

Juniper

Exam JN0-694

Enterprise Routing and Switching Support, Professional (JNCSP-
ENT)

Verson: Demo

[Total Questions: 10]

Question No : 1

-- Exhibit --

```
user@router> show route protocol bgp detail
```

```
inet.0: 20 destinations, 20 routes (19 active, 0 holddown, 1 hidden)
```

```
10.222.1.3/32 (1 entry, 1 announced)
```

```
*BGP Preference: 170/-101
```

```
Next hop type: Indirect
```

```
Address: 0x15ec944
```

```
Next-hop reference count: 3
```

```
Source: 1.1.1.1
```

```
Next hop type: Router, Next hop index: 536
```

```
Next hop: 1.1.1.1 via ge-0/0/1.0, selected
```

```
Protocol next hop: 1.1.1.1
```

```
Indirect next hop: 14081d0 262142
```

```
State:
```

```
Local AS: 65222 Peer AS: 65221
```

```
Age: 2:12 MetriC. 1 Metric2: 0
```

```
Task: BGP_65221.1.1.1.1+56417
```

```
Announcement bits (2): 0-KRT 4-Resolve tree 1
```

```
AS path: 65221 I
```

```
Communities: no-advertise
```

```
Accepted
```

```
Localpref: 100
```

```
Router ID. 10.222.1.1
```

-- Exhibit --

Click the Exhibit button.

You are troubleshooting a problem where an EBGP route is not being advertised to your local IBGP peers. You have received a 10.222.1.3/32 route from an EBGP peer as shown in the exhibit, but the route is not being advertised.


What is causing the problem?

- A. The route shows as a hidden route and cannot be advertised.
- B. The next hop for the route is indirect and prevents the route from being advertised.
- C. The community prevents the route from being advertised.
- D. The local preference value is too high for the route to be advertised.

Answer: C

Question No : 2

-- Exhibit --



```

user@R3> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table
inet.0          Tot Paths  Act Paths  Suppressed  History  Damp  State  Pending
Peer           AS         InPkt     OutPkt     OutQ     Flaps  Last  Up/Dwn
State|#Active/Received/Accepted/Damped...
172.22.0.1     65501      3         3         0        0      3 0/0/0/0  0/0/0/0
192.168.1.2    65502      8         6         0        0      1:52 0/5/5/0  0/0/0/0

user@R3> show route advertising-protocol bgp 172.22.0.1

user@R3> show route 100.100.0.0/16 terse

inet.0: 14 destinations, 19 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 100.100.0.0/24   O 150      0          >172.24.0.2
                          B 170      100         >172.24.0.2   I
* 100.100.1.0/24   O 150      0          >172.24.0.2
                          B 170      100         >172.24.0.2   I
* 100.100.2.0/24   O 150      0          >172.24.0.2
                          B 170      100         >172.24.0.2   I
* 100.100.3.0/24   O 150      0          >172.24.0.2
                          B 170      100         >172.24.0.2   I
* 100.100.4.0/24   O 150      0          >172.24.0.2
                          B 170      100         >172.24.0.2   I
    
```

-- Exhibit --

Click the Exhibit button.

You are asked to assist with a problem with a new EBGP peering between Site 1 and Site 2. Referring to the exhibit, Site 1 is not receiving the 100.100.0.0/16 routes from Site 2.

Which action will resolve the problem?

- A. Enable the advertise-inactive parameter for the EBGP peering.
- B. Enable the as-override parameter for the EBGP peering.
- C. Create an export policy to export the IBGP routes over the EBGP peering.
- D. Create a next-hop-self policy and apply it as an export policy to the EBGP peering.

Answer: A

Question No : 3

-- Exhibit --

```
user@router# show class-of-service

classifiers {
inet-precedence ipp-test {
import default;

forwarding-class best-effort {
loss-priority low code-points be;
}

forwarding-class expedited-forwarding {
loss-priority low code-points af21;
}

forwarding-class assured-forwarding {
loss-priority low code-points af11;
}

forwarding-class network-control {
```

```
loss-priority low code-points nc1;

}

}

}

interfaces {

ge-* {

scheduler-map map-test;

unit * {

classifiers {

inet-precedence ipp-test;

}

rewrite-rules {

inet-precedence ipp-rw-test;

inet-precedence default;

}

}

}

...

rewrite-rules {

inet-precedence ipp-rw-test {

forwarding-class best-effort {

loss-priority low code-point be;

loss-priority high code-point af21;

}

forwarding-class expedited-forwarding {
```

```
loss-priority high code-point af21;

loss-priority low code-point be;

}

forwarding-class assured-forwarding {

loss-priority low code-point af11;

loss-priority high code-point af11;

}

forwarding-class network-control {

loss-priority low code-point nc1;

loss-priority high code-point nc1;

}

}

}
```

```
user@router> show class-of-service
```

```
...
```

```
Code point type: inet-precedence
```

```
Alias Bit pattern
```

```
af11 001
```

```
af21 010
```

```
af31 011
```

```
af41 100
```

```
be 000
```

```
cs6 110
```

```
cs7 111
```

```
ef 101
```

```
nc1 110
```

nc2 111

-- Exhibit --

Click the Exhibit button.

