



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

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

```
[edit]
user@R2# show protocols bgp
traceoptions {
  file bgp.log;
  flag packets detail;
}
group ibgp {
  type internal;
  local-address 10.222.1.2;
  allow 10.222.0.0/24;
}

user@R2> show bgp summary
Groups: 1 Peers: 0 Down peers: 0
Table          Tot Paths  Act Paths  Suppressed    History  Damp State    Pending
inet.0
              0          0          0             0         0         0         0

user@R2> show bgp neighbor

user@R2> show bgp group
Group Type: Internal      AS: 65000          Local AS: 65000
Name: ibgp              Index: 0           Flags: <Export Eval>
Options: <Preference LocalAddress Refresh>
Local Address: 10.222.1.2 Holdtime: 90 Preference: 170
Total peers: 0           Established: 0
Allowed Unconfigured Peer Addresses:
10.222.0.0/24 orlonger
Trace options: detail packets
Trace file: /var/log/bgp.log size 0 files 10

Groups: 1 Peers: 0 External: 0 Internal: 0 Down peers: 0 Flaps: 0
Table          Tot Paths  Act Paths  Suppressed    History  Damp State    Pending
inet.0
              0          0          0             0         0         0         0
```

- 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 2

The exhibit shows part of the configuration for a router. You receive a complaint that the router is not correctly reclassifying all traffic to the best-effort forwarding class when the amount of IPv4 traffic exceeds 10 Mbps.



```
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        filter {
          input filter1;
        }
        policer {
          input policer1;
        }
        address 10.210.33.131/26;
      }
    }
  }
}
class-of-service {
  classifiers {
    inet-precedence ip_classifier_1 {
      forwarding-class best-effort {
        loss-priority low code-points [ 000 010 011 100 ];
      }
      forwarding-class assured-forwarding {
        loss-priority low code-points 001;
      }
      forwarding-class expedited-forwarding {
        loss-priority low code-points 101;
      }
      forwarding-class network-control {
        loss-priority low code-points 110;
        loss-priority high code-points 111;
      }
    }
  }
  interfaces {
    ge-0/0/0 {
      unit 0 {
        classifiers {
          inet-precedence ip_classifier_1;
        }
      }
    }
  }
}
firewall {
  policer policer1 {
    if-exceeding {
      bandwidth-limit 10m;
      burst-size-limit 2k;
    }
    then forwarding-class best-effort;
  }
  filter filter1 {
    term 1 {
      from {
        precedence b101;
      }
      then {
        count term1;
        forwarding-class expedited-forwarding;
      }
    }
    term 2 {
      from {
        forwarding-class-except best-effort;
      }
      then {
        policer policer1;
        count term2;
      }
    }
    term 3 {
      from {
        forwarding-class best-effort;
      }
      then count term3;
    }
  }
}
```

You have isolated the problem to traffic with the IP precedence bits set to the binary value 101. Which configuration is causing this behavior?

- A. the filter firewall filter's term 1
- B. the filter firewall filter's term 2
- C. the ip_classifier_1 classifier
- D. the policer1 policer

Correct Answer: A

QUESTION 3

You are monitoring a network that is configured with PIM sparse mode. An end user's PC (PC1) joins a multicast stream. The stream never switches from the rendezvous-point tree (RPT) to the shortest-path tree (SPT).

Which two statements explain this behavior? (Choose two.)

- A. An interface on the SPT is not configured for PIM.
- B. The designated router for PC1's LAN does not have a route to the multicast source.
- C. This is the normal operation of PIM sparse mode.
- D. This is a source-specific multicast (SSM) stream.



Correct Answer: AB

QUESTION 4

```
-- Exhibit -user@switch>show dot1x interface ge-0/0/1 detail ge-0/0/1.0 Role: Authenticator Administrative state: Auto
Supplicant mode: Multiple Number of retries: 3 Quiet period: 60 seconds Transmit period: 30 seconds Mac Radius:
Enabled Mac Radius Restrict: Enabled Reauthentication: Disabled Configured Reauthentication interval: 3600 seconds
Supplicant timeout: 30 seconds Server timeout: 30 seconds Maximum EAPOL requests: 2 Guest VLAN member:
Number of connected supplicants: 0 -- Exhibit -
```

Click the Exhibit button.

You are asked to troubleshoot an access control issue on your EX Series switch. The user connecting through port ge-0/0/1 indicates that their 802.1X supplicant is failing authentication and they are not able to connect to the network.

Referring to the exhibit, what is causing the 802.1X supplicant to fail?

- A. The supplicant mode is defined as multiple.
- B. The MAC restrict option is configured.
- C. The supplicant timeout value is set too low.
- D. The reauthentication interval is set too high.

Correct Answer: B

QUESTION 5

```
-- Exhibit
```



```

[edit]
user@R1# show protocols bgp
group ebgp {
  peer-as 65502;
  multipath;
  neighbor 172.22.0.1;
  neighbor 172.22.0.5;
}

user@R1> show route 100.100.0/16 terse

inet.0: 12 destinations, 17 routes (12 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   B 170    100           172.22.0.1      65502 I
                   >172.22.0.5
* 100.100.1.0/24   B 170    100           172.22.0.1      65502 I
                   >172.22.0.5
* 100.100.2.0/24   B 170    100           172.22.0.1      65502 I
                   >172.22.0.5
* 100.100.3.0/24   B 170    100           172.22.0.1      65502 I
                   >172.22.0.5
* 100.100.4.0/24   B 170    100           172.22.0.1      65502 I
                   >172.22.0.5
                   B 170    100           >172.22.0.1      65502 I

user@R1> show route forwarding-table | find 100.100.0.0/24
100.100.0.0/24   user      0 172.22.0.5      ucst  535    8 ge-0/0/10.0
100.100.1.0/24   user      0 172.22.0.5      ucst  535    8 ge-0/0/10.0
100.100.2.0/24   user      0 172.22.0.1      ucst  513    5 ge-0/0/1.0
100.100.3.0/24   user      0 172.22.0.5      ucst  535    8 ge-0/0/10.0
100.100.4.0/24   user      0 172.22.0.5      ucst  535    8 ge-0/0/10.0
...

```

-- Exhibit -Click the Exhibit button. Your network has two connections to your ISP. You have been asked to load-balance traffic across both

links that connect to your ISP. You have enabled multipath for this peer, but you are still not getting the expected load balancing.

Given the information shown in the exhibit, what else must you do?

- A. Configure and apply a load-balancing policy.
- B. Change the multipath parameter to multihop.
- C. Create a policy to manually change the next hops.
- D. Enable the keep all parameter.

Correct Answer: A

[Latest JN0-694 Dumps](#)

[JN0-694 VCE Dumps](#)

[JN0-694 Practice Test](#)