config-router)#default-metric 1RouterA(config-router)#redistribute igrp 100RouterA(config-router)#^Z

QUESTION 248:

You are using multiple routing protocols in different Autonomous Systems (AS) within the Certkiller network domain. You need to redistribute between the systems. You are using two-way redistribution but you want to ensure a loop free topology. Which action should help you avoid routing loop issues?

A. Manually configuring the administrative K-factor
B. Manually configuring the administrative distance
C. Manually configuring the static routes
D. Manually configuring the default gateway
E. None of the above

Answer: B

Explanation:

If a boundary router is running multiple IP routing protocols, then it may be possible that the router will learn about the same network from more than one routing protocol.
The administrative distance of a routing protocol provides a rate of trustworthiness as a source of routing information. Administrative distance is an integer from 0 to 255. The lowest administrative distance has the highest trust rating. An administrative distance of 255 means the routing information source cannot be trusted at all and should be ignored. An administrative distance of zero is reserved for directly connected interfaces and will always be preferred.
Specifying administrative distance values enables the Cisco IOS software to discriminate between sources of routing information. If two routes have the same network number, and possibly subnet information, the IOS software always picks the route whose routing protocol has the lowest administrative distance
distance command is used to modify the default administrative distance.
The distance command can also be applied with optional arguments to make changes to selected routes based on where they originate. The expanded syntax of the distance command is as follows:
Router(config-router)#distance weight [source-ip-address source-mask (access-list-number | name)]

QUESTION 249:

You need to implement route redistribution on a number of the Certkiller OSPF routers. Which three statements are true when configuring redistribution for OSPF? (Select three)

A. Subnets redistribute by default.
B. The default metric is 20.
C. The default metric is 10.
D. The default metric type is 1.
E. Subnets do not redistribute by default.
F. The default metric type is 2.

Answer: B, E, F

Explanation:
When redistributing into OSPF, it is not required to set a metric. If a metric is not stated, the default will be 20. It is a good idea to set the metric so all of the redistributed routes do not have the same metric.
Because OSPF is a classless routing protocol, it accepts routes that are not of classful boundaries. In order for OSPF to accomplish this, you need to append the keyword subnets to the end of the redistribution line. The subnets keyword allows for classless routes to be redistributed into OSPF. Without the subnets keyword, OSPF will accept only classful routes and will reject classless routes
OSPF defines two external metric types: Type 1 and Type 2. By default, routes are redistributed into OSPF as Type 2. From the perspective of the router computing the cost to a route external to the OSPF autonomous system, a Type 1 external route (marked with

E1 in the IP routing table) has a metric that is the sum of the internal OSPF cost, all the way back to the ASBR, and the external redistributed cost. From the viewpoint of the same router, a Type 2 external route (marked with E2 in the IP routing table) has a metric equal only to the redistributed cost that the ASBR originally advertised, with no premium for pathways back internal to the ASBR

QUESTION 250:

The Certkiller multi-protocol network is shown below:

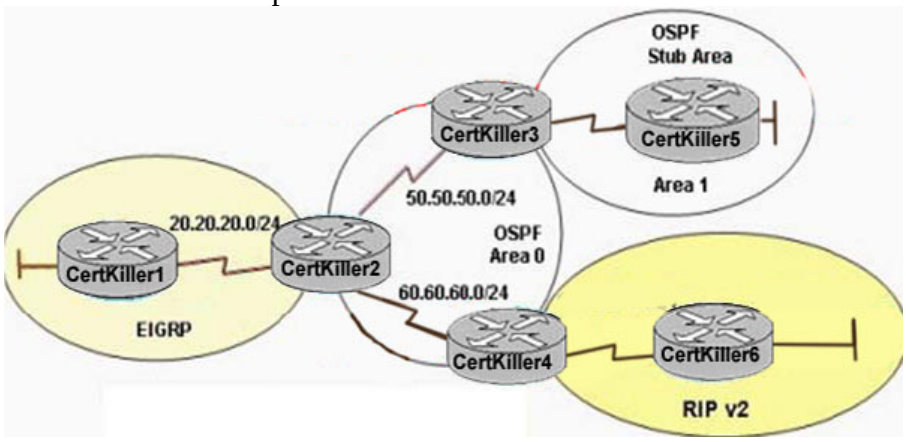

Configuration exhibit:

```
CertKiller2#show running-.config
!
router eigrp 100
  redistribute ospf 1 route-map CertKiller
  network 20.0.0.0
  default-metric 1500 1 255 1 1500
  no auto-summary
!
  router ospf 1
  network 50.50.50.0 0.0.0.255 area 0
  network 60.60.60.0 0.0.0.255 area 0
!
route-map CertKiller permit 10
  match route-type external
```

Refer to the exhibit and the partial configuration on router Certkiller 2. On router Certkiller 4 all RIP routes are redistributed into the OSPF domain. A second redistribution is configured on router Certkiller 2 using a route map. Based on the configuration on router Certkiller 2, which EIGRP external routes will be present in the routing table of Certkiller 1?

A. The routes originating from the RIP routing domain
B. All OSPF inter and intra-area routes
C. The routes originating from the OSPF stub area
D. All routes originating from RIP and OSPF routing domains
E. None of the above

Answer: A

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.
Route maps operate similar to access lists, by examining one line at a time and when a match is found, action is taken. Route maps are different from numbered access lists because they can be modified without changing the entire list. Each route map statement is given a number. If a sequence number is not specified, the first route map condition will automatically be numbered as ten (10). The second condition will automatically be numbered as 20, and so on. The optional sequence number can be used to indicate the position that a new route map is to have in the list of route maps already configured with the same name.
In Exhibit Access-list is created to deny from 50.50.50.0 and 60.60.60.0 (OSPF Domain) and that is called by route-map Certkiller .
While redistributing OSPF routes into EIGRP Certkiller rout-map is called so it denies to publish the RIP domain network into the EIGRP.

QUESTION 251:

You need to troubleshoot a route redistribution problem with the Certkiller multi-protocol network. Which three steps are most helpful in verifying proper route redistribution? (Select 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 routers not performing the route redistribution, use the "debug ip routing" command to verify the routing updates from the ASBR.
C. On the routers not performing the route redistribution, use the "show ip protocols" command to verify the routing information sources.
D. 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.
E. On the ASBR router performing the route redistribution, use the "show ip protocol" command to verify the redistribution configurations.

Answer: A, D, E

Explanation:
The following commands are used to verify the route redistribution operation:
1. Know the network topology, particularly where redundant routes exist.
2. Display the routing table of the appropriate routing protocol on a variety of routers in the internetwork using the show ip route command. For example, check the routing table on the ASBR as well as some of the internal routers in each AS.

3. Perform a traceroute on some of the routes that go across the ASs to verify that the shortest path is being used for routing. Make sure that traces are performed on networks where redundant routes exist.
4. If routing problems are encountered, use traceroute and debug commands to observe the routing update traffic on the ASBRs and internal routers.

QUESTION 252:

Router Certkiller 2 connects two different routing domains as shown below:

Certkiller 2 is configured as shown below:

```
router ispf 1
redistribute eigrp 100 metric 20 metric-type 1
network 172.16.0.0 0.0.255.255
```

During the redistribution process configured on Certkiller 2, 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? (Select two)

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

Answer: A, E

Explanation:
To support multiple routing protocols within the same internetwork efficiently, routing information must be shared among the different routing protocols. For example, routes learned from a RIP process may need to be imported into an IGRP process. This process of exchanging routing information between routing protocols is called route redistribution. Such redistribution can be one-way or two-way. One-way routes are where one protocol receives the routes from another. Two-way routes are where both protocols receive routes from each other. Routers that perform redistribution are called boundary routers because they border two or more autonomous systems or routing domains. This section examines route redistribution in detail, including the use of administrative distance, guidelines for redistribution implementation, and issues with redistribution configuration.
The syntax of the redistribute command is as follows:

Router(config-router)#redistribute protocol [process-id] {level-1 | level-1-2 | level-2} [metric metric-value][metric-type type-value] [match {internal | external 1 | external 2}][tag tag-value] [route-map map-tag][weight weight] [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.
Example:
Router A(config)#router ospf 109
Router A(config-router)#redistribute rip subnets
Router A(config-router)#network 130.10.62.0 0.0.0.255 area 0
Router A(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 253:

The Certkiller multi-protocol WAN is shown below:

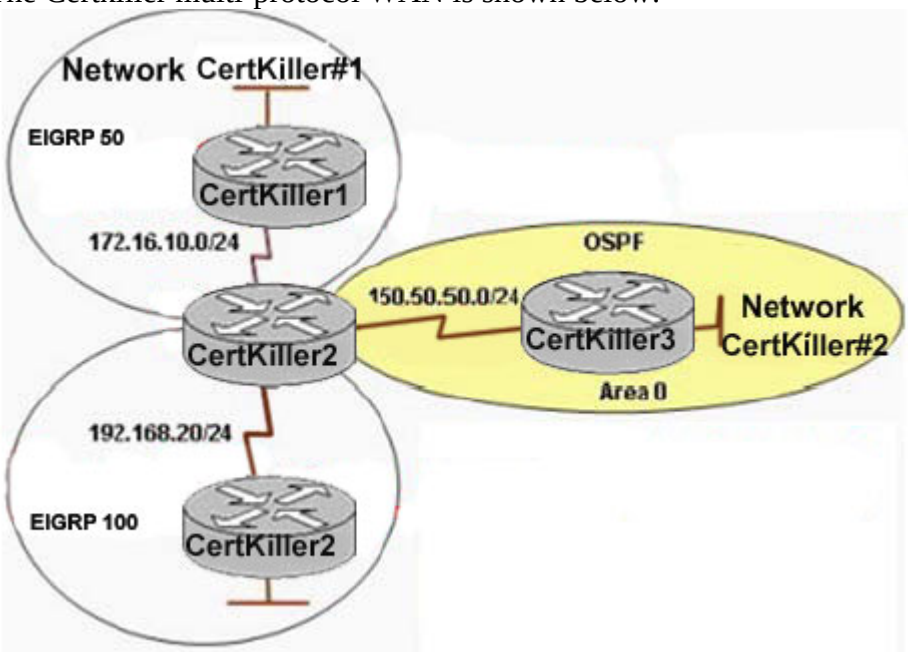

Certkiller 2 configuration exhibit:

```
CertKiller2#show running-config
!
router eigrp 50
 network 172.16.0.0
 no auto-summary
!
router eigrp 100
 redistribute ospf 1
 network 192.168.20.0
 default-metric 1500 1 255 1 1500
 no auto-summary
!
router ospf 1
 network 150.50.50.0 0.0.0.255 area 0
!
```

Study the exhibits above carefully. The routing protocols EIGRP and OSPF have

been configured as indicated in the exhibit. Given the partial configuration of router Certkiller 2, which network will be present in the routing table of Certkiller 4?

A. Network Certkiller #1 and Network Certkiller #2
B. neither Network Certkiller #1 nor Network Certkiller #2
C. Network Certkiller #2
D. Network Certkiller #1
E. None of the above

Answer: C

Explanation:
To support multiple routing protocols within the same internetwork efficiently, routing information must be shared among the different routing protocols. For example, routes learned from a RIP process may need to be imported into an IGRP process. This process of exchanging routing information between routing protocols is called route redistribution. Such redistribution can be one-way or two-way. One-way routes are where one protocol receives the routes from another. Two-way routes are where both protocols receive routes from each other. Routers that perform redistribution are called boundary routers because they border two or more autonomous systems or routing domains. This section examines route redistribution in detail, including the use of administrative distance, guidelines for redistribution implementation, and issues with redistribution configuration.
The syntax of the redistribute command is as follows:
Router(config-router)#redistribute protocol [process-id] {level-1 | level-1-2 | level-2} [metric metric-value][metric-type type-value] [match {internal | external 1 | external 2}][tag tag-value] [route-map map-tag][weight weight] [subnets]
In Exhibit ospf domain is redistributed into eigrp 100 domain so Network Certkiller 2 will present into Certkiller 4 router.

QUESTION 254:

Refer to the Certkiller network shown below:

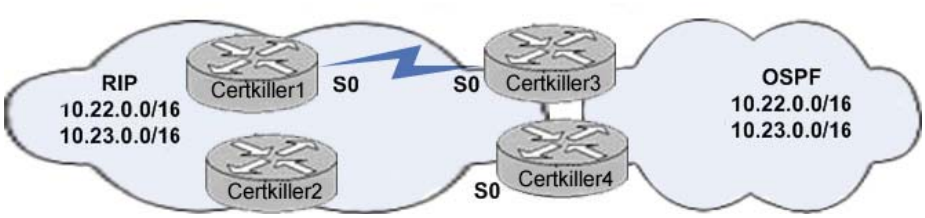

Certkiller 1 and Certkiller 2 belong to the RIP routing domain that includes the networks 10.20.0.0/16 and 10.21.0.0/16. Certkiller 3 and Certkiller 4 are performing two-way route redistribution between OSPF and RIP. A network administrator has discovered that Certkiller 2 is receiving OSPF routes for 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. Set the OSPF default metric to 20.
B. Apply an inbound ACL to the Certkiller 2 serial interface.
C. Configure distribute-lists on Certkiller 3 and Certkiller 4.
D. Change the RIP administrative distance on Certkiller 3 to 110.
E. Change the OSPF administrative distance on Certkiller 3 to 110.
F. None of the above

Answer: C

Explanation:
Use the distribute-list command to pick and choose which routing updates a router will send or receive. By referencing an access list, the distribute-list creates a route filter. This is a set of rules that precisely controls what routes a router will send or receive in a routing update. This command is available for all IP routing protocols and can be applied to either inbound or outbound routing updates. When applied to inbound updates, the syntax for configuring a route filter is as follows:
Router(config-router)#distribute-list access-list-number in [interface-name]
When applied to outbound updates, the syntax can be more complicated as shown in the following:
Router(config-router)#distribute-list access-list-number
out [interface-name | routing-process | as-number]
The routing-process and as-number options are invoked when exchanging routes between different routing protocols.

QUESTION 255:

Refer to the exhibits shown below:
Network topology exhibit:

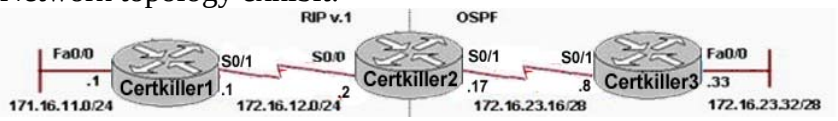

Certkiller 1 configuration exhibit:

```
Certkiller1#show ip route

172.16.0.0/24 is subnetted,2 subnets
C    172.16.12.0 is directly connected, Serial0/0
C    172.16.11.0 is directly connected, FastEthernet0/0
```

Certkiller 2 configuration exhibit:

```
Certkiller2#show run
!
route.rip.
network 172.16.0.0
redistribute ospf 1
default-metric 1
!
router ospf 1
network 172.16.23.16 0.0.0.15 area 0
redistribute rip subnets
default-metric 30
```

Certkiller 3 configuration exhibit:

```
Certkiller3# show ip route
    172.16.6.0/16 is variable connected,4 subnets2 masks
O E2  172.16.12.0 [110/20] via 172.16.23.17, 00:00:28, Serial0/1
O E2  172.16.11.0 [110/20] via 172.16.23.17, 00:00:28, Serial0/1
C     172.16.23.16/28 is directly connected, Serial0/1
C     172.16.23.32/28 is directly connected, FastEthernet0/0
```

RIP and OSPF are configured on the routers as shown in the exhibits. Certkiller 2 is configured with a two-way redistribution between RIP and OSPF domains. All routers can ping each other, but Certkiller 1 cannot see any of the OSPF routes in its routing table. What could the problem be?

A. The process of redistribution of RIP into OSPF does not require any metric conversion, so there is no need to define the metric using the default-metric command during the redistribution.
B. The metric for the OSPF routes that are redistributed into RIP is too low, a fact that prevents OSPF routes from being advertised into RIP.
C. OSPF and RIP use the same major network 172.16.0.0. Therefore, the keyword subnets is not required to redistribute protocols into OSPF.
D. Because OSPF has a longer mask for the same major network than RIP and because RIP version 1 is being used, none of the routes learned from OSPF will be advertised into RIP.
E. None of the above.

Answer: D

Explanation:
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:
Router A(config)#router ospf 109
Router A(config-router)#redistribute rip subnets
Router A(config-router)#network 130.10.62.0 0.0.0.255 area 0
Router A(config-router)#network 130.10.63.0 0.0.0.255 area 0

QUESTION 256:

A Certkiller router is configured for redistribution to advertise EIGRP routes into OSPF on a boundary router as shown below:
router ospf 1
redistribute eigrp 1 metric 25 subnets
Based on the configuration above, what is the function of the "25" in the redistribute command?

A. It specifies the administrative distance on the redistributed routes.
B. It specifies a new process-id to inject the EIGRP routes into OSPF.
C. It specifies the seed cost to be applied to the redistributed routes.
D. It specifies the metric limit of 25 subnets in each OSPF route advertisement.
E. None of the above

Answer: C

Explanation:
Use the distribute-list command to pick and choose which routing updates a router will send or receive. By referencing an access list, the distribute-list creates a route filter. This is a set of rules that precisely controls what routes a router will send or receive in a routing update. This command is available for all IP routing protocols and can be applied to either inbound or outbound routing updates. When applied to inbound updates, the syntax for configuring a route filter is as follows:
Router(config-router)#distribute-list access-list-number in [interface-name]
When applied to outbound updates, the syntax can be more complicated as shown in the following:
Router(config-router)#distribute-list access-list-number out [interface-name | routing-process | as-number]
The routing-process and as-number options are invoked when exchanging routes between different routing protocols.
The distribute-list command can filter any routes in either an outbound or an inbound update globally, or for a specific interface. The Cisco IOS permits one incoming and one outgoing global distribute-list for each routing process. It also permits one incoming and one outgoing distribute-list for each interface involved in a routing process. Keep track of which routing filters are applied globally and which are applied on specific interfaces with the show ip protocols command.

QUESTION 257:

When you're running an Interior Gateway Protocol (like OSPF or EIGRP); why would you want to redistribute these routes into BGP?

A. So BGP can propagate this information to other IGP neighbors.
B. So BGP can propagate this information to other IBGP neighbors.
C. So BGP can propagate this information to other EBGP neighbors.
D. So BGP can propagate this information to other OSPF neighbors.

Answer: C

Explanation:
When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). In order to advertise the IP subnets used within your network to other autonomous systems, these routes must be injected into BGP.
Note: Customer networks usually employ an Interior Gateway Protocol (IGP) such as RIP or OSPF for the exchange of routing information within their networks
Reference: Border Gateway Protocol
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

---

QUESTION 258:

Your network is running BGP as well as EIGRP. You are considering redistributing your BGP routes into your EIGRP. What factors must you consider before redistributing BGP routes to IGP?

A. IGPs are limited to 250 routes
B. A full BGP routing table may contain 100,000+ routes
C. Because of possible routing loops, Cisco router configuration does not allow BGP routes to be redistributed into an IGP.
D. Because BGP routes are not advertised unless they are known by the IGP, Cisco automatically redistributes routes into BGP.

Answer: B

Explanation:
The size of the BGP routing table must be taken into consideration. It is generally not recommended to redistribute BGP routes into your IGP due to the large number of routes that can be obtained from BGP.
Incorrect Answers:
A: There is no such limit.
C: BGP routes can very well be distributed into an IGP.
D: BGP routers are not redistributed automatically.

QUESTION 259:

Why is it sometimes necessary to redistribute IGP protocols like OSPF into BGP?

A. So that BGP can propagate this information to other IGP neighbors.
B. So that BGP can propagate this information to other IBGP neighbors.
C. So that BGP can propagate this information to other EBGP neighbors.
D. So that BGP can propagate this information to other OSPF neighbors.

Answer: C

Explanation:
Customer networks usually employ an Interior Gateway Protocol (IGP) such as RIP or OSPF for the exchange of routing information within their networks. In order to advertise these networks to other network across the Internet, they must be redistributed into BGP.
Reference: Border Gateway Protocol
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

---

QUESTION 260:

It isn't always in your best interests to redistribute dynamically learned routes from an interior gateway protocol to a border gateway protocol. Which of the following

are potential issues that can occur when doing so? (Select two)

A. Routing loops can occur.
B. The routes are automatically summarized.
C. External IGP learned routes might not necessarily have originated in this AS.
D. The BGP process will ignore the external IGP learned routes.

Answer: A, C

Explanation:
If redistribution is used, care must be taken that only local routes are redistributed. For example, routes learned from other autonomous systems (that were learned by redistributing BGP into the IGP) must not be sent out again from the IGP, or rooting loops could result.
Incorrect Answers:
B: With redistribution into exterior protocols such as BGP, no routes are automatically summarized.
D: All routes will be included in the redistribution, including externally learned routes.
Reference: Building Scalable Cisco Networks (Cisco Press) page 408.

QUESTION 261:

A Certkiller router was configured with BGP as shown below:
router bgp 6500
redistribute static
ip route 164.20.0.0 255.255.0.0 null 0
Based on this information, is this a recommended router configuration?

A. Yes. It allows BGP to advertise the 164.20.0.0 /16 network.
B. Yes. It results in all traffic for all subnets of 172.16.0.0 being dropped at this router.
C. No. Cisco prefers that you use the aggregate-address command to distribute IGP routes into BGP.
D. Yes. Cisco prefers this method of redistributing IGP routes into BGP over using the network command.

Answer: C

Explanation:
Redistribution of static routes configured to the null 0 interface into BGP is done to advertise aggregate routes rather than specific routes from the IP table. However, Cisco recommends the use of the aggregate-address-command instead.

QUESTION 262:

You are in the process of redistributing EIGRP into another routing protocol on your network. Which of the following commands would you execute if your

intention was to use altered administrative distance (AD) parameters for redistributing EIGRP into the second routing protocol?

A. default-metric eigrp metric
B. distance eigrp administrative-weight
C. distance eigrp internal-distance external-distance
D. distance eigrp external-distance internal-distance

Answer: C

Explanation:
The distance eigrp command is used to allow the use of two administrative distances---internal and external---that could be a better route to a node.
Syntax: distance eigrp internal-distance external-distance
Note: Use thedistance eigrp command if another protocol is known to be able to provide a better route to a node than was actually learned via external Enhanced IGRP or if some internal routes should really be preferred by Enhanced IGRP.
Incorrect Answers:
A: Incorrect usage of the default-metric command. Furthermore, this command would not be of use here.
B, D: This is the incorrect use of the distance eigrp command.

QUESTION 263:

Which three of the following IOS commands could an administrator use to verify route redistribution? (Select three)

A. debug
B. traceroute
C. show summary
D. show ip route
E. ipconfig

Answer: A, B, D

Explanation:
A: The debug command can be used to debug redistribution.
B: We can verify connectivity, and the presence of a route, with the traceroute command.
D: We can verify that the routes have been redistributed with the show ip route command. The routes will be shown.
Incorrect Answers
C: The show summary command displays a summary of relationships among owners, content rules, and services. It is not of help in this scenario.
E: This is a command used by PC hosts to verify that the IP stack is working properly.

QUESTION 264:

A Certkiller OSPF router has been configured in the following manner:
router ospf 1
redistribute eigrp 1 metric 33 subnets
Based on this information, what function does the 33 parameter in the redistribute command serve?

A. It specifies the metric cost to be applied to the redistributed routes.
B. It specifies the administrative distance on the redistributed routes.
C. It specifies the metric limit to 33 subnets in each OSPF route advertisement.
D. It specifies the process-id for the pseudo process that injects the EIGRP routes into OSPF.

Answer: A

Explanation: It specifies the metric cost to be applied to the redistributed routes. In this example, a metric of 33 will be applied to all of the EIGRP routes when redistributed into OSPF. The "subnets" keyword specifies that subnet mask information is to be preserved during this process.
Reference: Redistributing Routing Protocols
http://www.cisco.com/warp/public/105/redist.html#examples

QUESTION 265:

If you had a static route configured on a router, and that static route had to be advertised to other routers in the network; which one of the following statements would be true?

A. The router automatically advertises static routes to RIP routers.
B. You should configure redistribution using the redistribute command.
C. You should enable static advertisements using the static routes advertise command.
D. You should include the static route in a distribution list using the distribute-list command.

Answer: B.

Explanation:
If you want a router to advertise a static route in a routing protocol, you will need to redistribute it into a dynamic routing protocol.
Incorrect Answers:
A: Static routes are not automatically redistributed when they are configured. Static routes must be redistributed manually.
C: There is no static route advertise command.
D: This is used for filtering inbound and outbound routes.

Reference: Catherine Paquet and Diane Teare, "Building Scalable Cisco Networks" (Cisco Press 2001), p 465.

QUESTION 266:

The Certkiller network is displayed in the diagram below:

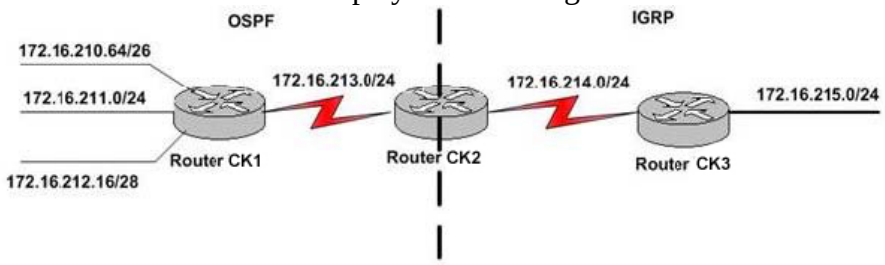

Assuming that CK2 is configured for mutual redistribution; which of the routes below should be present on Router CK3 's routing table? (Select four)

A. 172.16.211.0/24
B. 172.16.213.0/24
C. 172.16.214.0/24
D. 172.16.215.0/24
E. 172.16.210.64/26
F. 172.16.212.16/28

Answer: A, B, C, D

Explanation:
A: This is a classful route.
B: This route is redistributed from OSPF into IGRP.
C: This route is learned through IGRP.
D: This is a directly connected route.
Incorrect Answers:
E, F: These are subnetted routes. IGRP is a classful routing protocol. These routes will not be redistributed into IGRP.
Reference: http://www.cisco.com/warp/public/105/52.html
Redistributing Between Classful and Classless Protocols: EIGRP or OSPF into RIP or IGRP

QUESTION 267:

In order to enable a RIP network to communicate with an OSPF network, redistribution is configured on the Certkiller network. What kind of router would be used to redistribute RIP into an OSPF network?

A. ABR
B. ASBR
C. Internal router

D. Backbone router

Answer: B

Explanation:
External route summarization is specific to external routes that are injected into OSPF via redistribution. Only ASBRs can summarize external routes. These types of routes cannot be summarized by any other router type.

QUESTION 268:

One of the Certkiller routers is configured to redistribute the interior IP networks into BGP. What are two potential consequences that can occur when dynamically learned routes from IGP protocols get redistributed into BGP? (Select two)

A. Routing loops can occur.
B. The IGP routing table is reduced.
C. External IGP learned routes might not necessarily have originated in this AS.
D. Route processing is done using process switching instead of cache switching.

Answer: A, C

Explanation:
It is fairly common to redistribute IGP routes (such as Enhanced IGRP, IGRP, IS-IS, OSPF, and RIP routes) into BGP. However, precautions should be made when doing so. Some of your IGP routes might have been learned from BGP (C), so you need to use access lists to prevent the redistribution of routes back into BGP, or else routing loops can occur (A). In addition, when these IGRP routes are redistributed, all of the networks in the IGP routing table will be redistributed, including those that were learned externally.
Reference: Using the Border Gateway Protocol for Interdomain Routing
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/icsbgp4.htm

---

QUESTION 269:

What's the preferred method of route redistribution, when two routes of different protocols get exchanged? (Select two)

A. Use one way route redistribution when there is one path.
B. Use one way route distribution when there are multiple paths.
C. Use static routes when there are multiple paths.
D. Use two way route distribution when there is one path.
E. Use two way route redistribution where there are multiple paths.
F. Use static routes when there is one path.

Answer: B, C

Explanation:
B. One way redistribution- To avoid routing loops and problems with varying convergence times, allow routes to be exchanged in only one direction, not both directions. In the other direction, you should consider a default route.
C. When you want to prevent routing loops - Many companies have large enough networks that redundant paths are prominent. In some cases, for example, when a path to the same destination is learned from two different routing protocols, you may want to filter the propagation of one of the paths.
Reference: Building Scalable Cisco Networks (Cisco Press) page 472

QUESTION 270:

You are the administrator of a network and are planning on configuring route redistribution. The network you plan on working on is running different routing protocols, and the routers are connected together with redundant links. If your goal was to allow traffic to take the best route, what would you do first?

A. Define the default metric.
B. Identify the boundary router(s).
C. Determine which routing protocol is running in the core router(s).
D. Determine which routing protocol is running in the edge router(s).
E. None of the above

Answer: A

Explanation:
Because different protocols use different metrics for finding the best path to a destination, the first step in configuring mutual redistribution is to define the default metric that will be used as a baseline standard.

QUESTION 271:

You have a network with multiple routing protocols running in different Autonomous Systems, redistributed together with two-way redistribution. What could you do to prevent routing loops?

A. Manually configuring the static routes.
B. Manually configuring the default gateway.
C. Manually configuring the administrative k-value.
D. Manually configuring the administrative distance.

Answer: D

Explanation:

If you manually configure the administrative distance, you can easily decide which protocol the router will choose for each route, thus preventing the possibility of routing loops.
Incorrect Answers:
A: Manually configuring static routes could cause additional routing issues, since static routes are preferred over the dynamically learned routes they could override the information provided by the redistribution.
B: Manual routing configurations should be avoided when used with dynamically learned routes that are being redistributed.
C: These values are used only be IGRP and EIGRP, so they may not apply at all when redistributed into other routing protocols.

QUESTION 272:

Router CK1 is configured as shown below:
```
router igrp 100
network 197.135.20.0
network 197.135.24.0
network 197.135.27.0
redistribute rip
default-metric 10 100 255 1 1500
distance 140 0.0.0.0 255.255.255.255 9
access-list 9 permit 197.135.20.0
access-list 9 permit 197.135.24.0
access-list 9 permit 197.135.27.0
```
Which of the statements below correctly describe the configuration above? (Select two)

A. Networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 are allowed into the routing table.
B. The RIP learned routes to networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 will be assigned an administrative distance of 140.
C. The IGRP learned routes to networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 will be assigned an administrative distance of 140.
D. Changing the administrative distance to a number larger than the default value makes networks 197.135.20.0, 197.135.24.0, and 197.135.27.0 unreachable.

Answer: A, C

Explanation:
We are redistributing RIP into IGRP. The redistribute rip command specifies that routes learned via RIP will be advertised in the IGRP updates. All of the routes specified in the network statements will be allowed into the routing table. No route filtering is taking place with this command. The access lists are used only to define which networks will have their administrative distances changed when redistributed. In this example, the 3 networks in access list 9 will have their AD changed to 140 when redistributed.

We examine the following command:
distance 140 0.0.0.0 255.255.255.255 9
The 140 defines the administrative distance that specified routes will be assigned.
The 0.0.0.0 255.255.255.255 defines the source address of the router supplying the routing information, in this case any router.
The 9 defines the access-list to be used to filter incoming routing updates to determine which will have their administrative distance changed.
And one of the access-list statements:
access-list9 permit 197.135.27.0
9 is the access-list number.
permit allows all networks that match the address to be permitted, in this case to have their administrative distance changed.
197.135.27.0 A network to be permitted, in this case to have its administrative distance changed.
Incorrect Answers:
B: RIP routes are redistributed into IGRP; not vice versa. However, it is not clear that RIP even knows about this route, since the RIP configuration was omitted. We can, however, be certain than IGRP knows about it since these networks were added to the IGRP routing process.
D: In this case we are changing the AD to 140 for the specified networks. Although the AD has been changed to a value that is higher than the default, they are still reachable.

QUESTION 273:

DRAG DROP
You are a network administrator at Certkiller and you're in the process of migrating their networks from RIP to OSPF (you're using the private address of 10.0.0.0 for routing). You want to redistribute the RIP routes into OSPF and assign a metric of 60 to them. Drag the commands on the left onto the corresponding boxes on the right.

