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.pass4itsure.com/jn0-694.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

Referring to the exhibit, the prefix 3.3.3.3/32 is not in R1 \\'s routing table. Which two configuration changes on R2 would resolve the problem? (Choose two.)

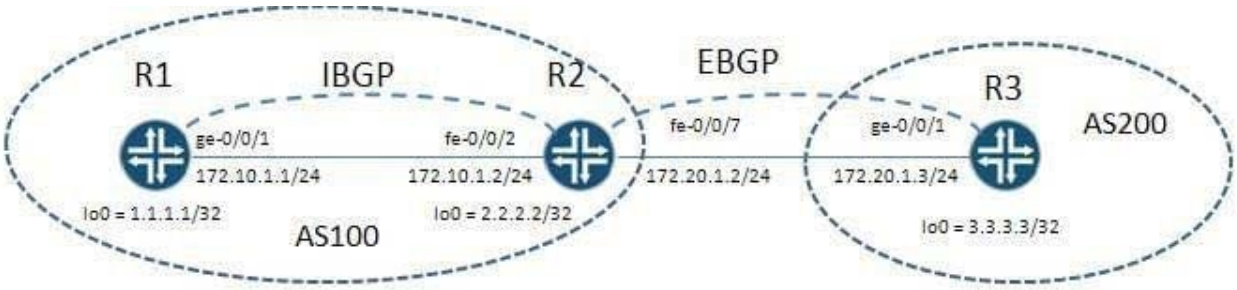


```
user@R1> show route receive-protocol bgp 2.2.2.2 all extensive

inet.0: 6 destinations, 6 routes (5 active, 0 holddown, 1 hidden)
  3.3.3.3/32 (1 entry, 0 announced)
     Accepted
     Nexthop: 172.20.1.3
     Localpref: 100
     AS path: 200 I

user@R1> show route 3.3.3.3 all

inet.0: 6 destinations, 6 routes (5 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

3.3.3.3/32          [BGP/170] 00:02:32, localpref 100, from 2.2.2.2
                      AS path: 200 I
                      Unusable
user@R1> show route

inet.0: 6 destinations, 6 routes (5 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[Direct/0] 00:47:57
                     > via lo0.0
2.2.2.2/32          *[OSPF/10] 00:04:18, metric 1
                     > to 172.10.1.2 via ge-0/0/1.0
172.10.1.0/24       *[Direct/0] 00:47:57
                     > via ge-0/0/1.0
172.10.1.1/32       *[Local/0] 00:47:57
                       Local via ge-0/0/1.0
224.0.0.5/32        *[OSPF/10] 00:47:58, metric 1
                       MultiRecv
```

A. Option A

B. Option B

C. Option C

D. Option D

Correct Answer: BC

QUESTION 2

You observe that a router is using an unusually high amount of CPU cycles. You determine that continuous SPF calculations in OSPF are being performed.

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

A. The wrong authentication keys between the OSPF neighbors are used.

B. The interface MTU is mismatched between the OSPF neighbors.

C. There are duplicate router IDs within the OSPF area.

D. An OSPF adjacency is flapping.

Correct Answer: CD

QUESTION 3

You are asked to troubleshoot the new IBGP peering session shown in the exhibit between R1 and R2. Which action will resolve the problem?

