

JN0-663.VCEplus.premium.exam.65q

Number: JN0-663  
Passing Score: 800  
Time Limit: 120 min  
File Version: 1.0



Website: <https://vceplus.com>

VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>

Facebook: <https://www.facebook.com/VCE.For.All.VN/>

Twitter : [https://twitter.com/VCE\\_Plus](https://twitter.com/VCE_Plus)

JN0-663

Service Provider Routing and Switching, Professional (JNCIP-SP)





## Exam A

### QUESTION 1

```
user@router> show ospf route 30.0.0.0/24
Topology default Route Table:

Prefix          Path Route    NH    Metric NextHop    Nexthop
                Type Type      Type                Interface Address/LSP
30.0.0.0/24     Ext2 Network IP           0 ge-0/0/1.0 5.0.0.1

user@router> show route protocol ospf 30.0.0.0/24

inet.0: 21 destinations, 23 routes, (21 active, 0 holddown, 0 hidden)
```

You notice an inconsistency between the routing table and the OSPF database, as shown in the exhibit.

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

- A. The LSA is a Type 4 LSA.
- B. An OSPF export policy is being applied to the route.
- C. An OSPF import policy is being applied to the route.
- D. The LSA is a Type 5 LSA.

**Correct Answer:** BC

**Section:** (none)

**Explanation**

**Explanation/Reference:**

### QUESTION 2





```
(65001)R1-----R2-----R3(65001)

[edit]
user@R2# run show route 11.11.11.0/24

inet.0 : 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

11.11.11.0/24      *[BGP/170] 00:04:55, localpref 100
                  AS path: 65001 I, validation-state: unverified
                  > to 172.16.1.1 via ge-0/0/0.0
                  [BGP/170] 00:10:33, localpref 100
                  AS path: 65001 65001 I, validation-state: unverified

[edit]
user@R2# show protocols bgp
group R1 {
    neighbor 172.16.1.1 {
        peer-as 65001;
    }
}
group R3 {
    neighbor 172.16.2.1 {
        peer-as 65001;
    }
}
local-as 65002;

[edit]
user@R2# show policy-options
policy-statement lb {
    then {
        load-balance per-packet;
    }
}
policy-statement prepend {
    term 1 {
        then as-path-prepend 65001;
    }
}

[edit]
user@R2# show routing-options
forwarding-table {
    export lb;
}
```



R2 is receiving the same route from R1 and R3. You must ensure that you can load balance traffic for that route.

Referring to the exhibit, which two configuration changes will allow load balancing? (Choose two.)

- A. Apply the `prepend` policy as an import policy under group R1.
- B. Configure `multipath` under the global BGP configuration.
- C. Configure `multipath` under group R1.
- D. Apply the `prepend` policy as an import policy under group R3.

**Correct Answer:** AB


**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 3**





```
[edit routing-instances CE-A1]
user@PE-1# show
instance-type vrf;
interface ge-0/0/9.0;
route-distinguisher 10.222.222.3:2;
vrf-target target:65511:101;
protocols {
  bgp {
    group CE-A1 {
      type external;
      peer-as 65111;
      neighbor 192.168.0.2;
    }
  }
}

[edit routing-instances CE-A2]
user@PE-2# show
instance-type vrf;
interface ge-0/0/9.0;
route-distinguisher 10.222.222.3:2;
vrf-target target:65511:101;
protocols {
  bgp {
    group CE-A2 {
      type external;
      peer-as 65111;
      neighbor 192.168.6.2;
    }
  }
}
```

Referring to the exhibit, hosts in Site 1 and Site 2 are unable to communicate with each other through the Layer 3 VPN.

What is the problem?

- A. The two sites are in the same AS.
- B. The two sites are using the same instance type.
- C. The two sites are using the same route target.
- D. The two sites are using the same route distinguisher.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 4**



```
[edit class-of-service]
user@router# show
classifiers {
    dscp classifierX {
        forwarding class low-priority {
            loss-priority low code-points 000000;
            loss-priority high code points 000001;
        }
        forwarding class medium-priority {
            loss-priority low code-points 000010;
            loss-priority high code points 000011;
        }
        forwarding class high-priority {
            loss-priority low code-points 000100;
            loss-priority high code points 000101;
        }
    }
}

forwarding-classes {
    class low-priority queue-num 0;
    class medium-priority queue-num 1;
    class high-priority queue-num 2;
    class network_control queue-num 3;
}
```

You manage an MX Series device which includes the configuration shown in the exhibit. Traffic marked with DSCP 000011 is entering the ge-1/0/4 interface at 102 Mbps. The traffic exits the device on the ge-1/0/5 interface. No other traffic is transiting the router.

