642-902 Q&As

Implementing cisco ip routing

Pass Cisco 642-902 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/642-902.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Cisco Official Exam Center

QUESTION 1

Refer to the exhibit. Router B and router C are performing mutual redistribution between OSPF and EIGRP, and their default metrics are configured the same. Router D has equal cost paths to networks where both paths are not really equal cost. For example, network 172.16.54.0 shows equal cost through both router B and router C, though in reality the cost is greater using router C. Other routers, though not shown, are connected to the 172.16.54.0 and 172.16.55.0 networks, and the same issues exist to those routers and the networks connected to them. What can be done so that data will be routed along the most optimal path in the network?

A. Redistribute connected interfaces on router B and router C.

B. Set the maximum number of equal cost paths to 1 in all routers.

C. When redistributing EIGRP into OSPF, set the external metric type to type E1.

D. Adjust the default metrics in router B and router C so that the values are different in each router.

E. None of these solutions will fix the problem. Migrate to a single dynamic routing protocol.

Correct Answer: E

From the output, we learn that all the External OSPF routes have metrics of 100 (the second parameters in [110/100]). This is not the default metric of OSPF Type 2 External route (the default value is 20) so the metrics of redistributed routes

have been modified. Maybe when redistributing into OSPF, the "metric" in the "redistribute" command or the "default-metric" command was used on router B and C to assign the metric of these routes. Something like this:

router ospf 1

redistribute eigrp 1 metric 100 subnets

or

router ospf 1

.....

default-metric 100

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

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

QUESTION 2

A network administrator recently redistributed RIP routes into an OSPF domain. However, the administrator wants to configure the network so that instead of 32 external type-5 LSAs flooding into the OSPF network, there is only one.

What must the administrator do to accomplish this?

A. Configure summarization on R1 with area 1 range 172.16.32.0 255.255.224.0

B. Configure summarization on R1 with summary-address 172.16.32.0 255.255.224.0

C. Configure area 1 as a stub area with area 1 stub

D. Configure area 1 as a NSSA area with area 1 stub nssa

Correct Answer: B

In many cases, the router doesn\\'t even need specific routes to each and every subnet (for example, 172.16.1.0/24). It would be just as happy if it knew how to get to the major network (for example, 172.16.0.0/16) and let another router take it

from there. In our telephone network example, the local telephone switch should only need to know to route a phone call to the switch for the called area code. Similarly, a router\\'s ability to take a group of subnetworks and summarize them as

one network (in other words, one advertisement) is called route summarization. Besides reducing the number of routing entries that a router must keep track of, route summarization can also help protect an external router from making

multiple changes to its routing table due to instability within a particular subnet. For example, let\\'s say that we were working on a router that connected to 172.16.2.0/24. As we were working on the router, we rebooted it several times. If we

were not summarizing our routes, an external router would see each time 172.16.2.0/24 went away and came back. Each time, it would have to modify its own routing table. However, if our external router were receiving only a summary route

(i.e., 172.16.0.0/16), then it wouldn\\'t have to be concerned with our work on one particular subnet. This is especially a problem for EIGRP, which can create stuck in active (SIA) routes that can lead to a network melt-down.

Summarization Example

We have the following networks that we want to advertise as a single summary route:

*

172.16.100.0/24

*

172.16.101.0/24

*

172.16.102.0/24

*

172.16.103.0/24

*

172.16.104.0/24

*

172.16.105.0/24

*

172.16.106.0/24

QUESTION 3

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

A. administrative distance

B. advertised successor

C. successor

D. feasible successor

Correct Answer: D

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

QUESTION 4

The company network is in the process of migrating the IP address scheme to use IPv6. Which of the following address types are associated with IPv6? (Select three)

A. Unicast

B. Private

C. Broadcast

D. Public

E. Multicast

F. Anycast

Correct Answer: AEF

Reference:

http://www.cisco.com/en/US/solutions/collateral/ns340/ns414/ns742/ns824/sbaBN_IPv6addr G.pdf (see page 3, IPv6 address types)

QUESTION 5

Based on the show ip bgp summary output,

```
R1# show ip bgp summary
  BGP table version is 8, main routing table version 8
  4 network entries (8/12 paths) using 832 bytes of memory
  5 BGP path attribute entries using 576 bytes of memory
  0 BGP route-map cache entries using 0 bytes of memory
  0 BGP filter-list cache entries using 0 bytes of memory
  2 received paths for inbound soft reconfiguration

  Neighbor        V    AS    MsgRcvd MsgSent   TblVer  InQ    OutQ    Up/Down      State/PfxRcd
  10.1.1.1        4    50001    80      81        8      0      0      04:15:10     2
  10.2.2.2        4    50002    79      81        0      0      0      00:00:15     Active
  10.3.3.3        4    50003    82      82        0      0      0      02:00:00     Idle
```

Which two statements are true? (Choose two.)

