



Free Questions for JN0-664 by vceexamstest

Shared by Briggs on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Your network is receiving the 203.0.113.0/24 network using EBGP from AS 64500 and AS 64501. Both of these advertisements have identical local-preference values, AS-path lengths, and BGP origin codes. You want to influence the way your AS sends traffic to the 203.0.113.0/24 network.

In this scenario, which attribute would you consider next when selecting the best path?

Options:

- A- router ID
- B- MED value
- C- peer IP address
- D- IGP metric

Answer:

A

Explanation:

as by default, the MED attribute is only compared for routes received from the same neighbouring AS. The next feasible tiebreaker in the BGP route selection algorithm would be Router ID.

Question 2

Question Type: MultipleChoice

You want to ensure that a single-area OSPF network will be loop free.

In this scenario, what are two requirements that satisfy this requirement? (Choose two.)

Options:

- A-** The DR/BDR ensures that each node within an area has the same information in their LSDBs.
- B-** The Shortest Path First algorithm must prune looped paths.
- C-** Nodes within an area must connect in a full mesh.
- D-** All nodes within an area must have the same information in their LSDBs.

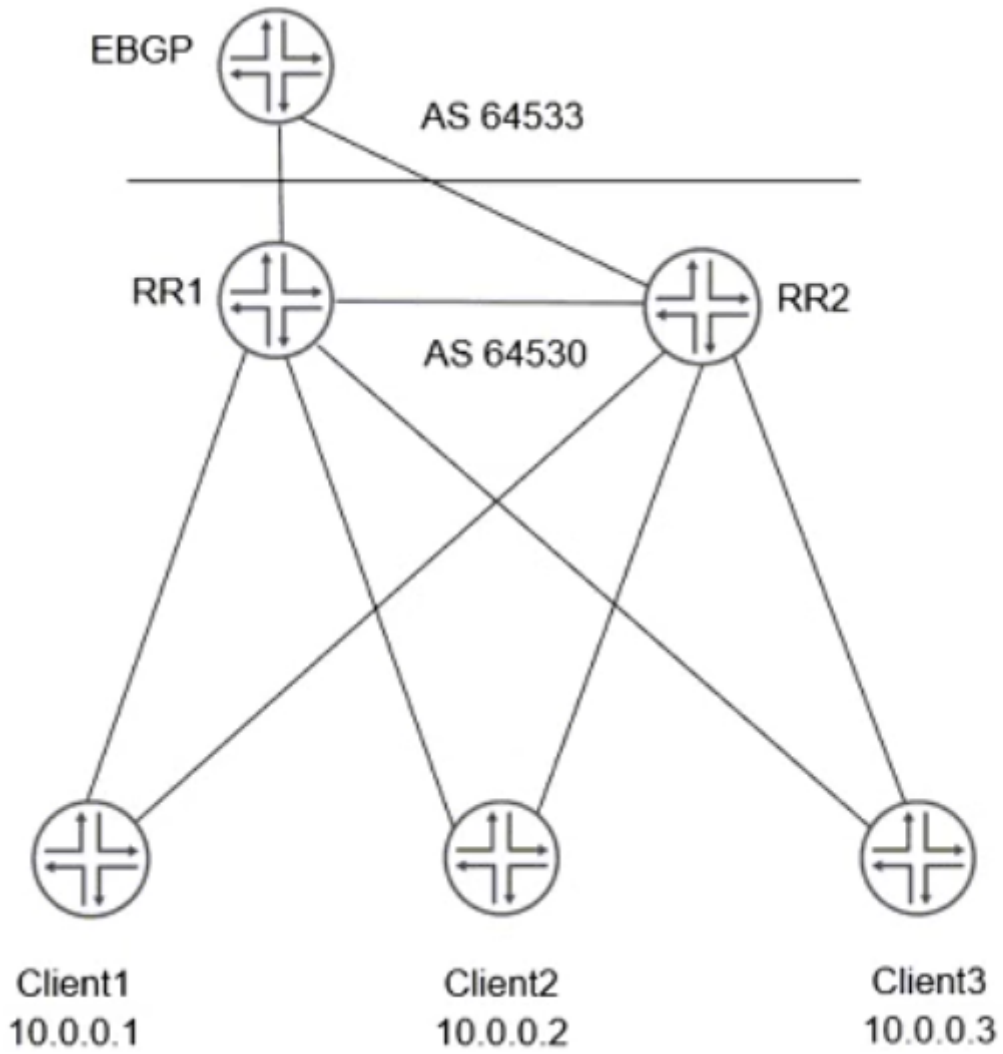
Answer:

B, D

Question 3

Question Type: MultipleChoice

Exhibit



Referring to the exhibit, which two statements are correct about the dual route reflectors within a cluster? (Choose two.)