In this scenario, what happens to traffic exceeding 100 Mbps?

- A. Traffic exceeding 100 Mbps is redirected to a rate limiter.
- B. Traffic exceeding 100 Mbps is buffered.
- C. Traffic exceeding 100 Mbps is dropped.
- D. Traffic exceeding 100 Mbps is forwarded.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 5** What is the purpose of the cluster-list attribute within a BGP route reflector group?

- A. to disable internal cluster re-advertisements
- B. to facilitate loop detection within the route reflector network
- C. to define the router that first advertised the route to the route reflector
- D. to override the router ID value within the cluster

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**



# QUESTION 6

```
user@host# show protocols ospf
area 0.0.0.6 {
  nssa {
    default-lsa {
      default-metric 10;
      metric-type 1;
      type-7;
    }
  }
  no-summaries;
  area-range 192.168.16.0/20;
}
```

Referring to the ABR configuration shown in the exhibit, which three statements are correct? (Choose three.)

- A. The ABR advertises a default route to the NSSA using a Type 7 LSA.
- B. The ABR advertises a single Type 3 summary LSA to the backbone area for all Type 1 and Type 2 LSAs in the 192.168.16.0/20 range.
- C. The ABR advertises a Type 5 external LSA to the backbone area for each Type 7 LSA in the NSSA.
- D. The ABR does not summarize any routes within the 192.168.16.0/20 range.
- E. The ABR advertises a single Type 5 external LSA to the backbone area for all Type 7 LSAs in the NSSA.

**Correct Answer:** ABC

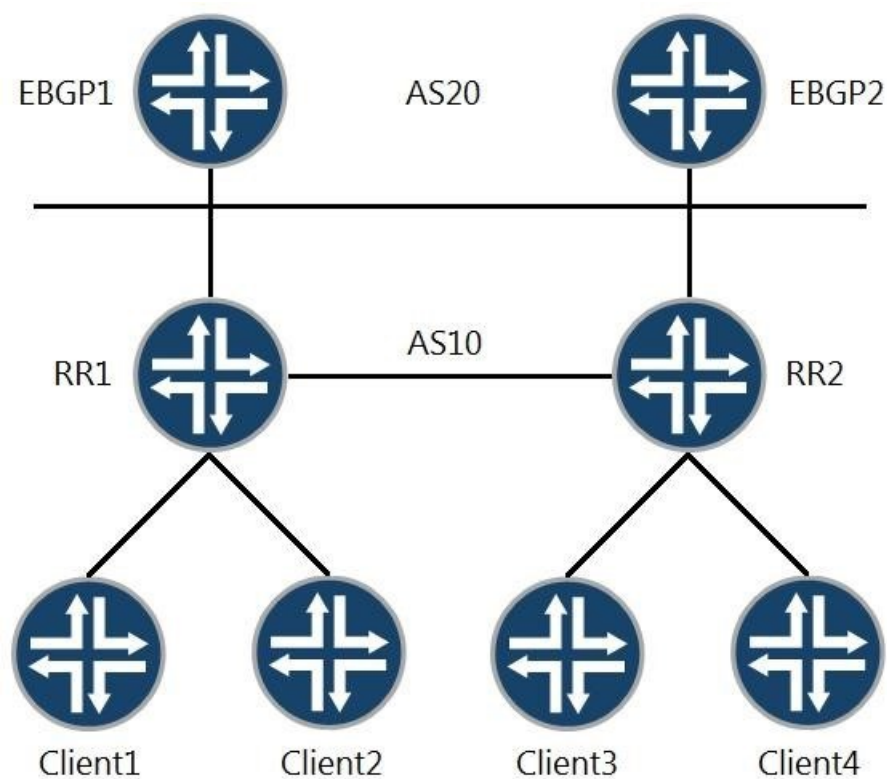
**Section:** (none)

**Explanation**

**Explanation/Reference:**



# QUESTION 7



Referring to the exhibit, which two statements about route reflectors are correct? (Choose two.)