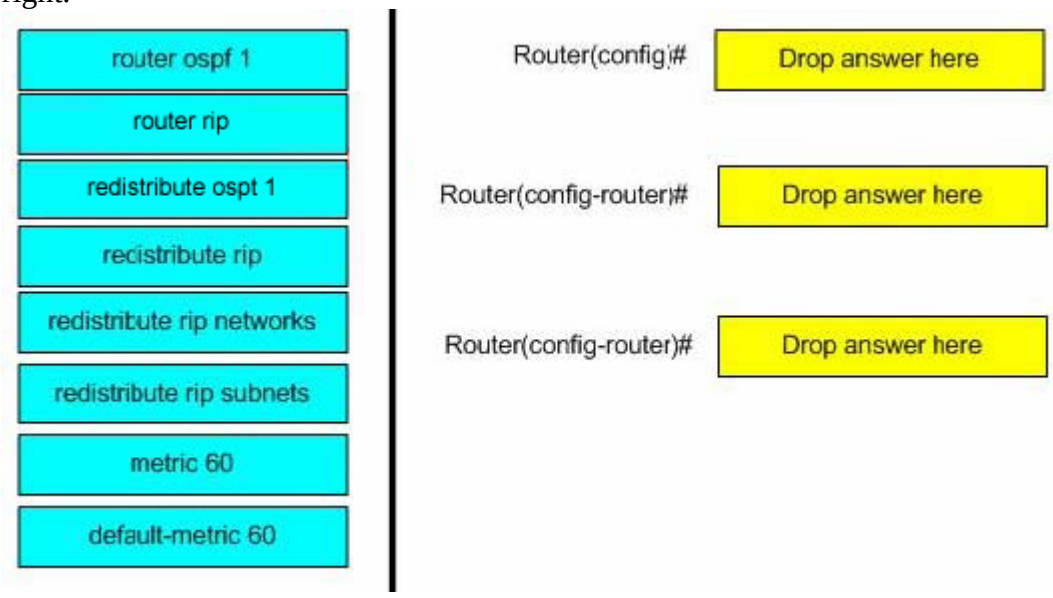

Answer:

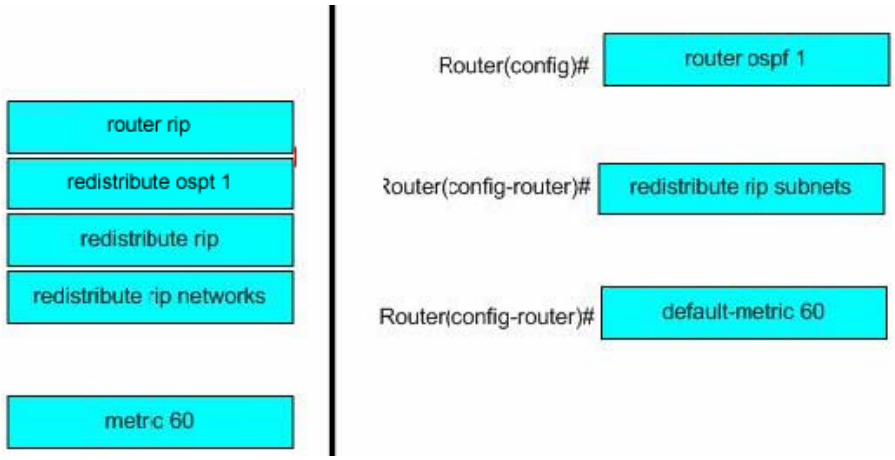

QUESTION 274:

A Certkiller router is configured for route redistribution as shown below:
ip route 30.0.0.0 255.0.0.0 172.16.1.2
ip route 192.168.1.0 255.255.255.0 172.10.1.2
router eigrp 100
redistribute ospf 100
network 172.19.1.0
router ospf 108
redistribute static
redistribute eigrp 100
network 172.16.1.0 00.0.25 area 0
distribute-list 5 out static
access-list 5 permit 30.0.0.0 0.255.255.255
In the above configuration, what is the function of the command line distribute-list 5 out static?

A. It denies the route to 30.0.0.0 via OSPF
B. It denies the route to 30.0.0.0 via EIGRP
C. It propagates the route to 30.0.0.0 via OSPF
D. It propagates the route to 30.0.0.0 via EIGRP

Answer: C

Explanation:
The "distribute-list 5 out static" command filters routes learned from static entries by using access list 5, before those routes are passed to the OSPF process. In this example, the static route to the 30.0.0.0/8 network matches access list 5, so it is permitted to be redistributed into the OSPF protocol.
Incorrect Answers:
A, B: This access-list permits, not denies, routes.
D: The route is propagated via OSPF, not via EIGRP.

QUESTION 275:

You are configuring redistribution to advertise EIGRP routes into OSPF on a boundary router. Given the configuration:
router ospf 1
redistribute eigrp 1 metric 25 subnets
Which is the function of the subnets in the redistribute command?

A. It specifies subnetted routes should be advertised into OSPF.
B. It specifies subnetted routes should be advertised out of OSPF.
C. It specified routes that will be summarized on the 25-bit boundary.
D. It specifies a limit of 25 subnets for each OSPF route advertisement.
E. None of the above.

Answer: A

Explanation:
The subnets keyword is used to specify that subnet mask information used by classless routing protocols should be preserved when redistributed into the other routing protocol. In this case, the EIGRP routes, including subnet masks, will be redistributed into OSPF.
Incorrect Answers:
B: Here the routes are being placed into the OSPF routing process, not advertised out of.
C, D: In this configuration, the value of "25" means that the metric will be 25 when placed into OSPF. It has nothing to do with the bit boundary or the number of subnets that can be advertised.

QUESTION 276:

The Certkiller network is in the process of changing the routing protocol from EIGRP to OSPF, as well as changing the IP network to the private 10.0.0.0/8 network. To do this, router CK1 has been configured as shown below:
router ospf 100
redistribute eigrp 100 metric 100 metric-type 1
network 172.16.0 0.0.0.255.255
All other configurations use the default values.
You want to successfully redistribute all networks and subnets on the Certkiller network. What can you do to accomplish this goal? (Choose two)

A. Change the OSPF process-id number from 100 to 1 in the router ospf 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. Add the network 10.0.0.0 0.255.255.255 command under router ospf 100.
F. Change the metric to an EIGRP compatible metric value for Bandwidth, Delay, Reliability, Load, MTU, (such as 64 1000 100 1 1500) in the redistribute command.

Answer: D, E

Explanation:
The use of the "subnets" keyword will ensure that the subnet mask information will be redistributed into the OSPF process. Since both EIGRP and OSPF support the use of VLSM, this is recommended.
Since the network is being migrated to OSPF, it would be best to begin advertising the network via the OSPF process, as shown in answer choice E.

QUESTION 277:

You are using multiple protocols in different Autonomous Systems (AS). You need to redistribute between the systems. You are using two-way redistribution.
Which action should help you avoid routing loop issues?

A. Manually configuring the static routes.
B. Manually configuring the default gateway.
C. Manually configuring the administrative K-factor.
D. Manually configuring the administrative distance.

Answer: D

Explanation:
If you manually configure the administrative distance, you can manually control which router will choose which route thus preventing the possibility of routing loops.

QUESTION 278:

SIMULATION
Network topology exhibit

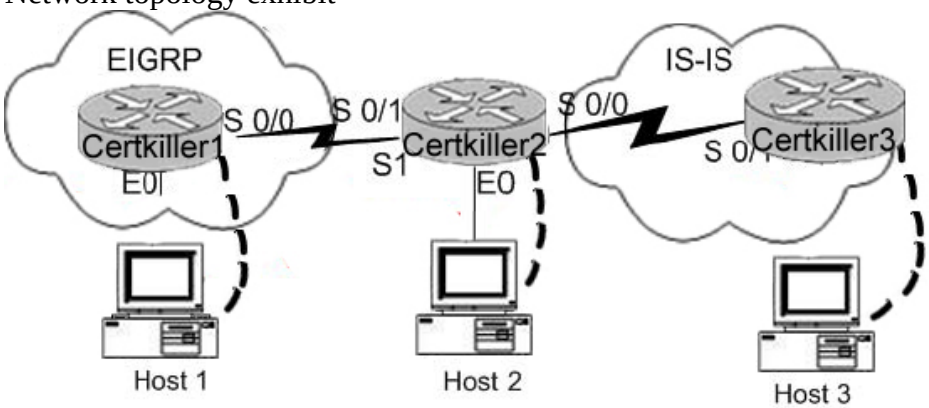

Certkiller .com recently completed a merger with Foo Inc. The two companies have been using separate routing protocols on their companies corporate networks, and an immediate solution is required for the two companies to begin sharing data. A boundary router, Certkiller 2, has been established to perform mutual redistribution of route information between the two networks. Configure route redistribution from

EIGRP into IS-IS and from IS-IS into EIGRP on the boundary router per the following requirements.
1. Seed metric for EIGRP must have the following characteristics:
Bandwidth=64 Kbps
2. Delay=100
3. Reliability=255
4. Load=1
5. MTU=1500
Seed metric for IS-IS must be set to 50.
Only redistribute Level-1 IS-IS routes into EIGRP.
Redistribute EIGPR routes into IS-IS as Level-1 routes.
Name: Certkiller 1
S0/0: 10.10.10.5/30
Loopback1: 1.1.1.1/32
Secret Password: Certkiller
Name: Certkiller 2
S0/0: 192.168.1.6/30
Loopback1: 2.2.2.2/32
Secret Password: Certkiller
Name: Certkiller 3
S0/0: 192.168.1.5/30
Loopback1: 3.3.3.3/32
Secret Password: Certkiller

Answer:

Explanation:
Certkiller 2> en
Certkiller 2# config t
Certkiller 2(config-router)# router isis
Certkiller 2(config-router)# redistribute eigrp 100 level-1 metric 50
Certkiller 2(config-router)# exit
Certkiller 2(config)# router eigrp 100
Certkiller 2(config-router)# redistribute isis level-1 metric 64 10 255 1 1500
Certkiller 2(config-router)# redistribute connected
Certkiller 2(config-router)# ^z
Certkiller 2# copy running-config start up-config.
Alternative #1:
1. Seed metric for EIGRP must have the following characteristics:
Bandwidth=64 Kbps
2. Delay=2 ms
3. Reliability=255
4. Load=1
5. MTU=1500
Name: Certkiller 1
S0/0: 10.197.197.5/30

Loopback1: 1.1.1.1/32
Secret Password: Certkiller
Name: Certkiller 2
S0/0: 192.168.191.6/30
Loopback1: 2.2.2.2/32
Secret Password: Certkiller
Name: Certkiller 3
S0/0: 192.168.191.5/30
Loopback1: 3.3.3.3/32
Secret Password: Certkiller
Alternative #2:
1. Seed metric for EIGRP must have the following characteristics:
Bandwidth=512 Kbps
2. Delay=5 ms
3. Reliability=255
4. Load=1
5. MTU=1500
6. Seed metric for IS-IS must be set to 24
Name: Certkiller 1
S0/0: 10.117.117.5/30
Loopback1: 1.1.1.1/32
Secret Password: Certkiller
Name: Certkiller 2
S0/0: 192.168.72.6/30
S0/1: 10.117.117.6/30
Secret Password: Certkiller
Name: Certkiller 3
S0/0: 192.168.72.5/30
Loopback1: 3.3.3.3/32
Secret Password: Certkiller

QUESTION 279:

Which three statements are true when configuring redistribution for OSPF? (Select 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

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.
For all redistributed OSPF routes, the default metric type is 2.
When there is a major net that is subnetted, you need to use the keyword subnet to redistribute protocols into OSPF. Without this keyword, OSPF only redistributes major nets that are not subnetted.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/technologies_tech_note09186a008009487e.shtml#ospf

---

QUESTION 280:

Router CK1 is configured with OSPF and EIGRP as shown in the diagram below:

```
router ospf 5
  network 10.0.0.0  0.0.255.255 area 0
!
router eigrp 100
   network 172.10.0.0
   redistribute ospf 5
```

You work as a network technician at Certkiller .com. You are troubleshooting a redistribution of OSPF routes into EIGRP. Study the exhibit carefully.
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: E

Explanation:
A default metric is required to redistribute a protocol into IGRP or EIGRP. Only connected routes and interface static routes can be redistributed without a default metric. If a default metric is not specified, it will default to a metric of 0 (infinity) and will not be advertised.
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_command_reference_chapter09186a008009

QUESTION 281:

Which three statements are true concerning redistributed routes when a default metric is not configured? (Select three)

A. RIP, IGRP, and EIGRP assign a default metric of 0 (infinity) to redistributed routes and will advertise these routes accordingly.
B. RIP, IGRP, and EIGRP assign a default metric 0 (infinity) to redistributed routes and will only advertise these routes if a valid seed metric is configured.
C. IS-IS assigns a default metric of 0 to redistributed routes.
D. IS-IS assigns a default metric of 10 to redistributed routes.
E. OSPF assigns a default metric of 1 for routes from all protocols except BGP, which gets a metric of 20.
F. OSPF assigns a default metric of 20 for routes from all protocols except BGP, which gets a metric of 1.

Answer: B, C, F

Explanation:
A default metric is required to redistribute a protocol into IGRP or EIGRP. Only connected routes and interface static routes can be redistributed without a default metric. If a default metric is not specified, it will default to a metric of 0 (infinity) and will not be advertised
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, as shown in the example above. If no metric is specified for the routes being redistributed into IS-IS, a metric value of 0 is used by default.
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/ CK3 65/technologies_tech_note09186a008009487e.shtml#isis

---

QUESTION 282:

The Certkiller WAN is shown in the following diagram:

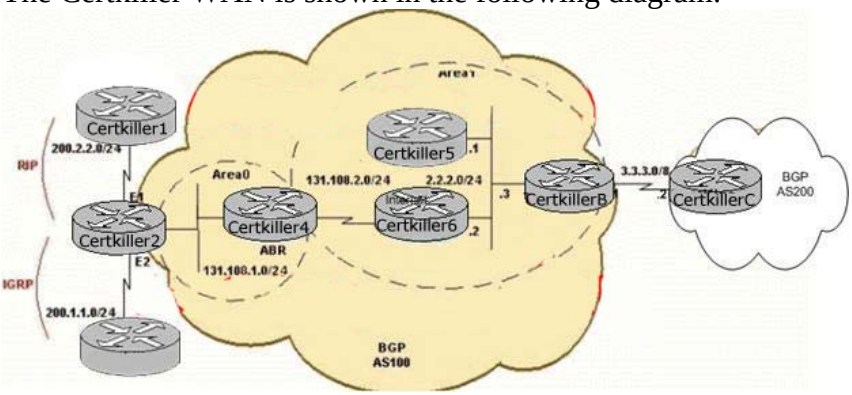

Exhibit, Router Configuration

```
Certkiller2 # show ip route
Gateway of last resort is not set

     2.0.0/24 is subnetted,1 subnets
C            2.2.2.0 is directly connected, Ethernet0/0
C    3.0.0.0/8 is directly connected, Serial1/0
O E2 200.1.1.0/24 [110/20] via 2.2.2.2, 00:16:17, Ethernet0/0
O E1 200.2.2.0/24 [110/104] via 2.2.2.2, 00:00:41, Ethernet0/0
     131.108.0.0/24 is subnetted, 2 subnets
O    131.108.2.0 [110/74] via 2.2.2.2, 00:16:17, Ethernet0/0
O IA 131.108.1.0 [110/84] via 2.2.2.2, 00:16:17, Ethernet0/0
```

Which command should be added to Certkiller B under router bgp 100 to allow only the external OSPF routes to be redistributed to Certkiller C?

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:
Redistribution of Only OSPF External (Type 1 and 2) Routes into BGP
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
Note: The exact scenario and network diagram shown in this question can be found in the reference link below.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/technologies_tech_note09186a00800943c5.shtml

---

QUESTION 283:

The Certkiller routed network is shown below:

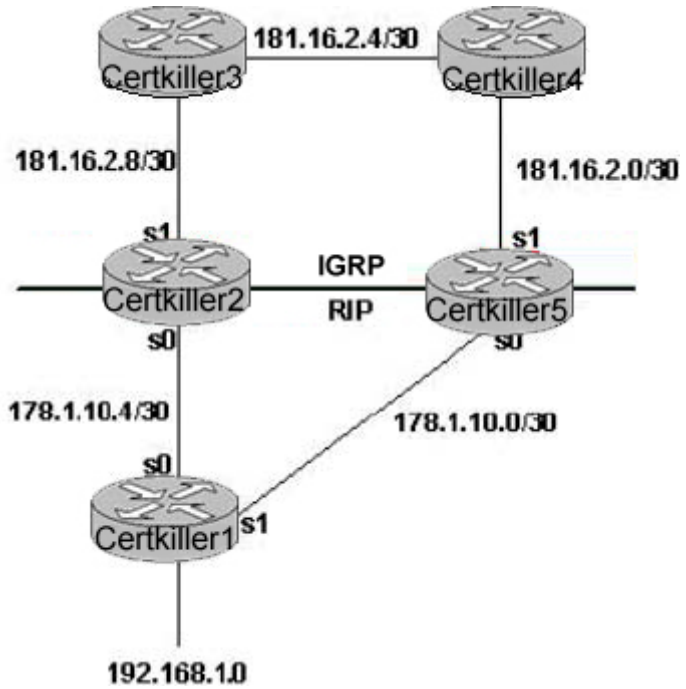

The relevant portion of the configuration of router Certkiller 2 is shown below:

```
hostname Certkiller2
!
router igrp 7
  network 181.16.0.0
  redistribute rip metric 1 1 1 1
  distribute-list 1 in serial 1
!
router rip
  network 178.1.0.0
  redistribute igrp 7 metric 2
!
access-list 1 deny 192.168.1.0
access-list 1 permit any
```

You work as a network technician at Certkiller .com. Study the exhibits carefully.
Both Certkiller 5 and Certkiller 2 are performing two-way IGRP/RIP route redistribution.
What behavior will Certkiller 2 use with the given Certkiller 2 configurations?

A. Updates about 192.168.1.0 that are learned through RIP will be ignored.
B. Updates about 192.168.1.0 will be advertised out the s0 interface.
C. Updates about 192.168.1.0 that are heard through the s1 interface will be ignored.
D. Updates about 192.168.1.0 that are learned through IGRP will be added to the routing table, but not advertised to Certkiller 1.

Answer: C

Explanation:
When a distribute list is defined, the access list that it is tied to is used to define the routing updates that are permitted. In our example, the access list specifies that only the 192.168.1.0 network is to be filtered, while all other network advertisements are allowed. Since the distribute list specifies that this filter is to be used for incoming updates on interface serial 1, choice C is correct.

QUESTION 284:

Router CK1 is configured as shown below:
```
router rip
network 10.0.0.0
!
router ospf 5
network 172.10.0.0 0.0.255.255 area 0
redistribute rip
```
Based on the exhibited configuration commands, which statement is true?

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 not be advertised.

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. The default metric type for OSPF is external, type 2.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/technologies_tech_note09186a008009487e.shtml#ospf

---

QUESTION 285:

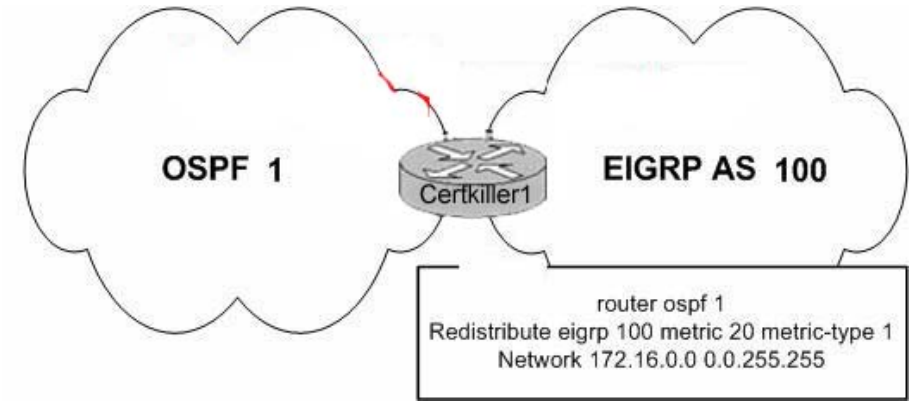

During the redistribution process configured on Certkiller A, some of the EIGRP routes, such as 10.1.1.0/24 and 10.2.2.0/24 are not being redistributed into the OSPF routing domain. Which two items could be a solution for this problem? (Choose two.)

A. Change 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 1 to 100 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, MTUs) in the redistribute command.

Answer: C, D

Explanation:
To support multiple routing protocols within the same internetwork efficiently, routing information must be shared among the different routing protocols. For example, routes learned from a RIP process may need to be imported into an IGRP process. This process of exchanging routing information between routing protocols is called route redistribution. Such redistribution can be one-way or two-way. One-way routes are where one protocol receives the routes from another. Two-way routes are where both protocols receive routes from each other. Routers that perform redistribution are called boundary routers because they border two or more autonomous systems or routing domains. This section examines route redistribution in detail, including the use of administrative distance, guidelines for redistribution implementation, and issues with redistribution configuration.
The syntax of the redistribute command is as follows:
Router(config-router)#redistribute protocol [process-id] {level-1 | level-1-2 | level-2} [metric metric-value][metric-type type-value] [match {internal | external 1 | external 2}][tag tag-value] [route-map map-tag][weight weight] [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.
Router A(config)#router ospf 109
Router A(config-router)#redistribute rip subnets
Router A(config-router)#network 130.10.62.0 0.0.0.255 area 0
Router A(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 286:

The Certkiller WAN is depicted below:

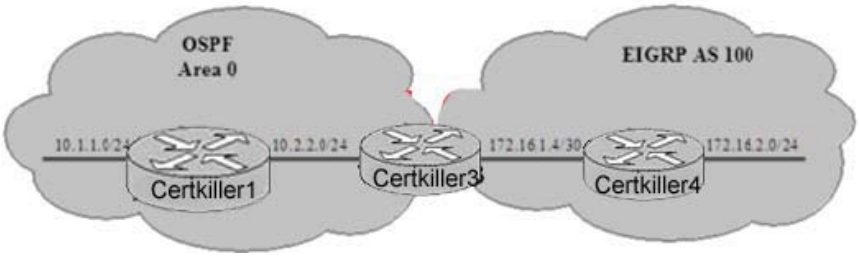

Certkiller 3 is redistributing the EIGRP routes into OSPF. What will the EIGRP routes appear in the routing table of Certkiller 1?

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

Answer: C

Explanation:
O E1 or O E2. The routes in this LSA are external to the autonomous system. They can be configured to have one of two values. E1 will include the internal cost to the ASBR added to the external cost reported by the ASBR. E2 does not compute the internal cost - it just reports the external cost to the remote destination.

QUESTION 287:

The Certkiller network is redistributing routing information from OSPF and EIGRP. Which three steps are most helpful in verifying proper route redistribution? (Select 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 ABSR, 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.
Incorrect Answers:
D: Issuing this command on a non-redistribution router will not tell us where and how the routes are originating from. This command will only be useful on the redistributing routers.
E: This command can not be used to verify the redistributed routes.

QUESTION 288:

The Certkiller network administrator has configured route filters on a router. What are three reasons to control routing updates via route filtering? (Select three)

A. For easier implementation
B. For simple security
C. To control network overhead on the wire
D. To hide certain networks from the rest of the organization
E. To prevent adjacencies from forming

Answer: B, C, D

Explanation:
Policy-based routing is a means by which administrators can implement routing that strays from the standard routing laid out in destination-based routing protocols. Destination-based routing protocols will route based on the shortest path to a destination. Policy-based routing allows an administrator to determine where they want to route traffic. Some reasons to implement policy-based routing are as follows. Keep in mind, these are not the only reasons administrators might base their decisions on:
* Traffic may need to be routed based on the type of traffic. For instance, non-interactive traffic, such as e-mail, may need to be routed over slower connections, and interactive traffic may need to be routed over faster connections.
* Traffic may need to be load-balanced in a way that differs from how the standard routing protocols would handle load balancing.
* Traffic may need to be prioritized based on the type of traffic it is.
* Traffic may need to be classified based on the source network or based on some other method. You would need to implement this type of policy- based routing when traffic from one network needs precedence over another network's traffic.

QUESTION 289:

Two Certkiller RIP routers are connected together as shown below:

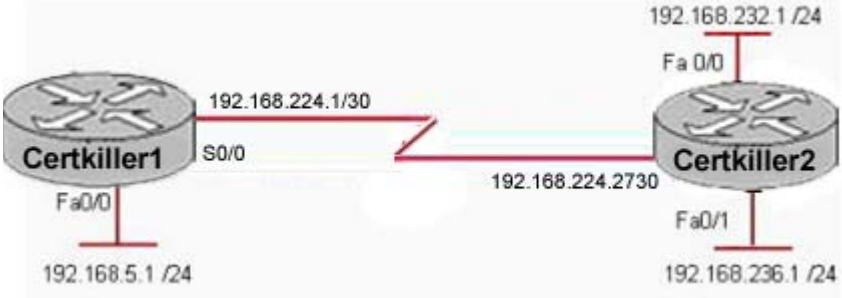

Certkiller 1 configuration exhibit:

```
Certkiller 1 (config)# router rip
Certkiller 1 (configure-router) # version 2
Certkiller 1 (configure-router) # network 192.168.224.0
Certkiller 1 (configure-router) # network 192.168.232.0
Certkiller 1 (configure-router) # network 192.168.236.0
```

Certkiller 2 configuration exhibit:

```
Certkiller2 (configu)# router rip
Certkiller2 (config-router)# version 2
Certkiller2 (config-route) # network 192.168.224.0
Certkiller2 (config-route) # network 192.168.5.0
```

Study the exhibit carefully. Which configuration command sequence would prevent the propagation of updates for network 192.168.236.0/24 to the Certkiller 1 router?

A. Certkiller 1(config)# router rip
Certkiller 1(config-router)# passive-interface s0/0
B. Certkiller 1(config)# access-list 1 deny 192.168.236.0
Certkiller 1(config)# access-list 1 permit any
Certkiller 1(config)# router rip
Certkiller 1(config-router)# distribute-list 1 out serial 0/0
C. Certkiller 1(config)# access-list 1 deny 192.168.236.0
Certkiller 1(config)# access-list 1 permit any
Certkiller 1(config)# interface serial 0/0
Certkiller 1(config-if)# ip access-group 1 in
D. Certkiller 2(config)# router rip
Certkiller 2(config-router)# passive-interface s0/0
E. Certkiller 2(config)# access-list 1 deny 192.168.236.0
Certkiller 2(config)# access-list 1 permit any
Certkiller 2(config)# interface serial 0/0
Certkiller 2(config-if)# ip access-group 1 out
F. Certkiller 2(config)# access-list 1 deny 192.168.236.0
Certkiller 2(config)# access-list 1 permit any
Certkiller 2(config)# router rip
Certkiller 2(config-router)# distribute-list 1 out serial 0/0

Answer: F

Explanation:
Use the distribute-list command to pick and choose which routing updates a router will send or receive. By referencing an access list, the distribute-list creates a route filter. This is a set of rules that precisely controls what routes a router will send or receive in a routing update. This command is available for all IP routing protocols and can be applied to either inbound or outbound routing updates. When applied to inbound updates, the syntax for configuring a route filter is as follows:
Router(config-router)#distribute-list access-list-number in [interface-name]
When applied to outbound updates, the syntax can be more complicated as shown in the following:
Router(config-router)#distribute-list access-list-number out [interface-name | routing-process | as-number]

QUESTION 290:

Study the exhibit below carefully. What is the correct configuration to enable router Certkiller 4 to exchange RIP routing updates with router Certkiller 1 but not with router Certkiller 3?

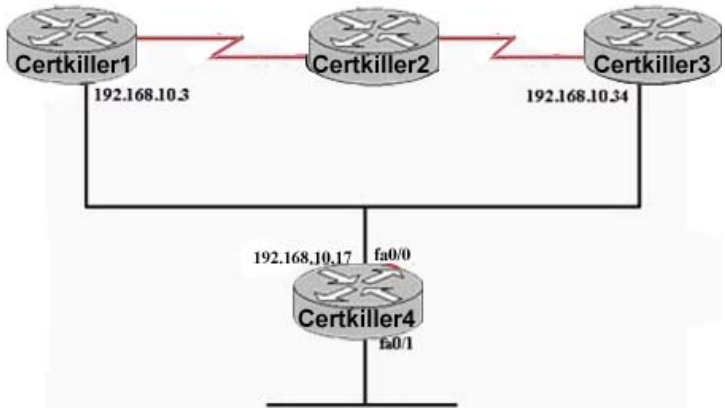

A. Certkiller 4(config)# router rip
Certkiller 4(config-router)# neighbor 192.168.10.3
Certkiller 4(config-router)# passive-interface fa0/0
B. Certkiller 4(config)# interface fa0/0
Certkiller 4(config-if)# neighbor 192.168.10.3
Certkiller 4(config-if)# passive-interface fa0/0
C. Certkiller 4(config)# router rip
Certkiller 4(config-router)# neighbor 192.168.10.34 no broadcast
Certkiller 4(config-router)# passive-interface fa0/0
D. Certkiller 4(config)# interface fa0/0
Certkiller 4(config-if)# neighbor 192.168.10.3
Certkiller 4(config-if)# passive-interface 192.168.10.34

Answer: A

Explanation:
Configure LAN interfaces as passive interfaces when enabling a routing protocol. A passive interface receives updates, but does not send them. The passive-interface command can be used with all IP interior gateway protocols. It can be used with RIP, IGRP, EIGRP, OSPF, and IS-IS. The syntax of the command is as follows:
Router(config-router)#passive-interface type number
The passive-interface command works differently with the different IP routing protocols that support it. In OSPF, the network address of the passive interface appears as a stub network. OSPF routing information is neither sent nor received by way of a passive interface. In EIGRP and OSPF, the router stops sending Hello packets on passive interfaces. When this happens, the router cannot form neighbor adjacencies. Therefore, the router cannot send and receive routing updates on the interface. It will be seen later in this module that the passive effect can be achieved for an EIGRP interface by using the distribute-list command. This can be done without preventing adjacency relationships.

QUESTION 291:

Router CK4 is configured to filter some routes when advertising to its neighbors. Which statement about route filtering is true when it is used to control routing updates?

A. Routes to be filtered are selected by using the distribute-group command.
B. Routes to be filtered are selected using only extended access lists.
C. Routes to be filtered are selected by using distribute lists.
D. Only outbound routes can be filtered.
E. None of the above.

Answer: C

Explanation:
Use the
distribute-list command to pick and choose which routing updates a router will send or receive. By referencing an access list, the distribute-list creates a route filter. This is a set of rules that precisely controls what routes a router will send or receive in a routing update. This command is available for all IP routing protocols and can be applied to either inbound or outbound routing updates. When applied to inbound updates, the syntax for configuring a route filter is as follows:
Router(config-router)#distribute-list access-list-number in [interface-name]
When applied to outbound updates, the syntax can be more complicated as shown in the following:
Router(config-router)#distribute-list access-list-number out [interface-name | routing-process | as-number]

QUESTION 292:

The Certkiller WAN is shown below

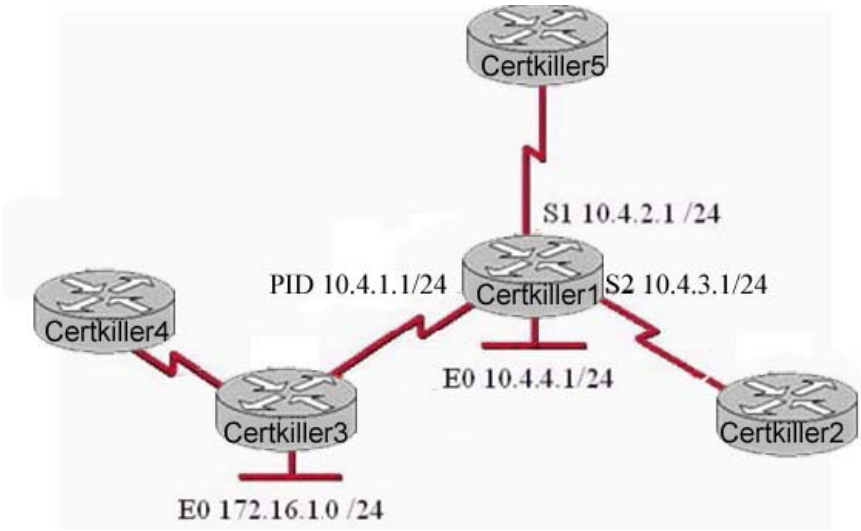

Certkiller 1 configuration exhibit:

```
Certkiller1(config)#router rip
Certkiller1(config-router)#network 10.0.0.0
Certkiller1(config-router)#distribute-list 44 in interface BRI0
Certkiller1 (config-router)#exit
Certkiller1 (config)#access-list44 deny 172.16.1.0 0.0.0.255
Certkiller1 (config)#access-list44 permit any
```

What are the effects of this RIP configuration on router Certkiller 1? (Select two)

A. User traffic from the 172.16.1.0 network is denied by access-list 44
B. No routing updates will be sent from router Certkiller 1 on interface BRI0 to router Certkiller 3
C. The routing table on router Certkiller 1 will be updated with the route to router Certkiller 4
D. The route to network 172.16.1.0 will not be entered into the routing table on router Certkiller 1
E. Router Certkiller 1 will not advertise the 10.0.0.0 network to router Certkiller 3

Answer: C, D

Explanation:
Use the distribute-list command to pick and choose which routing updates a router will send or receive. By referencing an access list, the distribute-list creates a route filter. This is a set of rules that precisely controls what routes a router will send or receive in a routing update. This command is available for all IP routing protocols and can be applied to either inbound or outbound routing updates. When applied to inbound updates, the syntax for configuring a route filter is as follows:
Router(config-router)#distribute-list access-list-number in [interface-name]
When applied to outbound updates, the syntax can be more complicated as shown in the following:
Router(config-router)#distribute-list access-list-number out [interface-name | routing-process | as-number]

The distribute-list command can filter any routes in either an outbound or an inbound update globally, or for a specific interface. The Cisco IOS permits one incoming and one outgoing global distribute-list for each routing process. It also permits one incoming and one outgoing distribute-list for each interface involved in a routing process. Keep track of which routing filters are applied globally and which are applied on specific interfaces with the show ip protocols command.

QUESTION 293:

Two Certkiller routers are connected as shown below:

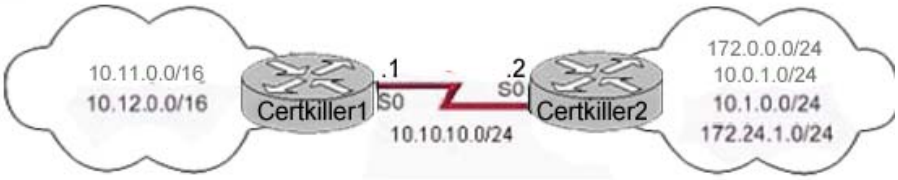

Router Certkiller 1 is configured as follows:

```
Hostname Certkiller1
!
router rip
 version 2
 network 10.0.0.0
distribute-list 10 in serial0
!
access-list 10 permit 10.0.0.0.0.0.255.255
```

Based on the information shown above, what is the effect of the "distribute-list" command in the Certkiller 1 configuration?

