



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.passapply.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 -(MSTI 2 regional root: 16386.2c:6b:f5:3e:f8:01)

{master:0}

user@switch> show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface Port ID Designated Designated Port State Role port ID bridge ID Cost

ge-0/0/6.0 128:519 128:519 16384.80711fbc 20000 BLK ALT ge-0/0/9.0 128:522 128:522

53248.2c6bf591a441 20000 FWD DESG ge-0/0/10.0 128:523 128:523 8192.80711fbe8110 20000 FWD

ROOT ge-0/0/12.0 128:525 128:525 49152.2c6bf53ef801 20000 BLK ALT

[...]

-- Exhibit -

Click the Exhibit button.

While troubleshooting an MSTP operation in your network, you see the output shown in the exhibit on one of your switches. You know that the MSTI 2 regional root bridge ID is 16386.2c:6b:f5:3e:f8:01.

Which port is attached to the root bridge of MSTI 2?

- A. ge-0/0/6
- B. ge-0/0/9
- C. ge-0/0/10
- D. ge-0/0/12

Correct Answer: D

QUESTION 2

-- Exhibit -(master:0)[edit ethernet-switching-options secure-access-port] user@switch# show

interface ge-0/0/1.0 {

static-ip 172.27.0.2 vlan v11 mac 00:0c:29:b5:89:7c;

no-dhcp-trusted;

}

vlan v11 {



```
arp-inspection;
```

```
}
```

```
interface ge-0/0/2.0 {
```

```
dhcp-trusted;
```

```
}
```

```
user@switch> show log messages | match arp
```

```
Feb 8 14:31:45 switch eswd[1280]: ESWD_DAI_FAILED. 3 ARP_REQUEST received, interface ge0/0/1.0[index 73],  
vlan v11[index 5], sender ip/mac 172.27.0.2/00:0c:29:b5:89:7d, receiver ip/mac
```

```
172.27.0.1/00:00:00:00:00:00
```

```
Feb 8 14:34:05 switch eswd[1280]: ESWD_DAI_FAILED. 3 ARP_REQUEST received, interface ge0/0/1.0[index 73],  
vlan v11[index 5], sender ip/mac 172.27.0.2/00:0c:29:b5:89:7d, receiver ip/mac
```

```
172.27.0.1/00:00:00:00:00:00 Feb 8 14:36:05 switch eswd[1280]: ESWD_DAI_FAILED. 3 ARP_REQUEST received,  
interface ge0/0/1.0[index 73], vlan v11[index 5], sender ip/mac 172.27.0.2/00:0c:29:b5:89:7d, receiver ip/mac  
172.27.0.1/00:00:00:00:00:00 -- Exhibit -
```

Click the Exhibit button.

You have been asked to troubleshoot a problem where a user is not able to send traffic through your switch. While troubleshooting, you see the log messages shown in the exhibit.

What is causing the problem?

- A. The eswd process has been corrupted.
- B. The receiver MAC in the packets is undefined.
- C. The defined MAC address is misconfigured.
- D. The static IP address is misconfigured.

Correct Answer: C

QUESTION 3

```
-- 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 4

-- 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.

Correct Answer: D

QUESTION 5

You are implementing Q-in-Q tunneling on an EX Series switch. You want the tunnel to support all C-VLANs; however, only some VLANs are able to send traffic across the tunnel. Switch-1 has the following configuration:

```
[edit vlans]
```

```
user@Switch-1# show
```

```
v100 {
```

```
vlan-id 100;
```



```
interface {  
ge-0/0/0.10;  
ge-0/0/1.20;  
}  
dot1q-tunneling {  
customer-vlans [ ];  
}  
}
```

What would solve this problem?

- A. Add family ethernet-switching to the tunnel-side interface on Switch-1.
- B. Implement RSTP.
- C. Q-in-Q tunneling will not work in this scenario; use a Layer 2 VPN instead.
- D. Remove the customer-vlans statement.

Correct Answer: C

[JN0-694 PDF Dumps](#)

[JN0-694 Practice Test](#)

[JN0-694 Study Guide](#)