

# Juniper JN0-351

**Enterprise Routing and Switching, Specialist**

**Questions And Answers PDF Format:**

**For More Information – Visit link below:**

**<https://www.certsgrade.com/>**


*Version = Product*



# Latest Version: 6.0

## Question: 1

Exhibit.



Which router will become the OSPF BDR if all routers are powered on at the same time?

- A. R4
- B. R1
- C. R3
- D. R2

**Answer: A**

Explanation:

OSPF DR/BDR election is a process that occurs on multi-access data links. It is intended to select two OSPF nodes: one to be acting as the Designated Router (DR), and another to be acting as the Backup Designated Router (BDR). The DR and BDR are responsible for generating network LSAs for the multi-access network and synchronizing the LSDB with other routers on the same network<sup>1</sup>.

The DR/BDR election is based on two criteria: the OSPF priority and the router ID. The OSPF priority is a value between 0 and 255 that can be configured on each interface participating in OSPF. The default priority is 1. A priority of 0 means that the router will not participate in the election and will never become a DR or BDR. The router with the highest priority will become the DR, and the router with the second highest priority will become the BDR. If there is a tie in priority, then the router ID is used as a tiebreaker.

The router ID is a 32-bit number that uniquely identifies each router in an OSPF domain. It can be manually configured or automatically derived from the highest IP address on a loopback interface or any active interface.

In this scenario, all routers have the same priority of 1, so the router ID will determine the outcome of the election. The router IDs are shown in the exhibit as RID values. The highest RID belongs to R4 (10.10.10.4), so R4 will become the DR. The second highest RID belongs to R3 (10.10.10.3), so R3 will become the BDR.

Reference:

1: OSPF DR/BDR Election: Process, Configuration, and Tuning 2: OSPF Designated Router (DR) and Backup Designated Router (BDR)

## Question: 2

Exhibit.

```
{master:0}[edit]
user@switch# run show interfaces terse
Interface           Admin  Link  Proto  Local          Remote
ge-0/0/0            up    up
gr-0/0/0            up    up
pfe-0/0/0           up    up
ge-0/0/1            up    up    up
ge-0/0/1.0          up    up    up    inet    172.23.11.10/24
                                   172.23.12.10/24
ge-0/0/2            up    up
ge-0/0/2.0          up    up    up    inet    172.23.11.100/24
ge-0/0/3            up    up
ge-0/0/3.0          up    up    up    inet    172.23.12.100/24
...
bme0                up    up
bme0.0              up    up    up    inet    128.0.0.1/2
                                   128.0.0.4/2
                                   128.0.0.16/2
                                   128.0.0.63/2
...
jsrv.1              up    up
lo0                 up    up
lo0.16385           up    up
lsi                 up    up
me0                 up    up
me0.0               up    up    up    inet    10.210.20.233/29
mtun                up    up
pimd                up    up
pime                up    up
tap                 up    up
vme                 up    down
```

What is the management IP address of the device shown in the exhibit?

- A. 10.210.20.233
- B. 172.23.12.100
- C. 128.0.0.1
- D. 172.23.11.10

**Answer: B**

Explanation:

The management IP address of a device is the IP address that is used to access the device for configuration and monitoring purposes. It is usually assigned to a dedicated management interface that is separate from the data interfaces. The management interface can be accessed via SSH, Telnet, HTTP, or other protocols.

In the exhibit, the list of interfaces and their statuses shows that the management interface is me0. This interface has an admin status of up, a protocol status of inet, a local address of 172.23.12.100/24, and a remote address of unspecified. This means that the me0 interface is active, has an IPv4 address assigned, and is not connected to another device.

Therefore, the management IP address of the device shown in the exhibit is 172.23.12.100.

Reference:

: [Management Interfaces Overview] : [Displaying Interface Status Information]

## Question: 3

Which three protocols support BFD? (Choose three.)

- A. RSTP
- B. BGP
- C. OSPF
- D. LACP
- F. FTP

**Answer: BCD**

Explanation:

BFD is a protocol that can be used to quickly detect failures in the forwarding path between two adjacent routers or switches. BFD can be integrated with various routing protocols and link aggregation protocols to provide faster convergence and fault recovery.

According to the Juniper Networks documentation, the following protocols support BFD on Junos OS devices<sup>1</sup>:

**BGP:** BFD can be used to monitor the connectivity between BGP peers and trigger a session reset if a failure is detected. BFD can be configured for both internal and external BGP sessions, as well as for IPv4 and IPv6 address families<sup>2</sup>.

**OSPF:** BFD can be used to monitor the connectivity between OSPF neighbors and trigger a state change if a failure is detected. BFD can be configured for both OSPFv2 and OSPFv3 protocols, as well as for point-to-point and broadcast network types<sup>3</sup>.

**LACP:** BFD can be used to monitor the connectivity between LACP members and trigger a link state change if a failure is detected. BFD can be configured for both active and passive LACP modes, as well as for static and dynamic LAGs<sup>4</sup>.

Other protocols that support BFD on Junos OS devices are:

**IS-IS:** BFD can be used to monitor the connectivity between IS-IS neighbors and trigger a state change if a failure is detected. BFD can be configured for both level 1 and level 2 IS-IS adjacencies, as well as for point-to-point and broadcast network types.