A. Certkiller 1 will filter the 10.1.0.0/24 and the 172.24.1.0/24 routes from the Certkiller 2 RIP updates.
B. Certkiller 1 will not filter any routes because there is no exact prefix match.
C. Certkiller 1 will filter only the 172.24.1.0/24 route from the Certkiller 2 RIP updates.
D. Certkiller 1 will permit only the 10.0.0.0/24 route in the Certkiller 2 RIP updates.
E. None of the above

Answer: A

Explanation:
Use the distribute-list command to pick and choose which routing updates a router will send or receive. By referencing an access list, the
distribute-list creates a route filter. This is a set of rules that precisely controls what routes a router will send or receive in a routing update. This command is available for all IP routing protocols and can be applied to either inbound or outbound routing updates.
When applied to inbound updates, the syntax for configuring a route filter is as follows:
Router(config-router)#distribute-list access-list-number in [interface-name]
When applied to outbound updates, the syntax can be more complicated as shown in the following:
Router(config-router)#distribute-list access-list-number out [interface-name | routing-process | as-number]

The distribute-list command can filter any routes in either an outbound or an inbound update globally, or for a specific interface. The Cisco IOS permits one incoming and one outgoing global distribute-list for each routing process. It also permits one incoming and one outgoing distribute-list for each interface involved in a routing process. Keep track of which routing filters are applied globally and which are applied on specific interfaces with the show ip protocols command.

QUESTION 294:

You are a network analyst at Certkiller and are currently in the process of analyzing two autonomous systems, each running a different routing protocol, but connected between each other with redundant paths. The junior administrator wants to prevent routing loops between the two autonomous systems and asks you if there's an IOS feature that could help him. How would you answer?

A. Route filtering.
B. Passive interfaces.
C. Static redistribution.
D. Two-way redistribution.

Answer: A

Explanation:
Multiple autonomous systems or routing domains can share route information through the redistribution process. Proper implementation of redistribution requires route filters to prevent feedback loops from forming. It is strongly recommended that redistribution between multiple ASs and multiple routing protocols be accompanied by route filters.
Reference: CCNP #640-503 Building Scalable Cisco Networks (Cisco Press), More EIGRP Scalability Rules

---

QUESTION 295:

Routing loops are becoming a problem on a particular network. What can an administrator do to reduce the occurrence of routing loops on a redistributed network?

A. Use multiple default gateways for redundancy.
B. Use one-way redistribution for greater stability.
C. Use two-way redistribution for greater stability.
D. Use overlapped routing protocols for redundancy.
E. None of the above.

Answer: B

Explanation:
One-way redistribution would help avoiding the routing loops problem.

Incorrect Answers:
A: Multiple default gateways would increase the risk of routing loops.
C: Two-way redistribution would increase the risk of routing loops.
D: Using several routing protocols would increase the risk of routing loops.

QUESTION 296:

Which one of the following BGP prefix lists would you use if you wanted to deny the default route 0.0.0.0?

A. ip prefix-list abc deny 0.0.0.0/0
B. ip prefix-list abc permit 0.0.0.0/32
C. ip prefix-list abc deny 255.255.255.255/0
D. ip prefix-list abc permit 255.255.255.255/32
E. None of the above

Answer: A

Explanation:
To deny the default route 0.0.0.0/0 use the following command:
ipprefix-list abc deny 0.0.0.0/0
Incorrect Answers:
B, D: We are required to deny it, not permit it.
C: This will effectively deny all routes, not just the default route.

QUESTION 297:

Which of the following prefix-list statements would you enter if you wanted to permit all prefixes that fall between /8 and /24 on the 192.0.0.0 network?

A. ip prefix-list abc permit 192.0.0.0/8 ge 8 le 24
B. ip prefix-list abc permit 192.0.0.0/8 ge 24 le 8
C. ip prefix-list abc permit 192.0.0.0/24 ge 24 le 8
D. ip prefix-list abc permit 192.0.0.0/24 ge 8 le 24

Answer: A

Explanation:
When used with prefix lists, the "ge" means "greater than or equal to" and the term "le" means "less than or equal to." Greater than (ge) 8 and less than (le) 24.
An 8 bit subnet mask should also be used, since we are looking to match all prefixes beginning on the entire 192.X.X.X range.

QUESTION 298:

You have a router that's been configured with multiple IP routing protocols, and

you're interested in checking out if and how inbound and outbound routing updates are being filtered. Which one of the following commands would list these filters?

A. show ip
B. show ip route
C. show ip protocols
D. show ip interface
E. show protocol filters

Answer: C

Explanation:
The command show ip protocols command is used to display the parameters and current state of the active routing protocol process. The information is presented on a routing protocol basis and includes applied inbound and outbound filters.
Incorrect Answers:
A: Show ip is an incomplete command.
B: The show ip route command displays active routes, not information on filters.
D: The show ip interface command lists a summary of an interface's IP information and status. However, it does not list the filters applied on a routing protocol basis.
E: This is an invalid command.

QUESTION 299:

DRAG DROP
You are a systems administrator of one of the Certkiller EIGRP networks. Your goal is to suppress Certkiller 's router from sending updates out of its Serial 0/0 interface without compromising its ability to form and maintain neighbor adjacencies on the same interface.
The router has already been configured as follows:
Router(config)#access-list 20 deny any
Router(config)#access-list 100 permit ip any any
Router(config)#router eigrp 1
Drag the correct answers on the left to the correct command line space on the right, to finish configuring the router.

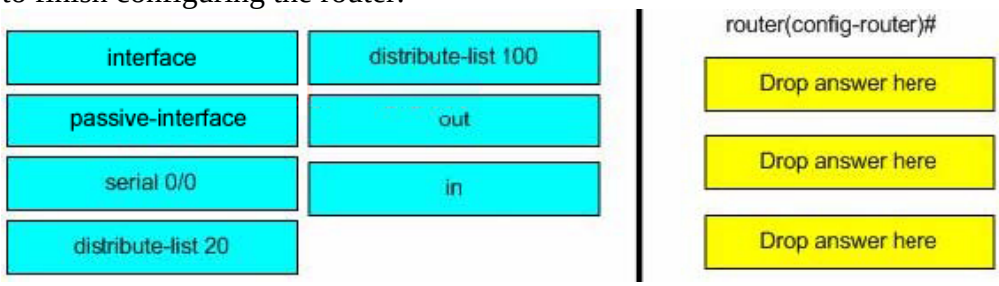

Answer:

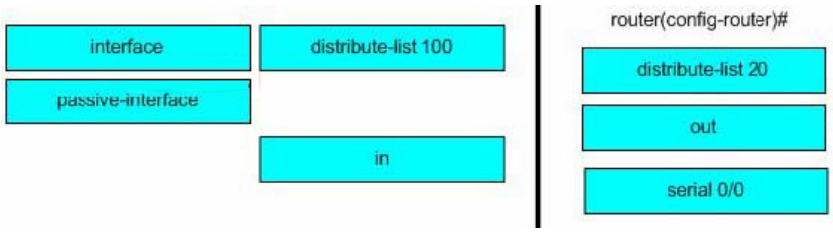

Explanation:
We do not want to allow the routing updates along serial 0/0 to go out. This distribute list will deny all IP networks from being advertised out the interface.

QUESTION 300:

Your OSPF router has one serial interface, and one Ethernet LAN interface. The sub-interface is configured in the following manner:
interface serial 0.122 point-to-point
ip address 192.168.1.1 255.255.255.0
encapsulation frame-relay
frame-relay interface-dlci 122
Your want to allow hosts on your LAN to send and receive data, but you don't want routing traffic to go through that interface. Which of the following commands should you use to complete your configuration?

A. interface serial 0.122 point-to-point
passive-interface ethernet 0
B. interface ethernet 0
ip address 192.168.12.1 255.255.255.0
passive-interface
C. router ospf 172
area 1 nssa
network 192.168.1.0 0.0.0.255 area 0
network 192.168.12.0 0.0.0.255 area 1
D. router ospf 172
passive-interface ethernet 0
network 192.168.1.0 0.0.0.255 area 0
network 192.168.12.0 0.0.0.255 area 1

Answer: D

Explanation:
We use the passive-interface command to configure the ethernet interface to be passive. The passive-interface router configuration command is used to disable sending routing updates on an interface.
Incorrect Answers:
A: We are not configuring the serial interface. Furthermore, the passive-interface command is a router configuration command, not an interface configuration command.
B: We cannot use the passive-interface command like this.

C: We should configure the Ethernet interface as passive, not the area as a not-so-stubby area (NSSA).

QUESTION 301:

DRAG DROP
You have a named Router CertK running EIGRP that has already been configured with the following:
Router CertK (config)#access-list 30 deny any
Router CertK (config)#access-list 40 permit ip any
Router CertK (config)#router eigrp 1
To complete your configuration you want to configure the serial 0/0 interface to prevent routing updates from going out that interface while still allowing the formation and upkeep of neighbor adjacencies on that interface. Drag the correct command phrase from the bottom and place it in the right command line sequence above.

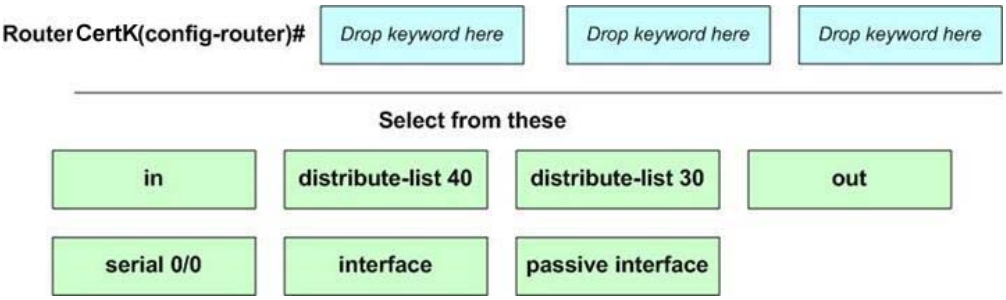

Answer:

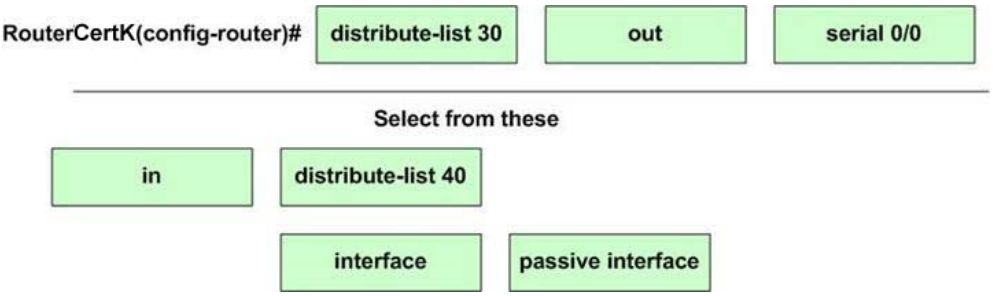

Explanation:
We use distribute list 30, which stops IP traffic, on outward traffic on the serial 0/0 interface.
RouterCK continues receiving routing updates from its neighbor, but the distribute-list prevents routes from being advertised out of serial 0. Furthermore, neighbor adjacencies are allowed to be formed between RouterCK and its neighbor on serial 0/0.
Incorrect Answers:
passiveinterface serial 0/0:
On EIGRP, passive interface causes the router to stop sending and receiving hello packets. This will prevent the interface from maintaining neighbor adjacencies.
distribute-list40 out serial 0/0:
We must stop outgoing traffic, not allow it.
distribute-list 30 in serial 0-0:
We must stop outgoing traffic, not incoming traffic.

Reference: How Does the Passive Interface Feature Work in EIGRP?
http://www.cisco.com/warp/public/103/16.html

QUESTION 302:

The Certkiller network consists of a main office and a single remote office. You need to configure the remote office in using the following guidelines:
1. The main office must learn all of its routes from the regional office.
2. The regional office must not learn routes from the main office.
3. The most scalable solution should be used.
Which of the following answer choices best describe what should be done to satisfy these requirements?

A. Configure static routes pointing to the network behind the central office router
B. Configure a default route pointing to the networks behind the central office router
C. Make the interface that is connected to the central office a passive interface to block incoming updates
D. Enable route update filtering on the interface that is connected to the central office to block incoming updates
E. None of the above will work

Answer: D

Explanation:
The only viable solution here is to enable the main office to dynamically learn about all of the regional office routes, while still not receiving any routes is by enabling incoming route filters.
Incorrect Answers:
A, B: Configuring static routes or default routes on the regional office will not enable the main office to learn about any of the regional office routes.
C: A passive interface would prevent all route updates to the main office. Passive interfaces do not filter incoming routing updates, so the regional office could still learn routes from the main office.
Note: Configuring the main office using a passive interface would satisfy the conditions here, but this question asks us how to configure the regional office, not the main office.

QUESTION 303:

If a router already has a route to the 172.27.0.0/16 network in the routing table, which of the following commands would cause EIGRP to flag that network as the candidate default route?

A. ip default-network 172.27.0.0
B. ip route 0.0.0.0 0.0.0.0 172.27.0.0
C. ip default-network 0.0.0.0
D. default-information originate

E. ip classless

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.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/ CK5 54/technologies_tech_note09186a0080094374.shtml

---

QUESTION 304:

With regard to controlling routing updates with the use of route filtering, which of the following statement is true?

A. Only inbound routes can be filtered.
B. Only outbound routes can be filtered.
C. Routes to be filtered are selected by using distribute lists.
D. Routes to be filtered are selected using only extended access lists.
E. Routes to be filtered are selected by using the distribute-group command.

Answer: C

Explanation:
The syntax for the distribute-list in/out command is:
distribute-listaccess-list-number in/out [interface-name]
where access-list-number is the standard IP access-list against which the contents of the incoming or outgoing routing update are matched. The [interface-name] argument is optional and specifies the interface on which the update is expected. It is important to note that the access-list referred to in access-list-number is applied to the contents of the update, not to the source or destination of the routing update packets. The router decides whether or not to include the contents in its routing table based on the access-lists. For example:
access-list 1 permit 1.0.0.0 0.255.255.255 router rip distribute-list 1 in!--- The distribute-list command is given!--- under the router configuration mode.
Incorrect Answers:
A, B: Distribute lists can be used to filter traffic based on incoming as well as outgoing routes.
D: Standard access-lists are used with distribute lists, not extended access lists.
E: This is an invalid command. The correct syntax is "distribute-list"

QUESTION 305:

Router CK1 has been configured to filter routes. Which of the following are 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 is used to filter out routing updates from other parts of the network, making certain networks unreachable from other parts of the organization. This can be done to simply decrease the amount of traffic overhead associated with routing updates over parts of the network, or for security reasons.
Incorrect Answers:
B: Route filtering requires a more complex router configuration, and can also lead to difficult troubleshooting.
E: Route filtering will only block certain routes from incoming or outgoing, but the routing protocol adjacencies will still be formed between neighboring routers.

QUESTION 306:

You are the network engineer at Certkiller . The router topology for the Certkiller network is shown in the following graphic:

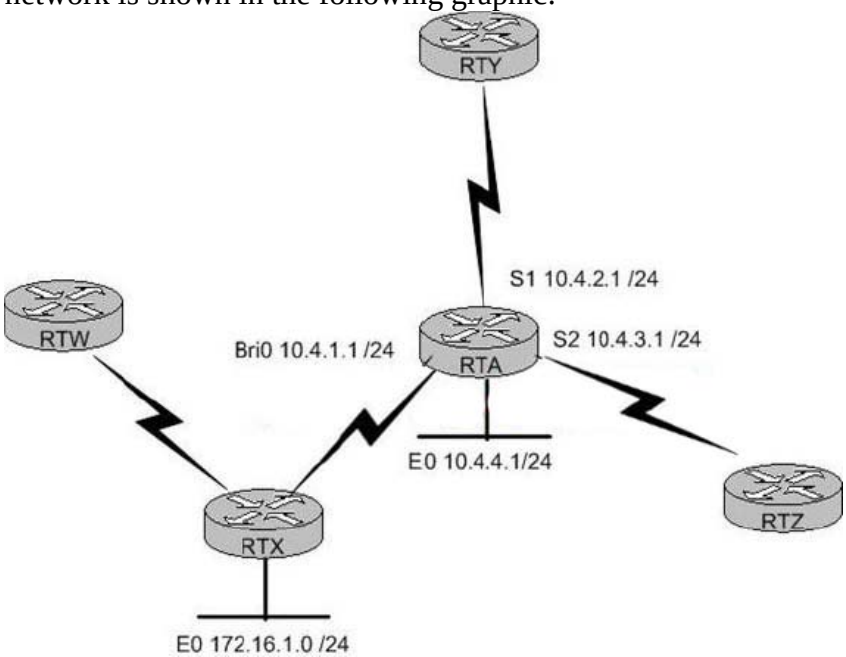

Router RTA is configured as follows:
RTA(config)#router rip
RTA(config-router)#network 10.0.0.0
RTA(config-router)#distribute-list 44 in interface BRI0
RTA(config-router)#exit

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 interface BRI0 on router RTA to router RTX.
B. Router RTA will not advertise the 10.0.0.0 network to router RTX.
C. The route 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:
Answer C is correct as the distribute list calls ACL 4 which denies network 172.16.1.0.
Answer E is correct because RTW networks are not being blocked.
Incorrect Answers:
A, B: Only incoming routes are filtered, not outgoing routes. All known routing entries will be advertised from router A.
D: This is incorrect because this ACL is being called by a distribute-list which only affects route table updates, not user traffic.

QUESTION 307:

The Certkiller WAN is depicted below:

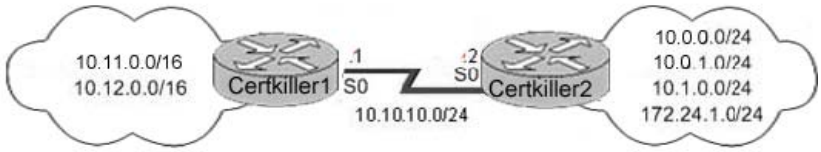

Exhibit

```
hostname Certkiller1
!
router rip
 version 2
 network 10.0.0.0
  distridute-list 10 in serial10
  !
access-list 10 permit 10.01.0.0.255.255
```

Based on the information above, what is the effect of the distribute list command in the Certkiller 1 configuration?

A. Certkiller 1 will filter only the 172.24.1.0/24 route from the Certkiller 2 RIP updates.
B. Certkiller 1 will permit only the 10.0.0.0/24 route in the Certkiller 2 RIP updates.
C. Certkiller 1 will filter 10.1.0.0/24 and the 172.24.1.0/24 routes from the Certkiller 2 RIP updates.
D. Certkiller 1 will not filter any routes because there is no exact prefix match.

Answer: C

Explanation:
The distribute list is used to define the routing updates to be permitted through the interface. In the example above, the filter list specifies that only the 10.0.1.0/16 networks will be permitted via incoming RIP updates through interface serial 0. All other incoming RIP updates, including the 10.1.0.0/24 and 172.24.1.0/24 will be filtered. Since the 10.0.1.0/24 is part of the 10.0.1.0/16 network, it will be allowed.

QUESTION 308:

Five Certkiller routers are connected as shown in the following exhibit:

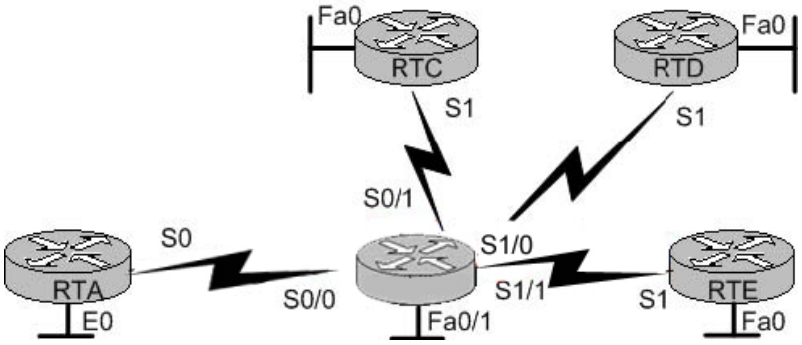

Refer to the exhibit. Assume all routers have been preconfigured to advertise their connected networks in the indicated IGP. Which three configurations would correctly propagate the RTA E0 LAN segment to RTC, but not to RTD or RTE? (Choose three)

A. RTB(config)# router rip
RTB(config-router)# version 1
RTB(config-router)# passive-interface default
RTB(config-router)# no passive-interface s0/0
RTB(config-router)# no passive-interface Fa0/0
B. RTB(config)# router rip
RTB(config-router)# version 2
RTB(config-router)# passive-interface s1/0
RTB(config-router)# passive-interface s1/1
C. RTB(config)# router igrp 123
RTB(config-router)# passive-interface default
RTB(config-router)# no passive-interface s0/1
D. RTB(config)# router eigrp 123
RTB(config-router)# passive-interface default
RTB(config-router)# no passive-interface s0/0
RTB(config-router)# no passive-interface s1/1
E. RTB(config)# router eigrp 123
RTB(config-router)# passive-interface default
RTB(config-router)# no passive-interface s0/0
RTB(config-router)# no passive-interface s0/1
F. RTB(config)# router ospf 1

RTB(config-router)# passive-interface s0/0
RTB(config-router)# passive-interface s0/1

Answer: B, C, E

Explanation:
B: In this answer choice, the connections to RTD and RTE are specified as passive, so no updates will be sent to these routers from router B. The link to RTC, however, is not listed as passive so the routes will be propagated to it.
C, D: In this answer choice, the passive-interface default statement is configured, meaning that any links not specified individually as not being passive will be. Since both of these choices leave the links to RTD and RTE as passive, and RTB explicitly specifies the link to RTC as not being passive, the route will propagate to RTC but not RTD and RTE.

QUESTION 309:

What is the purpose of configuring a router with the "IP Helper address" command?

A. IP Helper is used to direct BOOTP clients to a BOOTP server.
B. IP Helper is used to prevent the router form forwarding IP broadcasts.
C. IP Helper is used to allow IPX clients to communicate with IP-based servers.
D. IP Helper is used to accommodate compatibility routers using different IP routing protocols.

Answer: A

Explanation:
The ip helper-address command is used to have the Cisco IOS software forward User Datagram Protocol (UDP) broadcasts, including BOOTP, received on an interface. DHCP protocol information is carried inside of BOOTP packets. To enable BOOTP broadcast forwarding for a set of clients, configure a helper address on the router interface closest to the client. The helper address should specify the address of the DHCP server.
Note:
A DHCP server can be considered to be a BOOTP server, even though a DHCP server is more advanced.
Incorrect Answers:
B: Combined with the ip forward-protocol global configuration command, the ip helper-address command allows you to control which broadcast packets and which protocols are forwarded. However, the main purpose of the IP helper feature is not to prevent the router from forwarding IP broadcasts.
C: IP helper does not use IPX.
D: This is false.

QUESTION 310:

When you execute the "ip helper-address" command on a router, which three UDP ports get enabled automatically by default? (Select three)

A. 53 (DNS)
B. 69 (TFTP)
C. 515 (LPR)
D. 161 (SNMP)
E. 49 (TACACS)

Answer: A, B, E

Explanation:
To forward the BootP/DHCP request from the client to the DHCP server, the ip helper-address interface command is used. The IP helper-address can be configured to forward any UDP broadcast based on UDP port number. By default, the IP helper-address will forward the following UDP broadcasts:
1. DNS (port 53), time service (port 37)
2. Trivial File Transfer Protocol (TFTP) (port 69)
3. Terminal Access Control Access Control System (TACACS) service (port 49)
4. NetBIOS name server (port 137)
5. NetBIOS datagram server (port 138)
6. Boot Protocol (DHCP/BootP) client and server datagrams (ports 67 and 68)
7. IEN-116 name service (port 42)
Reference: Understanding and Troubleshooting DHCP in Catalyst Switch or Enterprise Networks
http://www.cisco.com/warp/public/473/100.html

QUESTION 311:

While troubleshooting an Internet related problem on the Certkiller network, you enable BGP packet debugging. Which attribute must exist in the BGP update packet?

A. LOCAL_ PREF
B. AGGREGATOR
C. AS_Path
D. Weight
E. None of the above

Answer: C

Explanation:
BGP updates are carried using TCP on port 179. In contrast, RIP updates use UDP port 520, while OSPF does not use a Layer 4 protocol. Because BGP requires TCP, IP

connectivity must exist between BGP peers. TCP connections must also be negotiated between them before updates can be exchanged. Therefore, BGP inherits those reliable, connection-oriented properties from TCP.
An AS_Path attribute is a well-known mandatory attribute, type code 2. It is the sequence of AS numbers that a route has traversed to reach a destination. The AS that originates the route adds its own AS number when sending the route to its external BGP peers. Thereafter, each AS that receives the route and passes it on to other BGP peers will prepend its own AS number to the list. Prepending is the act of adding the AS number to the beginning of the list. The final list has all the AS numbers that a route has traversed with the AS number of the AS that originated the route at the end of the list. This type of AS_Path list is called an AS_Sequence because all the AS numbers are ordered sequentially.
BGP uses the AS_Path attribute as part of the routing updates, update packet, to ensure a loop free topology on the Internet. Each route that gets passed between BGP peers will carry a list of all AS numbers that the route has already been through. If the route is advertised to the AS that originated it, that AS will see itself as part of the AS_Path attribute list and will not accept the route. BGP speakers prepend their AS numbers when advertising routing updates to other autonomous systems, external peers. When the route is passed to a BGP speaker within the same AS, the AS_Path information is left intact.

QUESTION 312:

There are numerous attributes associated with BGP. Which BGP attribute will not be advertised in routing updates to its neighboring routers?

A. Local preference
B. Next hop
C. Origin
D. Weight
E. AS_path
F. None of the above

Answer: D

Explanation:
The Weight attribute is similar to the Local Preference attribute in that it gives higher preference to the route that has a higher weight. The difference is that the weight parameter is local to the router and is not exchanged between routers. The weight parameter influences routes coming from different providers to the same router, one router with multiple connections to two or more providers. The weight parameter has a higher precedence than any other attribute. It is the most important attribute when determining route preference. The Weight attribute is a Cisco proprietary attribute.

QUESTION 313:

Synchronization has been disabled on a Certkiller BGP router. What is correct

about the BGP synchronization command? (Select two)

A. Synchronization must be enabled when implementing a multi-homed BGP connection to multiple ISPs.
B. Synchronization is necessary when peering with an EBGP neighbor.
C. Synchronization improves BGP routing convergence.
D. If it is turned ON, a prefix learned from IBGP neighbor is valid only if a non-bgp (IGP) route exists for that prefix.
E. Synchronization can be turned off if all the transit routers in an Autonomous system running full mesh IBGP.

Answer: D, E

Explanation:
The BGP synchronization rule states that a BGP router should not use, or advertise to an external neighbor, a route that is learned from IBGP unless that route is local or the router learns it from the IGP. In other words, BGP and the IGP must be synchronized before the networks learned from an IBGP neighbor can be used. If an AS passes traffic to another AS, BGP should not advertise a route before all routers in the AS have learned about the route via the IGP. A router learning a route via IBGP waits until the IGP has propagated the route within the AS and then advertises it to external peers. This rule ensures that all routers in the AS are synchronized and are able to route traffic that the AS advertises to other autonomous systems. This approach ensures consistency of routing information (avoids "black holes") within the AS. BGP synchronization is disabled by default in Cisco IOS Software Release 12.2(8)T and later; it was on by default in earlier Cisco IOS software releases. With the default of synchronization disabled, BGP can use and advertise to an external BGP neighbor routes learned from an IBGP neighbor that are not present in the local routing table BGP synchronization is unnecessary in some situations. It is safe to have BGP synchronization off only if all routers in the transit path in the AS are running full-mesh IBGP. Having synchronization disabled allows the routers to carry fewer routes in IGP and allows

---

QUESTION 314:

Certkiller uses BGP to advertise their IP blocks to the Internet. Which statement is true about EBGP?

A. EBGP requires a full mesh.
B. An internal routing protocol can be used to reach an EBGP neighbor.
C. A static route can be used to form an adjacency between neighbors.
D. The next hop does not change when BGP updates are exchanged between EBGP neighbors.
E. None of the above.

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. A static route can be used to form an adjacency between EBGP neighbors.

QUESTION 315:

All of the IBGP routers in the Certkiller network have peer relationships built with every other IBGP router in the Certkiller network. Why should iBGP sessions be fully meshed within a Transit AS?

A. Routes learned via EBGP are never propagated to other IBGP peers.
B. Routes learned via IBGP are never propagated to other IBGP peers.
C. A full mesh allows for optimal routing within the Transit AS.
D. Routes learned via IBGP are never propagated to other EBGP peers.
E. BGP requires redundant TCP sessions between IBGP peers.
F. None of the above.

Answer: B

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 316:

Routers CK1 , CK2 , and CK3 are all IBGP routers. Which statement is true about IBGP routers?

A. They do not need to be directly connected.
B. They must be fully meshed.
C. They can be in a different AS.
D. They must be directly connected.
E. None of the above.

Answer: A

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 317:

Router Certkiller 1 is a BGP router that has had the synchronization feature disabled. What is correct about the BGP synchronization command? (Select two)

A. Synchronization can be turned off if all the transit routers in an Autonomous system running full mesh IBGP.
B. If it is turned ON, a prefix learned from IBGP neighbor is valid only if a non-bgp (IGP) route exists for that prefix.
C. Synchronization improves BGP routing convergence.
D. Synchronization is necessary when peering with an EBGP neighbor.
E. Synchronization must be enabled when implementing a multi-homed BGP connection to multiple ISPs.

Answer: A, B

Explanation:
The BGP synchronization rule states that a BGP router should not use, or advertise to an external neighbor, a route that is learned from IBGP unless that route is local or the router learns it from the IGP. In other words, BGP and the IGP must be synchronized before the networks learned from an IBGP neighbor can be used. If an AS passes traffic to another AS, BGP should not advertise a route before all routers in the AS have learned about the route via the IGP. A router learning a route via IBGP waits until the IGP has propagated the route within the AS and then advertises it to external peers. This rule ensures that all routers in the AS are synchronized and are able to route traffic that the AS advertises to other autonomous systems. This approach ensures consistency of routing information (avoids "black holes") within the AS. BGP synchronization is disabled by default in Cisco IOS Software Release 12.2(8)T and later; it was on by default in earlier Cisco IOS software releases. With the default of synchronization disabled, BGP can use and advertise to an external BGP neighbor routes learned from an IBGP neighbor that are not present in the local routing table BGP synchronization is unnecessary in some situations. It is safe to have BGP synchronization off only if all routers in the transit path in the AS are running full-mesh IBGP. Having synchronization disabled allows the routers to carry fewer routes in IGP and allows

---

QUESTION 318:

Router CK1 is sending BGP information to all of its BGP peers. Which BGP

attribute will not be advertised in routing updates to its neighboring routers?

A. Weight
B. Origin
C. Local preference
D. AS_path
E. Next hop
F. None of the above

Answer: A

Explanation:
The Weight attribute is similar to the Local Preference attribute in that it gives higher preference to the route that has a higher weight. The difference is that the weight parameter is local to the router and is not exchanged between routers. The weight parameter influences routes coming from different providers to the same router, one router with multiple connections to two or more providers. The weight parameter has a higher precedence than any other attribute. It is the most important attribute when determining route preference. The Weight attribute is a Cisco proprietary attribute.

QUESTION 319:

EBGP is being used to advertise the Certkiller IP blocks to the Internet. Which statement is true about EBGP?

A. EBGP requires a full mesh.
B. A static route can be used to form an adjacency between neighbors.
C. The next hop does not change when BGP updates are exchanged between EBGP neighbors.
D. An internal routing protocol can be used to reach an EBGP neighbor.
E. None of the above

Answer: B

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 320:

Every routing protocol has a loop prevention mechanism. Which BGP attribute is used by BGP routers to prevent routing loops?

A. Origin
B. Weight
C. AS-path
D. Next-hop
E. Local preference
F. MED
G. None of the above

Answer: C

Explanation:
An AS_Path attribute is a well-known mandatory attribute, type code 2. It is the sequence of AS numbers that a route has traversed to reach a destination. The AS that originates the route adds its own AS number when sending the route to its external BGP peers. Thereafter, each AS that receives the route and passes it on to other BGP peers will prepend its own AS number to the list. Prepending is the act of adding the AS number to the beginning of the list. The final list has all the AS numbers that a route has traversed with the AS number of the AS that originated the route at the end of the list. This type of AS_Path list is called an AS_Sequence because all the AS numbers are ordered sequentially.
BGP uses the AS_Path attribute as part of the routing updates, update packet, to ensure a loop free topology on the Internet. Each route that gets passed between BGP peers will carry a list of all AS numbers that the route has already been through. If the route is advertised to the AS that originated it, that AS will see itself as part of the AS_Path attribute list and will not accept the route. BGP speakers prepend their AS numbers when advertising routing updates to other autonomous systems, external peers. When the route is passed to a BGP speaker within the same AS, the AS_Path information is left intact.

QUESTION 321:

The Certkiller BGP network is shown below:

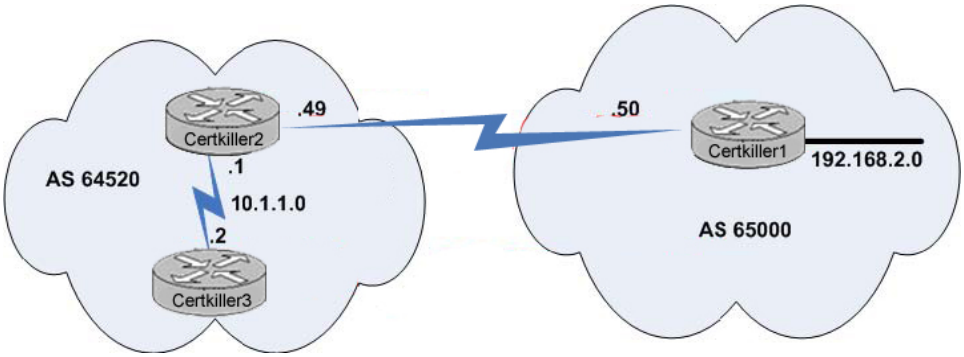

By default, when Certkiller 2 passes BGP advertisements from Certkiller 1 about network 192.168.2.0 to Certkiller 3, what address will be listed as the next-hop address?

