

JN0-694^{Q&As}

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

Pass Juniper JN0-694 Exam with 100% Guarantee

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

<https://www.leads4pass.com/jn0-694.html>

100% Passing Guarantee
100% Money Back Assurance

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

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

```
-- 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; } }
```

```
user@router# show firewall filter MF { term 1 { from { precedence 0; } then forwarding-class best-effort; } term 2 { from { precedence 5; } then forwarding-class expedited-forwarding; } term 3 { from { precedence 2; } then forwarding-class assured-forwarding; } term 4 { from { precedence 6; } then forwarding-class network-control; } term 5 { then accept; } }
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 IPP value af21 should be assigned to the expedited forwarding queue; however, this traffic is not being assigned to that queue.


Referring to the exhibit, what is causing this behavior?

- A. The af21 traffic is assigned to the assured forwarding queue because of the BA classifier.
- B. The af21 traffic is assigned to the assured forwarding queue because of the MF classifier.
- C. The af21 traffic is assigned to the best effort queue because of the MF classifier.
- D. The af21 traffic is assigned to the best effort queue because of the BA classifier.

Correct Answer: B

QUESTION 2

-- Exhibit



```
[edit]
user@R1# show protocols bgp
group ebgp {
  type external;
  multihop;
  local-address 1.1.1.1;
  peer-as 65222;
  neighbor 172.16.1.2;
}

[edit]
user@R1# show routing-options
static {
  route 172.16.1.2/32 next-hop 1.1.1.1;
}
router-id 10.222.1.1;
autonomous-system 65221;
```

Loopback IP:
172.16.1.2

-- Exhibit -

Click the Exhibit button.

You are asked to configure a multihop EBGP peering to a loopback address of your ISP. The peering does not establish, and the ISP has verified that the settings are correct on their side.

Referring to the exhibit, what is causing the problem?

- A. The peer-as parameter is misconfigured.
- B. The local-address parameter is misconfigured.
- C. The static route is misconfigured.
- D. The neighbor parameter is misconfigured.

Correct Answer: C

QUESTION 3

-- 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  Cksaum  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 0xd42  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  Cksaum  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.

Correct Answer: D

QUESTION 4

You have configured OSPF between two routers and the adjacency is not coming up. You confirm that the physical link between them is up and then run the commands shown in the exhibit on both routers. Which two configuration mistakes

apply? (Choose two.)

```
user@R1> show ospf interface
Interface          State      Area      DR ID      BDR ID      Nbrs
fe-0/0/1.0        DR        0.0.0.1   1.1.1.1    0.0.0.0     0
Type: LAN, Address: 10.50.10.26, Mask: 255.255.255.252, MTU: 1500, Cost: 1
DR addr: 10.50.10.26, 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 interface
Interface          State      Area      DR ID      BDR ID      Nbrs
fe-0/0/2.0        DR        0.0.0.2   1.1.1.2    0.0.0.0     0
Type: LAN, Address: 10.50.10.25, Mask: 255.255.255.252, MTU: 1500, Cost: 1
DR addr: 10.50.10.25, Priority: 128
Adj count: 0
Hello: 20, Dead: 80, ReXmit: 5, Not Stub
Auth type: None
Protection type: None
Topology default (ID 0) -> Cost: 1
```

- A. The hello timer is mismatched.
- B. The subnet is mismatched.
- C. The DR ID is mismatched.
- D. The area ID is mismatched.

Correct Answer: AD

QUESTION 5

-- 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.

Correct Answer: C

QUESTION 6

-- 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 andand 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

Correct Answer: A

QUESTION 7