Options:

- A- RR1 and RR2 must have the same duster ID to exchange routes learned from the client.
- B- RR1 and RR2 append the duster ID when advertising routes from dient to dient.
- C- RR1 and RR2 advertise routes learned from the clients to EBGp peers, using itself as the next hop.
- D- RR1 advertises routes from the client to RR2. using itself as the next hop.

Answer:

B, C

Question 4

Question Type: MultipleChoice

You are configuring anycast RP for load balancing and redundancy in your PIM-SM domain. You want to share active sources between RPs.

In this scenario, what are two solutions that will accomplish this task? (Choose two.)

Options:

- A- Configure MSDP on each RP router.
- B- Configure anycast PIM with the rp-set statement on each RP router.
- C- Configure anycast PIM with the rp-set statement on each source DR router.
- D- Configure MSDP on each source DR router.


Answer:

A, B

Question 5

Question Type: MultipleChoice

Exhibit



```

user@R1> show route forwarding-table matching 203.0.113.128/25
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index  NhRef Netif
203.0.113.128/28 user   0 10.1.1.1          ucst   576    11 ge-0/0/4.0
203.0.113.144/28 user   0 10.1.1.1          ucst   576    11 ge-0/0/4.0
203.0.113.160/28 user   0 10.1.1.1          ucst   576    11 ge-0/0/4.0
203.0.113.176/28 user   0 10.1.1.1          ucst   576    11 ge-0/0/4.0
203.0.113.192/28 user   0 10.1.1.1          ucst   576    11 ge-0/0/4.0
203.0.113.208/28 user   0 10.1.1.1          ucst   576    11 ge-0/0/4.0
203.0.113.224/28 user   0 10.1.1.1          ucst   576    11 ge-0/0/4.0
203.0.113.240/28 user   0 10.1.1.1          ucst   576    11 ge-0/0/4.0

```

You are troubleshooting the connection between AS 64496 and AS 64497 and notice that only one of the paths is being used for traffic forwarding.

Referring to the exhibit, which three actions will ensure that R1 is configured properly for load balancing BGP routes? (Choose three.)

Options:

- A-** Verify that the routing table on R1 has BGP routes for 203.0.113.128/25 with multiple next hops.
- B-** Verify that the multipath option is configured under protocols bgp on both R2 and R3.
- C-** Verify that there is a load balancing export policy under routing-options for the received BGP routes on R1.

D- Verify that the multipath option is configured under protocols bgp on R1.

E- Verify that an import load balancing policy exists under protocols bgp for the received BGP routes on R1.

Answer:

A, B, D

Question 6

Question Type: MultipleChoice

Exhibit

```
(65001)R1-----R2-----R3(65003)
[edit]
user@R2# run show route 11.11.11.0/24
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
11.11.11.0/24      *[BGP/170] 00:00:50, localpref 100
                   AS path: 65001 I, validation-state: unverified
                   > to 172.16.1.1 via ge-0/0/0.0
                   [BGP/170] 00:00:50, localpref 100
                   AS path: 65003 I, validation-state: unverified

[edit]
user@R2# show protocols bgp
group R1 {
    neighbor 172.16.1.1 {
        peer-as 65001;
    }
}
group R3 {
    neighbor 172.16.2.1 {
        peer-as 65003;
    }
}
local-as 65002;
[edit]
user@R2# show policy-options
policy-statement lb {
    then {
        load-balance per-packet;
    }
}
policy-statement prepend {
    term 1 {
        then as-path-prepend 65001;
    }
}
```

R2 is receiving the same route from R1 and R3. You must ensure that you can load balance traffic for that route.

Referring to the exhibit, which configuration change will allow load balancing?