A. 10.1.1.2
B. 10.1.1.1

C. 192.168.1.49
D. 192.168.1.50
E. None of the above

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:
1. For EBGP sessions, the next hop is the IP address of the neighbor that announced the route.
2. For IBGP sessions, where routes originated inside the AS, the next hop is the IP address of the neighbor that announced the route. For routes injected into the AS by way of EBGP, the next hop learned from EBGP is carried unaltered into IBGP. The next hop is the IP address of the EBGP neighbor from which the route was learned.
3. When the route is advertised on a multiaccess medium, such as Ethernet or Frame Relay, the next hop is usually the IP address of the interface of the router. This will be the interface connected to the media that originated the route.

QUESTION 322:

Which three of the statements below correctly describe the characteristics of Autonomous Systems in routed networks? (Select three)

A. Within an AS, all routers must run either BGP or IBGP.
B. An AS uses exterior gateway protocols (EGPs) to exchange information with other autonomous systems.
C. An AS is a group of routers under the same technical administration.
D. Within an AS, routes learned through BGP can be redistributed using interior gateway protocols.
E. Within an AS, routes learned through an interior protocol cannot be redistributed using BGP to other autonomous systems.

Answer: B, C, D

Explanation:
1. Exterior Gateway Protocol (EGP) routing protocol used to connect between autonomous systems.
2. The use of the term autonomous system in connection with BGP stresses the fact that the administration of an autonomous system appears to other autonomous systems to have a single coherent interior routing plan, and presents a consistent picture of those networks that are reachable through it.
3. BGP is used between autonomous systems

Incorrect Answers:
A: Within an AS, all routers need not run either interior BGP or exterior BGP. While all IBGP peers need to be fully meshed, all routers within the network do not need to run BGP.
E: One of the major reasons for running BGP is so that the interior public IP networks within an AS can be distributed to other systems in the Internet via BGP.
Reference: Building Scalable Cisco Networks (Cisco Press) page 313

QUESTION 323:

The Certkiller network is displayed in the following diagram:

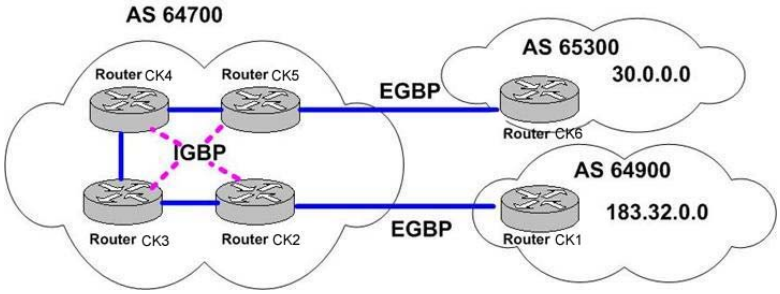

On the assumption that there is no IGP running in AS 64700 and synchronization is OFF, which router in AS 64700 is going to advertise the route to 183.32.0.0?

A. CK2 only
B. CK5 only
C. CK2 and CK5 only
D. CK2 , CK3 , and CK4 only
E. CK2 , CK3 , CK4 , and CK5

Answer: A

Explanation: We need to run IGP on AS 64700 to control the exit points from the AS. When router CK2 receives the EBGP route, it will advertise it to all IBGP routers. However, IBGP routers do not advertise the IBGP learned route unless it is also in their IGP table, or synchronization is turned off.
Reference: Using the Border Gateway Protocol for Interdomain Routing
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/icsbgp4.htm

---

QUESTION 324:

Which of the following statements about IBGP routers are true? (Select one.)

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:
Since BGP runs over TCP, the IBGP speakers need to only be able to reach the other IBGP speakers via a TCP connection across the network and do not need to be directly connected. Even if two IBGP neighbors are separated by other routers, they will still become neighbors as long as they are reachable via TCP.
Incorrect Answers:
A: Within an AS, all IBGP speakers must be configured as a full mesh. The only exceptions to this rule occur when router reflection or confederations are used.
B: IBGP routers need to be configured with the same AS number. The only exception to this rule is when BGP confederations are in use.
C: The routers only need to be reachable via a TCP connection.

QUESTION 325:

Given the following choices, what kind of BGP router advertises routes to other IBGP neighbors?

A. Client
B. EBGP peer
C. Route reflector
D. cluster of clients
E. None of the above

Answer: C

Explanation:
Ordinarily, with no route reflector, a full mesh of IBGP peers is required. Route reflectors modify the BGP split horizon rule by allowing the router configured as the route reflector to propagate routes learned by IBGP to other IBGP peers. Route reflectors reduce the number of BGP neighbor relationships in an AS.

QUESTION 326:

BGP has been configured on some of the Certkiller routers. The BGP routing process relies on two different types of tables. What are they? (Select two)

A. An IP routing table.
B. A BGP topology table.
C. A BGP attribute table.
D. A table that contains BGP information received from and sent to other routers.
E. A combined table that contains both IP routes and BGP information received from and sent to other routers.

Answer: B, C

Explanation:
BGP uses one table for topology and another for attribute. The attribute table assigns values to various attributes (weight, local preference, mult-exit discriminator, origin, AS path, next hop, & community) and the topology table matches the values of these attributes to the various routes it can select.

QUESTION 327:

Attributes are a means of sending additional route information over BGP. Which of the following statements are true regarding BGP attributes? (Select three)

A. MED is an optional attribute.
B. Origin is an optional attribute.
C. Next-hop is an optional attribute.
D. Local Preference is a discretionary attribute.
E. AS-Path is a well-known mandatory attribute.
F. Community is a well-known mandatory attribute.

Answer: A, D, E

Explanation:
A: The Multiple Exit Discriminator (MED) attribute is optional.
D: The local preference attribute is used to prefer an exit point from the local autonomous system (AS). It is a well-known discretionary attribute.
E: When a route advertisement passes through an autonomous system, the AS number is added to an ordered list of AS numbers that the route advertisement has traversed. The AS-Path attribute is mandatory.
Note: BGP metrics are called path attributes. Optional attributes are recognized by some implementations, but are expected not to be recognized by everyone. A well-known mandatory attributes must be present in all update messages.
The attributes defined by BGP include:
Well-known mandatory attributes:
1. AS-path
2. Next-hop
3. Origin
Well-known discretionary attributes:
1. Local preference
2. Atomic aggregate
Optional transitive attributes:
1. Aggregator
2. Communities
Optional non-transitive attribute:
1. Multi-Exit-Discriminator (MED)
Incorrect Answers:
B: The origin attribute indicates how BGP learned about a particular route. It can have

three values: IGP, EGP or incomplete. It is a required attribute.
C: The EBGP next-hop attribute is the IP address that is used to reach the advertising router. It is a required attribute.
F: The community attribute provides a way of grouping destinations, called communities, to which routing decisions can be applied. It is not mandatory.
Reference: Border Gateway Protocol
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

---

QUESTION 328:

The BGP community attribute tags route for the sake of consistency in route filtering. Which of the answer choices below correctly describe the community attribute?

A. Optional and transitive.
B. Optional and non-transitive.
C. Well-known and mandatory.
D. Well-known and discretionary.

Answer: A

Explanation:
The community attribute is an optional transitive attribute that can be in the range 0 to 4,294,967,200. Each network can be a member of more than one community.
The attributes defined by BGP include:
Well-known mandatory attributes:
1. AS-path
2. Next-hop
3. Origin
Well-known discretionary attributes:
1. Local preference
2. Atomic aggregate
Optional transitive attributes:
1. Aggregator
2. Communities
Optional non-transitive attribute:
1. Multi-Exit-Discriminator (MED)

QUESTION 329:

You are a guest lecturer at the Certkiller Academy teaching some CCNP hopefuls about BGP. One of the students, Bob asks you what a BGP community is. How would you respond?

A. Communities are tagged by default in outgoing updates.
B. Communities can only be used within one autonomous system.

C. Communities are a means of tagging routes to ensure consistent filtering.
D. Communities perform summarization of blocks of contiguous network prefixes.

Answer: C

Explanation:
A community is a group of destinations which share some common property. No tag is used by default. Communities are a means of tagging routes to ensure consistent filtering or route-selection policy.
Incorrect Answers
A: By default, all destinations belong to the general Internet community and are not tagged.
B: Each autonomous system administrator may define which communities a destination belongs to. Community information is passed on between different autonomous systems.
D: No summarization is performed by communities. Communications can be aggregated, however.
Reference: RFC 1997, BGP Communities Attribute

QUESTION 330:

The Certkiller IBGP routers are configured in a fully meshed fashion. Which of the following are reasons why IBGP routers should peer with each other within a single AS? (Select two)

A. IBGP routes are not propagated to other EBGP peers.
B. IBGP routes that a router originates are propagated to other IBGP peers.
C. IBGP routes are propagated to other IBGP speakers in the AS that are not peers.
D. IBGP routes that are learned from an IBGP neighbor are propagated to only EBGP peers.

Answer: A, B

Explanation:
IBGP routes are propagated to all IBGP peers and only the IBGP peers. They are not exchanged with EBGP peers, which is why all IBGP routers must be configured in a full mesh.
Note: You can configure Border Gateway Protocol (BGP) either within an autonomous system or between different autonomous systems. When run within an autonomous system, it's called internal BGP (IBGP). When run between different autonomous systems, it's called external BGP (EBGP).

QUESTION 331:

Routers CK1 and CK2 belong to the same BGP peer group. Which two of the statements below are true regarding this BGP peer group? (Select two)

A. The peer group name is passed to other routers in the peer group.
B. A peer group is a group of BGP neighbors with different update policies.
C. The peer group name is only local to the router on which it is configured.
D. A peer group allows options that affect outbound updates to be overridden.
E. A peer group is a more efficient way to update BGP than configuration individual neighbors.

Answer: C, E

Explanation:
C: The peer group name is only local to the router it is configured on, it is not passed to any other router.
E: Neighbors with the same update policies can be grouped into peer groups to simplify configuration and make update calculation more efficient.
Incorrect Answers:
A: A BGP peer group does not necessarily have a name, it might have a number.
B: BGP peers use the same update policy.
D: A peer group does not have this option.
Reference: RFC 1771, A Border Gateway Protocol 4 (BGP-4)

QUESTION 332:

Routers CK1 and CK2 are configured as BGP peers. Which of the following is true regarding this peering relationship? (Select two)

A. Periodic keepalives are used to verify connectivity.
B. Incremental keepalives are used to verify connectivity.
C. It provides a reliable connection between two BGP routers.
D. It provides a "best effort" connection between two BGP routers.

Answer: A, C

Explanation:
A: The default keepalive frequency is 60 seconds.
C: BGP peering provides a reliable connection between BGP routers. BGP utilizes the inherent reliability of TCP, since all BGP sessions are maintained over TCP port 179.
Incorrect Answers:
B: Keepalives messages are not incremental.
D: BGP peers form a reliable connection.
Reference: BGP Peer Groups
http://www.cisco.com/warp/public/459/29.html

QUESTION 333:

Router CK1 is configured as a route reflector. What would happen if router CK1 would receive an update from a peer in a different autonomous system?

A. It discards the update.
B. It sends the update to all IBGP peers.
C. It sends the update only to non clients.
D. It sends the update only to route reflector clients.
E. It sends the update to all routers in the autonomous system.

Answer: B

Explanation:
When any BGP router receives an update from an EBGP peer, it will forward this routing information to all IBGP peers. This is true for all EBGP routers, including route reflectors.
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1828/products_configuration_guide_chapter09186a00800c

---

QUESTION 334:

Which of the following statements is true about BGP peer communications?

A. Communication between BGP peers runs over RIP.
B. Communication between BGP peers runs over TCP.
C. Communication between BGP peers runs over UDP.
D. Communication between BGP peers runs over ICMP.

Answer: B

Explanation:
BGP communicate through a TCP connection. TCP port 179 is reserved for BGP sessions.
Reference: Border Gateway Protocol
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

---

QUESTION 335:

The Certkiller network is considered to be a single autonomous system (AS). Which three of the following statements are characteristics of an autonomous system? (Select three)

A. It uses only Interior Gateway Protocols (IGPs).
B. EGPs are used to connect different autonomous systems.
C. It is a set of routers under a single technical administration.
D. It uses EGPs to route packets to other autonomous systems and IGPs to route packets within the autonomous system.
E. It uses IGPs to route packets to other autonomous systems and EGPs to route packets

within the autonomous system.

Answer: B, C, D

Explanation:
A BGP autonomous system is a group of routers all administered as a single team working on mutual packet switching goals. EGP's (Exterior Gateway Protocol's) handle routing to outside autonomous systems which are on the exterior of that particular autonomous system, and IGP's (Interior Gateway Protocol's) route within, on the interior of autonomous systems.
Incorrect Answers:
A: This is incorrect because an autonomous system can use an exterior gateway protocol.
E: This is incorrect because the roles of IGP's and EGP's are reversed in this answer choice.

QUESTION 336:

Which of the following statements is true about the BGP synchronization command? (Select two)

A. Synchronization must be enabled when implementing a multi-homed BGP connection to multiple ISPs.
B. If it is turned ON, a prefix learned from IBGP neighbor is valid only if a non-BGP (IGP) route exists for that prefix.
C. Synchronization is necessary when peering with an EBGP neighbor.
D. Synchronization improves BGP routing convergence.
E. Synchronization can be turned off if all the transit routers in an Autonomous system are running a fully meshed IBGP.

Answer: B, E

Explanation:
If your autonomous system will be passing traffic through it from another autonomous system to a third autonomous system, it is very important that your autonomous system be consistent about the routes that it advertises. For example, if your BGP were to advertise a route before all routers in your network had learned about the route through your IGP, your autonomous system could receive traffic that some routers cannot yet route. To prevent this from happening, BGP must wait until the IGP has propagated routing information across your autonomous system. This causes BGP to be synchronized with the IGP. Synchronization is enabled by default.
Only if all routers in the transit path in the AS are running BGP it is safe to turn synchronization off.
Reference: Building Scalable Cisco Networks (Cisco Press) page 33
http://www.cisco.com/en/US/products/sw/iosswrel/ps1826/products_configuration_guide_chapter09186a008008

QUESTION 337:

The Certkiller BGP network is shown in the diagram below:

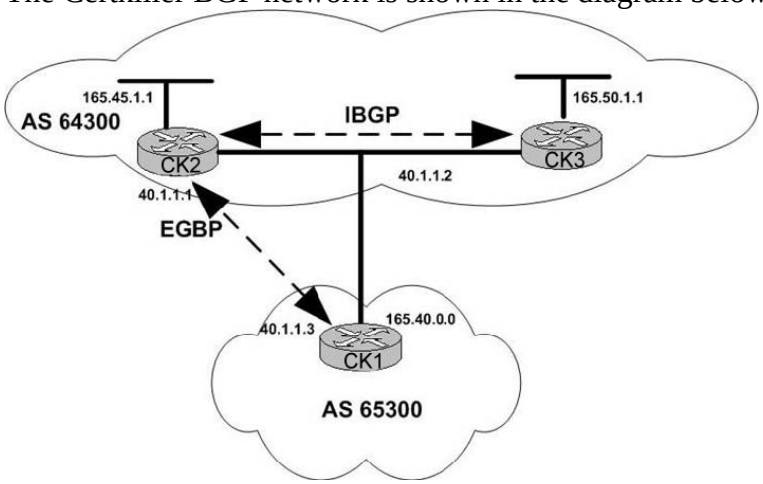

Router CK1 receives BGP routing advertisements from router CK2 about the network 165.50.0.0. By default, what is the value of the next hop attribute?

A. 40.1.1.1
B. 40.1.1.2
C. 40.1.1.3
D. Router CK1 does not accept the advertisement from Router CK2 because Router CK1 is not peering with Router CK3 via BGP.
E. Router CK2 does not advertise network 165.50.0.0 to Router CK1 because the network is not directly connected to Router CK2 .

Answer: A

Explanation:
The BGP next-hop attribute is a well-known mandatory attribute that indicates the next hop IP address that is to be used to reach a destination. For EBGP, the next hop is the IP address of the neighbor specified who sent the update, Router CK2 in this scenario.
Incorrect Answers:
B: Router CK2 learned this route through IBGP with the next-hop of 40.1.1.2. This value will not be used instead of 40.1.1.1 because the next hop information is not preserved across AS 64300 when being sent to a different AS.
C: This is the local interface on Router CK1 itself.
D: The advertisement is accepted.
E: The route is advertised.
Reference: Reference: Cisco Press's BSCI 2nd Edition pg 529

QUESTION 338:

The Certkiller network is shown in the following exhibit:

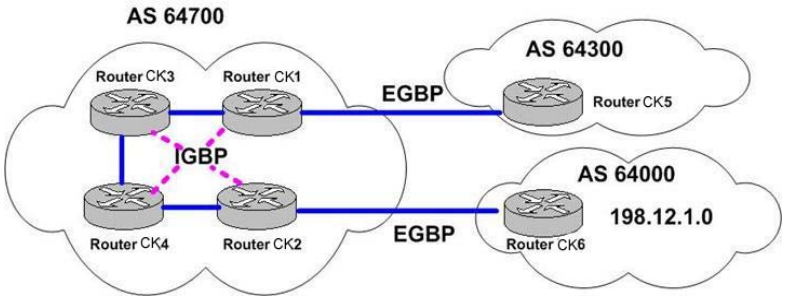

1. all routers are using BGP
2. synchronization is OFF in AS 647000
Regarding the route to 198.12.1.0; which routers will this route propagate to?

A. Router CK6
B. Routers CK2 and CK6
C. Routers CK2 , CK4 , and CK6
D. Routers CK2 , CK3 , CK4 , and CK6
E. Routers CK1 , CK2 , CK3 , CK4 , and CK6
F. Routers CK1 , CK2 , CK3 , CK4 , CK5 , and CK6
G. It will not be propagated to any other router

Answer: F

Explanation:
Router CK 6 will advertise the 192.12.1.0 route to Router CK2 using EBGP. This route would propagate to Router CK1 , CK3 , and CK4 through the fully meshed IGBP sessions. Since synchronization is turned off Router CK1 will advertise the internally learned route to Router CK5 using EBGP.

QUESTION 339:

The Certkiller BGP network is shown below:

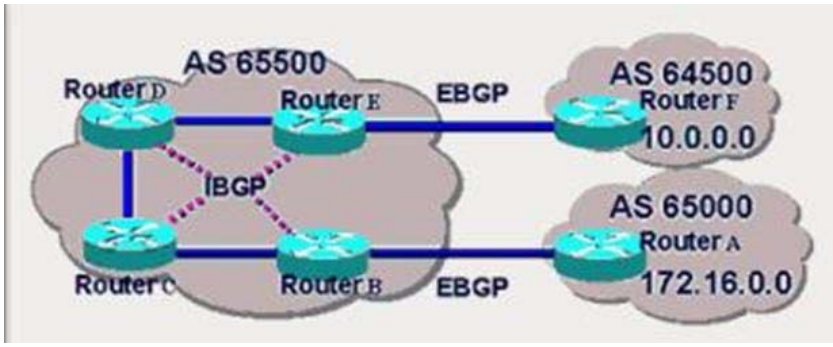

1. synchronization is OFF
2. NO IGP is running in AS 65500
Which router in AS 65500 is going to end up advertising a route to network 172.16.0.0?

A. B only

B. E only
C. B and E only
D. B, C and D only
E. B, C, D, and E

Answer: C

Explanation:
Router B learns about the 172.16.0.0 network route via the EBGP session from router A. With synchronization turned off, it will pass this route to all of its fully meshed IBGP peer routers. Although all routers will receive the route, only routers B and E will advertise this route to other routers. Router B will advertise the route to all of the fully meshed IBGP peers, and router E will advertise it to AS 64500.
Reference: Using the Border Gateway Protocol for Interdomain Routing
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/icsbgp4.htm

---

QUESTION 340:

The Certkiller network administrator is planning to a use
private AS number. Which of the following are choices for a private AS number?

A. 10080
B. 48512
C. 64128
D. 64524

Answer: D

Explanation:
This autonomous system designator is a 16-bit number, with a range of 1 to 65535. RFC 1930 provides guidelines for the use of AS numbers. A range of AS number, 64512 through 65535, is reserved for private use, much like the private Internet Protocol (IP) addresses defined in RFC 1918.
Reference: Building Scalable Cisco Networks (Cisco Press) page 312

QUESTION 341:

The router that is being used for the Internet is low on memory. Because multiple ISP's are being used, BGP is needed for Internet routing. Which of the following types of routes should be received from the Internet providers if the goal is to consume as little memory as possible?

A. Only default routes.
B. Only partial BGP routes.
C. Only internal BGP routes.
D. Only redistributed routes.

Answer: A

Explanation:
Today, the Internet consists of over 100,000 routes, so accepting the full routing table from an ISP can consume a great deal of router memory. By only accepting default routes from the ISPs ensures a minimal impact on the system. There is no risk of having your autonomous system (AS) becoming a transit AS. By receiving only default routes, the entire Internet will be reachable while only using a small amount of memory for the routing table.
Note: One recommendation of Cisco, not listed here, is to use AS_PATH filters for accepting only routes originated from an ISP and its directly connected Autonomous Systems (AS), instead of receiving the full BGP routing table from an ISP.
Reference: Achieve Optimal Routing and Reduce BGP Memory Consumption
http://www.cisco.com/warp/public/459/41.shtml

QUESTION 342:

Routers CK1 , CK2 , and CK3 are all IBGP peers within the Certkiller network. Why should these peer sessions be fully meshed within the AS?

A. Because BGP requires redundant TCP sessions between iBGP peers.
B. Because a full mesh allows for optimal routing within the Transit AS.
C. Because routes learned via iBGP are never propagated to other eBGP peers.
D. Because routes learned via iBGP are never propagated to other iBGP peers.
E. Because routes learned via eBGP are never propagated to other iBGP peers.

Answer: D

Explanation:
Since BGP does not relay routing traffic from one Interior BGP session to another (only from an Exterior BGP session to an IBGP session or another EBGP session), BGP speakers must be fully meshed. When a BGP routing update is received from a neighboring AS, it must be relayed directly to all other BGP speakers in the AS. Do not expect to relay BGP paths from one router, through another, to a third, all within the same AS.
Incorrect Answers:
A: BGP only requires a single TCP session to each peer.
B: The full mesh is needed so that all routers receive the routing information, but it does not provide for any level of route optimization.
C, E: Routes from EBGP peers are propagated to all IBGP peers, and routes learned via IBGP are sent to EBGP peers by default.

QUESTION 343:

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:
BGP uses path information to determine routing loops throughout the network. A simplified view of AS_PATH is that it is the list of Autonomous Systems that a route goes through to reach its destination. Loops are detected and avoided by checking for your own AS number in AS_PATH's received from neighboring Autonomous Systems. If your AS number appears in the AS path the second time, then a loop has occurred.

QUESTION 344:

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:
BGP relies on TCP to provide for a reliable connection between peer routers. The statement in E is false, as the notification message is used to terminate the BGP session, not to maintain or establish the peering session. BGP systems send notificationmessages when an error condition is detected. After the message is sent, the BGP session and the TCP connection between the BGP systems are closed. Notification messages consist of the BGP header plus the error code and subcode, and data that describes the error.
Incorrect Answers:
A: This statement is true.
B: BGP uses TCP port 179. TCP is a reliable, connection oriented protocol.
C: This statement is true. The route must actually exist in the routing table before it will be advertised using the network BGP configuration command.
D: These are both true statements regarding BGP. BGP relies on TCP to maintain sessions.

QUESTION 345:

Which BGP state is a router in when it has found the IP address in the neighbor statement and has created and sent out a BGP packet but the router has not received a response back?

A. Update
B. Connect
C. Active
D. Established

Answer: B

Explanation:
The following table describes the various states that a BGP router can be in:
BGPStateMachine
BGP goes through six states to establish an adjacency.
Idle - incoming connections are refused, and the system gets ready to start speaking BGP. After this is done (by way of a Start event), move to Connect.
Connect - a connection is made to the peer. Send a BGP OPEN message, and go to OpenSent.
Active - a connection comes in from a peer. Send a BGP OPEN message, and go to OpenSent.
OpenSent - Wait for an OPEN message from the peer. When received, send a KEEPALIVE and go to OpenConfirm.
OpenConfirm - Wait for the KEEPALIVE from the peer, then move to Established.
Established - Bidirectional communication is established. Start sending UPDATE and KEEPALIVE messages as required

---

QUESTION 346:

Which BGP attribute will not be advertised in routing updates to its neighboring routers?

A. weight
B. local preference
C. origin
D. AS_path
E. next hop

Answer: A

Explanation:
Weight is proprietary to Cisco and is used in route selection. It is local to the router, and because it is not propagated to other routers, there is no problem with compatibility. When there are multiple paths, it selects a path to a destination with different next hops to

the same destination. Note that the weight attribute has no code. Because it is a local attribute and is not propagated to other routers, no code is needed.
Reference: Cisco Press 642-801 Study Guide p.527

QUESTION 347:

DRAG DROP
Place the BGP attributes in the correct order used for determining a route.

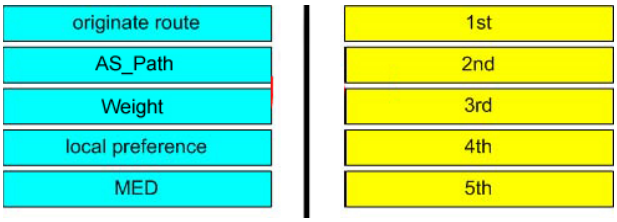

Answer:

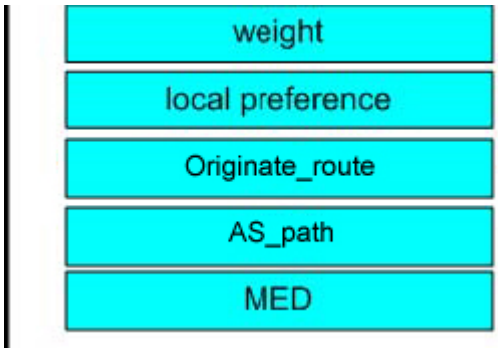

Explanation:
1. Weight: The Weight attribute is similar to the Local Preference attribute in that it gives higher preference to the route that has a higher weight. The difference is that the weight parameter is local to the router and is not exchanged between routers. The weight parameter influences routes coming from different providers to the same router, one router with multiple connections to two or more providers. The weight parameter has a higher precedence than any other attribute. It is the most important attribute when determining route preference. The Weight attribute is a Cisco proprietary attribute.
2. Local Preference: Local Preference is a well-known discretionary attribute, type code 5. The Local Preference attribute is a degree of preference given to a route for comparison with other routes for the same destination. Higher Local Preference values are preferred. Local Preference, as indicated by the name, is local to the AS and is exchanged between IBGP peers only. Local Preference is not advertised to EBGP peers. Routers within a multihomed AS may learn that they can reach the same destination network by way of neighbors in two or more different autonomous systems. In effect, there could be two or more exit points from the local AS to any given destination. Use the Local Preference attribute to force the BGP routers to prefer one exit point instead of another when routing to a particular destination network. Because this attribute is communicated within all BGP routers inside the AS, all BGP routers will have a common view on how to exit the AS.
3. Originate Route: The Origin attribute is a well-known mandatory attribute, type code 1, that indicates the origin of the routing update. BGP allows the following three types of

origins:
1. IGP - The prefix is internal to the originating AS.
2. EGP - The prefix was learned by way of some EGP, such as BGP.
3. Incomplete - The prefix was learned by some other means, probably redistribution.
4. AS_Path:
An AS_Path attribute is a well-known mandatory attribute, type code 2. It is the sequence of AS numbers that a route has traversed to reach a destination. The AS that originates the route adds its own AS number when sending the route to its external BGP peers. Thereafter, each AS that receives the route and passes it on to other BGP peers will prepend its own AS number to the list. Prepending is the act of adding the AS number to the beginning of the list. The final list has all the AS numbers that a route has traversed with the AS number of the AS that originated the route at the end of the list. This type of AS_Path list is called an AS_Sequence because all the AS numbers are ordered sequentially.
BGP uses the AS_Path attribute as part of the routing updates, update packet, to ensure a loop free topology on the Internet. Each route that gets passed between BGP peers will carry a list of all AS numbers that the route has already been through. If the route is advertised to the AS that originated it, that AS will see itself as part of the AS_Path attribute list and will not accept the route. BGP speakers prepend their AS numbers when advertising routing updates to other autonomous systems, external peers. When the route is passed to a BGP speaker within the same AS, the AS_Path information is left intact.
5. MED : 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.

QUESTION 348:

Which attribute must exist in the BGP update packet?

A. LOCAL_PREF
B. AGGREGATOR
C. AS_Path
D. Weight

Answer: C

Explanation:
The AS PATH is a well known mandatory attribute. Mandatory means that it is required by all routers. These attributes are required and are therefore recognized by all BGP implementations.

QUESTION 349:

Which three conditions can cause BGP neighbor establishment to fail? (Select three)

A. There is an access list blocking all TCP traffic between the two BGP neighbors.
B. The EBGP neighbor is not directly connected, and the ebgp-multi-hop option is set to the default value.
C. The IBGP neighbor is not directly connected.
D. BGP synchronization is enabled in a Transit AS with fully-meshed IBGP neighbors.
E. The BGP update interval is different between the two BGP neighbors.
F. The BGP neighbor is referencing an incorrect AS number in its neighbor statement.

Answer: A, B, F

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 350:

You need to configure EBGP on a Certkiller router. 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.
E. None of the above

Answer: C

Explanation:
Only choice C is correct. To reach a EBGP peer, a static route can be used with the EBGP multi-hop command to establish a neighbor adjacency. Unlike interior routing protocols, EBGP neighbors do not need to physically be connected to each other. All that is needed is a path to the peer IP address, and a valid TCP port 179 connection between the BGP routers.

QUESTION 351:

A problem was reported that the 10.10.10.0/24 prefix was not injected into the local BGP table on a Certkiller router named CK1 . The following information is available from this router:
CK1 Configuration:
router bgp 65001
network 10.0.0.0
neighbor 172.16.1.1 remote-as 65002
no auto-summary
Routing table information:
show ip route | include 10
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 of the CK1 ?

A. The 172.16.1.1 neighbor is down.
B. The prefix 10.10.10.0/24 is not a 'connected' route.
C. This route is not a BGP learned route.
D. The network command is wrong.
E. None of the above

Answer: D

Explanation:
The network command is used with IGPs, such as RIP, to determine the interfaces on which to send and receive updates. The command also indicates which directly connected networks to advertise. However, when configuring BGP, the network command does not affect what interfaces BGP runs on. Therefore, configuring just a network statement will not establish a BGP neighbor relationship. This is a major difference between BGP and IGPs. The network statement follows this syntax:
Router(config-router)#network network-number [mask network-mask]
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.

QUESTION 352:

Observe the diagram shown below:

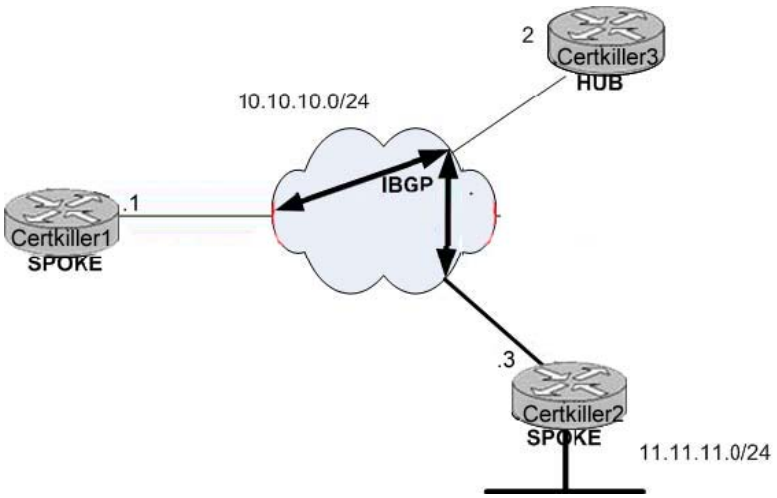

In this network, Certkiller 3 is the hub router and Certkiller 1 and Certkiller 2 are the spokes. There are no virtual circuits between the spoke locations. What is needed to successfully route traffic to the 11.11.11.0/24 network from Certkiller 1?

A. The neighbor 10.10.10.1 next-hop-self command on Certkiller 2.
B. The neighbor 10.10.10.1 next-hop-self command on Certkiller 1.
C. The neighbor 10.10.10.1 next-hop-self command on Certkiller 3.
D. Nothing is required. This is the default behavior on this topology.
E. None of the above.

Answer: C

Explanation:
It is quite common for a BGP speaker to use the same update policies for its peers. An update policy consists of the same outbound route maps, distribute lists, filter lists, update source, and so on. Having to configure the same update policy on a BGP speaker for all of its neighbors can become strenuous, and more important, when making modifications to the update policy, there is room for errors. There is a way around this, and it's known as peer groups. Peer groups allow you to group all of a BGP speaker's neighbors that need to use the same policy into one group. The update policy is then applied to all members of that peer group.