- A. RR1 advertises routes learned from Client1 to RR2 with itself as the next hop.
- B. RR2 advertises routes learned from Client3 to EBG2 with itself as the next hop.
- C. RR1 and RR2 need the same cluster ID to exchange routes learned from their clients.
- D. RR2 adds its cluster ID when advertising routes from Client4 to Client3.


**Correct Answer:** BD

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 8



You are provisioning Layer 2 circuits between sites CE1, CE2, and CE3.

Referring to the exhibit, which statement is true?

- A. A point-to-multipoint LSP must be created between sites.
- B. Each site must have only one VLAN configured to the PE.
- C. Site PE1 must have a point-to-multipoint link configured towards the core.
- D. Two VLANs must be configured from PE 1 to CE 1.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 9

What information is stored in a VRF table for a BGP Layer 2 VPN? (Choose two.)

- A. Layer 2 encapsulation
- B. local site ID
- C. remote interface of local CE device



- D. logical interlace provisioned to local CE device
- E. label-switched path

**Correct Answer:** BD

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 10** Which two statements about wide and narrow metrics used in IS-IS are correct? (Choose two.)

- A. Wide metrics are enabled with the `wide-metrics-only` parameter under `protocols isis hierarchy`.
- B. Narrow metrics are enabled by default and use 8 bits in TLVs to send information.
- C. Wide metrics are sent by default and use 24 bits in TLVs to send information.
- D. Disabling narrow metrics results in external routes being leaked from L1 to L2 areas automatically.

**Correct Answer:** AB

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 11**



```
[edit]
user@R4# run show pim rps
Instance: PIM.master

address-family INET
RR address      Type      Mode      Holdtime Timeout Groups Group prefixes
22.22.22.22     bootstrap sparse    150      108      0      224.0.0.0/4
33.33.33.33     bootstrap sparse    150      108      2      224.1.0.0/16

[edit]
user@R4# run show route 22.22.22.22

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

22.22.22.22/32    *[IS-IS/18] 00:32:27, metric 10
                  > to 10.1.1.2 via ge-0/0/0.0

inet.2: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0        *[Static/5] 00:13:55
                  > to 10.1.1.6 via ge-0/0/1.0

[edit]
user@R4# run show route 33.33.33.33

inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

33.33.33.33/32    *[IS-IS/18] 00:32:43, metric 10
                  > to 10.1.1.6 via ge-0/0/1.0

inet.2: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0        *[Static/5] 00:14:25
                  > to 10.1.1.6 via ge-0/0/1.0

[edit]
user@R2# run show protocols pim
rp {
  bootstrap {
    family inet {
      priority 200;
    }
  }
  local {
    address 22.22.22.22;
    group-ranges {
      224.0.0.0/4;
    }
  }
}
interface all;

[edit]
user@R3# show protocols pim
rp {
  bootstrap {
    family inet {
      priority 210;
    }
  }
  local {
    address 33.33.33.33;
    group-ranges {
      224.1.0.0/16;
    }
  }
}
interface all;
```



R4 is directly connected to both RPs (R2 and R3). R4 is currently sending all joins upstream to R3 but you want to load balance the joins between both RPs.

Referring to the exhibit, which configuration change will solve this issue?



- A. Configure the `join-load-balance` parameter under PIM on R4.
- B. Configure the default route in `inet.2` on R4 from R3 as the next hop to both R3 and R2.
- C. Configure the `group-range` parameter to be the same on R2 and R3.
- D. Configure the bootstrap priority on R2 to be the same as R3.