**RIP:** BFD can be used to monitor the connectivity between RIP neighbors and trigger a route update if a failure is detected. BFD can be configured for both RIP version 1 and version 2 protocols, as well as for IPv4 and IPv6 address families.

**VRRP:** BFD can be used to monitor the connectivity between VRRP routers and trigger a priority change if a failure is detected. BFD can be configured for both VRRP version 2 and version 3 protocols, as well as for IPv4 and IPv6 address families.

The protocols that do not support BFD on Junos OS devices are:

**RSTP:** RSTP is a spanning tree protocol that provides loop prevention and rapid convergence in layer 2 networks. RSTP does not use BFD to detect link failures, but relies on its own hello mechanism that sends BPDU packets every 2 seconds by default.

**FTP:** FTP is an application layer protocol that is used to transfer files between hosts over a TCP connection. FTP does not use BFD to detect connection failures, but relies on TCP's own retransmission and timeout mechanisms.

Reference:

1: [Configuring Bidirectional Forwarding Detection] 2: [Configuring Bidirectional Forwarding Detection for BGP] 3: [Configuring Bidirectional Forwarding Detection for OSPF] 4: [Configuring Bidirectional Forwarding Detection for Link Aggregation Control Protocol] : [Configuring Bidirectional Forwarding Detection for IS-IS] : [Configuring Bidirectional Forwarding Detection for RIP] : [Configuring Bidirectional Forwarding Detection for VRRP] : [Understanding Rapid Spanning Tree Protocol] : [Understanding FTP]

## Question: 4

Exhibit.

The exhibit displays a Juniper CLI session on a PE-1 device. The user enters the command `show route table ISPI.inet.0`, which returns no output. The user then enters `configure` and configures a routing instance named `ISPI` of type `forwarding`. The configuration includes a static route `0.0.0.0/0` with next-hop `203.0.113.2` and an import policy `ISPI-import`. The user then shows the policy options for `ISPI-import`, which are `from instance master;` and `then accept;`. To the right of the CLI output is a network diagram showing a central device labeled `PE-1` with two interfaces: `ge-0/0/1` connected to a box labeled `inet.0` and `ge-0/0/2` connected to a box labeled `ISPI.inet.0`.

The `isp_i_inet.0` route table has currently no routes in it.

What will happen when you commit the configuration shown on the exhibit?

- A. The `inet.0` route table will be completely overwritten by the `isp_i_inet.0` route table.
- B. The `inet.0` route table will be imported into the `isp_i_inet.0` route table.
- C. The `ISPI.inet.0` route table will be completely overwritten by the `inet.0` route table.
- D. The `ISPI.inet.0` route table will be imported into the `inet.0` route table.

**Answer: B**

Explanation:

The configuration shown in the exhibit is an example of a routing instance of type `virtual-router`. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters that create a separate routing domain on a Juniper device<sup>1</sup>. A `virtual-router` routing instance allows administrators to divide a device into multiple independent virtual routers, each with its own routing table<sup>2</sup>. The configuration also includes a `rib-group` statement, which is used to import routes from one routing table to another. A `rib-group` consists of an `import-rib` statement, which specifies the source routing table, and an `export-rib` statement, which specifies the destination routing table.

In this case, the rib-group name is inet-to-isp, and the import-rib statement specifies inet.0 as the source routing table. The export-rib statement specifies ispi.inet.0 as the destination routing table. This means that the routes from inet.0 will be imported into ispi.inet.0. Therefore, the correct answer is B. The inet.0 route table will be imported into the ispi.inet.0 route table.

Reference:

1: Routing Instances Overview 2: Virtual Routing Instances : [rib-group (Routing Options)]

## Question: 5

Which statement is correct about graceful Routing Engine switchover (GRES)?

- A. The PFE restarts and the kernel and interface information is lost.
- B. GRES has a helper mode and a restarting mode.
- C. When combined with NSR, routing is preserved and the new master RE does not restart rpd.
- D. With no other high availability features enabled, routing is preserved and the new master RE does not restart rpd.

**Answer: C**

Explanation:

The Graceful Routing Engine Switchover (GRES) feature in Junos OS enables a router with redundant Routing Engines to continue forwarding packets, even if one Routing Engine fails<sup>1</sup>. GRES preserves interface and kernel information, ensuring that traffic is not interrupted<sup>1</sup>. However, GRES does not preserve the control plane<sup>1</sup>.

To preserve routing during a switchover, GRES must be combined with either Graceful Restart protocol extensions or Nonstop Active Routing (NSR)<sup>1</sup>. When GRES is combined with NSR, nearly 75 percent of line rate worth of traffic per Packet Forwarding Engine remains uninterrupted during GRES<sup>1</sup>. Any updates to the primary Routing Engine are replicated to the backup Routing Engine as soon as they occur<sup>1</sup>.

Therefore, when GRES is combined with NSR, routing is preserved and the new master RE does not restart rpd<sup>1</sup>.

For More Information – **Visit link below:**  
<https://www.certsgrade.com/>

## PRODUCT FEATURES

-  **100% Money Back Guarantee**
-  **90 Days Free updates**
-  **Special Discounts on Bulk Orders**
-  **Guaranteed Success**
-  **50,000 Satisfied Customers**
-  **100% Secure Shopping**
-  **Privacy Policy**
-  **Refund Policy**

16 USD Discount Coupon Code: **NB4XKTMZ**



Visit us at: <https://www.certsgrade.com/pdf/jn0-351>