Option	Description
advertise-map	Specifies the route map for conditional advertisement.
advertisement-interval	Sets the minimum interval between sending eBGP routing updates.
default-originate	Originates the default route to this neighbor.
description	Creates a neighbor-specific description.

inbound/outbound distribute-list	Filters updates to/from this neighbor.
ebgp-multihop	Allows eBGP neighbors not on directly connected networks.
inbound/outbound filter-list	Establishes BGP filters.
maximum-prefix	Sets the maximum number of prefixes accepted from this peer.
next-hop-self	Disables the next-hop calculation for this neighbor.
password	Sets a password.
inbound/outbound prefix-list	Applies a prefix list to a neighbor.
remote-as	Specifies a BGP neighbor.
remove-private-AS	Removes a private AS number from outbound updates.
inbound/outbound route-map	Applies a route map to a neighbor.
route-reflector-client	Configures a neighbor as a route reflector client.
send-community	Sends the COMMUNITY attribute to this neighbor.
soft-reconfiguration	Per neighbor soft reconfiguration.
timers	Sets BGP per neighbor time rs.
unsuppress-map	Route map to selectively unsuppress suppressed routes.
update-source	Source of routing updates.
version	Sets the BGP version to match a neighbor.
weight	Sets the default weight for routes from this neighbor.

QUESTION 353:

Certkiller has a BGP network and a BGP route of 197.22.129.0/24 that should be propagated to all of the devices. The route is currently not in any of the routing tables. The Certkiller administrator determines that an access list is the cause of the problem. The administrator changes the access list to allow this route, but the route still does not appear in any of the routing tables. What should be done to propagate this route?

A. Use the service-policy command to adjust the QOS policy to allow the route to propagate.

B. Clear the BGP session.
C. Use the release BGP routing command.
D. Change both the inbound and outbound policy related to this route.
E. None of the above

Answer: B

Explanation:
When configuring BGP, changes made to an existing configuration may not appear immediately. In order to force BGP to clear its table and reset BGP sessions, use the clear ip bgp * command :
Router#clear ip bgp *
The asterisk (*) is a wildcard that matches all table entries. Therefore, all BGP routes are lost while the neighbor relationships are reset. This is expedient and very useful in a lab situation, but caution should be exercised when issuing this command on a production router. On an Internet backbone router, it may be more appropriate to use this command with a specific IP address, as shown in the following:
Router#clear ip bgp 192.168.0.0

QUESTION 354:

Study the Certkiller network shown below carefully:

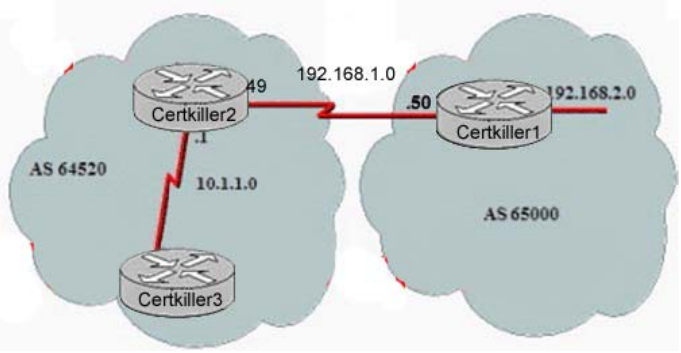

By default, when Certkiller 2 passes BGP advertisements from Certkiller 1 about network 192.168.2.0 to Certkiller 3, what address will be listed as the next-hop address?

A. 192.168.1.50
B. 192.168.1.49
C. 10.1.1.1
D. 10.1.1.2
E. None of the above

Answer: A

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:
1. For EBGP sessions, the next hop is the IP address of the neighbor that announced the route.
2. For IBGP sessions, where routes originated inside the AS, the next hop is the IP address of the neighbor that announced the route. For routes injected into the AS by way of EBGP, the next hop learned from EBGP is carried unaltered into IBGP. The next hop is the IP address of the EBGP neighbor from which the route was learned.
3. When the route is advertised on a multiaccess medium, such as Ethernet or Frame Relay, the next hop is usually the IP address of the interface of the router. This will be the interface connected to the media that originated the route.

QUESTION 355:

The Certkiller BGP network is shown below:

Configuration exhibit:

```
hostname Certkiller1
!
interface loopback 0
ip address 10.10.10.2  255.255.255.255
!
interface serial 0/1
ip address 172.16.10.2 255.255.255.252
!
interface serial 0/1
ip address 172.16.10.5  255.255.255.252

route bgp 65444
neighbor 10.10.10.1 remote-as 65111
neighbor 10.10.10.1 update-source loopback 0
neighbor 10.10.10.1 ebgp-multihop
```

Study the exhibits above carefully. Routers Certkiller 1 and Certkiller 2 are running BGP but the session is active. What command(s) needs to be added to establish the BGP session?

A. neighbor 10.10.10.1 next-hop-self
B. network 10.10.10.0
C. no synchronization
D. ip route 10.10.10.1 255.255.255.255 s0/0

ip route 10.10.10.1 255.255.255.255 s0/1
E. None of the above

Answer: D

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 356:

The Certkiller WAN is shown in the following network topology exhibit:

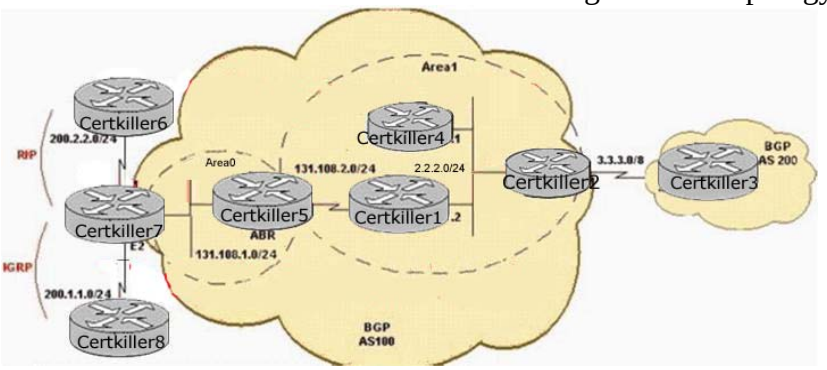

Configuration exhibit:

```
Certkiller1 show ip route
Gateway of last resort is not set

     2.0.0.0/24 is subnetted. 1 subnetted
C          2.2.2.0 is directly connected, Ethernet0/0
C    3.0.0.0/8 is directly connected, Serial1/0
O E2 200.1.1.0/24 [110/20] via 2.2.2.2, 00:16:17, Ethernet0/0
O E1 200.2.2.0/24 [110/104] via 2.2.2.2, 00:00:41, Ethernet0/0
     31.108.0.0/24 is subnetted 2 subnets
O    131.108.1.0 [110/84] via 2.2.2.2, 00:16:17, Ethernet0/0
O IA 131.108.1.0 [110/84] via 2.2.2.2, 00:16:17, Ethernet0/0
```

Based on the information shown above, which command should be added to Certkiller 2 under the "router bgp 100" to allow only the external OSPF routes to be redistributed to Certkiller 3?

A. redistribute ospf 1 match external 1
B. redistribute ospf 1 match external 1 external 2
C. redistribute ospf 1 match external 2
D. redistribute ospf 1

Answer: B

Explanation:
To support multiple routing protocols within the same internetwork efficiently, routing information must be shared among the different routing protocols. For example, routes learned from a RIP process may need to be imported into an IGRP process. This process of exchanging routing information between routing protocols is called route redistribution. Such redistribution can be one-way or two-way. One-way routes are where one protocol receives the routes from another. Two-way routes are where both protocols receive routes from each other. Routers that perform redistribution are called boundary routers because they border two or more autonomous systems or routing domains
The syntax of the redistribute command is as follows:
Router(config-router)#redistribute protocol [process-id] {level-1 | level-1-2 | level-2} [metric metric-value][metric-type type-value] [match {internal | external 1 | external 2}][tag tag-value] [route-map map-tag][weight weight] [subnets]
Example:
Router A(config)#router rip
Router A(config-router)#redistribute ospf 109 match internal external 1 external 2
Router A(config-router)#default-metric 10
Note that the redistribute command includes the OSPF process ID, 109. The other keywords, match internal external 1 and external 2, instruct RIP to redistribute internal OSPF routes, as well as external Type 1 and Type 2 routes. This is the default for OSPF redistribution. These keywords are required only if its behavior is to be modified.

QUESTION 357:

A problem was reported that the 10.10.10.0/24 prefix was not injected into the local BGP table on Router Certkiller 3. The following information is available from Router Certkiller 3:
Configuration:
router bgp 65001
network 10.0.0.0
neighbor 172.16.1.1 remote-as 65002
no auto-summary
Routing table information:
show ip route | include 10
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. The 172.16.1.1 neighbor is down.
B. The network command is wrong.
C. This route is not a BGP learned route.
D. The prefix 10.10.10.0/24 is not a connected route.
E. None of the above

Answer: B

Explanation:

The network command is used with IGPs, such as RIP, to determine the interfaces on which to send and receive updates. The command also indicates which directly connected networks to advertise. However, when configuring BGP, the network command does not affect what interfaces BGP runs on. Therefore, configuring just a network statement will not establish a BGP neighbor relationship. This is a major difference between BGP and IGPs. The network statement follows this syntax:
Router(config-router)#network network-number [mask network-mask]
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.

QUESTION 358:

A Certkiller BGP router is configured as shown below:

```
Certkiller1 #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   rip/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
```

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
The 192.168.0.0 network is not being propagated throughout the network. Based on the information shown above what is the reason the 192.168.0.0 route is not being advertised?

A. The network 10.0.0.0 statement is missing mask 0.255.255.255.
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 192.168.0.0 statement is missing mask 255.255.0.0
E. The auto-summary configuration is missing.

Answer: D

Explanation:
The network command is used with IGPs, such as RIP, to determine the interfaces on which to send and receive updates. The command also indicates which directly connected networks to advertise. However, when configuring BGP, the network command does not affect what interfaces BGP runs on. Therefore, configuring just a network statement will not establish a BGP neighbor relationship. This is a major

difference between BGP and IGPs. The network statement follows this syntax:
Router(config-router)#network network-number [mask network-mask]
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.

QUESTION 359:

You need to troubleshoot a BGP neighbor issue between routers CK1 and CK2 .
Which three conditions can cause BGP neighbor establishment to fail? (Select three)

A. There is an access list blocking all TCP traffic between the two BGP neighbors.
B. BGP synchronization is enabled in a transit autonomous system with fully-meshed IBGP neighbors.
C. The BGP neighbor is referencing an incorrect autonomous system number in its neighbor statement.
D. The BGP update interval is different between the two BGP neighbors.
E. The EBGP neighbor ebgp-multihop option is set to the default value.
F. The IBGP neighbor is not directly connected.

Answer: A, C, E

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.
When BGP neighbors first establish a connection, they exchange all candidate BGP routes. After this initial exchange, incremental updates are sent as network information changes. As discussed in earlier modules, incremental updates are more efficient than complete table updates. This is especially true with BGP routers, which may contain the complete Internet routing table.
Peers advertise destinations that are reachable through them by using update messages. These messages contain route prefix, AS path, path attributes such as the degree of preference for a particular route, and other properties

---

QUESTION 360:

Based on the "show ip bgp summary" output shown below on router Certkiller 1, which two statements are true? (Select two)

A. The BGP session to the 10.2.2.2 neighbor is established.

B. The router is attempting to establish a BGP peering session with the 10.1.1.1 neighbor.
C. The router is attempting to establish a BGP peering session with the 10.2.2.2 neighbor.
D. The BGP session to the 10.3.3.3 neighbor is established.
E. The BGP session to the 10.1.1.1 neighbor is established.
F. 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.

Answer: C, E

Explanation:
Show ip bgp summary command displays the summary of all BGP connections.
The six states of the BGP FSM are described as follows:
* Idle - Idle is the first state of a BGP connection. BGP is waiting for a start event. It is normally initiated by an administrator or a network event. At the start event, BGP initializes its resources and resets a connect retry timer. Then it starts listening for a TCP notice that BGP can transition back to Idle from any other state in case of errors.
* Connect
- In the Connect state, BGP is waiting for the TCP connection to be completed. If the TCP connection is successful, the state transitions to OpenSent. If the TCP connection fails, the state transitions to the Active state, and the router tries to connect again. If the connect retry timer expires, the state remains in the Connect state, the timer is reset, and a TCP connection is initiated. In case of any other event, initiated by the system or the administrator, the state returns to Idle.
* Active - In the Active state, BGP is trying to acquire a peer by initiating a TCP connection. If it is successful, it transitions to OpenSent. If the connect retry timer expires, BGP restarts the connect timer and returns to the Connect state. While active, BGP is still listening for a connection that may be initiated from another peer. The state may go back to Idle in case of other events, such as a stop event initiated by the system or the operator.
In general, a neighbor state that is switching between "Connect" and "Active" is an indication that something is wrong and that there are problems with the TCP connection. It could be because of many TCP retransmissions, or the incapability of a neighbor to reach the IP address of its peer.
* OpenSent - In the OpenSent state, BGP is waiting for an open message from its peer. The open message is checked for correctness. In case of errors, such as an incompatible version number or an unacceptable AS, the system sends an error notification message and goes back to idle. If there are no errors, BGP starts sending keepalive messages and resets the keepalive timer. At this stage, the hold time is negotiated and the smaller value is taken. If the negotiated hold time is zero (0), the hold timer and the keepalive timer are not restarted.
At the OpenSent state, BGP recognizes whether the peer belongs to the same AS or to a different AS. BGP does this by comparing its AS number to the AS number of its peer. A same AS is an IBGP peer and a different AS is an EBGP peer.
When a TCP disconnect is detected, the state falls back to Active. For any other errors, such as an expiration of the hold timer, BGP sends a notification message with the corresponding error code. Then it returns to the Idle state.

* OpenConfirm - While in OpenConfirm state, BGP is waiting for a keepalive or notification message. If a keepalive message is received, the state goes to the Established state, and the neighbor negotiation is complete. If the system receives an update or keepalive message, it restarts the hold time, assuming that the negotiated hold time is not zero. If a notification message is received, the state falls back to Idle. The system sends periodic keepalive messages at the rate set by the keepalive timer. In the case of any TCP disconnect or in response to any stop event, initiated by the system or the administrator, the state returns to Idle. In response to any other event, the system sends a notification message with an FSM error code and returns to the Idle state.
* Established - Established is the final state in the neighbor negotiation. BGP starts exchanging update packets with its peers. If it is non-zero, the hold timer is restarted at the receipt of an update or keepalive message.

QUESTION 361:

One of the Certkiller routers is unable to form a neighbor relationship with an external BGP peer. Which two statements are true about external BGP neighbor relationships? (Select two)

A. EBGP neighbors use TCP port 179 to exchange BGP routing tables.
B. The BGP split-horizon rule specifies that routes learned via IBGP are never propagated to other EBGP peers.
C. EBGP neighbors must be in different autonomous systems.
D. Static routes or an interior gateway protocol is required between EBGP neighbors.
E. When an EBGP neighbor receives an update from another EBGP neighbor, it should not forward the update to other EBGP neighbors.
F. Loopback addresses should be used between EBGP neighbors.

Answer: A, C

Explanation:
BGP updates are carried using TCP on port 179. In contrast, RIP updates use UDP port 520, while OSPF does not use a Layer 4 protocol. Because BGP requires TCP, IP connectivity must exist between BGP peers. TCP connections must also be negotiated between them before updates can be exchanged. Therefore, BGP inherits those reliable, connection-oriented properties from TCP.
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 362:

You need to troubleshoot an issue on a Certkiller BGP routers and need to verify the neighbor configuration. Which command displays the IBGP and EBGP neighbors

that are configured?

A. show ip bgp
B. show ip bgp peers
C. show ip bgp summary
D. show ip bgp paths
E. None of the above

Answer: C

Explanation:
Show ip bgp summary command displays the summary of all BGP connections.

QUESTION 363:

A Certkiller BGP router needs to advertise networks to the upstream Internet providers. Which two methods advertise internal networks to external ISPs via BGP? (Select two)

A. Disabling synchronization
B. Forcing the next-hop address
C. Defining routes via the network statement
D. Using aggregate routes

Answer: C, 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 364:

The following output was seen on a Certkiller router:

```
Certkiller1# 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        1    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
```

Based on the show "ip bgp summary" output above, which two statements are true? (Select two)

A. The router is attempting to establish a BGP peering session with the 10.1.1.1 neighbor.
B. The BGP session to the 10.3.3.3 neighbor is established.
C. 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.
D. The router is attempting to establish a BGP peering session with the 10.2.2.2 neighbor.
E. The BGP session to the 10.1.1.1 neighbor is established.
F. The BGP session to the 10.2.2.2 neighbor is established.

Answer: D, E

Explanation:
Show ip bgp summary command displays the summary of all BGP connections.
The six states of the BGP FSM are described as follows:
* Idle - Idle is the first state of a BGP connection. BGP is waiting for a start event. It is normally initiated by an administrator or a network event. At the start event, BGP initializes its resources and resets a connect retry timer. Then it starts listening for a TCP notice that BGP can transition back to Idle from any other state in case of errors.
* Connect - In the Connect state, BGP is waiting for the TCP connection to be completed. If the TCP connection is successful, the state transitions to OpenSent. If the TCP connection fails, the state transitions to the Active state, and the router tries to connect again. If the connect retry timer expires, the state remains in the Connect state, the timer is reset, and a TCP connection is initiated. In case of any other event, initiated by the system or the administrator, the state returns to Idle.
* Active - In the Active state, BGP is trying to acquire a peer by initiating a TCP connection. If it is successful, it transitions to OpenSent. If the connect retry timer expires, BGP restarts the connect timer and returns to the Connect state. While active, BGP is still listening for a connection that may be initiated from another peer. The state may go back to Idle in case of other events, such as a stop event initiated by the system or the operator.
In general, a neighbor state that is switching between "Connect" and "Active" is an indication that something is wrong and that there are problems with the TCP connection. It could be because of many TCP retransmissions, or the incapability of a neighbor to reach the IP address of its peer.
* OpenSent - In the OpenSent state, BGP is waiting for an open message from its peer. The open message is checked for correctness. In case of errors, such as an incompatible version number or an unacceptable AS, the system sends an error notification message and goes back to idle. If there are no errors, BGP starts sending keepalive messages and resets the keepalive timer. At this stage, the hold time is negotiated and the smaller value is taken. If the negotiated hold time is zero (0), the hold timer and the keepalive timer are not restarted.
At the OpenSent state, BGP recognizes whether the peer belongs to the same AS or to a different AS. BGP does this by comparing its AS number to the AS number of its peer. A same AS is an IBGP peer and a different AS is an EBGP peer.

When a TCP disconnect is detected, the state falls back to Active. For any other errors, such as an expiration of the hold timer, BGP sends a notification message with the corresponding error code. Then it returns to the Idle state.
* OpenConfirm - While in OpenConfirm state, BGP is waiting for a keepalive or notification message. If a keepalive message is received, the state goes to the Established state, and the neighbor negotiation is complete. If the system receives an update or keepalive message, it restarts the hold time, assuming that the negotiated hold time is not zero. If a notification message is received, the state falls back to Idle. The system sends periodic keepalive messages at the rate set by the keepalive timer. In the case of any TCP disconnect or in response to any stop event, initiated by the system or the administrator, the state returns to Idle. In response to any other event, the system sends a notification message with an FSM error code and returns to the Idle state.
* Established - Established is the final state in the neighbor negotiation. BGP starts exchanging update packets with its peers. If it is non-zero, the hold timer is restarted at the receipt of an update or keepalive message.

QUESTION 365:

Refer to the exhibits below:
Network Topology exhibit:

Certkiller 1 configuration exhibit:

```
hostname Certkiller1
!
interface loopback 0
ip address 10.10.10.2 255.255.255.255

interface serial 0/0
ip address 172.16.10.2 255.255.255.252
!
interface serial 0/1
ip address 172.16.10.5 255.255.255.252
!
router bgp 65444
neighbor 10.10.10.1 remote-as 65111
neighbor 10.10.10.1 update-source loopback 0
neighbor 10.10.10.1 ebgp-multihop
```

In the network shown above, routers Certkiller 1 and Certkiller 2 are running BGP but the session is 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. neighbor 10.10.10.1 next-hop-self
C. no synchronization
D. network 10.10.10.0
E. None of the above

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 366:

A Certkiller BGP router is configured as shown below:

```
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
```

Examine the above configuration. What does the route map named test accomplish in this configuration?

A. It permits only the 10.0.0.0/8 prefix to be received from the 10.1.1.1 neighbor
B. It marks all prefixes received from the 10.1.1.1 neighbor with a MED of 200
C. It marks the 10.0.0.0/8 prefix received from the 10.1.1.1 neighbor with a MED of 200
D. It marks the 10.0.0.0/8 prefix advertised to the 10.1.1.1 neighbor with a MED of 200
E. It marks all prefixes advertised to the 10.1.1.1 neighbor with a MED of 200
F. It permits only the 10.0.0.0/8 prefix to be advertised to the 10.1.1.1 neighbor

Answer: D

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 367:

The command "bgp always-compare-med" was added to router CK1 . What will this command accomplish?

A. It forces CK1 to compare the communities of routes from different autonomous systems
B. It forces CK1 to compare the weight of routes from different autonomous systems
C. It forces CK1 to compare metrics of routes from different autonomous systems
D. It forces CK1 to compare the local preference of routes from different autonomous systems
E. None of the above

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
If the bgp always-compare-med router configuration command is enabled, all paths are fully comparable, including those from other autonomous systems in the confederation, even if the bgp deterministic med command is also enabled.

QUESTION 368:

The "show ip bgp" command was issued on a Certkiller router as shown below:

```
Certkiller1#show ip bgp
BGP table version is 16 local router ID is 10.200.200.13
0.0.0.0/24 hayhayhahaaliahajjaghfh history, * valid, > best, i - internal
Origin codes i-is.16    BGP.7-  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
```

Based on the Certkiller 2 output, which statement is true?

A. The best path to reach the 192.168.11.0 prefix is via 10.200.200.11.
B. 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.
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 best path to reach the 192.168.11.0 prefix is via 10.200.200.12.
E. None of the above.

Answer: A

Explanation:
Local Preference is a well-known discretionary attribute, type code 5. The Local Preference attribute is a degree of preference given to a route for comparison with other routes for the same destination. Higher Local Preference values are preferred. Local Preference, as indicated by the name, is local to the AS and is exchanged between IBGP peers only. Local Preference is not advertised to EBGP peers.
Routers within a multihomed AS may learn that they can reach the same destination network by way of neighbors in two or more different autonomous systems. In effect, there could be two or more exit points from the local AS to any given destination. Use the Local Preference attribute to force the BGP routers to prefer one exit point instead of another when routing to a particular destination network. Because this attribute is communicated within all BGP routers inside the AS, all BGP routers will have a common view on how to exit the AS.
In exhibit, 10.200.200.11 has higher local value then 10.200.200.12 So, best path for 192.168.11.0 is via 10.200.200.11

QUESTION 369:

According to Cisco's official recommendations, what is the preferred method for populating the BGP table on your router with IGP routes?

A. Use the network command.
B. Redistribute EBGP routes into BGP.
C. Redistribute dynamic routes into BGP.
D. Redistribute static routes into the IGP.

Answer: A

Explanation:
For BGP the network command allows BGP to advertise a network that is already in the IP table.
Note: An alternate solution, not listed here, is to redistribute the IGP routes into BGP.
Incorrect Answers:
B: We want to distribute internal routes into BGP, not external routes.
C: Redistribution from the IGP is NOT advised because there is a great reliance on the IGP table. It could cause instability.
D: We want to populate the BGP table, not the IGP table.

QUESTION 370:

While troubleshooting a BGP problem on the Certkiller network you notice that the 10.10.10.0/24 prefix is not being injected into the local BGP table of CK1 . The relevant configuration of router CK1 is shown below:
router bgp 65001
network 10.0.0.0
neighbor 172.16.1.1 remote-as 65002
no auto-summary
Routing table information:
show ip route | include 10
O 10.10.10.0/24 [110/11] via 192.168.1.1, 2d00h, Ethernet0/0
Why doesn't the local BGP table have the prefix?

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:
The correct syntax should be "network 10.10.10.0 mask 255.255.255.0" under the BGP routing process. Without the correct subnet mask specified, the route will not get injected into the BGP routing table, even if it is learned via an IGP. In this case, the route is known via OSPF.

QUESTION 371:

What would happen if a full mesh of BGP sessions were configured within a single large autonomous system? (Select two)

A. Many UDP sessions will be created.
B. More memory and CPU will be consumed.
C. This configuration is not permitted by default.

D. A significant amount of bandwidth for BGP updates and retransmissions can be used.
E. Permanent Virtual Circuits (PVCs) must be created to link the fully meshed BGP sessions.

Answer: B, D

Explanation:
The number of TCP connections that are required for a full mesh grows exponentially. The formula for the total number of BGP sessions is N(N-1)/2. For example, for 10 IBGP routers to become fully meshed 45 peering sessions will need to be configured.
B: More memory and CPU resources are required on the routers to support the full mesh BGP.
D: There will be many TCP sessions and a significant amount of bandwidth will be required for BGP traffic.
Incorrect Answers:
A: Many TCP, not UDP, sessions will be created.
C: It is permitted.
E: PVCs are not required.

QUESTION 372:

Regarding the use of BGP prefix lists, which one of the statements below is true?

A. They start at 10 by default.
B. They automatically increment by ten by default.
C. They are displayed using the show ip prefix-list command.
D. They are displayed using the show bgp prefix-list command.
E. None of the above.

Answer: C

Explanation:
The show ip prefix-list command is used to display information about a prefix list or prefix list entries.
Note: Filtering by prefix list involves matching the prefixes of routes with those listed in the prefix list, similar to using access lists.
Incorrect Answers:
A, B: By default, the entries of a prefix list will have sequence values of 5, 10, 15 etc. They start at 5 and increment by 5.
D: This is an invalid command.

QUESTION 373:

BGP is being used on the Certkiller network for Internet routing. In EBGP, which of the following configuration lines advertises the subnet 154.2.1.0 255.255.255.0 to EBGP neighbors?

A. Router (config-router)#network 154.2.1.0
B. Router (config-router)#network 164.2.1.0
C. Router (config-router)#network-advertise 154.2.1.0
D. Router (config-router)#network 154.2.1.0 mask 255.255.255.0

Answer: D

Explanation:
When configuring BGP, we specify the mask keyword also along with the network command. If the network mask is omitted, the IP network is taken at the network class boundary.

QUESTION 374:

Routers A, B, and C are all part of the same autonomous system and are configured for BGP routing. The relevant configurations for RTR A and RTR B are shown below:

```
Hostname RTR A
!
interface serial 0/0
ip address 140.140.13.2 255.255.255.252
description Connection to ISP A
!
interface 10/0
ip address 10.10.10.2 255.255.0
description Connection to RTR B
!
interface f0/1
ip address 172.16.30.1 255.255.255.0
description Connection to RTR C
router bgp 5201
neighbor 10.10.10.3 remote-as 5201
neighbor 140.140.13.2 remote-as 5203
Hostname RTR B
!
interface 10/0
ip address 10.10.10.3 255.255.255.0
description Connection to RTR A
!
interface 10/1
ip address 10.10.20.2 255.255.255.0
description Connection to RTR C
!
router bgp 5201
neighbor 10.10.10.2 remote-as 5201
```

neighbor 10.10.20.1 remote-as 5201

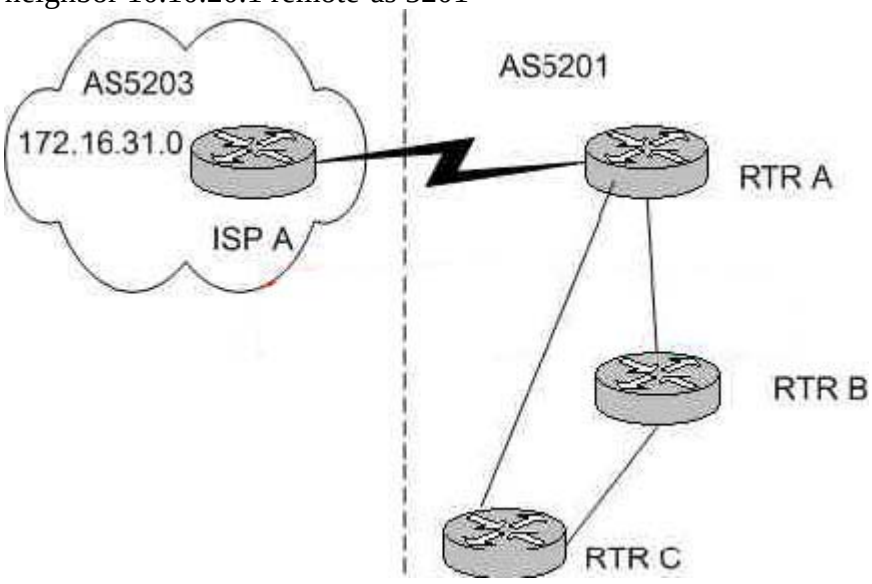

Based on the information provided, what will RTR A do with the updates it receives from ISP A?

A. Send the update to RTR B and C.
B. Send the update to RTR B only.
C. Ignore and discard the packet.
D. Update its tables and discard the packet.

Answer: C

Explanation:
When specifying the remote external AS in the BGP peer statement, the IP address of the neighbor router should be used as the peer. In this example, the remote peer specified in RTR A is 140.140.13.2, but as we can see from the configuration file this is the IP address of its own serial interface. The correct statement should have been "neighbor 140.140.13.1 remote-as 5203" under the BGP process.
Because of this, the EBGP peering session will not be established and all BGP updates from ISP A will be ignored.

QUESTION 375:

Your network has 2 separate T1 connections with two different providers. TO support this, BGP is being configured. The required information regarding the ISPs is shown below:
Your network: 164.67.36.0/24
Your AS number: 300
AS number of ISP1: 1005
AS number of ISP2: 1010
Based on the information provided, which of the following commands would you enter to advertise your network to ISP1 and ISP2 via BGP?

A. router bgp 1005
network 164.67.36.0 mask 255.255.255.0
neighbor 15.1.1.1 remote-as 1005
neighbor 25.1.1.1 remote-as 1010
B. router bgp 300
network 164.67.36.0 mask 255.255.255.0
neighbor 15.1.1.1 remote-as 1005
neighbor 25.1.1.1 remote-as 1010
C. router bgp 300
network 164.67.36.0
neighbor 15.1.1.1 remote-as 1005
neighbor 25.1.1.1 remote-as 1010
D. router bgp 1010
network 164.67.36.0
neighbor 15.1.1.1 remote-as 1005
neighbor 25.1.1.1 remote-as 1010

Answer: B

Explanation:
Step 1: router bgp 300
The router bgp command is used to activate the BGP protocol and identify the local autonomous system.
Step 2: network 164.67.36.0 mask 255.255.255.0
The network command controls which networks are originated by this router.
Syntax: network network-number network-mask
Step 3:
neighbor 15.1.1.1 remote-as 1005
neighbor 25.1.1.1 remote-as 1010
The neighbor remote-as command to identify a peer router with which the local router will establish a session.
Incorrect Answers:
A, D: We must specify the local autonomous system in the router bgp command. We must use AS 300, not AS 1005 or AS 1010
C: Both the network number and the network mask must be specified with the network command.

QUESTION 376:

The Certkiller Internet router is configured as follows::
router bgp 64000
network 17.0.0.0
neighbor 178.5.1.1 remote-as 64000
neighbor 197.4.1.2 remote-as 64100
Based on this router configuration, which of the following statements are true?

(Select two)

A. The router with IP address 178.5.1.1 runs IBGP with RouteCK.
B. The router with IP address 178.5.1.1 runs EBGP with RouterCK.
C. The router with IP address 197.4.1.2 runs EBGP with RouterCK.
D. The router with IP address 178.5.1.1 runs as a community member with RouterCK.
E. The router with IP address 197.4.1.2 runs as a peer group member with RouteCK.

Answer: A, C

Explanation:
The router bgp command is used to activate the BGP protocol and identify the local autonomous system.
The neighbor command activates a BGP session with another router using either IBGP or EBGP.
Syntax: neighbor { ip-address|peer-group-name}remote-as autonomous-system
The value placed in the autonomous system field of the neighbor command determines whether the communication with the neighbor is an EBGP or an IBGP session.
A: If the autonomous system field configured in the router bgp command is identical to the field in the neighbor remote-as command, then BGP will initiate an internal session (IBGP). Here both the local AS and the neighbor 178.5.1.1 are in the AS 64000.
C: If the field values are different, then BGP will initiate an external session (EBGP). The network 197.4.1.2 has a different AS number.

QUESTION 377:

You wish to permit all BGP prefixes between /10 and /18 for the 207.0.0.0 network. How should a BGP prefix list be configured to do this?

A. ip prefix-list 207.0.0.0/8 ge 10 le 18
B. ip prefix-list 207.0.0.0/8 ge 18 le 10
C. ip prefix-list 207.0.0.0/24 ge 10 le 18
D. ip prefix-list 207.0.0.0/24 ge 18 le 10
E. None of the above

Answer: A

Explanation:
With prefix lists, the "ge" means "greater than or equal to" while the le means "less than or equal to." Here, we want to specify all prefixes greater than or equal to 10 and less than or equal to 18 as specified in choice
A. In addition, we must specify an 8 bit network
mask, not a 24 bit network mask.

QUESTION 378:

Routers CK1 and CK2 are configured for BGP as shown below:
RTR CK1
router bgp 200
neighbor 183.215.22.1 remote-as 200
neighbor 183.215.22.1 update-source loopback 1
RTR CK2
router bgp 200
neighbor 147.229.1.1 remote-as 200
These two routers currently have an active BGP peering session between them and they are able to pass routes to each other. Based on these configuration files, which of the following are true? (Select three)

A. RTR CK1 and RTR CK2 are running IBGP inside AS 200
B. The IP address of RTR CK1 's Loopback 1 interface is 147.229.1.1.
C. The IP address of RTR CK1 's Loopback 1 interface is 183.215.22.1.
D. RTR CK1 and RTR CK2 are running EBGP between the autonomous systems.
E. RTR CK1 has forced BGP to use the loopback IP address as the source in the TCP neighbor connection.

