

Juniper

JN0-664 Exam

Service Provider Routing and Switching, Professional



Thank you for Downloading JN0-664 exam PDF Demo

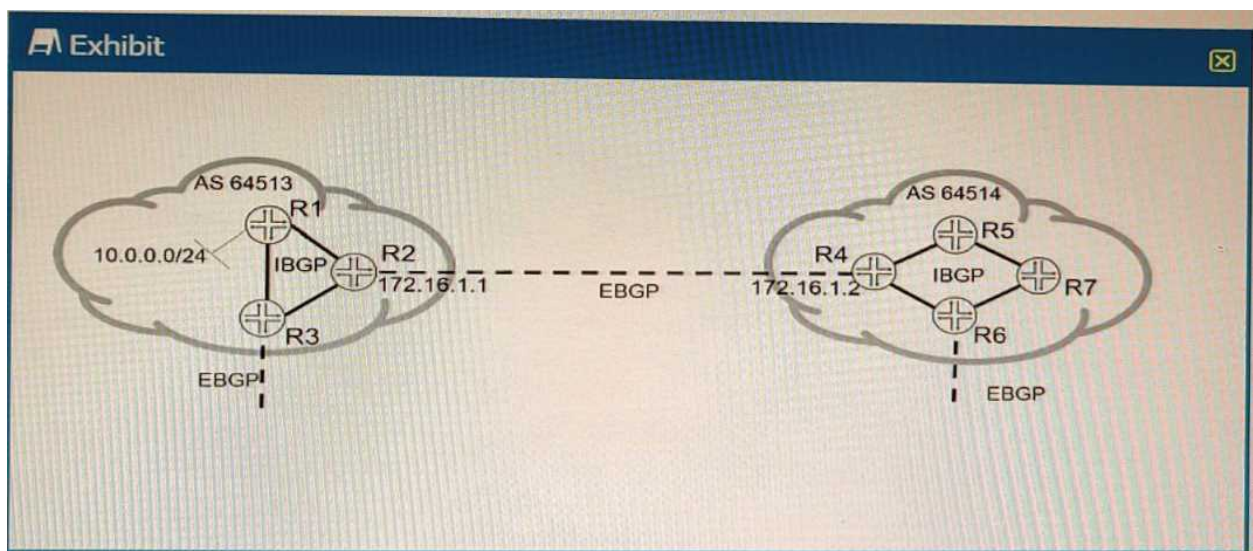
You can Buy Latest JN0-664 Full Version Download

<https://www.certkillers.net/Exam/JN0-664>

Version: 6.0

Question: 1

Exhibit.



Referring to the exhibit; the 10.0.0.0/24 EBGP route is received on R5; however, the route is being hidden.

What are two solutions that will solve this problem? (Choose two.)

- A. On R4, create a policy to change the BGP next hop to itself and apply it to IBGP as an export policy
- B. Add the external interface prefix to the IGP routing tables
- C. Add the internal interface prefix to the BGP routing tables.
- D. On R4, create a policy to change the BGP next hop to 172.16.1.1 and apply it to IBGP as an export policy

Answer: AB

Explanation:

the default behavior for iBGP is to propagate EBGP-learned prefixes without changing the next-hop. This can cause issues if the next-hop is not reachable via the IGP. One solution is to use the next-hop self command on R4, which will change the next-hop attribute to its own loopback address. This way, R5 can reach the next-hop via the IGP and install the route in its routing table.

Another solution is to add the external interface prefix (120.0.4.16/30) to the IGP routing tables of R4 and R5. This will also make the next-hop reachable via the IGP and allow R5 to use the route. According to [2](#), this is a possible workaround for a pure IP network, but it may not work well for an MPLS network.

The reason why the route is being hidden is that R5 cannot reach the BGP next hop 10.0.0.1, which is the address of R1. R5 does not have a route to 10.0.0.0/24 in its routing table, and neither does R4. Therefore, R5 cannot resolve the BGP next hop and marks the route as hidden.

There are two solutions that will solve this problem:

Option A: On R4, create a policy to change the BGP next hop to itself and apply it to iBGP as an export policy. This way, R5 will receive the route with a next hop of 172.16.1.2, which is reachable via the IGP. [This solution is also known as next-hop-self1.](#)

Option B: Add the external interface prefix to the IGP routing tables. This way, R4 and R5 will learn a route to 10.0.0.0/24 via the IGP and be able to resolve the BGP next hop. [This solution is also known as recursive lookup2.](#)

Option C is not correct because adding the internal interface prefix to the BGP routing tables will not help R5 reach the BGP next hop 10.0.0.1.