Traffic with the IP precedence value af21 ingresses the router and should be rewritten with the same value as it egresses; however, this traffic is rewritten to a different value.

Referring to the exhibit, what is the source of this problem?

- A. The BA classifier is assigning the traffic to the best-effort queue with a high loss priority.
- B. The BA classifier is assigning the traffic to the best-effort queue with a low loss priority.
- C. The BA classifier is assigning the traffic to the expedited forwarding queue with a high loss priority.
- D. The BA classifier is assigning the traffic to the expedited forwarding queue with a low loss priority.

Answer: D

Question No : 4

-- Exhibit --

```
user@R1> show class-of-service interface ge-0/0/0
```

```
Physical interface: ge-0/0/0, Index: 134
```

```
Queues supported: 8, Queues in use: 4
```

```
Scheduler map: , Index: 2
```

```
Congestion-notification: Disabled
```

```
Logical interface: ge-0/0/0.0, Index: 69
```

```
Object Name Type Index
```

```
Classifier ipprec-compatibility ip 13
```

-- Exhibit --

Click the Exhibit button.

You are sending traffic to the ge-0/0/0 interface on R1 with the expedited forwarding (101) IP precedence bits. However, the counters on the router show that it is not processing any traffic in the expedited forwarding queue.

Referring to the exhibit, what is causing the problem?

- A. The classifier does not contain a mapping for this precedence value.
- B. The classifier is assigning the traffic a high loss priority.
- C. The classifier is assigning the traffic to the best-effort forwarding class.
- D. The classifier is dropping the traffic.

Answer: C

Question No : 5

Your Junos device is dropping certain traffic flows, while allowing other traffic flows to pass through the device unaffected.

Which CoS component is causing this problem?

- A. BA classification
- B. RED
- C. MF classification
- D. Rewrite rules

Answer: D

Question No : 6

You use static routes for connectivity to the ISP. Your ISP recently switched to using different links for multicast and unicast traffic. Following the change, users in your company were unable to receive multicast traffic through the ISP.

What must you configure on your router to reestablish multicast connectivity to your ISP?

- A. Add a static default route to the ISP in the inet.2 routing table.
- B. Add the default-rpf-interface parameter under the [edit routing-options multicast] hierarchy.
- C. Add the upstream-interface parameter under the [edit protocols pim] hierarchy.
- D. Disable PIM on the interface used for unicast traffic.

Answer: A

Question No : 7

-- Exhibit --

```
user@R1> show ospf neighbor
```

```
Address Interface State ID Pri Dead
```

```
10.222.0.2 ge-0/0/1.0 Init 10.222.1.2 128 32
```

```
user@R1> show ospf interface detail
```

```
Interface State Area DR ID BDR ID Nbrs
```

```
ge-0/0/1.0 DR 0.0.0.0 10.222.1.1 0.0.0.0 1
```

```
Type: LAN, Address: 10.222.0.1, Mask: 255.255.255.252, MTU: 1500, Cost: 1
```

```
DR addr: 10.222.0.1, Priority: 128
```

```
Adj count: 0
```

```
Hello: 10, DeaD. 40, ReXmit: 5, Not Stub
```

```
Auth type: MD5, Active key ID. 10, Start time: 1970 Jan 1 00:00:00 UTC
```

```
Protection type: None
```

```
Topology default (ID 0) -> Cost: 1
```

```
lo0.0 DR 0.0.0.0 10.222.1.1 0.0.0.0 0
```

```
Type: LAN, Address: 10.222.1.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
```

```
DR addr: 10.222.1.1, Priority: 128
```

```
Adj count: 0
```

```
Hello: 10, DeaD. 40, ReXmit: 5, Not Stub
```