Answer: A, B, E

Explanation:
A: The two statements that define a remote AS (autonomous) use the same AS number. Both routers must belong to the same AS and there IBGP is used.
B: The IP address of the CK1 loopback address is 147.229.1.1 as specified in the neighbor statement of the CK2 router: neighbor 147.229.1.1 remote-as 200
E: The "neighbor 183.215.22.1 update-source loopback 1" command issued at CK1 makes CK1 use this loopback for the BGP connection to CK2 . See Note 2 below.
Note 1: Syntax of neighbor command:
Router(config-router)#neighbor {ip-address| peer-group-name}
remote-asas-number This command specifies a BGP neighbor.
Note 2: For iBGP, you might want to allow your BGP connections to stay up regardless of which interface is used to reach a neighbor. To enable this configuration, you first configure a loopback interface and assign it an IP address (neighbor 183.215.22.1 update-source loopback 1). Next, configure the BGP update source to be the loopback interface (we have to assume this step - it is not indicated by the exhibit). Finally, configure your neighbor to use the address on the loopback interface (neighbor 147.229.1.1 remote-as 200).
Incorrect Answers:
C: 183.215.22.1 is the IP address of CK2 , not CK1 .
D: There is only one AS and the routers belong to that AS. Within an AS IBGP is used, not EBGP.

QUESTION 379:

The Certkiller BGP router is configured as shown below:

router bgp 64000
neighbor 172.16.1.1 remote-as 64000
neighbor 10.1.1.2 remote-as 64550
network 200.52.1.192 mask 255.255.255.224
no synchronization
aggregate-address 200.52.1.0 255.255.255.0
The BGP peers 172.16.1.1 and 10.1.1.2 in the above configuration are active, as well as the interface with IP address 200.52.1.192. Based on this information, which statement below is true about your configuration?

A. Router CertK has an IBGP connection with neighbor 10.1.1.2.
B. Router CertK has an EBGP connection with the neighbor 172.16.1.1.
C. Router CertK advertises only a route 200.52.1.0 255.255.255.0 in BGP.
D. Router CertK advertises only a route 200.52.1.192 255.255.255.224 in BGP.
E. Router CertK advertises both the routes to 200.52.1.0 255.255.255.0 and 200.52.1.192 255.255.255.224 in BGP.

Answer: E

Explanation:
The aggregate route, 200.52.1.0 255.255.255.0, and the more specific route, 200.52.1.192 255.255.255.224, will both be advertised. To only advertise the aggregate route the summary-only option of the aggregate-address command must be used.
Note: The aggregate-address command is used to create an aggregate, or summary, entry in the BGP table.
Incorrect Answers:
A: The neighbor 10.1.1.2 belongs to another autonomous system and is a EBGP neighbor, not an IBGP neighbor.
B: Router A and neighbor 172.16.1.1 belong to the same autonomous system and is an IBGP neighbor, not an EBGP neighbor.
C: The summary-only option of the aggregate-address command is used to only advertise the summary and not the specific routes.
D: The aggregate route, 200.52.1.0 255.255.255.0, will be advertised as well.

QUESTION 380:

The Certkiller network consists of two separate autonomous systems as shown below:

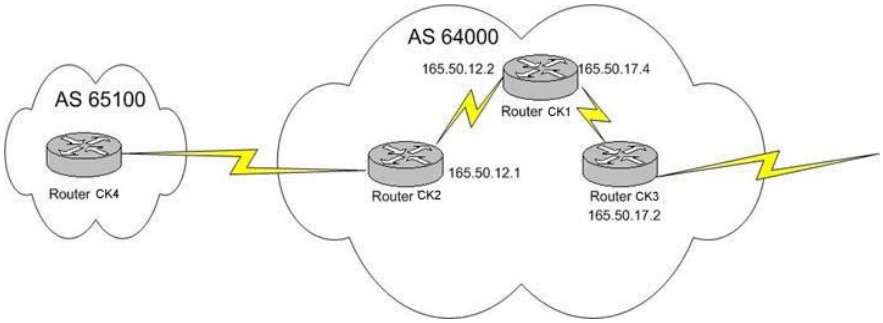

You need to configure Router CK2 as a BGP route reflector and Router CK1 as the client. Assuming that Router CK3 isn't running BGP, which two of the commands below would you enter on CK2 to satisfy your goals? (Select two)

A. neighbor 165.50.12.1 remote-as 65100
B. neighbor 165.50.12.2 remote-as 64000
C. neighbor 165.50.12.1 route-reflector-client
D. neighbor 165.50.12.2 route reflector-client

Answer: B, D

Explanation:
B: Router CK2 (config-router)# neighbor 165.50.12.2 remote-as 64000
We configure router CK1 (165.50.12.2) as a neighbor in AS 64000.
D: Router CK2 (config-router)# neighbor 165.50.12.2 route-reflector-client
Configures the router CK2 as a BGP route reflector and configures the specified neighbor CK1 (165.50.12.2) as its client.
Incorrect Answers:
A: We must specify router CK1 as neighbor, not CK2 itself (165.50.12.1). Furthermore, we should use the local AS (64000), not the remote AS 65100.
C: We must specify router CK1 as route reflector client, not CK2 itself (165.50.12.1).

QUESTION 381:

The Certkiller network consists of a series of routers that are all configured for IBGP. Which one of the following IBGP characteristics is true?

A. The IBGP routers must always be fully meshed.
B. The IBGP routers can be in a different AS.
C. The IBGP routers must be directly connected.
D. The IBGP routers do not need to be directly connected.
E. None of the above are true.

Answer: D

Explanation:
The IBGP routers do not have to be directly connected. The remote IBGP peers need only be reachable via a TCP connection. For example, if the network is also running an interior routing protocol such as EIGRP or OSPF, the remote IBGP router could be many hops away, as long as it is reachable via the IGP that is being used.
Incorrect Answers:
A: Using route reflectors or confederations a full mesh topology is not necessary.
B: The IBGP routers must be placed in the same AS. Peers that are in different autonomous systems are using EBGP, not IBGP.
C: The IBGP routers do not have to be directly connected.

QUESTION 382:

Which of the following statements regarding BGP peer groups are true? (Select two)

A. Peer members inherit all options of the peer group.
B. Peer groups can be used to simplify BGP configurations.
C. Peer groups are optional non-transitive attributes for BGP.
D. A peer group allows options that affect outbound updates to be overridden.
E. A common name should be used on all routers because this information is passed between neighbors.

Answer: A, B

Explanation:
A BGP peer group is a group of BGP neighbors with the same update policies.
A: By default members of the peer group inherit all of the configuration options of the peer group.
Note: Members can also be configured to override these options
B: Peer groups are normally used to simplify router configurations when many neighbors have the same policy.
Note: BGP neighbors who share the same outbound policies can be grouped together in what is called a BGP peer group. Instead of configuring each neighbor with the same policy individually, Peer group allows to group the policies which can be applied to individual peer thus making efficient update calculation along with simplified configuration.
Incorrect Answers:
C: The only Optional non-transitive attribute in BGP is MED.
D: Does not apply.
E: Using a common router name provides no BGP benefits at all.
Reference: BGP Peer Groups, http://www.cisco.com/warp/public/459/29.html

---

QUESTION 383:

A BGP router is configured as shown below:

```
interface ethernet 0
ip address 10.10.10.1 255.255.0.0
!
int serial 0
ip address 172.16.1.1 255.255.255.252
!
router bgp 65001
neighbor 192.168.1.1 remote-as 65002
```

Based on the above configuration, which of the following BGP statements would inject the 10.10.0.0/16 prefix into the BGP routing table?

A. network 10.0.0.0
B. network 10.10.0.0 mask 255.255.0.0
C. network 10.10.10.1 mask 255.255.255.255
D. network 10.10.10.0 mask 255.255.255.0
E. network 10.0.0.0 mask 255.255.0.0

Answer: B

Explanation:
The /16 mask is equal to 255.255.0.0, so answer choice B matches the address and the mask. To specify the route as classless, the mask keyword should be included or the network will be summarized at the network boundary.

QUESTION 384:

Why would a systems administrator enter in the network
command when configuring BGP?

A. Local routes matching the network command are filtered from the BGP routing table.
B. Local routes matching the network command can be installed into the BGP routing table.
C. Routes matching the network command will be filtered from BGP routing updates.
D. External routes matching the network command will be installed into the BGP routing table.
E. None of the above.

Answer: B

Explanation:
Use the network router configuration command to permit BGP to advertise a network if it is present in the IP routing table.
Reference: Building Scalable Cisco Networks (Cisco Press) page 342

QUESTION 385:

The Certkiller network is depicted in the following diagram:

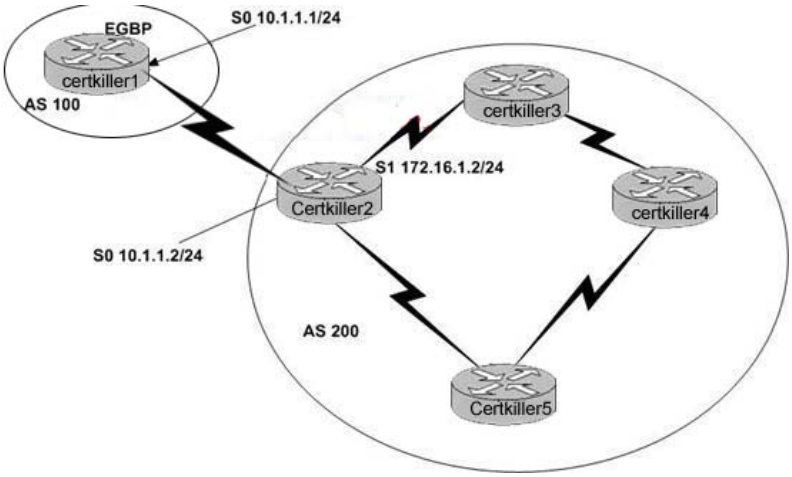

Which of the following command sets would you use if you want Certkiller 1 and Certkiller 2 to exchange BGP routing information? (Select two)

A. Certkiller 2(config)#router bgp 200
Certkiller 2(config-router)#neighbor 10.1.1.1 remote-as 100
B. Certkiller 1(config)#router bgp 100
Certkiller 1(config-router)#exit
Certkiller 1(config)#interface S0
Certkiller 1(config-if)#neighbor 10.1.1.2 remote-as 200
C. Certkiller 1(config)#router bgp 100
Certkiller 1(config-router)#neighbor 10.1.1.2 remote-as 200
D. Certkiller 2(config)#router bgp 100
Certkiller 2(config-router)exit
Certkiller 2(config)#interface S0
Certkiller 2(config-if)#neighbor remote-as 100
E. Certkiller 1(config)#router bgp 100
Certkiller 1(config-router)#network 10.0.0.0
Certkiller 1(config-router)#neighbor 10.1.1.2 remote-as 100
F. Certkiller 2(config)#router bgp 200
Certkiller 2(config-router)#network 10.0.0.0
Certkiller 2(config-router)#network 10.1.1.1 remote-as 200

Answer: A, C

Explanation:
Choices A and C give the correct syntax for configuring EBGP peering sessions.
Incorrect Answers:
B, D: The BGP peering configurations are placed under the BGP routing process, not in interface configuration mode.
E, F: The network 10.0.0.0 command is not required here, since this is an EBGP peering session using the directly connected interface. Since each neighbor shares the 10.0.0.0 network they each already know how to reach this network.

QUESTION 386:

Router CK-1 is configured for BGP routing as shown below:
router bgp 65300
network 27.0.0.0
neighbor 192.23.1.1 remote-as 65300
From the perspective of router CK-1, what kind of router is the router with IP address 192.23.1.1?

A. A peer router running IBGP
B. A peer router running EBGP
C. A community member running IBGP
D. A peer group member running IBGP
E. A peer group member running EBGP

Answer: A

Explanation:
Both the local and remote router is configured with the same autonomous system number so they are peer routers running IBGP.

QUESTION 387:

When you're configuring BGP on a Cisco router, what is true of the command " network"?

A. Local routes matching the network command are filtered from the BGP routing table.
B. Local routes matching the network command can be installed into the BGP routing table.
C. Sending and receiving BGP updates is controlled by using a number of different filtering methods.
D. The route to a neighbor autonomous system must have the correct MED applied to be installed into BGP routing table.
E. None of the above.

Answer: B

Explanation:
The network command allows BGP to advertise an IGP route if it is already in the IP table. A matching route must exist in the routing table before the network is announced. The network command is used to permit BGP to advertise a network if it is present in the IP routing table.

QUESTION 388:

Router CK1 needs to be configured to advertise a specific network. Which of the

following commands would you use if you wanted to advertise the subnet 154.2.1.0 255.255.255.0 to the EBGP neighbors on your subnet?

A. Router (config-router)#network 154.2.1.0
B. Router (config-router)#network 164.2.1.0
C. Router (config-router)#network-advertise 154.2.1.0
D. Router (config-router)#network 154.2.1.0 mask 255.255.255.0
E. None of the above

Answer: D

Explanation:
The network command is used to specify the networks to be advertised by the Border Gateway Protocol (BGP) and multiprotocol BGP routing processes.
Syntax: network network-number [masknetwork-mask] [route-map map-name]
Mask and route-map are optional. If the
mask keyword is configured, then an exact match must exist in the routing table.
Incorrect Answers:
A: If we do not specify the subnet mask then additional networks are allowed to be advertised. The classful subnet mask of 154.2.1.0 is 255.255.0.0 - a Class B network.
B: This is using the incorrect IP address, as well as a missing subnet mask.
C: The network-advertise is an invalid command.

QUESTION 389:

You are the administrator of a company with BGP connections to multiple ISP's. How could you configure BGP to make it favor one particular ISP for outbound traffic?

A. Configure weight
B. Enable route reflector
C. Create a distribute list
D. Enable the Longer Autonomous System path option.
E. All of the above.

Answer: A

Explanation:
If the router learns about more than one route to the same destination, the route with the highest weight will be preferred. Weight is a Cisco BGP parameter that is local to the router. When terminating multiple ISP connections into the same router, weight can be used to affect which path is chosen for outbound traffic.
Incorrect Answers:
B: A route reflector cannot be used to influence outbound traffic. A route reflector modifies the BGP split horizon rule by allowing the router configured as the route reflector to propagate routes learned by IBGP to other IBGP peers. This saves on the

number of BGP TCP sessions that must be maintained, and also reduces the BGP routing traffic.
C: Distribute lists restrict the routing information that the router learns or advertises. By itself a distribute list cannot make routes from one ISP be preferred to routers from another ISP.
D: This choice describes ASD path pre-pending, which would be used to influence the path that incoming traffic takes, not outgoing.

QUESTION 390:

What are two solutions to overcome the full mesh requirement with iBGP? (Choose two)

A. BGP groups
B. BGP local preference
C. Route reflector
D. Confederation
E. Aggregate-address

Answer: C, D

Explanation:
In general, all IBGP peers must be configured to be fully meshed. If they are not, then all of the IBGP routers will not have the updated information from the external BGP routers. There are two ways to overcome the scalability issues of a full IBGP mesh: route reflectors and confederations. With route reflectors, internal BGP routers peer only with the route reflector. With confederations, the AS is broken up into smaller, more manageable sub autonomous systems.

QUESTION 391:

An ISP is running a large IBPG network with 25 routers. The full mesh topology that is currently in place is inefficiently using up bandwidth from all of the BGP traffic.
What can the administrator configure to reduce the number of BGP neighbor relationships within the AS?

A. Route reflectors
B. Route maps
C. Route redistribution
D. Peer groups
E. Aggregate addresses

Answer: A

Explanation:

In general, all IBGP peers must be configured to be fully meshed. If they are not, then all of the IBGP routers will not have the updated information from the external BGP routers. There are two ways to overcome the scalability issues of a full IBGP mesh: route reflectors and confederations. With route reflectors, internal BGP routers peer only with the route reflector, and then the route reflectors connect with each other. This can considerably reduce the number of IBGP sessions. Another solution to the scalability problem of IBGP is the use of confederations. With confederations, the AS is broken up into smaller, more manageable sub autonomous systems.

QUESTION 392:

DRAG DROP
Arrange the BGP attributes on the left in order or priority.
Your boss at certkiller.com asks you to place the BGP attributes in the correct order used for determining a route.

BGP Attributes, select from these	Place here
originate route	Place first here
AS_path	Place second here
weight	Place third here
Local preference	Place fourth here
MED	Place firth here

Answer:
Your boss at Certkiller.com asks you to place the BGP attributes in the correct order used for determining a route.

BGP Attributes, select from these	Place here
	weight
	Local Preference
	Originate route
	AS_Path
	MED

Explanation:
How the Best Path Algorithm Works:
BGP assigns the first valid path as the current best path. It then compares the best path with the next path in list, until it reaches the end of the list of valid paths. The following is a list of rules used to determine the best path.

1. Prefer the path with the highest WEIGHT.
Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's configured.
2. Prefer the path with the highest LOCAL_PREF. Note the following:
3. 1. Path without LOCAL_PREF is considered as having the value set with the bgp default local-preference command, or 100 by default.
2. Prefer the path that was locally originated via a network or aggregate BGP subcommand, or through redistribution from an IGP. Local paths sourced by network or redistribute commands are preferred over local aggregates sourced by the aggregate-address command.
3. Prefer the path with the shortest AS_PATH. Note the following:
4. 1. The AS_CONFED_SEQUENCE and AS_CONFED_SET are not included in the AS_PATH length.
2. Prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE.
3. Prefer the path with the lowest multi-exit discriminator (MED). Note the following:
4. 1. This comparison is only done if the first (neighboring) AS is the same in the two paths; any confederation sub-ASs are ignored. In other words, MEDs are compared only if the first AS in the AS_SEQUENCE is the same for multiple paths. Any preceding AS_CONFED_SEQUENCE is ignored.
2. If bgp always-compare-med is enabled, MEDs are compared for all paths. This option needs to be enabled over the entire AS, otherwise routing loops can occur.
3. If bgp bestpath med-confed is enabled, MEDs are compared for all paths that consist only of AS_CONFED_SEQUENCE (paths originated within the local confederation).
4. Paths received from a neighbor with a MED of 4,294,967,295 will have the MED changed to 4,294,967,294 before insertion into the BGP table.
5. Paths received with no MED are assigned a MED of 0, unless bgp bestpath missing-as-worst is enabled, in which case they are assigned a MED of 4,294,967,294.
6. The bgp deterministic med command can also influence this step as demonstrated in the How BGP Routers Use the Multi-Exit Discriminator for Best Path Selection.
7. Prefer external (eBGP) over internal (iBGP) paths. If bestpath is selected, go to Step 9 (multipath).
Note:Paths containing AS_CONFED_SEQUENCE and AS_CONFED_SET are local to the confederation, and therefore treated as internal paths. There is no distinction between Confederation External and Confederation Internal.
8. Prefer the path with the lowest IGP metric to the BGP next hop. Continue, even if bestpath is already selected.
9. Check if multiple paths need to be installed in the routing table for BGP Multipath. Continue, if bestpath is not selected yet.
10. When both paths are external, prefer the path that was received first (the oldest one). This step minimizes route-flap, since a newer path will not displace an older one, even if it would be the preferred route based on the next decision criteria (Steps 11, 12, and 13). Skip this step if any of the following is true:
1. The bgp best path compare-routerid command is enabled.
Note:This command was introduced in Cisco IOS(r) Software Releases 12.0.11S, 12.0.11SC, 12.0.11S3, 12.1.3, 12.1.3AA, 12.1.3.T, and 12.1.3.E.

10. 1. The router ID is the same for multiple paths, since the routes were received from the same router.
2. There is no current best path. An example of losing the current best path occurs when the neighbor offering the path goes down.
3. Prefer the route coming from the BGP router with the lowest router ID. The router ID is the highest IP address on the router, with preference given to loopback addresses. It can also be set manually using the bgp router-id command.
Note:If a path contains route-reflector (RR) attributes, the originator ID is substituted for the router ID in the path selection process.
12. If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length. This will only be present in BGP route-reflector environments. It allows clients to peer with RRs or clients in other clusters. In this scenario, the client must be aware of the RR-specific BGP attribute.
13. Prefer the path coming from the lowest neighbor address. This is the IP address used in the BGP neighbor configuration, and corresponds to the remote peer used in the TCP connection with the local router.

QUESTION 393:

The Certkiller BGP routing table consists of the following network routes:

192.168.12.0/24
192.168.13.0/24
192.168.14.0/24
192.168.15.0/24

What is the correct command to summarize these prefixes into a single summary prefix of 192.168.12.0/22 while also allowing for the advertisement of the more specific prefixes?

A. network 192.168.12.0 mask 255.255.252.0
B. network 192.168.12.0 mask 0.0.3.255
C. network 192.168.12.0
D. aggregate-address 192.168.12.0 255.255.252.0
E. aggregate-address 192.168.12.0 255.255.252.0 summary-only
F. aggregate-address 192.168.12.0 255.255.252.0 as-set

Answer: D

Explanation:
To summarize BGP prefixes into one aggregated route, use the "aggregate-address" command. When used alone, this will advertise the aggregate route, along with the individual specific routing entries. To advertise only the aggregated route, use the "summary-only" keyword, as specified in choice E.

QUESTION 394:

The Certkiller BGP network is displayed below:

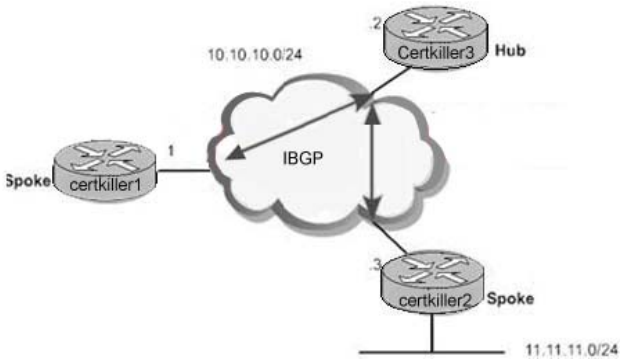

Certkiller 3 is the hub router and Certkiller 1 and Certkiller 2 are the spokes. There are no virtual circuits between the spoke locations. Each router is in a separate AS.
What is needed to successfully route traffic to the 11.11.11.0/24 network from Certkiller 1?

A. The neighbor 10.10.10.1 next-hop-self command on Certkiller 1.
B. The neighbor 10.10.10.1 next-hop-self command on Certkiller 2.
C. The neighbor 10.10.10.1 next-hop-self command on Certkiller 3.
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)

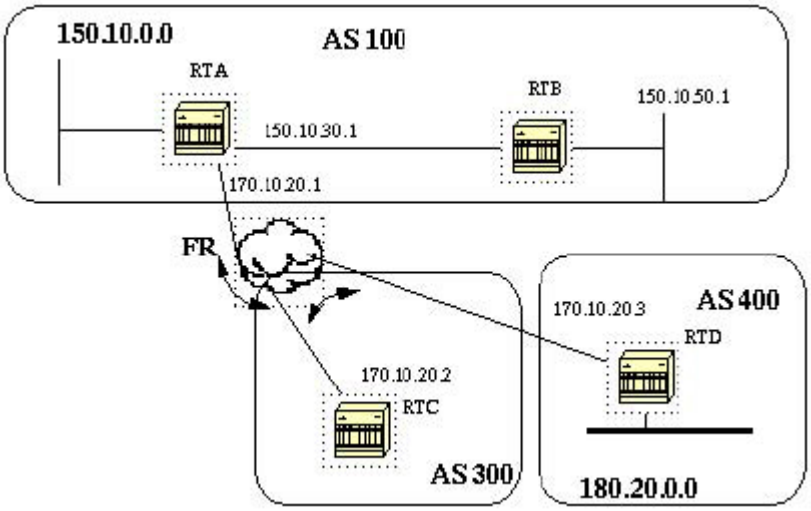

If the common media as you see in the shaded area above is a frame relay or any NBMA cloud then the exact behavior will occur as if we were connected via Ethernet. RTC will advertise 180.20.0.0 to RTA with a next hop of 170.10.20.3.
The problem is that RTA does not have a direct PVC to RTD, and cannot 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 395:

The following was seen on router CK1 :

```
Certkiller1#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
```

Based on the show ip bgp summary output in the exhibit, which two statements are true? Select 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 is attempting to establish a BGP peering to 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 BGP peering with the 10.2.2.2 neighbor.

Answer: A, F

Explanation:
In BGP, established valid neighbors will have a number under the state/pfxRcd heading. Neighbors that show an active state mean that the router is actively trying to establish a peering relationship. Idle peers are not active, and not actively trying to peer. In this case, only the peer to 10.1.1.1 is currently operational with two route prefixes received, while the peer to 10.2.2.2 is trying to become established.

QUESTION 396:

While verifying BGP operation on the Certkiller router, you issue the "show ip bgp" command as shown below:
routerCK>show ip bgp
BGP table version is 1046033, local router ID is 198.32.162.100
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
* > 143.16.0.0 128.214.63.2 0 400 0 200 1
* 143.16.0.0 192.208.10.5 0 300 0 300 1
* 143.16.0.0 143.16.63.5 0 100 0 200 1
* 143.16.0.0 203.250.13.41 0 100 0 500 1
From the information above, which path will the network 143.16.0.0 prefer to take to exit the AS?

A. 128.214.63.2
B. 192.208.10.5
C. 128.213.63.5
D. 203.250.13.41
E. All of the above will be used in a round robin fashion.

Answer: A

Explanation:
Local preference (LocPref) is a well-known discretionary attribute that provides an indication to routers in the AS about which path is preferred to exit the AS. A path with a higher local preference is more preferred. In this scenario the following entry has the highest local preference value of 400.
Network Next Hop Metric LocPrf Weight Path
* > 128.213.0.0 128.214.63.2 0 400 0 200 1
The preferred exit path of the AS is therefore 128.214.63.2, as noted by the">" which refers to the best path for this destination.

QUESTION 397:

The IP BGP table of router P1R3 is displayed in the following diagram:

```
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
```

Based on the P1R3 output shown above, which statement is true?

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.
E. None of the above

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. Within a router, the path with the highest weight will be preferred.

QUESTION 398:

Which IOS command could you use to verify if a BGP router is or isn't a route reflector?

A. show bgp neighbor
B. show running-config
C. show route-reflector
D. show route-reflector-client
E. None of the above

Answer: A

Explanation:
The show bgp neighbor command indicates if a particular neighbor is a route reflector client.
Note: Route reflectors modify the BGP split horizon rule by allowing the router configured as the route reflector to propagate routes learned by IBGP to other IBGP peers.
Incorrect Answers:
B: Displaying the current configuration would not be useful in finding information in the route reflector status.
C, D: There are no such commands.

QUESTION 399:

Which IOS command would you enter if you wanted to view a list of IBGP and EBGP neighbor relationships that are configured?

A. show ip bgp

B. show ip bgp paths
C. show ip bgp peers
D. show ip bgp summary
E. show ip bgp protocols

Answer: D

Explanation:
The show ip bgp summary command displays the status of all BGP connections.
Neighbors with corresponding AS values will be listed; both interior and external.
Incorrect Answers:
A: The show ip bgp command displays routes in the BGP routing table, not the neighbors.
B: The show ip bgp paths command is used to display all the BGP paths in the database. However, it does not list the neighbors.
C, E: There is no such command.
Reference:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_r/1rprt1/1rbgp.htm

---

QUESTION 400:

The 192.168.0.0/16 network is not being propagated throughout the network via BGP as expected. Observe the BGP configuration commands from the advertising router shown below.

```
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
```

What is the reason the 192.168.0.0/16 route is not being advertised?

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:
The network 192.168.0.0 statement is missing mask 255.255.0.0. Without the mask command used in a network statement, the route may not get properly injected into the BGP routing process.

QUESTION 401:

Exhibit:
Certkiller A# 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, Ethernet0
P 192.168.1.32/28, 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 192.168.1.77 (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
Based on the above show ip eigrp topology output, which three statements are true? (Choose three.)

A. Certkiller A is in AS 200
B. Certkiller A will balance between three paths to reach the 192.168.1.48/28 prefix, because all three paths have the same AD of 40512000.
C. The best path for Certkiller A to reach the 192.168.1.48/28 prefix is via 192.168.1.66.
D. 40512000 is the AD to reach the 192.168.1.48/28 prefix.
E. All of 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 Certkiller A 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 402:

There are numerous attributes used to determine the best path in BGP. Which BGP path attribute is Cisco proprietary?

A. MED

B. Local preference
C. Next-hop
D. AS-path
E. Origin
F. Weight
G. None of the above

Answer: F

Explanation:
The Weight attribute is similar to the Local Preference attribute in that it gives higher preference to the route that has a higher weight. The difference is that the weight parameter is local to the router and is not exchanged between routers. The weight parameter influences routes coming from different providers to the same router, one router with multiple connections to two or more providers. The weight parameter has a higher precedence than any other attribute. It is the most important attribute when determining route preference. The Weight attribute is a Cisco proprietary attribute.

QUESTION 403:

DRAG DROP
Your work as a network technician at Certkiller .com. You need to explain to your trainee, Miss Certkiller, the order for which BGP attributes are used for determining a route.

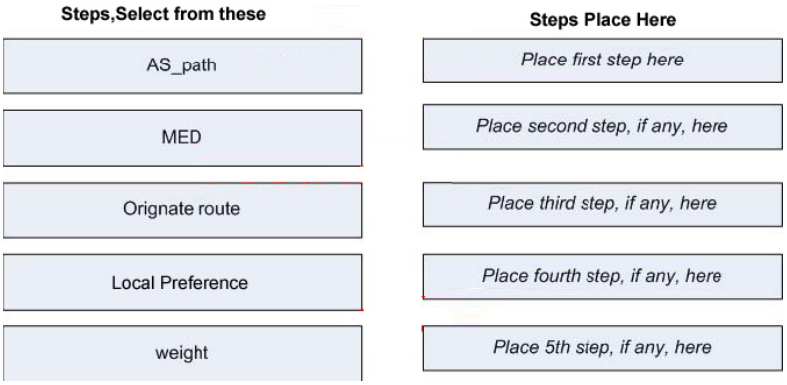

Answer:

Explanation:
1. Weight: The Weight attribute is similar to the Local Preference attribute in that it gives higher preference to the route that has a higher weight. The difference is that the weight parameter is local to the router and is not exchanged between routers. The weight parameter influences routes coming from different providers to the same router, one router with multiple connections to two or more providers. The weight parameter has a higher precedence than any other attribute. It is the most important attribute when determining route preference. The Weight attribute is a Cisco proprietary attribute.
2. Local Preference: Local Preference is a well-known discretionary attribute, type code 5. The Local Preference attribute is a degree of preference given to a route for comparison with other routes for the same destination. Higher Local Preference values are preferred. Local Preference, as indicated by the name, is local to the AS and is exchanged between IBGP peers only. Local Preference is not advertised to EBGP peers. Routers within a multihomed AS may learn that they can reach the same destination network by way of neighbors in two or more different autonomous systems. In effect, there could be two or more exit points from the local AS to any given destination. Use the Local Preference attribute to force the BGP routers to prefer one exit point instead of another when routing to a particular destination network. Because this attribute is communicated within all BGP routers inside the AS, all BGP routers will have a common view on how to exit the AS.
3. Originate Route: The Origin attribute is a well-known mandatory attribute, type code 1, that indicates the origin of the routing update. BGP allows the following three types of origins:
1. IGP - The prefix is internal to the originating AS.
2. EGP - The prefix was learned by way of some EGP, such as BGP.
3. Incomplete - The prefix was learned by some other means, probably redistribution.
4. AS_Path: An AS_Path attribute is a well-known mandatory attribute, type code 2. It is the sequence of AS numbers that a route has traversed to reach a destination. The AS that

originates the route adds its own AS number when sending the route to its external BGP peers. Thereafter, each AS that receives the route and passes it on to other BGP peers will prepend its own AS number to the list. Prepending is the act of adding the AS number to the beginning of the list. The final list has all the AS numbers that a route has traversed with the AS number of the AS that originated the route at the end of the list. This type of AS_Path list is called an AS_Sequence because all the AS numbers are ordered sequentially.
BGP uses the AS_Path attribute as part of the routing updates, update packet, to ensure a loop free topology on the Internet. Each route that gets passed between BGP peers will carry a list of all AS numbers that the route has already been through. If the route is advertised to the AS that originated it, that AS will see itself as part of the AS_Path attribute list and will not accept the route. BGP speakers prepend their AS numbers when advertising routing updates to other autonomous systems, external peers. When the route is passed to a BGP speaker within the same AS, the AS_Path information is left intact.
5. MED : 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.

QUESTION 404:

Router CK1 is a BGP router and it 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 origin code
B. Lowest neighbor IP address
C. Shortest AS-path
D. Lowest MED
E. Highest local preference
F. None of the above

Answer: E

Explanation:
Local Preference is a well-known discretionary attribute, type code 5. The Local Preference attribute is a degree of preference given to a route for comparison with other routes for the same destination. Higher Local Preference values are preferred. Local Preference, as indicated by the name, is local to the AS and is exchanged between IBGP peers only. Local Preference is not advertised to EBGP peers.
Routers within a multihomed AS may learn that they can reach the same destination network by way of neighbors in two or more different autonomous systems. In effect,

there could be two or more exit points from the local AS to any given destination. Use the Local Preference attribute to force the BGP routers to prefer one exit point instead of another when routing to a particular destination network. Because this attribute is communicated within all BGP routers inside the AS, all BGP routers will have a common view on how to exit the AS.