Option D is not correct because changing the BGP next hop to 172.16.1.1 on R4 will not help R5 either, since R5 does not have a route to 172.16.1.1 in its routing table.

[Reference: 1: Configuring Next-Hop-Self for iBGP Peers 2: Understanding Recursive Lookup](#)

Question: 2

You are responding to an RFP for a new MPLS VPN implementation. The solution must use LDP for signaling and support Layer 2 connectivity without using BGP. The solution must be scalable and support multiple VPN connections over a single MPLS LSP. The customer wants to maintain all routing for their Private network.

In this scenario, which solution do you propose?

- A. circuit cross-connect
- B. BGP Layer 2 VPN
- C. LDP Layer 2 circuit
- D. translational cross-connect


Answer: C

Explanation:

AToM (Any Transport over MPLS) is a framework that supports various Layer 2 transport types over an MPLS network core. One of the transport types supported by AToM is LDP Layer 2 circuit, which is a point-to-point Layer 2 connection that uses LDP for signaling and MPLS for forwarding. LDP Layer 2 circuit can support Layer 2 connectivity without using BGP and can be scalable and efficient by using a single MPLS LSP for multiple VPN connections. The customer can maintain all routing for their private network by using their own CE switches.

Question: 3

Exhibit.

 Exhibit

```
user@R1# show interfaces
ge-1/2/3 {
  unit 0 {
    description to-R2;
    family inet {
      address 10.1.1.1/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.16.1/32;
    }
    family iso {
      address 49.0001.1921.6801.6001.00;
    }
  }
}
user@R1# show protocols
isis {
  interface ge-1/2/3.0 {
    level 2 disable;
  }
}
...
user@R2# show interfaces
ge-1/2/3 {
  unit 0 {
    description to-R1;
    family inet {
      address 10.1.1.2/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.16.2/32;
    }
    family iso {
      address 49.0001.1921.6801.6002.00;
    }
  }
}
```

Referring to the exhibit, what must be changed to establish a Level 1 adjacency between routers R1 and R2?

- A. Change the level 1 disable parameter under the R1 protocols isis interface lo0.0 hierarchy to the level 2 disable parameter.
- B. Remove the level 1 disable parameter under the R2 protocols isis interface lo0.0 configuration hierarchy.
- C. Change the level 1 disable parameter under the R2 protocols isis interface ge-1/2/3.0 hierarchy to the level 2 disable parameter
- D. Add IP addresses to the interface ge-1/2/3 unit 0 family iso hierarchy on both R1 and R2.

Answer: B

Explanation:

IS-IS routers can form Level 1 or Level 2 adjacencies depending on their configuration and network topology. Level 1 routers are intra-area routers that share the same area address with their neighbors. Level 2 routers are inter-area routers that can connect different areas. Level 1-2 routers are both intra-area and inter-area routers that can form adjacencies with any other router.

In the exhibit, R1 and R2 are in different areas (49.0001 and 49.0002), so they cannot form a Level 1 adjacency. However, they can form a Level 2 adjacency if they are both configured as Level 1-2 routers. R1 is already configured as a Level 1-2 router, but R2 is configured as a Level 1 router only, because of the level 1 disable command under the lo0.0 interface. This command disables Level 2 routing on the loopback interface, which is used as the router ID for IS-IS.

Therefore, to establish a Level 1 adjacency between R1 and R2, the level 1 disable command under the R2 protocols isis interface lo0.0 hierarchy must be removed. This will enable Level 2 routing on R2 and allow it to form a Level 2 adjacency with R1.

Question: 4

You are asked to protect your company's customers from amplification attacks. In this scenario, what is Juniper's recommended protection method?

- A. ASN prepending
- B. BGP FlowSpec
- C. destination-based Remote Triggered Black Hole
- D. unicast Reverse Path Forwarding

Answer: C

Explanation:

amplification attacks are a type of distributed denial-of-service (DDoS) attack that exploit the characteristics of certain protocols to amplify the traffic sent to a victim. For example, an attacker can send a small DNS query with a spoofed source IP address to a DNS server, which will reply with a much larger response to the victim. This way, the attacker can generate a large amount of traffic with minimal resources.