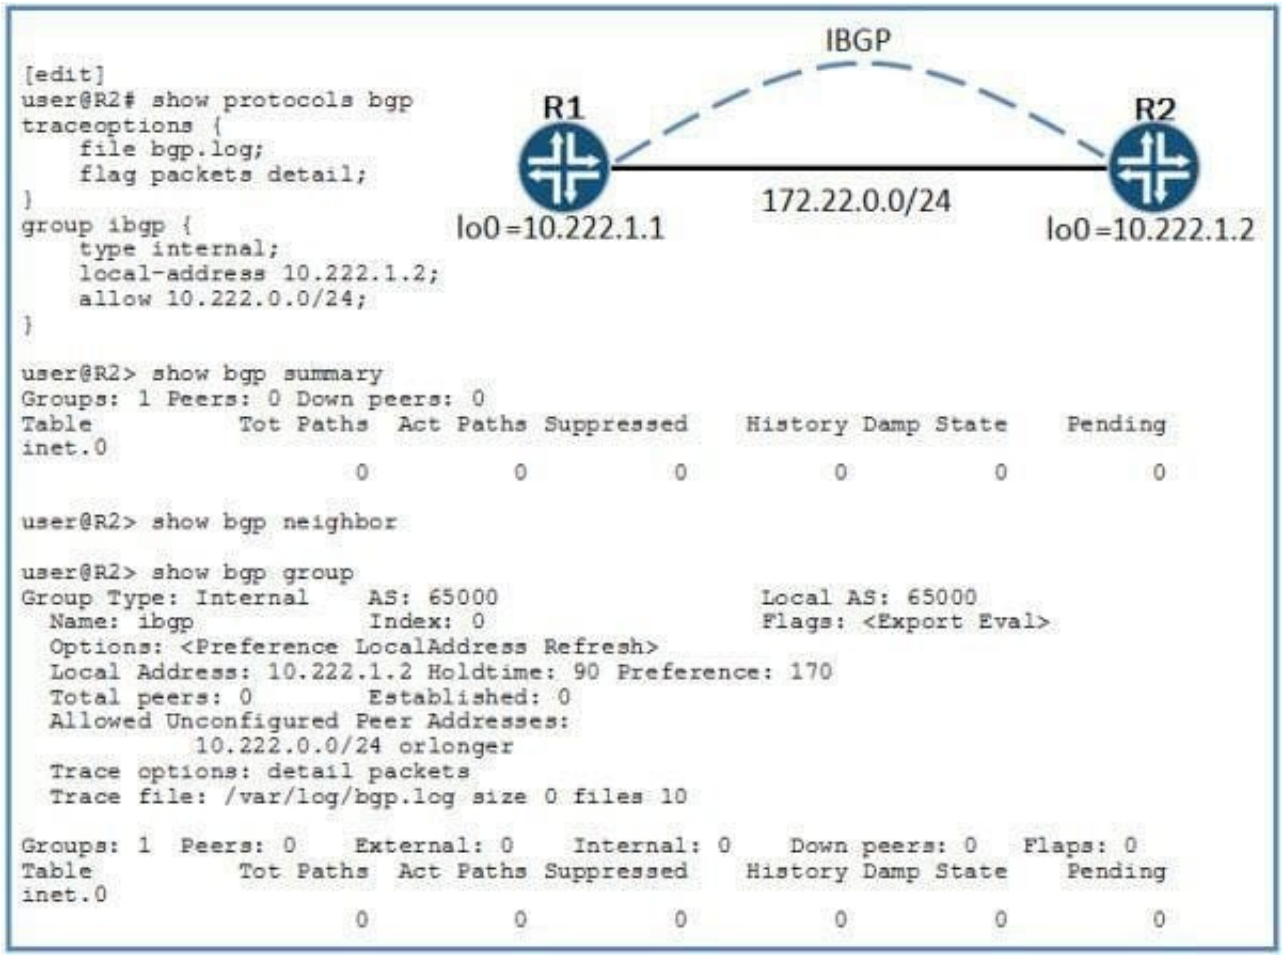

A. Configure the multihop option.

B. Configure the accept-remote-nexthop option.

C. Change the allowed peer range to 10.222.1.0/24.

D. Change the allowed peer range to 172.22.0.0/24.

Correct Answer: C

QUESTION 4

There is a lot of traffic marked with IP precedence values af2l and af3l that ingresses the router. The af3l 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.

Correct Answer: C

QUESTION 5

-- 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.

Correct Answer: C