Auth type: None

Protection type: None

Topology default (ID 0) -> Cost: 0

user@R2> show ospf neighbor

user@R2> show ospf interface detail

Interface State Area DR ID BDR ID Nbrs

ge-0/0/1.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 0

Type: P2P, Address: 10.222.0.2, Mask: 255.255.255.252, MTU: 1500, Cost: 1

Adj count: 0

Hello: 10, DeaD. 40, ReXmit: 5, Not Stub

Auth type: MD5, Active key ID. 10, Start time: 1970 Jan 1 00:00:00 UTC

Protection type: None

Topology default (ID 0) -> Cost: 1

lo0.0 DR 0.0.0.0 10.222.1.2 0.0.0.0 0

Type: LAN, Address: 10.222.1.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0

DR addr: 10.222.1.2, Priority: 128

Adj count: 0

Hello: 10, DeaD. 40, ReXmit: 5, Not Stub

Auth type: None

Protection type: None

Topology default (ID 0) -> Cost: 0

-- Exhibit --

Click the Exhibit button.

You are trying to establish an OSPF adjacency between R1 and R2, but the adjacency does not establish.

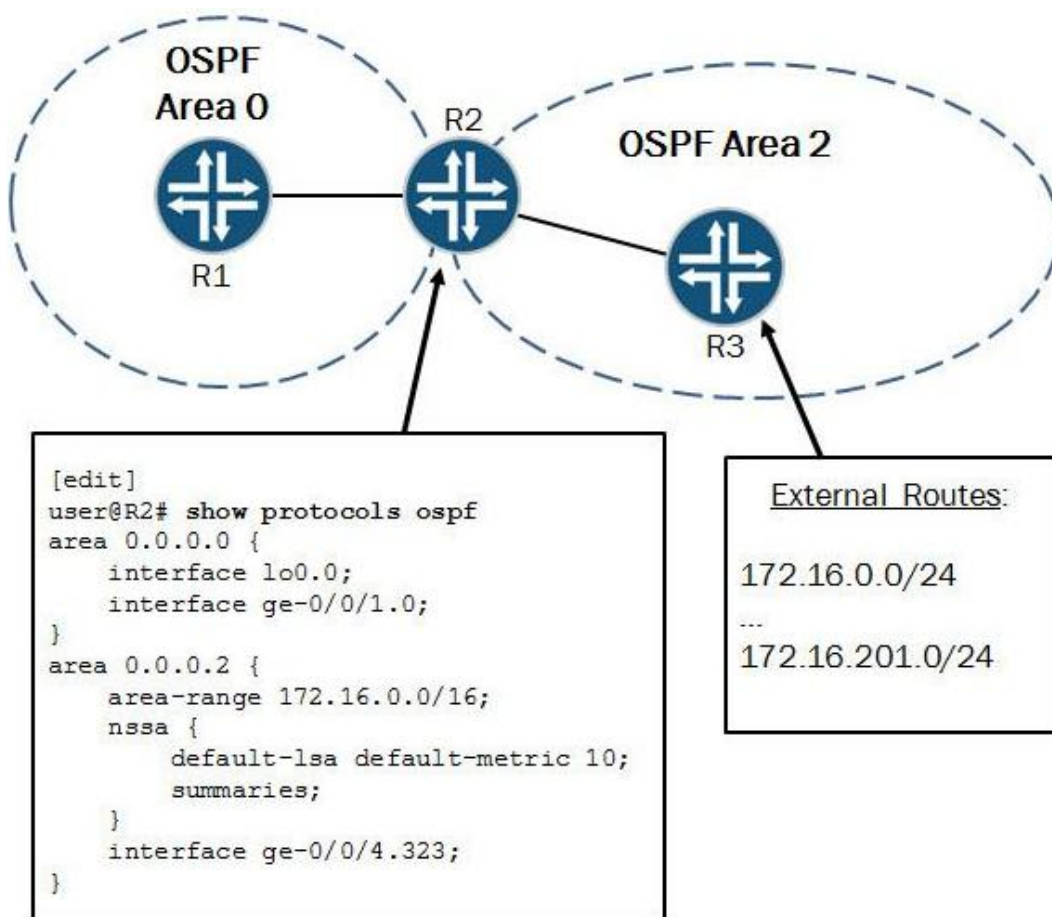
Referring to the exhibit, what is causing the adjacency to fail?

- A. The MD5 key ID values are mismatched between R1 and R2.
- B. R1 has both family inet and family iso configured on the link toward R2.
- C. The IP subnet mask is mismatched between R1 and R2.
- D. The interface type is mismatched between R1 and R2.

Answer: D

Question No : 8

You are troubleshooting a problem where external routes are not being summarized into the OSPF backbone.



Referring to the exhibit, what needs to be done to resolve this problem?