Options:

- A-** Configure the multipath parameter under the global BGP configuration.
- B-** Apply the prepend policy as an import policy under group R1.
- C-** Configure the multipath multiple-as parameter under the global BGP configuration.
- D-** Apply the prepend policy as an import policy under group R3.

Answer:

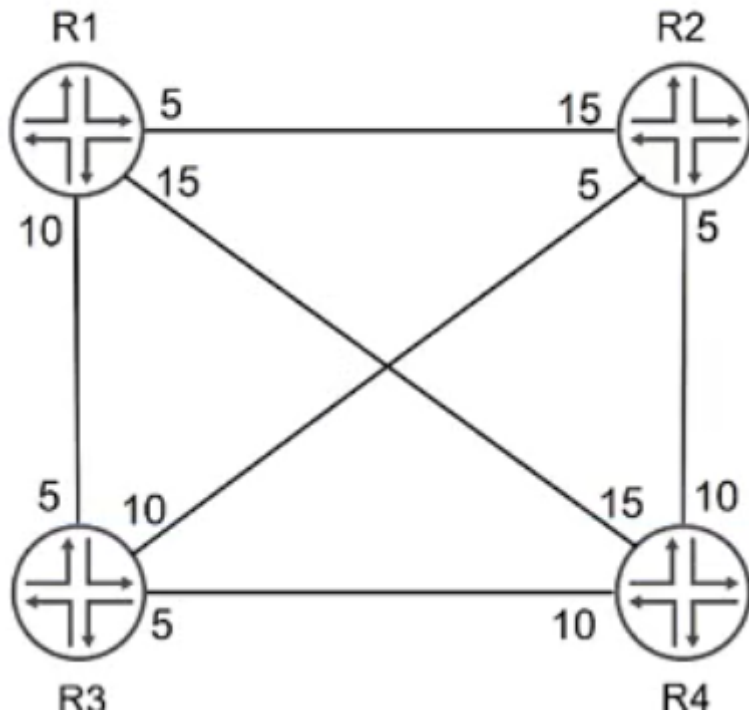
C

Question 7

Question Type: MultipleChoice

Exhibit.

OSPF Area 0



Referring to the exhibit, which path would traffic passing through R1 take to get to R4?

Options:

A- R1 -> R3 -> R4

B- R1 -> R2 -> R3 -> R4

C- R1 -> R2 -> R4

D- R1 -> R4

Answer:

C

Explanation:

The OSPF cost is carried in the LSAs that are exchanged within an OSPF area. When a router calculates the cost to a destination it uses the cost of the exit interface of each router in the path to the destination.

Question 8

Question Type: MultipleChoice

Which two statements are correct regarding bootstrap messages that are forwarded within a PIM sparse mode domain? (Choose two.)

Options:

- A- Bootstrap messages are forwarded only to routers that explicitly requested the messages within the PIM sparse-mode domain
- B- Bootstrap messages distribute RP information dynamically during an RP election.
- C- Bootstrap messages are used to notify which router is the PIM RP
- D- Bootstrap messages are forwarded to all routers within a PIM sparse-mode domain.

Answer:

B, D

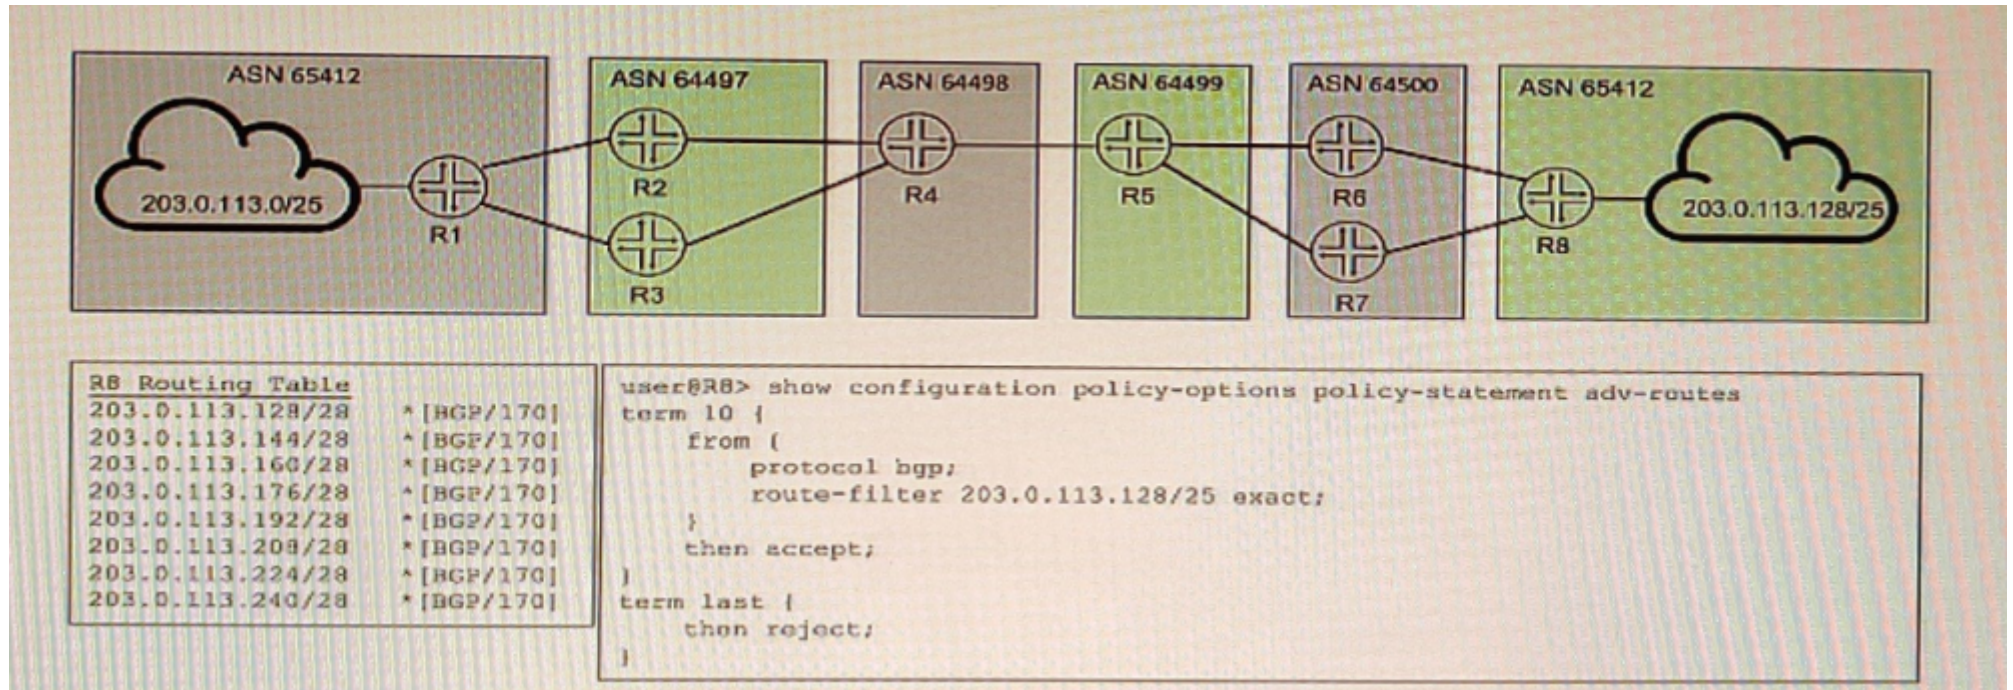
Explanation:

Bootstrap messages are PIM messages that are used to distribute rendezvous point (RP) information dynamically during an RP election. Bootstrap messages are sent by bootstrap routers (BSRs), which are routers that are elected to perform the RP discovery function for a PIM sparse-mode domain. Bootstrap messages contain information about candidate RPs and their multicast groups, as well as BSR priority and hash mask length. Bootstrap messages are forwarded to all routers within a PIM sparse-mode domain using hop-by-hop flooding.

Question 9

Question Type: MultipleChoice

Exhibit



You are attempting to summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500. You implement the export policy shown in the exhibit and all routes from the routing table stop being advertised.

In this scenario, which two steps would you take to summarize the route in BGP? (Choose two.)

Options:

- A- Remove the from protocol bgp command from the export policy.
- B- Add the set protocols bgp family inet unicast add-path command to allow additional routes to the RIB tables. -
- C- Add the set routing-options static route 203.0.113.123/25 discard command.
- D- Replace exact in the export policy with orlonger.