QUESTION 405:

Router CK1 is a BGP router used to find the best path to traffic destined to the Internet. 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. Next-hop
B. MED
C. Origin
D. Weight
E. Local preference
F. AS-path
G. None of the above

Answer: D

Explanation:
The Weight attribute is similar to the Local Preference attribute in that it gives higher preference to the route that has a higher weight. The difference is that the weight parameter is local to the router and is not exchanged between routers. The weight parameter influences routes coming from different providers to the same router, one router with multiple connections to two or more providers. The weight parameter has a higher precedence than any other attribute. It is the most important attribute when determining route preference. The Weight attribute is a Cisco proprietary attribute.

QUESTION 406:

A Certkiller router was configured as shown below:

```
router bgp 65003
network 10.10.20.0 mask 255.255.255.0
 neighbor 12.113.176.25 remote-as 7018
 neighbor 12.113.176.25 timers 15 45
 neighbor 12.113.176.25 send-community
 neighbor 12.113.176.25 route-map Advertised_Customer_Networks out
 neighbor 12.113.176.25 ote-as 2386
 neighbor 12.10.20.2 next-hop-self
 neighbor 12.10.20.2 route-map iBGP_Advertised_Routes out
 no auto-summary
!
ip prefix-list Local100 seq 5 permit 10.110.110.110/32
ip prefix-list Local100 seq 10 permit 10.10.20.0/24
!
ip prefix-list Local90 seq 5 permit 10.109.109.109/32
!
route-map Advertised_Customer_Networks permit 10
 match ip address prefix-list Local100
 set community 7018:100
!
route-map Advertised_Customer_Networks permit 20
 match ip address prefix-list Local90
 set community 7018:90
!
route-map iBGP_Advertised_Routes deny 10
 match route-type local
!
route-map iBGP_Advertised_Routes permit 20
 match route-type external
```

Based on the configuration shown above, which router command is responsible for sending the community attributes of the router to its EBGP neighbor?

A. neighbor 12.10.20.2 route-map iBGP_Advertised_Routes out
B. neighbor 12.113.176.25 send-community
C. neighbor 12.113.176.25 route-map Advertised_Customer_Networks out
D. neighbor 12.10.20.2 next-hop-self
E. None of the above

Answer: B

Explanation:
A community is a group of destinations that have some common attribute. Destinations can be added to a community by setting their COMMUNITY attribute. Routing policies can then be enforced based on using the COMMUNITY attribute to affect routing decisions. Destinations can be grouped into a single community or multiple communities regardless of their physical location and autonomous system. By default, all routes belong to the Internet, a well-known community.
neighbor peer-address send-community
peer-address - the address used by a BGP peer for the BGP session.
Once communities have been configured for a network, you can use the communities to filter and manipulate the routes belonging to the community. In order to accomplish this, you first need to create a community list. The community list contains all of the communities that you want the policy to affect.

QUESTION 407:

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:
How the Best Path Algorithm Works:
BGP assigns the first valid path as the current best path. It then compares the best path with the next path in list, until it reaches the end of the list of valid paths. The following is a list of rules used to determine the best path.
1. Prefer the path with the highest WEIGHT.
Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's

configured.
2. Prefer the path with the highest LOCAL_PREF.
3. Prefer the path that was locally originated via a network or aggregate BGP subcommand, or through redistribution from an IGP. Local paths sourced by network or redistribute commands are preferred over local aggregates sourced by the aggregate-address command.
4. Prefer the path with the shortest AS_PATH. Note the following:
5. Prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE.
6. Prefer the path with the lowest multi-exit discriminator (MED). Note the following:
7. Prefer external (eBGP) over internal (iBGP) paths. If bestpath is selected, go to Step 9 (multipath).
8. Prefer the path with the lowest IGP metric to the BGP next hop. Continue, even if bestpath is already selected.
9. Check if multiple paths need to be installed in the routing table for BGP Multipath. Continue, if bestpath is not selected yet.
10. When both paths are external, prefer the path that was received first (the oldest one). This step minimizes route-flap, since a newer path will not displace an older one, even if it would be the preferred route based on the next decision criteria (Steps 11, 12, and 13).
11. Prefer the route coming from the BGP router with the lowest router ID. The router ID is the highest IP address on the router, with preference given to loopback addresses. It can also be set manually using the bgp router-id command.
12. If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length. This will only be present in BGP route-reflector environments. It allows clients to peer with RRs or clients in other clusters. In this scenario, the client must be aware of the RR-specific BGP attribute.
13. Prefer the path coming from the lowest neighbor address. This is the IP address used in the BGP neighbor configuration, and corresponds to the remote peer used in the TCP connection with the local router.
Reference:
http://www.cisco.com/en/US/tech/ CK3 65/technologies_tech_note09186a0080094431.shtml

---

QUESTION 408:

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:
How the Best Path Algorithm Works:
BGP assigns the first valid path as the current best path. It then compares the best path with the next path in list, until it reaches the end of the list of valid paths. The following is a list of rules used to determine the best path.
1. Prefer the path with the highest WEIGHT.
Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's configured.
2. Prefer the path with the highest LOCAL_PREF. Note the following:
3. 1. Path without LOCAL_PREF is considered as having the value set with the bgp default local-preference command, or 100 by default.
2. Prefer the path that was locally originated via a network or aggregate BGP subcommand, or through redistribution from an IGP. Local paths sourced by network or redistribute commands are preferred over local aggregates sourced by the aggregate-address command.
3. Prefer the path with the shortest AS_PATH. Note the following:
4. 1. The AS_CONFED_SEQUENCE and AS_CONFED_SET are not included in the AS_PATH length.
2. Prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE.
3. Prefer the path with the lowest multi-exit discriminator (MED).
4. Prefer external (eBGP) over internal (iBGP) paths. If bestpath is selected, go to Step 9 (multipath).
5. Prefer the path with the lowest IGP metric to the BGP next hop. Continue, even if bestpath is already selected.
6. Check if multiple paths need to be installed in the routing table for BGP Multipath. Continue, if bestpath is not selected yet.
7. When both paths are external, prefer the path that was received first (the oldest one). This step minimizes route-flap, since a newer path will not displace an older one, even if it would be the preferred route based on the next decision criteria (Steps 11, 12, and 13).
8. Prefer the route coming from the BGP router with the lowest router ID. The router ID is the highest IP address on the router, with preference given to loopback addresses. It can also be set manually using the bgp router-id command.
9. If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length. This will only be present in BGP route-reflector environments. It allows clients to peer with RRs or clients in other clusters. In this scenario, the client must be aware of the RR-specific BGP attribute.
10. Prefer the path coming from the lowest neighbor address. This is the IP address used in the BGP neighbor configuration, and corresponds to the remote peer used in the TCP connection with the local router.

QUESTION 409:

The Certkiller network admin needs to configure IP multicast on a new router.
Which well-defined routing protocol would a network administrator configure on multicast routers when member routers are widely dispersed?

A. Distance Vector Multicast Routing Protocol (DVMRP)
B. Protocol Independent Multicast Dense Mode (PIM-DM)
C. Protocol Independent Multicast Sparse Mode (PIM-SM)
D. Core-Based Trees (CBT)
E. Multicast Open Shortest Path First (MOSPF)
F. None of the above

Answer: C

QUESTION 410:

IP multicast was implemented on a Certkiller LAN segment. Which two multicast IP addresses can be represented by the multicast MAC address 0100.5e0A.0A07? (Select two)

A. 229.11.10.7
B. 228.10.10.7
C. 229.138.10.8
D. 228.10.138.7
E. 229.138.10.7
F. 228.10.10.8

Answer: B, E

QUESTION 411:

Certkiller uses IGMP version 2 in their IP multicast network. How is the designated querier elected in IGMPv2?

A. The first router to appear on a subnet is designated.
B. The router with the lowest IP address on a subnet is designated.
C. The host that responds first to the election query is designated.
D. The host with the lowest MAC address on a segment is designated.
E. None of the above.

Answer: B

QUESTION 412:

Study the exhibit below carefully. Switch Certkiller 3 is receiving IGMP frames only on interface FastEthernet 0/3. Given IGMP snooping is enabled, out of which port or ports will switch Certkiller 3 forward multicast traffic?

```
Certkiller3 # show vlan

VLAN Name                Status     ports
----------------------------------------------------------------
1    default             active     Fa0/1  Fa0/2, Fa0/3, Fa0/4
10   VLAN0010            active     Fa0/5, Fa0/6, Fa0/9,Fa0/8
20   VLAN0020            active     Fa0/9, Fa0/10, Fa0/11,Fa0/12
30   VLAN0030            active     Fa0/13, Fa0/14,Fa0/15, Fa0/16,
40   VLAN0040            active     Fa0/17, Fa0/18,Fa0/19, Fa0/20,
50   VLAN0050            active     Fa0/21, Fa0/22,Fa0/23, Fa0/24

1002  Fddi-default                  act/unsup
1003 token-ring-default             act/unsup
1004 fddinet-default                act/unsup
1005 trnet-default                  act/unsup
```

A. FastEthernet 0/1, 0/2, 0/3, 0/4
B. FastEthernet 0/1, 0/2, 0/4
C. All ports
D. FastEthernet 0/1, 0/2, 0/4, 0/5, 0/6, 0/7, 0/8
E. FastEthernet 0/3
F. None of the ports

Answer: E

QUESTION 413:

Router CK1 is configured as a RP in the Certkiller multicast network. What is the purpose of a rendezvous point (RP)?

A. It acts as the designated router for a broadcast segment when multicast routing is enabled
B. It is used in PIM dense and sparse mode to create a database of all multicast sources
C. It acts as a meeting place for sources and receivers of multicast traffic
D. It is used in PIM dense mode to create a database of all multicast sources
E. None of the above

Answer: C

QUESTION 414:

The IGMP snooping feature has been enabled on the Certkiller multicast network. Which two statements are true about Internet Group Management Protocol (IGMP) snooping? (Select two)

A. IGMP snooping and Cisco Group Membership Protocol (CGMP) were developed to help Layer 3 switches make intelligent forwarding decisions on their own.
B. IGMP snooping is enabled with the ip multicast-routing global configuration command.
C. IGMP snooping and Cisco Group Membership Protocol (CGMP) can be used simultaneously on a switch.
D. IGMP snooping examines IGMP join/leave messages so that multicast traffic is forwarded only to hosts that sent an IGMP message toward the router.
E. IGMP snooping is an IP multicast constraining mechanism for Layer 2 switches.

Answer: D, E

QUESTION 415:

You have been tasked with configuring IP multicast on a new Certkiller router.
Which three IP multicast related statements are true? (Select three)

A. Multicast addresses 224.0.0.5 and 224.0.0.6 are limited scoped addresses that have been reserved for OSPF.
B. Multicast addresses 224.0.1.0 through 238.255.255.255 are called globally scoped addresses. They are used to multicast data between organizations and across the Internet.
C. Multicast addresses 239.0.0.0 through 239.255.255.255 are called limited scope addresses. They are constrained to a local group or organization.
D. Multicast addresses 224.0.0.0 through 224.0.0.255 are used for network protocols on local LAN segments. Because they are always transmitted with a Time to Live (TTL) of 1, they are never forwarded by a router.
E. The multicast address 224.0.0.1 is a globally scoped address that has been reserved for the Network Time Protocol (NTP) by the IANA.

Answer: B, C, D

QUESTION 416:

IP multicast needs to be configured on the Certkiller LAN switches. Which technology manages multicast traffic at Layer 2 by configuring Layer 2 LAN interfaces dynamically to forward multicast traffic only to those interfaces that want to receive it?

A. IGMP
B. PIM-DM
C. DVMRP
D. MOSPF
E. IGMP snooping
F. None of the above

Answer: E

QUESTION 417:

The Certkiller network administrator assigns an address of 239.255.8.5 to an application running on a device with an Ethernet MAC address of 01.b2.7d.05.f1.80.
Which Layer 2 multicast address will this device use?

A. 01.b2.7d.05.f1.80
B. 01.b2.7d.0a.08.05

C. 01.00.5e.05.f1.80
D. 01.00.5e.7F.08.05
E. ff.ff.ff.ff.ff.ff
F. Non of the above

Answer: D

QUESTION 418:

Study the exhibit of the Certkiller network carefully:

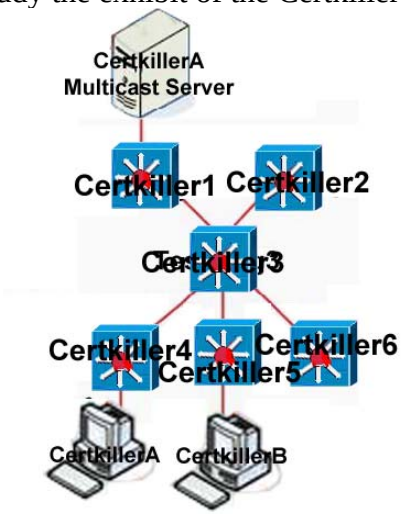

All multilayer switches are running PIM dense mode. Recipient A and Recipient B are sending IGMPv2 join messages to their respective multilayer switches. Based on this information, which statement is true?

A. Switch Certkiller 3 is the root of the multicast tree. Switches Certkiller 1, Certkiller 3, Certkiller 4, and Certkiller 5 will participate in the multicast tree once pruning has taken place.
B. The Multicast Server is the root of the multicast tree. Switches Certkiller 1, Certkiller 3, Certkiller 4, and Certkiller 5 will participate in the multicast tree once pruning has taken place.
C. Switch Certkiller 3 is the root of the multicast tree. Switches Certkiller 3, Certkiller 4, and Certkiller 5 will participate in the multicast tree once pruning has taken place.
D. The Multicast Server is the root of the multicast tree. Switches Certkiller 4 and Certkiller 5 will participate in the multicast tree once pruning has taken place.
E. Switch Certkiller 1 is the root of the multicast tree. Switches Certkiller 1, Certkiller 3, Certkiller 4, and Certkiller 5 will participate in the multicast tree once pruning has taken place.
F. Switch Certkiller 1 is the root of the multicast tree. Switches Certkiller 1, Certkiller 4, and Certkiller 5 will participate in the multicast tree once pruning has taken place.
G. None of the above.

Answer: B

QUESTION 419:

The Certkiller network is implementing IP multicasting to support a new application. Which two statements about multicast addressing are true? (Select two)

A. To calculate the Layer 2 multicast address, the host maps the last 23 bits of the IP address into the last 24 bits of the MAC address. The high-order bit is set to 1.
B. To calculate the Layer 2 multicast address, the host maps the last 24 bits of the IP address into the last 24 bits of the MAC address. The high-order bit is set to 0.
C. The last 3 bytes (24 bits) of the multicast MAC address are 0x01-00-5E. This is a reserved value that indicates a multicast application.
D. The first 3 bytes (24 bits) of the multicast MAC address are 0x01-00-5E. This is a reserved value that indicates a multicast application.
E. The first 23 bits of the multicast MAC address are 0x01-00-5E. This is a reserved value that indicates a multicast application.
F. To calculate the Layer 2 multicast address, the host maps the last 23 bits of the IP address into the last 24 bits of the MAC address. The high-order bit is set to 0.

Answer: D, F

QUESTION 420:

You wish to represent some IPv6 networks in a shorter form. What number is a valid representation for the 200F:0000:0000:AB00:0000:0000:0000:0000/56 IPv6 prefix?

A. 200F:0:0:AB/56
B. 200F::AB00/56
C. 200F::AB/56
D. 200F:0:0:AB00::/56
E. None of the above

Answer: D

QUESTION 421:

IGMP has been configured on a number of devices within the Certkiller IP multicast network. Which three statements are true about the Internet Group Management Protocol (IGMP)? (Select three)

A. IGMP version 3 enables a multicast receiving host to specify to the router which sources it should forward traffic from.
B. There are three IGMP modes: dense mode, sparse mode, and sparse-dense mode.
C. IGMP is a multicast routing protocol that makes packet-forwarding decisions independent of other routing protocols such as EIGRP.
D. IGMP snooping runs on Layer 3 routers.

E. IGMP messages are IP datagrams with a protocol value of 2, destination address of 224.0.0.2, and a TTL value of 1.
F. IGMP is used to register individual hosts with a multicast group.

Answer: A, E, F

QUESTION 422:

The IGMP snooping feature has been enabled on the Certkiller IP Multicast network. Which two statements are true about Internet Group Management Protocol (IGMP) snooping? (Select two)

A. IGMP snooping and Cisco Group Membership Protocol (CGMP) were developed to help Layer 3 switches make intelligent forwarding decisions on their own.
B. IGMP snooping is enabled with the ip multicast-routing global configuration command.
C. IGMP snooping and Cisco Group Membership Protocol (CGMP) can be used simultaneously on a switch.
D. IGMP snooping examines IGMP join/leave messages so that multicast traffic is forwarded only to hosts that sent an IGMP message toward the router.
E. IGMP snooping is an IP multicast constraining mechanism for Layer 2 switches.

Answer: D, E

QUESTION 423:

The Certkiller IP Multicast network has been upgraded in order to support IGMPv3. Which enhancement was added to IGMP version 3?

A. Destination filtering
B. Membership report message
C. Membership query message
D. Leave group message
E. Source filtering
F. None of the above

Answer: E

QUESTION 424:

You need to assign an IP multicast address to a new Certkiller application. Which three IP addresses are valid multicast addresses? (Select three)

A. 249.1.2.3
B. 169.254.23.59
C. 223.254.255.254

D. 227.43.34.2
E. 225.1.1.1
F. 238.3.3.3
G. 255.255.255.255

Answer: D, E, F

QUESTION 425:

A Certkiller device wants to participate in an IP multicast application. Which type of IGMP message is sent when a network client wants to join a multicast group?

A. Host membership report
B. Host membership notification
C. Host membership query
D. Host membership status
E. None of the above

Answer: A

QUESTION 426:

You need to verify the IP multicast configuration on a Certkiller router. Which three statements are true about IP multicast configuration? (Select three)

A. PIM sparse mode and PIM sparse-dense mode require an RP on the network.
B. PIM dense mode interfaces are always added to the multicast routing table in a router.
C. PIM sparse mode and PIM dense mode require an RP on the network.
D. PIM sparse-dense mode and PIM dense mode require an RP on the network.
E. PIM sparse-dense mode acts as PIM dense mode if an RP is not known.
F. PIM sparse mode interfaces are always added to the multicast routing table in a router.

Answer: A, B, E

QUESTION 427:

Study the exhibit below carefully. What are two important facts in interpreting the output of the "show ip pim" interface command issued on the Certkiller router below? (Select two)

```
Certkiller1# show ip pim interface
     Address          Interface     Version/Mode    Nbr    Query     DR
                                                    Count  Intvl
     192.168.10.1     Ethernet0     v2/Sparse-Dense 1      30      192.168.10.2
     192.168.9.3      Ethernet1     v2/Sparse-Dense 1      30      192.168.9.5
```

A. Multiaccess segments elect a DR based on lowest IP address.
B. Multiaccess segments elect a DR based on highest IP address.
C. Multiaccess multicast segments do not elect a DR.
D. The RP is only seen in version 2 of Sparse-Dense mode.

E. Point-to-Point links do not display DR information.

Answer: B, E

QUESTION 428:

Study the exhibit below carefully. IP multicast configuration changes have been made on several routers. However, the IP multicast table displayed in the exhibit still does not reflect the changes.

```
CertKillerA# show ip route
Ip Multicast Routing Table
Files of Mhash is sto Files of Mhash is sto    ocal, P - Pruned
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.0.255.3), uptime 5:29:15, RP is 198.92.37.2, flags: SC
  Incoming interface: FastEthernet 1/1, RPF neighbor 10.3.35.1, Dvmrp
  Outgoing interface list:
    Ethernet0, Forward/Sparse, 5:29:15/0:02:57

(198.92.46.1, 224.0.255.3), uptime 5:29:15, expires 0:02:59, flags: C
  Incoming interface: FastEthernet 1/1, RPF neighbor 10.3.35.1
  Outgoing interface list:
    Ethernet Forward/Sparse, 5:29:15/0:02:57

<Output omitted>
```

Based on the information shown above, what should be done to display the new routing table information?

A. Issue the no ip mroute-cache privileged EXEC command.
B. Issue the clear ip route * privileged EXEC command.
C. Issue the clear ip igmp group privileged EXEC command.
D. Issue the clear ip mroute privileged EXEC command.
E. None of the above.

Answer: D

QUESTION 429:

The following debug output was seen on a Certkiller IP multicast router:

```
CertKiller3 # debug ip mrouting 224.2.0.1
MRT: Create (',224.2.0.1) if_input NULL
MRT: Create (224.69.15.0/24, 225.2.2.4), if_input Ethernet0, RPF nbr 224.69.61.15
MRT: Create (224.69.39.0/24, 225.2.2.4), if_input Ethernet1, RPF nbr 224.0.0.0
MRT: Create (10.9.0.0/16.224.2.0.1), if_input Ethernet1, RPF nbr 0.0.0.0
MRT: Create (10.10.0.0/16.224.2.0.1), if_input Ethernet1, RPF nbr 0.0.0.0
```

Given the output shown above, which two statements are true? (Select two)

A. Multicast route to 224.69.15.0/24 was added to the mroute table and created by a source directly connected to the router.
B. The reverse path forwarding (RPF) for the route 224.2.0.1 failed to find the interface on which the multicast packet was received.

C. Multicast route to 10.16.0.0/16 was added to the mroute table and created by a source directly connected to the router.
D. This router received an IGMP host report from a group member or a PIM join message.
E. The route to 224.69.15.0/24 will be out Ethernet 0.

Answer: C, D

QUESTION 430:

Study the exhibit carefully regarding the configuration of router Certkiller A:

```
CertKillerA#show run
    <output omitted>
    ip multicast-routing

    interface GigabitEthernet0/0
     ip address 172.16.4.16 255.255.255.0
     ip pim sparse-dense-mode
```

What statement is true based upon the configuration in the exhibit shown above?

A. A rendezvous point must be configured in order for the interface to operate in sparse mode.
B. The rendezvous point is IP address 172.16.4.16.
C. If a rendezvous point is configured, the interface cannot operate in dense mode.
D. IGMP version 2 is being used.
E. The switch will use sparse mode first, dense mode second.
F. CGMP version 2 is being used.
G. None of the above

Answer: A

QUESTION 431:

Study the exhibit shown below:

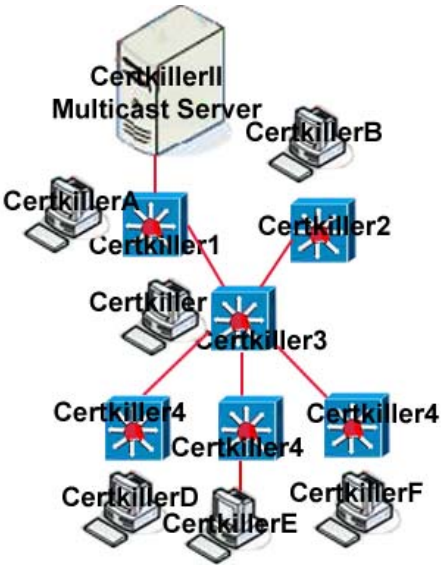

In this Certkiller multicast network all multilayer switches are running PIM sparse mode. Host Certkiller B and Host Certkiller F are sending IGMPv2 join messages to their respective multilayer switches. Based on this information, which statement is true?

A. Switch Certkiller 1 is the rendezvous of the multicast tree.
B. Switches Certkiller 2 and Certkiller 6 will participate in the multicast tree once pruning has taken place.
C. Switches Certkiller 1, Certkiller 2, Certkiller 3, and Certkiller 6 will participate in the multicast tree.
D. The multicast server is the rendezvous point of the multicast tree.
E. Switches Certkiller 1, Certkiller 2, Certkiller 3, and Certkiller 6 will participate in the multicast tree once pruning has taken place.

Answer: C

QUESTION 432:

You have been tasked with setting up IP multicast routing on some new Certkiller routers. Which three statements are true about IP multicast configuration? (Select three)

A. PIM sparse mode interfaces are always added to the multicast routing table in a router.
B. PIM sparse mode and PIM sparse-dense mode require an RP on the network.
C. PIM sparse-dense mode and PIM dense mode require an RP on the network.
D. PIM sparse mode and PIM dense mode require an RP on the network.
E. PIM dense mode interfaces are always added to the multicast routing table in a router.
F. PIM sparse-dense mode acts as PIM dense mode if an RP is not known.

Answer: B, E, F

QUESTION 433:

You need to configure PIM on a Certkiller IP multicast router. Which two tasks are required to configure PIM for IP multicast routing? (Select two)

A. Enable CGMP.
B. Join a multicast group.
C. Enable IP multicast routing.
D. Enable PIM on an interface.
E. Configure the TTL threshold.

Answer: C, D

QUESTION 434:

Study the exhibit below carefully:

```
CertKiller1#debug ip igmp
12:32:51:065: IGMP: Send v2 Query on Ethernet1 to 224.0.0.1
12:32:51:069· IGMP: Set report delay time to 9.4 seconds for 224.0.1.40 on Ethernet1
12:32:36:909: IGMP: Received v1 Report from 192.168.9.1 (Ethernet1) for 239.255.0.1
12:32:56:917: IGMP: Starting old host present timer for 239.255.0.1 on Ethernet1
12:32:01:065: IGMP: Send v2 Report for 224.0.1.40 on Ethernet1
12:32:01:069: IGMP: Received v2 Report from 192.168.9.4 (Ethernet1) for 224.0.1.40
12:33:51:065: IGMP: Send v2 Query on Ethernet1 to 224.0.0.1
```

Which two conclusions can be drawn from the "debug ip igmp" output shown above on router Certkiller 1? (Select two)

A. The IP multicast groups are 224.0.0.1, 224.0.1.40, and 239.255.0.1.
B. Reverse Path Forwarding (RPF) is enabled for 224.0.1.40.
C. The router sent an IGMP version 2 query out interface Ethernet1 at multicast address 224.0.0.1.
D. IP PIM RP mapping is static.
E. Router Certkiller 1 received an IGMP report version 1 from host 192.168.9.1.
F. Reverse Path Forwarding (RPF) is enabled for 192.168.9.4.

Answer: C, E

QUESTION 435:

The following output was seen on a Certkiller router:

```
CertKiller1# show ip pim interface
Address          Interface     Version/Mode     Nbr   Query     DR
                                                Count Intvl
192.168.10.1     Ethernet0     v2/Sparse-Dense  1     30     192.168.10.2
192.168.9.3      Ethernet1     v2/Sparse-Dense  1     30     192.168.9.5
```

What are two important facts in interpreting the output of this "show ip pim interface" command? (Select two)

A. Multiaccess segments elect a DR based on lowest IP address.
B. Multiaccess, multicast segments do not elect a DR.
C. Multiaccess segments elect a DR based on highest IP address.
D. Point-to-Point links do not display DR information.

E. The RP is only seen in version 2 of Sparse-Dense mode.

Answer: C, D

QUESTION 436:

Study the following exhibit carefully:

```
CertKiller2# show vlan

VLAN Name              Status    Ports
---- -------------    ---------- ----------------------------
1    default           active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
10   VLAN0010          active    Fa0/5, Fa0/6, Fa0/9,Fa0/8
20   VLAN0020          active    Fa0/9, Fa0/10, Fa0/11,Fa0/12
30   VLAN0030          active    Fa0/13, Fa0/14,Fa0/15, Fa0/16,
40   VLAN0040          active    Fa0/17, Fa0/18,Fa0/19, Fa0/20,
50   VLAN0050          active    Fa0/21, Fa0/22,Fa0/23, Fa0/24

1002 iddl-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
```

Switch Certkiller 2 is receiving IGMP frames only on interface FastEthernet 0/3. Given IGMP snooping is enabled, out of which port or ports will switch Certkiller 2 forward multicast traffic?

A. All ports
B. FastEthernet 0/1, 0/2, 0/4, 0/5, 0/6, 0/7, 0/8
C. FastEthernet 0/1, 0/2, 0/4
D. None of the ports
E. FastEthernet 0/1, 0/2, 0/3, 0/4
F. FastEthernet 0/3

Answer: F

QUESTION 437:

A Certkiller IP multicast router was configured as shown below:

```
ip muticast-routing
!
interface ethernet 0
  ip pim sparse-dense-mode
!
interface ethernet 1
  ip pim sparse-dense-mode
!
```

On the basis of the partial configuration provided above, what additional configuration is required to allow the router to properly participate in a PIM sparse-dense mode scenario?

A. IGMP needs to be enabled on the router.
B. The PIM dense-mode state, refresh interval needs to be configured.
C. The RP configuration needs to be added for the router to participate in dense mode network areas.
D. The RP configuration needs to be added for the router to participate in sparse mode network areas.
E. None of the above

Answer: D

QUESTION 438:

You need to configure router CK1 with the IP address of the RP in the Certkiller IP multicast network. Which command sequence is required to configure the IP address of the rendezvous point (RP) located at IP address 10.10.10.1?

A. Certkiller 2(config)# ip pim rp-candidate gigabitethernet0/2
B. Certkiller 2(config)# ip pim accept-rp 10.10.10.1 1
C. Certkiller 2(config)# ip igmp join-group 10.10.10.1
D. Certkiller 2(config)# ip pim rp-address 10.10.10.1
E. None of the above

Answer: D

QUESTION 439:

PIM has been implemented in the Certkiller multicast network. When IP multicast is enabled via PIM, which mode uses the flood and prune method?

A. Bidir-PIM
B. PIM-RP
C. PIM sparse-dense
D. PIM-SM
E. PIM-DM
F. None of the above

Answer: E

QUESTION 440:

You want to integrate IP multicast into the Certkiller network as efficiently as possible. When building an IP multicast domain using PIM which mode assumes that other routers do not want to forward multicast packets for the group?

A. PIM-DM

B. PIM-RP
C. IGMP snooping
D. CGMP
E. PIM-SM
F. None of the above

Answer: E

QUESTION 441:

The Certkiller network uses the autoconfiguration feature in their IPv6 network. Which statement best describes autoconfiguration for IPv6?

A. A router connected to the local link sends its 64 bit prefix address and the host appends its MAC address to create its unique IPv6 address.
B. A router connected to the local link sends a broadcast periodically to advertise the subnet and the host appends its locally configured IPv4 address to create a unique IPv6 address.
C. A host will contact its local router to receive its unique IPv6 address.
D. The local DHCP server will detect a new host has been added to the LAN segment and automatically send a unique IPv6 address to the new host.
E. None of the above.

Answer: A

QUESTION 442:

IPv6 has been implemented on the Certkiller WAN. Which address type does the IPv6 address FF05:0:0:0:0:0:0:2 specify?

A. Site local unicast
B. Multicast
C. Link local
D. Aggregable global unicast
E. Unspecified
F. None of the above

Answer: B

Explanation:
All IPv6 multicast addresses start with the first 8 bits of the address set to 1. Thus all multicast addresses start with the hexadecimal notation FF (1111 1111). The multicast range is as follows:
FF00::/8
FFFF::/8
The second octet, following to the first octet of FF, identifies both the scope and the

lifetime of the multicast address. In this way, IPv6 has millions of group multicast addresses to use in current and emerging technologies.
Reference: CCNP Self-Study CCNP BSCI Exam Certification Guide p.112

QUESTION 443:

Certkiller 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
A. For a full
explanation of the IPv6 address representation, see the reference link below:
Reference: http://www.ietf.org/rfc/rfc1924.txt?number=1924

QUESTION 444:

You need to design an addressing scheme for the Certkiller IPv6 network. What is the difference between the IPv6 addresses ::/0 and ::/128?

A. ::/0 is the unspecified address, and ::/128 is the multicast address.
B. ::/0 is the unicast address, and ::/128 is the multicast address.
C. ::/0 is the anycast address, and ::/128 is the multicast address.
D. ::/0 is the default route, and ::/128 is the unspecified address.
E. ::/0 is the unicast address, and ::/128 is the anycast address.
F. ::/0 is the anycast address, and ::/128 is the default address.
G. None of the above

Answer: D

QUESTION 445:

An IPv6 packet contains some differences to an IPv4 packet. Which three are characteristics of IPv6? (Select three)

A. An IPv6 header contains the next header field.
B. An IPv6 header contains the protocol field.
C. An IPv6 header contains the header checksum field.
D. An IPv6 address is 128 bits long.
E. An IPv6 header is 20 bits long.
F. IPv6 routers send RA messages.

Answer: A, D, F

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 446:

IPv6 addresses can be usually be represented in a much sorter form. Which IPv6 address correctly compresses the IPv6 unicast address 2001:0:0:0:0DB8:0:0:417A?

A. 2001::0DB8::417A
B. 2001:0DB8:0:0:417A
C. 2001::DB8:0:0:417A
D. 2001:::0DB8::417A
E. 2001:0DB8:417A
F. 2001:::0DB8:0:0:417A
G. None of the above

Answer: C

QUESTION 447:

Hosts on the Certkiller LAN are IPv6 enabled. When an IPv6 enabled host boots, it sends a router solicitation (RS) message to the Certkiller router. This IPv6 router responds with a router advertisement (RA). Which two items are contained in the RA? (Select two)

A. Any route advertisements it has received
B. Request for the local host IP address
C. Keepalive timers
D. Prefixes for the link
E. Lifetime of the prefix
F. IPv6 address for the host

Answer: D, E

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 448:

The Certkiller 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

Explanation:
IP version 6 introduced the concept of anycasts. The three IP address types used for IPv6 are:
1. Unicast- An IPv6 unicast address is an identifier for a single interface, on a single node. A packet that is sent to a unicast address is delivered to the interface identified by that address.
2. Anycast- An anycast address is an address that is assigned to a set of interfaces that typically belong to different nodes. A packet sent to an anycast address is delivered to the closest interface-as defined by the routing protocols in use-identified by the anycast address.
3. Multicast - An IPv6 multicast address is an IPv6 address that has a prefix of FF00::/8 (11111111). An IPv6 multicast address is an identifier for a set of interfaces that typically belong to different nodes.
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide_chapter09186a0080110dd2.h

---

QUESTION 449:

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:
Due to the method of allocating certain styles of IPv6 addresses, it will be common for addresses to contain long strings of zero bits. In order to make writing addresses

containing zero bits easier a special syntax is available to compress the zeros. The use of "::" indicates multiple groups of 16-bits of zeros. The "::" can only appear once in an address. The "::" can also be used to compress the leading and/or trailing zeros in an address.
For example the following addresses:
1080:0:0:0:8:800:200C:417A a unicast address
FF01:0:0:0:0:0:0:43 a multicast address
0:0:0:0:0:0:0:1 the loopback address
0:0:0:0:0:0:0:0 the unspecified addresses
may be represented as:
1080::8:800:200C:417A a unicast address
FF01::43 a multicast address
::1 the loopback address
:: the unspecified addresses
In our example, the trailing zero's can be compressed into ::, giving us 200F:0:AB00:: as the address.