A. The BGP session to the 10.1.1.1 neighbor is established.

B. The BGP session to the 10.2.2.2 neighbor is established.

C. The BGP session to the 10.3.3.3 neighbor is established.

D. The router is attempting to establish a BGP peering session with the 10.1.1.1 neighbor.

E. The BGP session to the 10.3.3.3 neighbor is established, but the router has not received any BGP routing updates from the 10.3.3.3 neighbor.

F. The router is attempting to establish a BGP peering session with the 10.2.2.2 neighbor.

Correct Answer: AF

The main point of this question is the "State/PfxRcd" column, which shows the BGP states. Below is the list of BGP states in order, from startup to peering:

1 Idle: the initial state of a BGP connection. In this state, the BGP speaker is waiting for a BGP start event, generally either the establishment of a TCP connection or the re- establishment of a previous connection. Once the connection is

established, BGP moves to the next state.

2 Connect: In this state, BGP is waiting for the TCP connection to be formed. If the TCP connection completes, BGP will

move to the OpenSent stage; if the connection cannot complete, BGP goes to Active

3 Active: In the Active state, the BGP speaker is attempting to initiate a TCP session with the BGP speaker it wants to peer with. If this can be done, the BGP state goes to OpenSent state.

4 OpenSent: the BGP speaker is waiting to receive an OPEN message from the remote BGP speaker 5 OpenConfirm: Once the BGP speaker receives the OPEN message and no error is detected, the BGP speaker sends a KEEPALIVE

message to the remote BGP speaker

6 Established: All of the neighbor negotiations are complete. You will see a number (2 in this case), which tells us the number of prefixes the router has received from a neighbor or peer group.

QUESTION 6

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

A. eigrp stub

B. eigrp stub connected

C. eigrp stub summary

D. eigrp stub connected static

E. eigrp stub receive-only

Correct Answer: A

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

QUESTION 7

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

A. In interface configuration mode, the OSPFv3 area ID combination assigns interfaces to OSPFv3 areas.

B. In router configuration mode, the network wildcard area ID combination assigns networks to OSPFv3 areas.

C. In interface configuration mode, the IPv6 OSPF process area ID combination assigns interfaces to OSPFv3 areas.

D. In router configuration mode, the IPv6 OSPF interface area ID combination assigns interfaces to OSPFv3 areas.

Correct Answer: C

Reference:

http://www.hh.se/download/18.4cf286ee134f03ddb7b800015/1326882212358/Chapter3_VT 2012.pdf (slide 42)

QUESTION 8

What does the command clear ipv6 ospf process accomplish?

A. The OSPF adjacencies are cleared and initiated again.

B. The route table is cleared. Then the OSPF neighbors are reformed.

C. The shortest path first (SPF) algorithm is performed on the LSA database.

D. The OSPF database is repopulated. Then the shortest path first (SPF) algorithm is performed.

Correct Answer: D

The command "clear ipv6 ospf" will clear the present routing table and force the OSPFv3 process to build a new one. This command is often used when something in the network was changed or for debugging purpose. When the "process" keyword is added, which means "clear ipv6 ospf process", the OSPF database is cleared and repopulated then the SPF algorithm is performed.

QUESTION 9

Refer to the exhibit.