One of the methods to protect against amplification attacks is destination-based Remote Triggered Black Hole (RTBH) filtering. This technique allows a network operator to drop traffic destined to a specific IP address or prefix at the edge of the network, thus preventing it from reaching the victim and consuming bandwidth and resources. RTBH filtering can be implemented using BGP to propagate a special route with a next hop of 192.0.2.1 (a reserved address) to the edge routers. Any traffic matching this route will be discarded by the edge routers.

Question: 5

Exhibit


```

user@router> show l2vpn connections
Layer-2 VPN connections:
Legend for connection status (St)
EI -- encapsulation invalid          NC -- interface encapsulation not
CCC/TCC/VPLS                         CCC/TCC/VPLS
EM -- encapsulation mismatch         WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down        NP -- interface hardware not present
CM -- control-word mismatch          -> -- only outbound connection is up
CN -- circuit not provisioned        <- -- only inbound connection is up
OR -- out of range                   Up -- operational
OL -- no outgoing label              Dn -- down
LD -- local site signaled down       CF -- call admission control failure
RD -- remote site signaled down      SC -- local and remote site ID collision
LN -- local site not designated      LM -- local site ID not minimum designated
RN -- remote site not designated     RM -- remote site ID not minimum designated
XX -- unknown connection status     IL -- no incoming label
MM -- MTU mismatch                  MI -- Mesh-Group ID not available
BK -- Backup connection              ST -- Standby connection
FF -- Profile parse failure          PB -- Profile busy
RS -- remote site standby            SN -- Static Neighbor
LB -- Local site not best-site       RB -- Remote site not best-site
VM -- VLAN ID mismatch              HS -- Hot-standby Connection
Legend for interface status
Up -- operational
Dn -- down

Instance: vpn-A
Edge protection: Not-Primary
Local site: CE1-2 (2)
connection-site Type St      Time last up      # Up trans
1               rmt  Up      Apr 11 14:35:27 2020      1
Remote PE: 172.17.20.1, Negotiated control-word: Yes (Null)
Incoming label: 21, Outgoing label: 22
Local interface: ge-0/0/6.610, Status: Up, Encapsulation: VLAN
Flow Label Transmit: No, Flow Label Receive: No

```

Which two statements about the output shown in the exhibit are correct? (Choose two.)

- A. The PE is attached to a single local site.
- B. The connection has not flapped since it was initiated.
- C. There has been a VLAN ID mismatch.
- D. The PE router has the capability to pop flow labels

Answer: AD

Explanation:

[According to 1 and 2](#), BGP Layer 2 VPNs use BGP to distribute endpoint provisioning information and

set up pseudowires between PE devices. BGP uses the Layer 2 VPN (L2VPN) Routing Information Base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path.

In the output shown in the exhibit, we can see some information about the L2VPN RIB and the pseudowire state. Based on this information, we can infer the following statements:

The PE is attached to a single local site. This is correct because the output shows only one local site ID (1) under the L2VPN RIB section. A local site ID is a unique identifier for a site within a VPLS domain. If there were multiple local sites attached to the PE, we would see multiple local site IDs with different prefixes.

The connection has not flapped since it was initiated. This is correct because the output shows that the uptime of the pseudowire is equal to its total uptime (1w6d). This means that the pseudowire has been up for one week and six days without any interruption or flap.

There has been a VLAN ID mismatch. This is not correct because the output shows that the remote and local VLAN IDs are both 0 under the pseudowire state section. A VLAN ID mismatch occurs when the remote and local VLAN IDs are different, which can cause traffic loss or misdelivery. If there was a VLAN ID mismatch, we would see different values for the remote and local VLAN IDs.

The PE router has the capability to pop flow labels. This is correct because the output shows that the flow label pop bit is set under the pseudowire state section. The flow label pop bit indicates that the PE router can pop (remove) the MPLS flow label from the packet before forwarding it to the CE device. The flow label is an optional MPLS label that can be used for load balancing or traffic engineering purposes.

Thank You for trying JN0-664 PDF Demo

To Buy New JN0-664 Full Version Download visit link below

<https://www.certkillers.net/Exam/JN0-664>

Start Your JN0-664 Preparation

Use Coupon “**CKNET**” for Further discount on the purchase of Full Version Download. Test your JN0-664 preparation with actual exam questions.