```
-- Exhibit -user@R1> show bgp neighbor 172.10.3.201 Peer: 172.10.3.201+54714 AS 64512 Local: 172.10.3.202+179 AS 64513 Type: External State: Established Flags: Last State: OpenConfirm Last Event: RecvKeepAlive Last Error: None Export: [ export ] Options: Local Address: 172.10.3.202 Holdtime: 90 Preference: 170 Local AS: 64513 Local System AS: 0 Number of flaps: 0 Peer ID. 10.247.194.254 Local ID. 10.247.24.6 Active Holdtime: 90 Keepalive Interval: 30 Peer index: 0 BFD. disabled, down Local Interface: ge-0/0/0.500 NLRI for restart configured on peer: inet-unicast NLRI advertised by peer: inet-unicast NLRI for this session: inet-unicast Peer supports Refresh capability (2) Restart time configured on the peer: 120 Stale routes from peer are kept for: 300 Restart time requested by this peer: 120 NLRI that peer supports restart for: inet-unicast NLRI that restart is negotiated for: inet-unicast NLRI of received end-of-rib markers: inet-unicast NLRI of all end-of-rib markers sent: inet-unicast Peer supports 4 byte AS extension (peer-as 64512) Peer does not support Addpath Table inet.0 Bit: 30000 RIB State: BGP restart is complete RIB State: VPN restart is complete Send state: in sync Active prefixes: 7 Received prefixes: 7 Accepted prefixes: 7 Suppressed due to damping: 0 Advertised prefixes: 30 Last traffic (seconds): Received 5 Sent 18 Checked 8 Input messages: Total 40 Updates 3 Refreshes 0 Octets 877 Output messages: Total 55 Updates 13 Refreshes 0 Octets 1764 Output Queue[2]: 0 -- Exhibit -
```

Click the Exhibit button.

A customer reports that BGP graceful restart is not working on R1. After a Routing Engine failover, R1 did not set the restart state bit in its Open message. The customer provides the BGP neighbor output shown in the exhibit.

Referring the exhibit, what is causing this problem?

- A. BGP graceful restart is not enabled on R1.
- B. BGP graceful restart is not enabled on the peer device.
- C. The restart duration time is not configured on R1.
- D. The restart duration time is not configured on the peer device.

Correct Answer: A

QUESTION 8

Your switch is experiencing a problem where a port that should have only one host connected occasionally shows that multiple MAC addresses are being learned.

Which configuration setting would ensure that no extra hosts can join the network using this switch port?

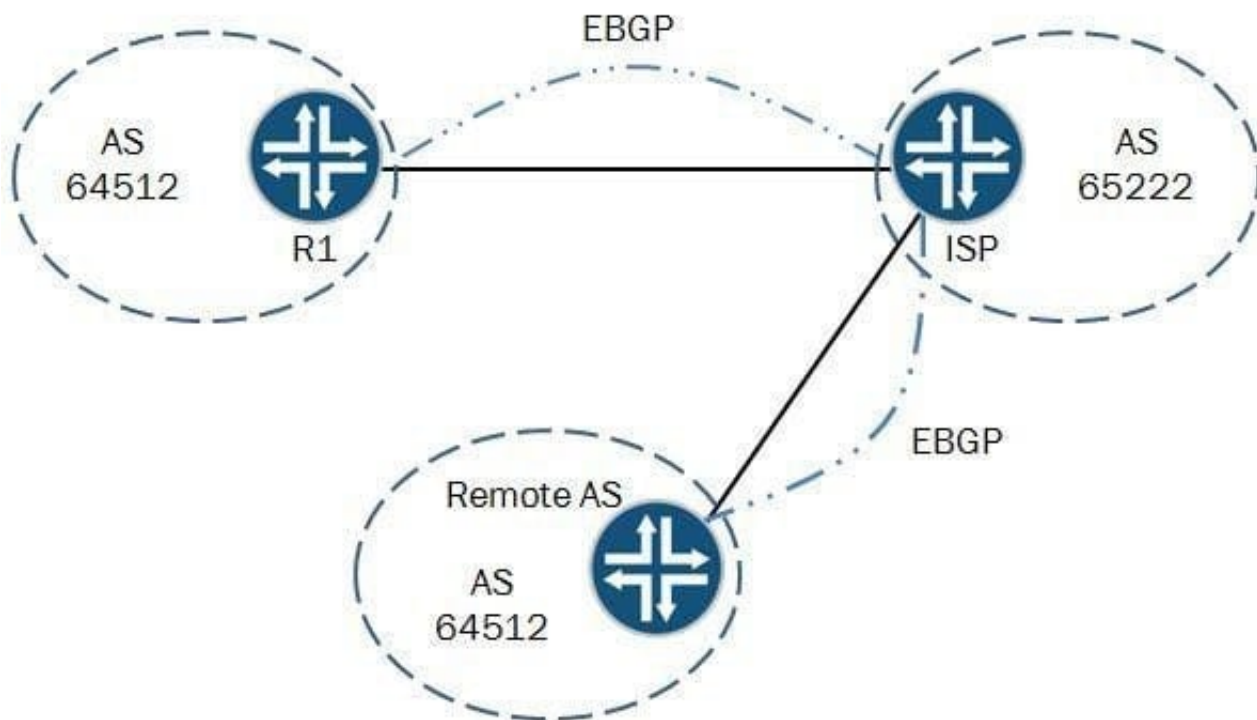
- A. mac-limit
- B. no-mac-learning
- C. persistent-learning
- D. bpdu-block-on-edge

Correct Answer: D

QUESTION 9

You want to receive routes from a remote EBGP peer as shown in the exhibit. The remote site has the same AS number as your own. You have successfully established the EBGP peering, but are not receiving any routes.

Which BGP group configuration parameter will enable you to receive routes from the remote EBGP peer?



- A. advertise-inactive
- B. multipath
- C. as-override
- D. remove-private

Correct Answer: C

QUESTION 10

You recently deployed two Anycast RPs. Multicast clients in the network are reporting that they are receiving traffic from some, but not all, multicast sources.

What are two solutions for this problem? (Choose two.)

- A. Configure MSDP between the Anycast RPs.
- B. Configure rp-set for the Anycast RPs.
- C. Configure multicast BGP between the Anycast RPs.
- D. Configure the network to always use the RPT and not switch over to the SPT.

Correct Answer: AB

[Latest JN0-694 Dumps](#)

[JN0-694 VCE Dumps](#)

[JN0-694 Study Guide](#)