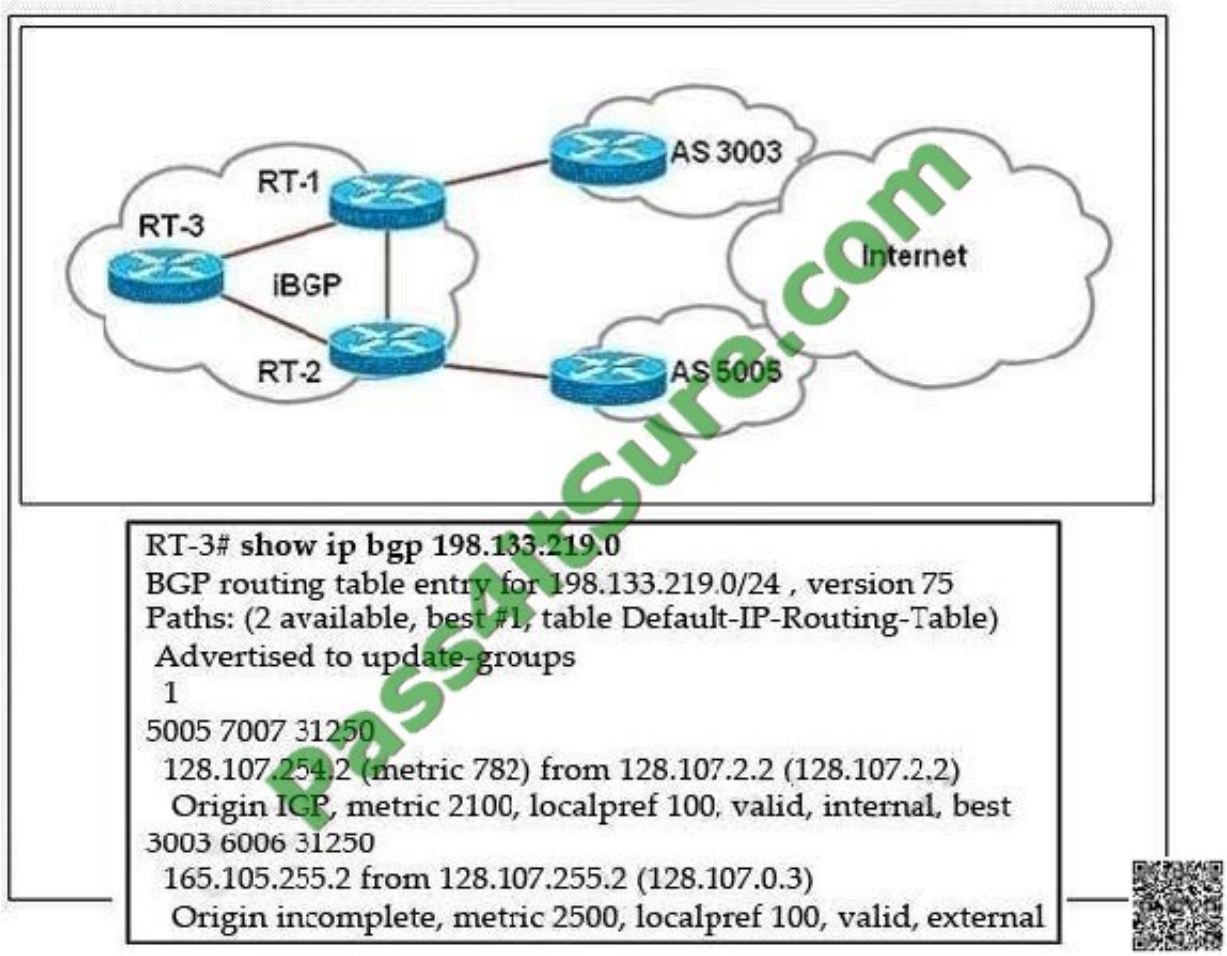

Router RT-1 chooses one path to network 198.133.219.0/24. Indicate the reason Router RT-1 chooses this "best" path.

A. In making its decision about the best path, RT-1 gives precedence to the origin code.

B. In making its decision about the best path, RT-1 gives precedence to the BGP MED values.

C. IP address 128.107.2.2 is lower than 128.107.255.2.

D. In making its decision about the best path, RT-1 prefers the IGP metrics.

E. RT-1 prefers internal BGP routes.

F. IP address 128.107.254.2 is lower than 128.107.255.2.

Correct Answer: A

After BGP receives updates about different destinations from different autonomous systems, it chooses the best path to reach a specific destination. The following process summarizes how BGP chooses the best route on a Cisco router.

Prefer the route with the highest weight. (The weight attribute is proprietary to Cisco and is local to the router only.)

If multiple routes have the same weight, prefer the route with the highest local preference value. (The local preference is used within an autonomous system.) If multiple routes have the same local preference, prefer the route that the local

router originated. A locally originated route has a next hop of 0.0.0.0 in the BGP table. If none of the routes were locally originated, prefer the route with the shortest autonomous system path. If the autonomous system path length is the same,

prefer the lowest origin code (IGP

autonomous system is the same for all routes considered, unless the bgp always-compare-med command is enabled If the routes have the same MED, prefer external paths to internal paths. If synchronization is disabled and only internal

paths remain, prefer the path through the closest IGP neighbor, which means that the router prefers the shortest internal path within the autonomous system to reach the destination (the shortest path to the BGP next hop).

For EBGP paths, select the oldest route to minimize the effect of routes going up and down (flapping). Prefer the route with the lowest neighbor BGP router ID value. If the BGP router IDs are the same, prefer the router with the lowest

neighbor IP address.

QUESTION 10

Which statement is correct based upon the following output from the show command on RT1?