Answer:

C, D

Explanation:

To summarize routes from the 203.0.113.128/25 IP block on R8 to AS 64500, you need to do the following:


Add the set routing-options static route 203.0.113.128/25 discard command. This creates a static route for the summary prefix and discards any traffic destined to it. This is necessary because BGP can only advertise routes that are present in the routing table.

Replace exact in the export policy with orlonger. This allows R8 to match and advertise any route that is equal or more specific than the summary prefix. The exact term only matches routes that are exactly equal to the summary prefix, which is not present in the routing table.

Question 10

Question Type: MultipleChoice

Exhibit



You want to use both links between R1 and R2. Because of the bandwidth difference between the two links, you must ensure that the links are used as much as possible.

Which action will accomplish this goal?

Options:

- A-** Define a policy to tag routes with the appropriate bandwidth community.
- B-** Disable multipath.

- C-** Ensure that the metric-out parameter on the Gigabit Ethernet interface is higher than the 10 Gigabit Ethernet interface.
- D-** Enable per-prefix load balancing.

Answer:

A

Explanation:

<https://www.juniper.net/documentation/us/en/software/junos/sampling-forwarding-monitoring/bgp/topics/concept/bgp-multipath-unequal-understanding.html>

Question 11

Question Type: MultipleChoice

In which two ways does OSPF prevent routing loops in multi-area networks? (Choose two.)

Options:

- A- All areas are required to connect as a full mesh.
- B- The LFA algorithm prunes all looped paths within an area.
- C- All areas are required to connect to area 0.
- D- The SPF algorithm prunes looped paths within an area.

Answer:

C, D

Explanation:

OSPF is an interior gateway protocol that uses link-state routing to exchange routing information among routers within a single autonomous system. OSPF prevents routing loops in multi-area networks by using two methods: area hierarchy and SPF algorithm. Area hierarchy is the concept of dividing a large OSPF network into smaller areas that are connected to a backbone area (area 0). This reduces the amount of routing information that each router has to store and process, and also limits the scope of link-state updates within each area. All areas are required to connect to area 0 either directly or through virtual links². SPF algorithm is the method that OSPF uses to calculate the shortest path to each destination in the network based on link-state information. The SPF algorithm runs on each router and builds a shortest-path tree that represents the topology of the network from the router's perspective. The SPF algorithm prunes looped paths within an area by choosing only one best path for each destination³.

Question 12

Question Type: MultipleChoice

Exhibit

```

user@R4> show pim rps
Instance: PIM.master
address-family INET
RP address      Type      Mode    Holdtime Timeout Groups Group prefixes
10.1.255.2      bootstrap sparse   150     118     0 224.1.1.0/24
10.1.255.3      bootstrap sparse   150     118     2 224.1.1.0/28
user@R4> show route 10.1.255.2
inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.1.255.2/32    *[IS-IS/18] 00:32:27, metric 10
                 > to 10.1.1.2 via ge-0/0/0.0
inet.2: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0       *[Static/5] 00:13:55
                 > to 10.1.1.6 via ge-0/0/1.0
user@R4> show route 10.1.255.3

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.1.255.3/32    *[IS-IS/18] 00:32:43, metric 10
                 > to 10.1.1.6 via ge-0/0/1.0
inet.2: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0       *[Static/5] 00:14:25
                 > to 10.1.1.6 via ge-0/0/1.0

[edit]
user@R2# show protocols pim
rp {
  bootstrap {
    family inet {
      priority 200;
    }
  }
}

```

R4 is directly connected to both RPs (R2 and R3) R4 is currently sending all ,o,ns upstream to R3 but you want all joins to go to R2 instead Referring to the exhibit, which configuration change will solve this issue?

Options:

- A-** Change the bootstrap priority on R2 to be higher than R3
- B-** Change the default route in inet.2 on R4 from R3 as the next hop to R2
- C-** Change the local address on R2 to be higher than R3.
- D-** Change the group-range to be more specific on R2 than R3.

Answer:

D

To Get Premium Files for JN0-664 Visit

<https://www.p2pexams.com/products/jn0-664>

For More Free Questions Visit

<https://www.p2pexams.com/juniper/pdf/jn0-664>

