

Juniper

Exam JN0-694

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

Version: 6.1

[Total Questions: 52]

Question No : 1

-- Exhibit --

```
user@router# run show log bgp-test
```

...

```
Jun 10 23:50:43.056697 BGP SEND 192.168.133.1+179 -> 192.168.133.0+64925
```

```
Jun 10 23:50:43.056739 BGP SEND message type 3 (Notification) length 23
```

```
Jun 10 23:50:43.056760 BGP SEND Notification code 2 (Open Message Error) subcode 7  
(unsupported capability)
```

```
Jun 10 23:50:43.056781 BGP SEND Data (2 bytes): 00 04
```

```
Jun 10 23:50:52.215104 advertising receiving-speaker only capability to neighbor  
::192.168.133.0 (External AS 300)
```

```
Jun 10 23:50:52.215173 bgp_send. sending 59 bytes to ::192.168.133.0 (External AS 300)
```

```
Jun 10 23:50:52.215200
```

```
Jun 10 23:50:52.215200 BGP SEND ::192.168.133.1+179 -> ::192.168.133.0+57107
```

```
Jun 10 23:50:52.215233 BGP SEND message type 1 (Open) length 59
```

```
Jun 10 23:50:52.215256 BGP SEND version 4 as 23456 holdtime 90 id 10.200.1.1 parmlen  
30
```

```
Jun 10 23:50:52.215276 BGP SEND MP capability AFI=2, SAFI=1
```

```
Jun 10 23:50:52.215294 BGP SEND Refresh capability, code=128
```

```
Jun 10 23:50:52.215312 BGP SEND Refresh capability, code=2
```

```
Jun 10 23:50:52.215332 BGP SEND Restart capability, code=64, time=120, flags=
```

```
Jun 10 23:50:52.215353 BGP SEND 4 Byte AS-Path capability (65), as_num 2123456789
```

```
Jun 10 23:50:52.216018
```

```
Jun 10 23:50:52.216018 BGP RECV ::192.168.133.0+57107 -> ::192.168.133.1+179
```

```
Jun 10 23:50:52.216058 BGP RECV message type 3 (Notification) length 21
```

```
Jun 10 23:50:52.216079 BGP RECV Notification code 2 (Open Message Error) subcode 2  
(bad peer AS number)
```

```
Jun 10 23:51:15.058112 advertising receiving-speaker only capability to neighbor
```

192.168.133.0 (External AS 300)

Jun 10 23:51:15.058192 bgp_senD. sending 59 bytes to 192.168.133.0 (External AS 300)

Jun 10 23:51:15.058217

Jun 10 23:51:15.058217 BGP SEND 192.168.133.1+50083 -> 192.168.133.0+179

Jun 10 23:51:15.058250 BGP SEND message type 1 (Open) length 59

Jun 10 23:51:15.058273 BGP SEND version 4 as 65001 holdtime 90 id 10.200.1.1 parmlen 30

Jun 10 23:51:15.058294 BGP SEND MP capability AFI=1, SAFI=128

Jun 10 23:51:15.058312 BGP SEND Refresh capability, code=128

Jun 10 23:51:15.058331 BGP SEND Refresh capability, code=2

Jun 10 23:51:15.058386 BGP SEND Restart capability, code=64, time=120, flags=

Jun 10 23:51:15.058416 BGP SEND 4 Byte AS-Path capability (65), as_num 65001

Jun 10 23:51:15.058651 bgp_pp_rcv:3140: NOTIFICATION sent to 192.168.133.0 (External AS 300): code 6 (Cease) subcode 7 (Connection collision resolution), Reason: dropping 192.168.133.0 (External AS 300), connection collision prefers 192.168.133.0+53170 (proto)

Jun 10 23:51:15.058680 bgp_senD. sending 21 bytes to 192.168.133.0 (External AS 300)

Jun 10 23:51:15.058702

Jun 10 23:51:15.058702 BGP SEND 192.168.133.1+50083 -> 192.168.133.0+179

Jun 10 23:51:15.058735 BGP SEND message type 3 (Notification) length 21

Jun 10 23:51:15.058755 BGP SEND Notification code 6 (Cease) subcode 7 (Connection collision resolution)

Jun 10 23:51:15.059557 advertising receiving-speaker only capability to neighbor 192.168.133.0 (External AS 300)

Jun 10 23:51:15.059594 bgp_senD. sending 59 bytes to 192.168.133.0 (External AS 300)

Jun 10 23:51:15.059617

Jun 10 23:51:15.059617 BGP SEND 192.168.133.1+179 -> 192.168.133.0+53170

Jun 10 23:51:15.059649 BGP SEND message type 1 (Open) length 59

Jun 10 23:51:15.059671 BGP SEND version 4 as 65001 holdtime 90 id 10.200.1.1 parmlen

30

Jun 10 23:51:15.059691 BGP SEND MP capability AFI=1, SAFI=128

Jun 10 23:51:15.059709 BGP SEND Refresh capability, code=128

Jun 10 23:51:15.059727 BGP SEND Refresh capability, code=2

Jun 10 23:51:15.059747 BGP SEND Restart capability, code=64, time=120, flags=

Jun 10 23:51:15.059768 BGP SEND 4 Byte AS-Path capability (65), as_num 65001

Jun 10 23:51:15.060383 bgp_process_caps: mismatch NLRI with 192.168.133.0 (External AS 300): peer: (1) us: (4)

Jun 10 23:51:15.060445 bgp_process_caps:2578: NOTIFICATION sent to 192.168.133.0 (External AS 300): code 2 (Open Message Error) subcode 7 (unsupported capability) value 4

Jun 10 23:51:15.060470 bgp_send. sending 23 bytes to 192.168.133.0 (External AS 300)

Jun 10 23:51:15.060492

Jun 10 23:51:15.060492 BGP SEND 192.168.133.1+179 -> 192.168.133.0+53170

Jun 10 23:51:15.060556 BGP SEND message type 3 (Notification) length 23

Jun 10 23:51:15.060578 BGP SEND Notification code 2 (Open Message Error) subcode 7 (unsupported capability)

Jun 10 23:51:15.060600 BGP SEND Data (2 bytes): 00 04

-- Exhibit --

Click the Exhibit button.