Choose one from below:

A. OSPFv3 establishes neighbor adjacencies by using Link-local addresses.

B. OSPFv3 establishes neighbor adjacencies by using IPv4 addresses.

C. OSPFv3 establishes neighbor adjacencies by using global IPv6 addresses.

D. RT1 owns a subnet mask of 64 bits.

Correct Answer: A

The point of this question is the OSPFv3.

As shown information, we notice that the router with both ipv6 address and ipv4 address and both of them are in ospf prosess, so in OSPFv3, we can find that, the DR BDR address are all link local address, so in OSPFv3 it establist neighbor

with link local address.

QUESTION 11

Refer to the exhibit.

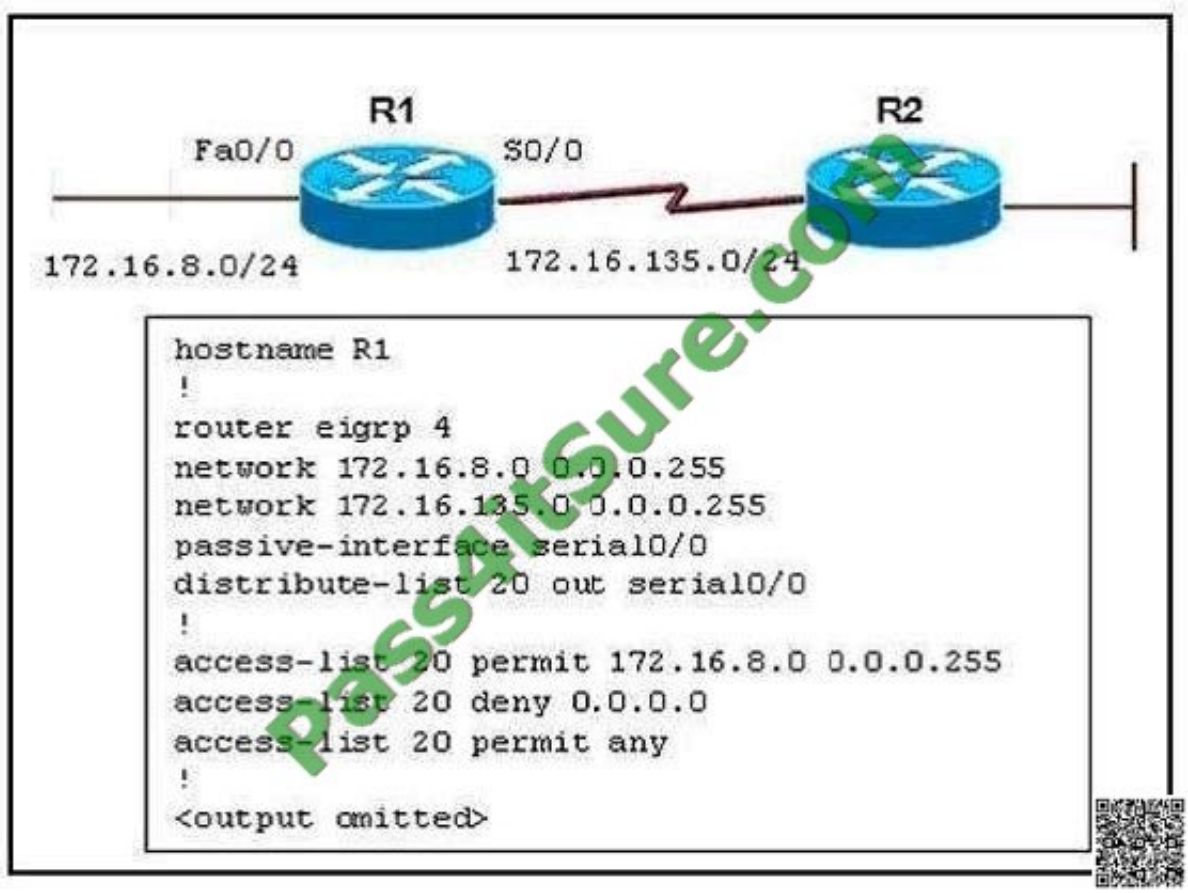

Routers R1 and R2 are running EIGRP and have converged. On the basis of the information that is presented, which statement is true?

A. All outgoing routing updates from router R1 to router R2 will be suppressed, but the inbound updates will continue to be received.

B. All incoming routing updates from R2 will be suppressed, but the outgoing updates will continue to be sent.

C. Both outgoing and incoming routing updates on R1 will be stopped because of the passive- interface Serial0/0 configuration statement.

D. Both outgoing and incoming routing updates on R1 will be permitted because the distribute-list 20 out Serial0/0 command cannot be used with association with the outgoing interface.