- A. The area-range parameter needs to be under Area 0.
- B. The area-range parameter needs to be under the nssa hierarchy.

- C. The summaries parameter needs to be removed under the/issa hierarchy.
- D. The area-range parameter must include the override-metric parameter.

Answer: B

Question No : 9

-- Exhibit --

```

user@R1> show route
inet.0: 5 destinations, 5 routes (5 active, 0
holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32      *[Direct/0] 00:01:10
                > via lo0.0
2.2.2.2/32      *[OSPF/10] 00:00:13, metric 1
                > to 172.10.1.2 via ge-0/0/1.0
172.10.1.0/24   *[Direct/0] 00:01:10
                > via ge-0/0/1.0
172.10.1.1/32   *[Local/0] 00:01:10
                Local via ge-0/0/1.0
224.0.0.5/32    *[OSPF/10] 00:01:10, metric 1
                MultiRecv

user@R1> show ospf database
Jun 12 03:33:34
  OSPF database, Area 0.0.0.0
  Type ID Adv Rtr Seq Age Opt Cksum Len
  Router 2.2.2.2 2.2.2.2 0x80000005 30 0x22 0xeb10 60
  Router *200.200.200.200 200.200.200.200 0x80000009 7 0x22 0x1d42 48
  Network *172.10.1.1 200.200.200.200 0x80000005 2 0x22 0xcc62 32
  Network *172.20.1.3 200.200.200.200 0x80000004 3600 0x22 0x42e1 32

user@R1> show ospf database
Jun 12 03:33:46
  OSPF database, Area 0.0.0.0
  Type ID Adv Rtr Seq Age Opt Cksum Len
  Router 2.2.2.2 2.2.2.2 0x80000005 42 0x22 0xeb10 60
  Router *200.200.200.200 200.200.200.200 0x8000000d 3 0x22 0x1546 48
  Network *172.10.1.1 200.200.200.200 0x80000006 6 0x22 0xca63 32
  Network *172.20.1.3 200.200.200.200 0x80000005 3600 0x22 0x40e2 32

user@R1> show ospf interface ge-0/0/1.0 detail
Interface State Area DR ID BDR ID Nbrs
ge-0/0/1.0 DR 0.0.0.0 200.200.200.200 2.2.2.2 1
Type: LAN, Address: 172.10.1.1, Mask: 255.255.255.0,
MTU: 1500, Cost: 1
DR addr: 172.10.1.1, BDR addr: 172.10.1.2, Priority:
128

user@R1> show ospf neighbor detail
Address Interface State ID Pri Deac
172.10.1.2 ge-0/0/1.0 Full 2.2.2.2 128 31
  
```

-- Exhibit --

Click the Exhibit button.

Referring to the exhibit, you are configuring an OSPF network. All OSPF adjacencies come up and stay stable. But neither R1 nor R2 has the prefix 200.200.200.200/32 in its routing table.

What is causing this problem?

- A. R2 does not have the export policy for prefix 200.200.200.200/32.
- B. R1 does not have routes to network 172.10.1.0/24.

C. R2 is BDR on both network 172.10.1.0/24 and 172.20.1.0/24.

D. The router ID of R1 is the same as the router ID of R3.

Answer: D

Question No : 10

-- Exhibit --

[edit routing-instances]

user@router# show vr1 routing-options

instance-import [vr1 vr2];

[edit routing-instances]

user@router# show vr2 routing-options

instance-import [vr1 vr2];

[edit routing-instances]

user@router# top show policy-options policy-statement vr1

term 1 {

from instance vr1;

then accept;

}

term 2 {

then reject;

}

[edit routing-instances]

user@router# top show policy-options policy-statement vr2

term 1 {

from instance vr2;

```
then accept;
```

```
}
```

```
term 2 {
```

```
then reject;
```

```
}
```

```
-- Exhibit --
```

Click the Exhibit button.

A network engineer wants to leak routes between routing instances vr1 and vr2. No routes from vr2 are showing up in vr1.

Which change should the engineer make to accomplish this task?

A. [edit routing-instances]

```
user@router# delete vr1 routing-options instance-import
```

```
[edit routing-instances]
```

```
user@router# set vr1 routing-options instance-import (vr1 || vr2)
```

B. [edit routing-instances]

```
user@router# delete vr1 routing-options instance-import
```

```
[edit routing-instances]
```

```
user@router# set vr1 routing-options instance-import (vr1 && vr2)
```

C. [edit routing-instances]

```
user@router# set vr1 routing-options auto-export
```

D. [edit routing-instances]

```
user@router# set vr1 routing-options interface-routes rib-group vr2
```

Answer: A