Referring to the exhibit, what is causing the IPv4 BGP peering to stay in an active state?

- A. The peer AS is incorrect.
- B. The peer does not support 4-byte AS values.
- C. The peer has an NLRI mismatch.
- D. The peer has an incorrect IP address.

Answer: C

-- 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.

Answer: B

Question No : 3

There is a lot of traffic marked with IP precedence values af21 and af31 that ingresses the router. The af31 traffic should be using the expedited forwarding queue, but the traffic is

much lower than expected and there are no drops seen on the egress interface.

```
user@router# show class-of-service
interfaces {
  ge-* {
    scheduler-map map-test;
  }
}
scheduler-maps {
  map-test {
    forwarding-class best-effort scheduler be;
    forwarding-class expedited-forwarding scheduler ef;
    forwarding-class assured-forwarding scheduler af;
    forwarding-class network-control scheduler nc;
  }
}
schedulers {
  be {
    transmit-rate percent 70;
    priority high;
  }
  ef {
    transmit-rate percent 15;
    priority low;
  }
  af {
    transmit-rate percent 10;
    priority strict-high;
  }
  nc {
    transmit-rate percent 5;
    priority high;
  }
}

user@router# show firewall
policer ef {
  if-exceeding {
    bandwidth-limit 8k;
    burst-size-limit 15k;
  }
  then forwarding-class best-effort;
}
policer as {
  if-exceeding {
    bandwidth-limit 5m;
    burst-size-limit 15k;
  }
  then forwarding-class best-effort;
}
policer nc {
  if-exceeding {
    bandwidth-limit 5m;
    burst-size-limit 15k;
  }
}
<<cont next column>>
```



```
    then forwarding-class best-effort;
}
filter MF {
  term 1 {
    from {
      precedence 3;
    }
    then {
      policer ef;
      forwarding-class expedited-forwarding;
    }
  }
  term 2 {
    from {
      precedence 2;
    }
    then {
      policer as;
      forwarding-class assured-forwarding;
    }
  }
  term 3 {
    from {
      precedence 6;
    }
    then {
      policer nc;
      forwarding-class network-control;
    }
  }
  term 4 {
    then {
      forwarding-class best-effort;
      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
```

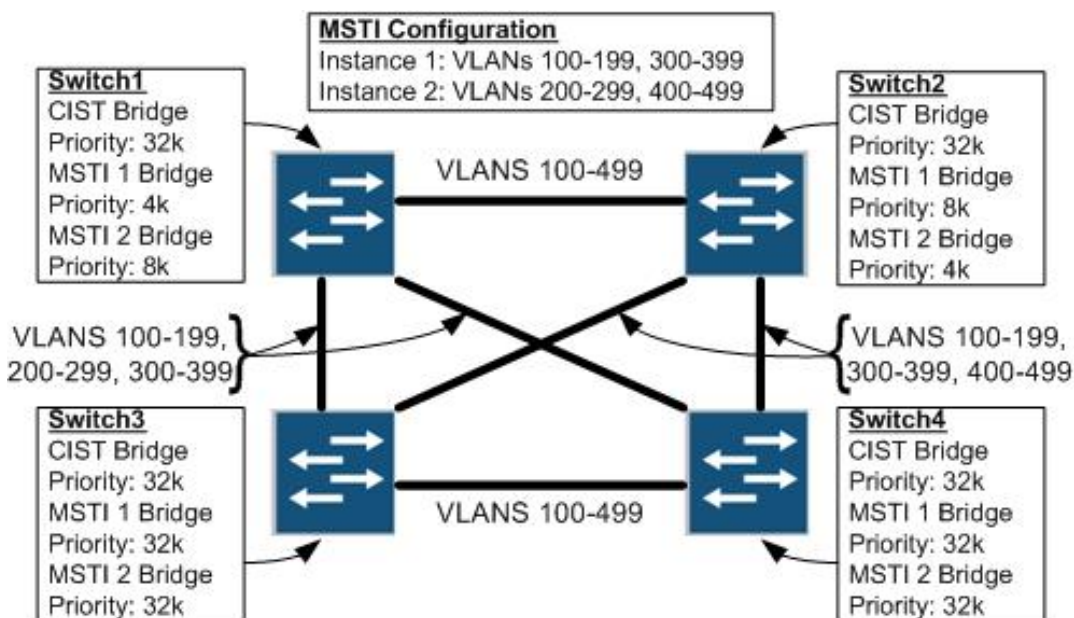
Referring to the exhibit, what is causing the problem?

- A. The assured forwarding queue has a strict high priority and is starving the expedited forwarding queue.
- B. The expedited forwarding queue has a low priority value; therefore the traffic is not serviced.
- C. The MF classifier is forwarding most of the af3l traffic to the best-effort queue.
- D. The MF classifier is does not match on af3l and therefore the traffic is being dropped.

Answer: C

Question No : 4

-- Exhibit --



-- Exhibit --

Click the Exhibit button.

The exhibit shows a small switched network, some details about the MSTP configuration in the network, and the VLANs that are trunked over each link. When Switch2 reboots, users in VLAN 400 on Switch3 report that they lose connectivity to resources in VLAN 400 on Switch4.

What is the cause of this problem?