Correct Answer: C

You can use the passive-interface command to control the advertisement of routing information. The command enables the suppression of routing updates over some interfaces while it allows updates to be exchanged normally over other interfaces. With most routing protocols, the passive-interface command restricts outgoing advertisements only. However, when used with Enhanced Interior Gateway Routing Protocol (EIGRP), the effect is slightly different. With EIGRP running on a network, the passive- interface command stops both outgoing and incoming routing updates, since the effect of the command causes the router to stop sending and receiving hello packets over an interface. Reference: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f0a.shtml

QUESTION 12

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

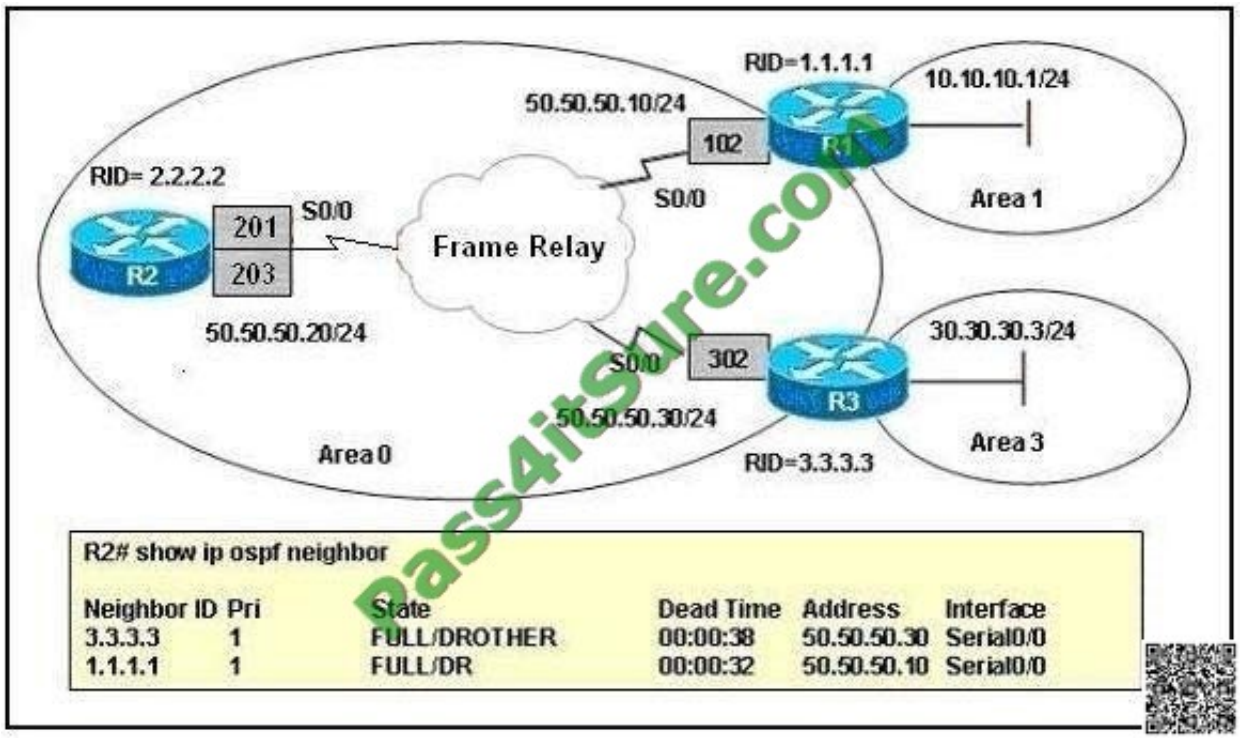

A. The ip ospf network non-broadcast command should be configured on each Frame Relay interface.

B. The ip ospf network broadcast command should be configured on each Frame Relay interface.

C. The neighbor command should be configured under the OSPF routing process on all routers.

D. The ip ospf priority value on the hub router should be set to 0.

E. The ip ospf priority value on the spoke routers should be set to 0.

Correct Answer: E

In an NBMA network topology, neighbors are not discovered automatically. OSPF tries to elect a DR and a BDR due to the multi-access nature of the network, but the election fails since neighbors are not discovered because NBMA environment doesn\\'t forward broadcast and multicast packets. Neighbors must be configured manually to overcome these problems. Also, additional configuration is necessary in a hub and spoke topology to make sure that the hub routers, which have connectivity with every other spoke router, are elected as the DR and BDR. You must set the spoke interfaces to an OSPF priority of zero, this ensures that the spokes will not become the DR or BDR.