QUESTION 450:

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

Explanation:
To support the processing and forwarding of IP version 6 packets, the "ipv6 unicast-routing "command must be enabled in global configuration mode. The following is a configuration example:
ipv6 unicast-routing
interface ethernet0/0
description connection to enterprise
ipv6 address 2002:C0A8:2101::1/128

QUESTION 451:

Certkiller is in the process of migrating from IP version 4 to version 6. Which two statements are true about using IPv4 and IPv6 simultaneously on a network segment? (Select two)

A. IPv6 provides for more host IP addresses but IPv4 provides for more network addresses.

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. IPv4 and IPv6 addresses can be simultaneously assigned to a host but not to a router interface.
D. 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.
E. Hosts can be configured to receive both IPv4 and IPv6 addresses via DHCP.

Answer: D, E

QUESTION 452:

The Certkiller WAN is shown below:

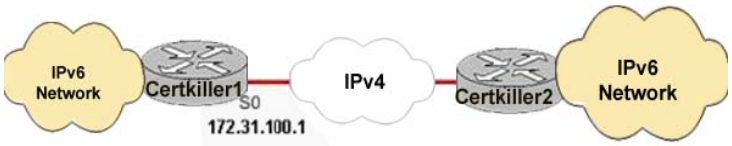

Based on the network above, if Certkiller 1 is configured for 6to4 tunneling, what will be the prefix of its IPv6 network?

A. AC1F:6401::/32
B. 1723:1100:1::/48
C. FFFF:AC1F:6401::/16
D. 2002:AC1F:6401::/48
E. 3FFE:AC1F:6401::/32

Answer: D

QUESTION 453:

A Certkiller router is configured as shown below:

```
<output omitted>
!
ipv6 ueic storouting
!
interface Ethernet0
  ip address 192.168.99.1 255.255.255.0
!
interface Tunnel2002
 no ip address
 no ip redirects
 ipv6 address 2002:C0A8:6301::1/128
 tunnel source ethernet0
 tunnel made ipv6ip 6to4
!
ipv6 route 2002::/16 Ethernet2002
ipv6 route ::/0 2002:C0A8:2101::1
!
<output omitted>
```

Given the partial configuration shown above, which statement regarding IPv6 is true?

A. The configuration is an example of a 6to4 tunnel.
B. The configuration is an example of a 4to6 tunnel.
C. The configuration is an example of an encrypted IPv6 VPN tunnel.
D. The configuration is an example of a one to one IPv6 tunnel.
E. None of the above.

Answer: A

QUESTION 454:

Certkiller is migrating from IPv4 to IPv6. What happens when an IPv6 enabled router running version 6 to 4 must send a packet to a remote destination and the next hop is the address of 2002::/16?

A. The IPv6 packet is encapsulated in an IPv4 packet using an IPv4 protocol type of 41.
B. The IPv6 packet has its header removed and replaced with an IPv4 header.
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.
E. None of the above

Answer: A

QUESTION 455:

Certkiller is implementing IPv6 within their network. Which two statements are true concerning IPv6? (Select two)

A. Like IPv4, IPv6 broadcasts are sent to all nodes on a LAN segment.
B. Mobile IP is built into IPv6 by default.
C. The leading zeros in an address format are mandatory.
D. IPSec is mandatory and built into IPv6.

Answer: B, D

QUESTION 456:

Certkiller 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

Explanation:
The four key strategies for deploying IPv6 are as follows:
1. Deploying IPv6 over IPv4 tunnels: These tunnels encapsulate the IPv6 traffic within the IPv4 packets, and are primarily for communication between isolated IPv6 sites or connection to remote IPv6 networks over an IPv4 backbone. The techniques include using manually configured tunnels, generic routing encapsulation (GRE) tunnels, semiautomatic tunnel mechanisms such as tunnel broker services, and fully automatic tunnel mechanisms such as IPv4-compatible and 6to4.
2. Deploying IPv6 over dedicated data links: This technique enables isolated IPv6 domains to communicate by using the same Layer 2 infrastructure as for IPv4, but with IPv6 using separate Frame Relay or ATM PVCs, separate optical links, or dense Wave Division Multiplexing (dWDM).
3. Deploying IPv6 over MPLS backbones: This technique allows isolated IPv6 domains to communicate with each other, but over an MPLS IPv4 backbone. Multiple techniques are available at different points in the network, but each requires little change to the backbone infrastructure or reconfiguration of the core routers because forwarding is based on labels rather than the IP header itself.
4. Deploying IPv6 using dual-stack backbones: This technique allows IPv4 and IPv6 applications to coexist in a dual IP layer routing backbone. All routers in the network need to be upgraded to be dual-stack with IPv4 communication using the IPv4 protocol stack and IPv6 communication using the IPv6 stack.
In addition to the strategies for deploying IPv6 within your IPv4 environment, you also need protocol translation mechanisms (for example, a NAT-PT device to connect IPv6-only web browsers to IPv4-only web servers) or dual-stack servers (for example, an e-mail server that handles IPv4-only and IPv6-only mail clients) to allow communication between applications using IPv4 and applications using IPv6. These mechanisms become increasingly important as IPv6 deployment moves from the testing to the actual usage phase, and more relevant as application developers decide that continuing to support IPv4 is not cost-effective.
Reference:
http://www.cisco.com/en/US/tech/ CK8 72/technologies_white_paper09186a00800c9907.shtml#1098273

QUESTION 457:

Certkiller is migrating from IPv4 to IPv6 in the network shown below. If Certkiller 1 is configured for 6to4 tunneling, what will the prefix of its IPv6 network be?

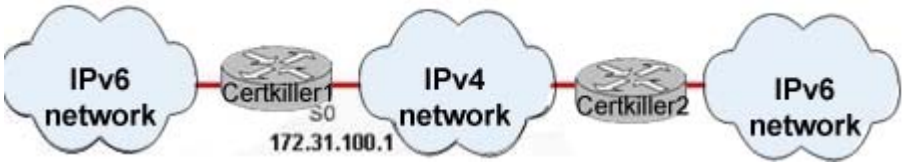

A. 3FFE:AC1F:6401::/32
B. 2002:AC1F:6401::/48
C. 1723:1100:1::/48
D. FFFF:AC1F:6401::/16
E. AC1F:6401::/32
F. None of the above

Answer: B

QUESTION 458:

The Certkiller network will use 6to4 tunneling to migrate from IPv4 to IPv6. Which statement is true concerning 6to4 tunneling?

A. Hosts and routers inside a 6to4 site will need a special code.
B. An edge router must use IPv6 address of 2002::/16 in its prefix.
C. IPv4 traffic is encapsulated with an IPv6 header.
D. The edge routers can use any locally configured IPv6 address.
E. None of the above

Answer: B

QUESTION 459:

The following output was seen on an existing Certkiller router:

```
Routing Process "ospfv3 1" with ID 172.16.3.3
 It is an autonomous system boundary router
  Redistribution Initial Retres from,
     status
 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
 Number of external LSA 1. Checksum Sum 0x218D
 Number of areas inethrs router is 1. 1 normal 0 stub 0 nssa
    Area 1
         Number of interface sign is rate is
        SPF algorithm executed   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
```

Based on the information shown above, what IOS command produces this output?

A. show ipv6 ospf interface
B. show ipv6 ospf
C. show ip ospf
D. show ip ospf interface
E. None of the above

Answer: B

QUESTION 460:

The command "clear ipv6 ospf process" was issued on a Certkiller router. What does this command accomplish?

A. The route table is cleared. Then the OSPF neighbors are reformed.
B. The OSPF adjacencies are cleared and initiated again.
C. The OSPF database is repopulated and then the shortest path first (SPF) algorithm is performed.
D. The shortest path first (SPF) algorithm is performed on the LSA database.
E. None of the above

Answer: C

QUESTION 461:

Study the exhibit shown below carefully.

```
CertKiller3 # 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
FF80::A8BB:CCFF:FE00:6E00
  Timerinterva is configured Hello  0  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)
```

On the basis of the output from the "show ipv6 interface command" shown on router Certkiller 3 , what two statements must be true? (Select two)

A. This is the designated router (DR) on the Ethernet 0/0 link.
B. The IP address of the designated router (DR) is FE80::A8BB:CCFF:FE00:6E00.
C. OSPF version 3 is enabled to support IPv6.
D. Interface Ethernet 0/0 has been configured with the ipv6 ospf authentication ipsec spi 500 md5 command.
E. Interface Ethernet 0/0 was configured with the ipv6 ospf 1 area 1 command.

Answer: C, D

QUESTION 462:

You have been tasked with setting up OSPF on an existing Certkiller router using IPv6. Which command enables OSPF for IPv6?

A. ipv6 router ospf process-id
B. router ospf process-id
C. router ospf ipv6 process-id
D. ipv6 ospf process-id
E. None of the above

Answer: D

QUESTION 463:

Certkiller uses OSPF as the routing protocol in their IPv6 network. Which three statements about configuring OSPF in an IPv6 network are true? (Select three)

A. OSPF version 3 will support IPv6.
B. Multiple instances of OSPF for IPv6 can be run on a link.
C. IPv4 addresses cannot be used as the router ID in OSPF for IPv6.
D. Networks must be explicitly configured using the network command in router OSPF configuration mode.
E. OSPF version 2 will support 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: A, B, F

Certkiller .com, Scenario

Configuration exhibit:

```
Certkiller3 #show ip eigrp topology all-links
IP EIGRP Topology table for AS(65001) / ID (172.20.10.1)

Codes: P- passive A-AActive U- Update Q-Query R-Reply.
    R- Reply Static S-Via Static

P 172.20.3.128/25, 2 successors, FD is 30720, serno 9
      via 172.20.10.2 (30720/28160), FastEthernet0/1
      via 172.20.3.2 (30720/28160), FastEthernet0/3
P 10.69.0.0/24, 1 successors, FD is 156160, serno 16
      via 172.20.3.2 (156160/128256), FastEthernet0/3
      via 172.20.10.2 (158720/156160), FastEthernet0/1
P 172.20.10.0/24, 1 successors, FD is 28160, serno 1
      via Connected, FastEthernet0/1
P 172.20.10.0/30. 1 successors., FD is 20514560, serno 15
      via 172.20.1.1 (20514560/205122000), FastEthernet0/2
      via 172.20.10.2 (20517120/20514560), FastEthernet0/1
P 172.20.1.0/24, 1 successors, FD is 28160, serno 2
      via Connected, FastEthernet0/2
P 172.20.2.0/24, 1 successors, FD is 30720, serno 8
      via 172.20.10.2 (30720/28160), FastEthernet0/1
      via 172.20.3.2 (33280/30720), FastEthernet0/3
P 172.20.3.0/25, 1 successors, FD is 28160, serno 3
      via Connected, FastEthernet0/3
```

You work as a network administrator at Certkiller .com. Certkiller .com is large company with offices in USA, Canada and Mexico. As far as network routing protocol EIGRP is used exclusively. Your boss at Certkiller .com, Miss Certkiller, has asked you to study the EIGRP convergence
You use the show ip eigrp topology all-links command. Please refer to the exhibit for the output.
Use the scenario information to answer the three scenario questions below.

Certkiller.com Questions (3 Questions)

QUESTION 464:

Note: Please refer to the Certkiller .com scenario.
Which three networks does the Switch Certkiller 3 device have feasible successors for? (Select three)

A. 172.26.1.0/24
B. 172.26.3.128/25
C. 172.26.3.0/25
D. 10.151.0.0/24
E. 172.26.2.0/24
F. 172.26.0.0/30

Answer: D, E, F

QUESTION 465:

Note: Please refer to the Certkiller .com scenario.
Which three EIGRP routes will be installed for the 172.20.3.128/25 and 172.20.2.0/24 networks? (Select three)

A. D 172.20.3.128/25 [90/28160] via 172.20.1.2, 01:26:35. FastEthernet0/2
B. D 172.20.2.0/24 [90/28160] via 172.20.10.2, 02:10:11. FastEthernet0/1
C. D 172.20.3.128/25 [90/30720] via 172.20.10.2, 01:26:35. FastEthernet0/1
D. D 172.20.2.0/24 [90/30720] via 172.20.10.2, 02:10:11. FastEthernet0/1
E. D 172.20.2.0/24 [90/33280] via 172.20.3.2, 02:10:11. FastEthernet0/3
F. D 172.20.3.128/25 [90/30720] via 172.20.3.2, 01:26:35. FastEthernet0/3

Answer: C, D, F

QUESTION 466:

Note: Please refer to the Certkiller .com scenario.
Which three networks is the router at 172.26.10.2 directly connected to? (Select three)

A. 172.26.10.0/24
B. 172.26.3.0/25
C. 172.26.2.0/24
D. 172.26.0.0/30
E. 172.26.3.128/25
F. 172.26.1.0/24

Answer: A, C, E

Certkiller .com #2, Scenario

Network topology exhibit:

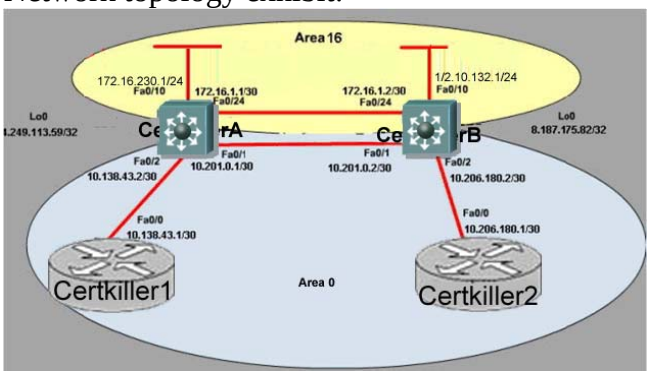

You work as a network administrator at Certkiller .com. Certkiller .com is large company with offices in Europe and Asia. As far as network routing protocol OSPF is used exclusively. The company and the company network have both been growing rapidly. Multiple adds, moves and changes have been applied to the network. Your boss at Certkiller .com, Miss Certkiller, 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 4 questions:

QUESTION 467:

Note: Please refer to the Certkiller .com #2 scenario.
Examine the following excerpt from the "show ip ospf" command on switch Certkiller A:
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 32
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 causing the different missing routes throughout the network?

A. Area 32 is configured with authentication.
B. Area 32 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 32 has been configured as a total stub network
E. Area 32 has been configured as a stub network
F. Area 0 is discontiguous.
G. None of the above

Answer: F

QUESTION 468:

Note: Please refer to the Certkiller .com #2 scenario.
Which configuration command on Certkiller A (with a similar command on Certkiller B) will provide an immediate solution to the missing route problem?

A. no area 32 stub
B. no area 32 authentication message-digest
C. area 32 virtual-link 9.152.105.122
D. area 32 virtual-link 172.16.4.2
E. no area 32 stub no-summary
F. network 172.16.0.0.0.0.255.255 area 32
G. None of the above

Answer: C

QUESTION 469:

Note: Please refer to the Certkiller .com #2 scenario.
The log of Certkiller A 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 FasstEthernet0/1 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

QUESTION 470:

Note: Please refer to the Certkiller .com #2 scenario.
The Certkiller 1rRouter has lost connectivity to Certkiller 1. The following is Certkiller 1's current route table:
172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
O IA 172.16.240.0/24 [110/11] via 10.218.67.1, 00:00:03, FastEthernet0/0

O IA 172.16.209.0/24 [110/12] via 10.218.67.1, 00:00:03, FastEthernet0/0
O IA 172.16.4.0/30 [110/2] via 10.218.67.1, 00:00:03, FastEthernet0/0
10.0.0.0/30 is subnetted, 1 subnets
C 10.218.67.0 is directly connected, FastEthernet0/0
Which expected route is missing from Certkiller 1's route table based on the topology during the maintenance period?

A. o 10.249.0.0 [110/2] via 10.218.67.1, 00:00:09, FastEthernet0/0
B. o IA 9.152.105.122 [110/3] via 10.218.67.1, 00:00:09, FastEthernet0/0
C. o IA 10.191.213.0 [110/3] via 10.218.67.1, 00:00:09, FastEthernet0/0
D. o IA 10.249.0.0 [110/2] via 10.218.67.1, 00:00:09, FastEthernet0/0
E. o IA 3.42.36.239 [110/2] via 10.218.67.1, 00:00:09, FastEthernet0/0
F. o 10.191.213.0 [110/3] via 10.218.67.1, 00:00:09, FastEthernet0/0

Answer: F

Mixed Questions (13 Questions)

QUESTION 471:

Exhibit:

```
Routing Process "ospfv3 1" with ID 172.16.3.3
 It is an autonomous system boundary router
 Redistribution external routers from
    statics
 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
 Number 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
```

You work as a network technician at Certkiller .com. You trainee Certkiller shows you the IOS command output displayed in the exhibit.
What command did Jack use to produce this output?

A. show ip RIP
B. show ipv6 ospf
C. show ip ospf
D. show ip ospf interface
E. show ipv6 ospf interface

Answer: C

Explanation:
To display general information about Open Shortest Path First (OSPF) routing processes, use the show ip ospf command in EXEC mode.
The following is sample output from the show ip ospf command when entered without a specific OSPF process ID:
Router# show ip ospf Routing Process "ospf 201" with ID 10.0.0.1 and Domain ID 10.20.0.1Supports only single TOS(TOS0) routesSupports opaque LSASPF schedule delay 5 secs, Hold time between two SPFs 10 secsMinimum LSA interval 5 secs. Minimum LSA arrival 1 secsLSA group pacing timer 100 secsInterface flood pacing timer 55 msecsRetransmission pacing timer 100 msecsNumber of external LSA 0. Checksum Sum 0x0Number of opaque AS LSA 0. Checksum Sum 0x0Number of DCbitless external and opaque AS LSA 0Number of DoNotAge external and opaque AS LSA 0Number of areas in this router is 2. 2 normal 0 stub 0 nssaExternal flood list length 0Area BACKBONE(0)Number of interfaces in this area is 2Area has message digest authenticationSPF algorithm executed 4 timesArea ranges areNumber of LSA 4. Checksum Sum 0x29BEBNumber of opaque link LSA 0. Checksum Sum 0x0Number of DCbitless LSA 3Number of indication LSA 0Number of DoNotAge LSA 0Flood list length 0Area 172.16.26.0Number of interfaces in this area is 0Area has no authentication SPF algorithm executed 1 timesArea ranges are192.168.0.0/16 Passive AdvertiseNumber of LSA 1. Checksum Sum 0x44FDNumber of opaque link LSA 0. Checksum Sum 0x0 Number of DCbitless LSA 1Number of indication LSA 1Number of DoNotAge LSA 0 Flood list length 0

QUESTION 472:

Exhibit:

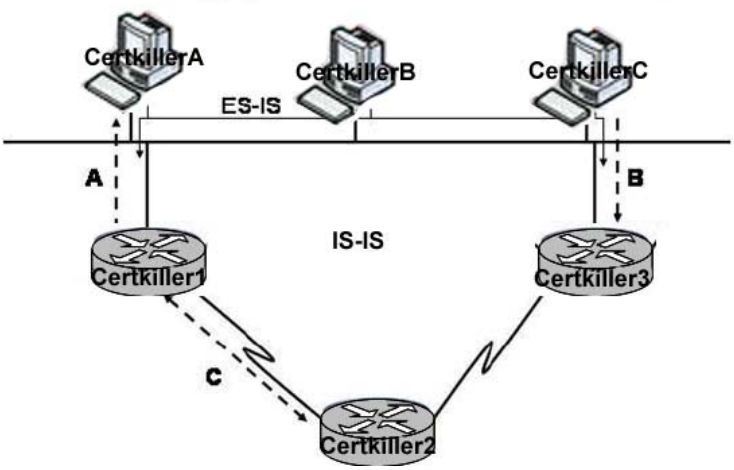

You work as a network technician at Certkiller .com. You are examining the network displayed in the exhibit. Your boss, Miss Certkiller, asks you what correctly matches the letter to the IS-IS hello PDU type. Select three.

A. A=ISH
B. A=IIH
C. A=ESH

D. B=ISH
E. B=IIH
F. B=ESH
G. C=ISH
H. C=IIH
I. C=ESH

Answer: A, F,H,

Explanation:
Hello PDU (ESH, ISH, IS-IS Hello [IIH]) - Used to establish and maintain adjacencies.
ESHs are sent from ESs to ISs. ISHs are sent from ISs to ESs. IIHs are sent between ISs.
Note that ESH and ISH PDUs are ES-IS PDUs, not IS-IS PDUs.

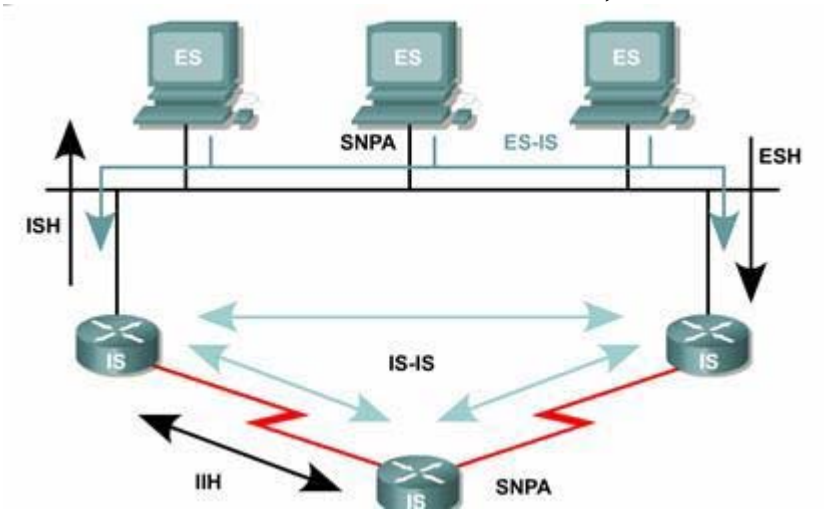

QUESTION 473:

Network topology exhibit:

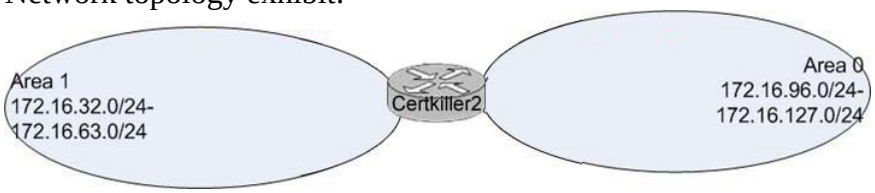

You work as a network technician at Certkiller .com. You are examining the network displayed in the exhibit. You want to summarize all routes from area 0 to area 1. Which command need to be issued on Certkiller 2?

A. summary-address 172.16.96.0 255.255.224.0
B. summary-address 172.16.96.0 0.0.63.255
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. area 0 range 172.16.96.0 255.255.224.0
F. area 0 range 172.16.96.0 255.255.255.0
G. None of the other alternatives provides the required result.

Answer: D

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:
1. Interarea route summarization - Interarea route summarization is done on ABRs and applies to routes from within each area. It does not apply to external routes injected into OSPF by way of redistribution. To take advantage of summarization, network numbers within areas should be contiguous.
2. External route summarization - External route summarization is specific to external routes that are injected into OSPF by way of redistribution. Here again, it is important to ensure that external address ranges that are being summarized are contiguous. Summarization of overlapping ranges from two different routers could cause packets to be sent to the wrong destination. Only ASBRs can summarize external routes.
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 474:

Study the exhibit shown below carefully.

```
<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)/1:0.0)
  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)
```

On the basis of the output shown on router Certkiller 3 , what two statements must be true? (Select two)

A. The output was generated by the show ip interface command.
B. OSPF Version 2 has been enabled to support IPv6
C. This is the designated router (DR) on the FastEthernet 0/0 link.
D. Interface FastEthernet 0/0 was configured with the ipv6 ospf 1 area 1 command.
E. The router was configured with the commands:
router ospf 1
network 172.16.6.0 0.0.0.255 area 1
F. The IP address of the backup designated router (BDR) is FE80:205:5FFF:FED3:5808

Answer: A,F

Explanation:
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 **ipv6 unicast-routing** command.
Step 3	Enable IPv6 on the interface using the **ipv6 ospf area** command.
Step 4	(Optional) Configure OPSFv3 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] interfacee [interface] command.

QUESTION 475:

Network topology exhibit:

Certkiller 1 configuration exhibit:

```
Certkiller1#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
```

Certkiller 2 configuration exhibit:

```
Certkiller2#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
```

Certkiller 3 configuration exhibit:

```
Certkiller3 # show running-config
    <Output omitted>
    key chain schenectady
     key 1
      key-string 0987654321
     key 2
      key-string 1234576890
    <output omitted>
    interface GigabitEthernet0/
     ip address 172.16.1.2 255.
     ip authentication mode eigr
     ip authentication key-chain
     <Output omitted>
     router eigrp 1
     network 172.16.0.0
     no auto-summary
```

What can be said about the network in this scenario? Please refer to the exhibits.

A. Because the key strings do not match, router Certkiller 1 will not be able to ping router Certkiller 2 and router Certkiller 3.
B. Router Certkiller 1 will be able to ping router Certkiller 2 and Certkiller 3
C. Because the key names do not match, router Certkiller 1 will not be able to ping router Certkiller 2 and router Certkiller 3.
D. Because autosummarization needs to be turned on for EIGRP on router Certkiller 2, router Certkiller 1 will be able to ping router Certkiller 2 but will not be able to ping router Certkiller 3

Answer: D

QUESTION 476:

Network topology exhibit:

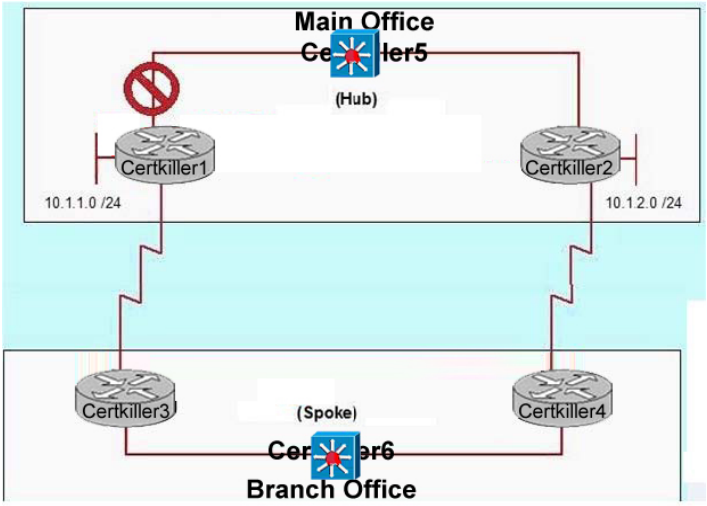

You work as a network administrator at Certkiller .com. The network topology is displayed in the exhibit. You have set up a hub and spoke topology with redundant connections using EIGRP. However, your boss, Mrs. Certkiller, is worried that a network outage between Router Certkiller 1 and Router Certkiller 2 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 you use to configure the spoke routers as EIGRP stub routers that will not advertise connected network, static routers, or summary addresses?

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

Answer: C

QUESTION 477:

What is the IPv6 address FF02:2 used for?

A. all hosts in a particular multicast group
B. all routers in an autonomous system
C. all routers in a local segment
D. all host in a local segment

Answer: C

QUESTION 478:

Your Certkiller trainee asks you, in the context of IPv6 and OSPF, what best describes a type 9 LSA?

A. link LSA
B. interarea prefix LSA for ABRs
C. router LSA
D. switch LSA
E. interarea prefix LSA

Answer: E

QUESTION 479:

Exhibit:

```
Certkiller3#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>
```

You work as a network administrator at Certkiller .com. Please study the exhibit carefully.
Router Certkiller 3 is connected to networks 172.16.1.0/26 and 172.16.1.24/27. On the basis of the partial output in the exhibit, which statement is true?

A. Router Certkiller 3 will deny the 172.16.1.0/26 route while permitting the 172.16.1.64/27 route to be advertised.
B. Router Certkiller 3 should be configured with an ACL instead of an ip prefix-list command
C. Router Certkiller 3 will advertise both routes.
D. Router Certkiller 3 will deny the 172.16.1.0/27 route while permitting the 172.16.1.64/26 route to be advertised.

Answer: D

QUESTION 480:

Network topology exhibit:

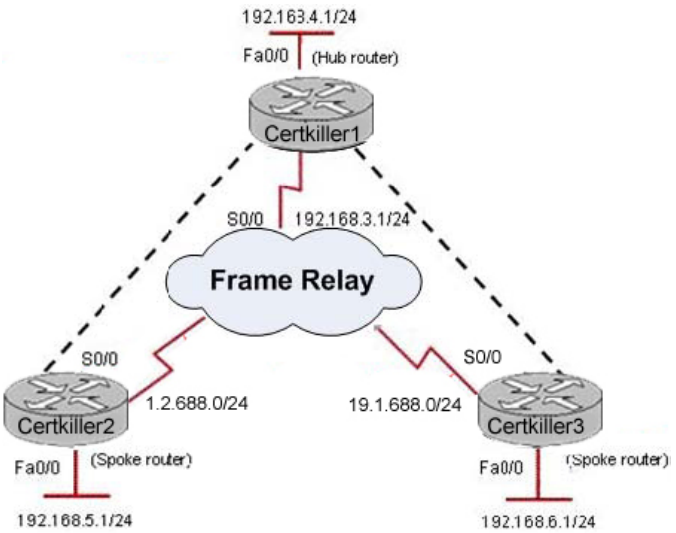

You work as a network administrator at Certkiller .com. Please study the exhibit carefully.
Router Certkiller 1 is the hub router while the Certkiller 2 and Certkiller 2 routers are spoke router for Certkiller 1. The Frame Relay network is configured with EIGRP, and the entire network is in autonomous system 1. However, router Certkiller 2 and Certkiller 3 are not receiving each other's routes.
What is the solution?

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

Answer: B

QUESTION 481:

Network topology exhibit:

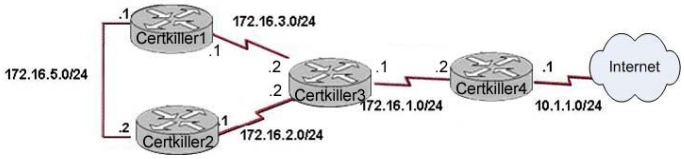

Certkiller 4 configuration exhibit:

```
hostname Certkiller4
!
<output omitted>

router eigrp 4
  network 10.0.0.0
  network 172.16.0.0
  no auto-summary
!
```

You work as a network administrator at Certkiller .com. Please study the exhibits carefully.
All routers in the network have been configured with EIGRP.
You are required to fine tune the configuration of Certkiller 4. In particular you want to advertise the default route to the neighbors of Certkiller 4.
What command should you use?

A. Certkiller 4(config-router)#default-information originate
B. Certkiller 4(config)#ip route 10.0.0.0 255.0.0.0 10.1.1.1
C. Certkiller 4(config)#ip route 0.0.0.0 0.0.0.0 10.1.1.1
D. Certkiller 4(config-router)#ip default-network 10.0.0.0

Answer: D

QUESTION 482:

You work as a network administrator at Certkiller .com. You trainee asks you what metric would be assigned to redistribute BGP routes in the following scenario.
Scenario:
A metric is not specified for routes that are redistributed into OSPF.
The default metric that is assigned to the routes is 20, with the exception of BGP routes.

A. 200
B. 20
C. 1
D. 5
E. 0
F. 10
G. 100

Answer: C

QUESTION 483:

Network topology exhibit.

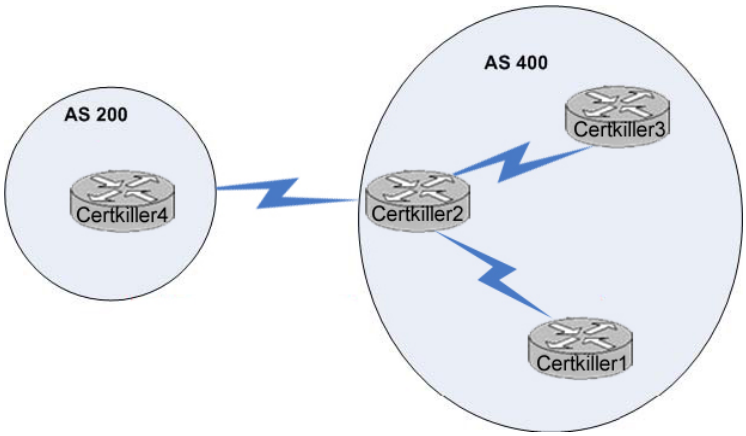

You work as a network administrator at Certkiller .com. Please study the exhibit carefully.
You have configured all routes for BGP. The EBGP routes received on Certkiller 2 are showing up in BGP tables of router Certkiller 3 and Certkiller 1 respectively. However, these EBGP routers do show up in their IP tables. What could be the problem?

A. The BGP routers in the autonomous system 400 are not logically fully-meshed
B. Synchronization in autonomous system 400 is turned is on.
C. EBGP multihop is not configured on routers Certkiller 3 and Certkiller 1.
D. Synchronization in autonomous system 400 is turned is off.
E. Routers Certkiller 3 and Certkiller 1 do not receive the same routes via an IGP.

Answer: E