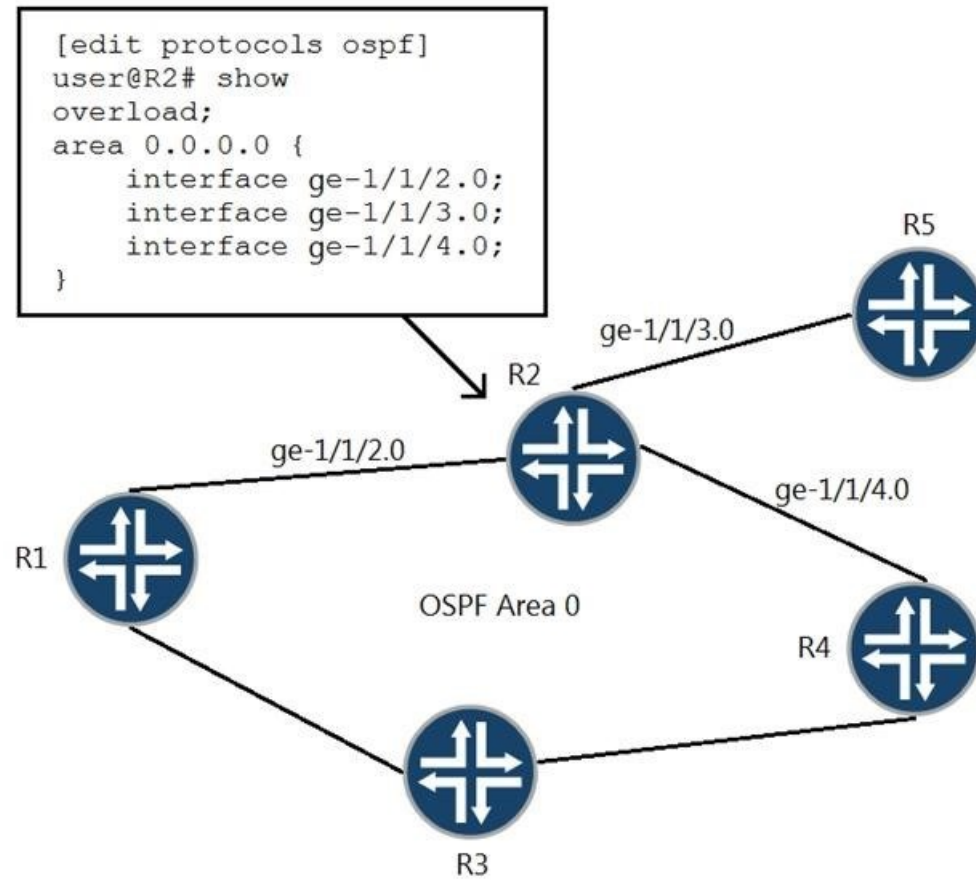
**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 12



Referring to the exhibit, which two statements are correct? (Choose two.)

- A. R2 stops sending LSAs into the network.
- B. The OSPF interface metrics on R2 are all set to 65535.
- C. R1 will never forward transit traffic through R2.
- D. Transit traffic from R1 to R4 will traverse R3.

**Correct Answer:** BD

**Section:** (none)


**Explanation**

**Explanation/Reference:**

#### QUESTION 13



```
user@R2# show protocols isis
level 1 disable;
interface ge-0/0/0.0;
interface ge-0/0/1.0 {
    level 2 metric 300;
}
```



```
user@R1# show protocols isis
level 1 disable;
interface ge-0/0/0.0;

user@R3# show protocols isis
level 1 disable;
interface ge-0/0/1.0;
```

#### AREA 49.0001

Referring to the exhibit, what will the IS-IS cost be for R1 to reach R3?

- A. 73
- B. 20
- C. 301
- D. 310

**Correct Answer:** D


**Section:** (none)

**Explanation**

**Explanation/Reference:**



#### QUESTION 14



Referring to the exhibit, traffic sent from CE-A2 to PE3 does not loop back to CE-A2 through PE2.

Winch two EVPN functions accomplish this task? (Choose two.)

- A. multicast ingress replication
- B. aliasing
- C. split horizon
- D. designated forwarder election



**Correct Answer:** CD

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 15** Which two types of LSAs have an area scope? (Choose two.)

- A. Type 2
- B. Type 5
- C. Type 11
- D. Type 7

**Correct Answer:** AC

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 16**

```
user@R1> show ospf3 interface
Interface      State      Area      DR ID      BDR ID      Nbrs
ge-0/0/0.0     DR         0.0.0.0    172.16.1.2  172.16.1.1    1
ge-0/0/0.0     PtToPt     0.0.0.1    0.0.0.0    0.0.0.0      1
ge-0/0/1.0     BDR        0.0.0.1    172.16.1.1  172.16.1.2    1

user@R1> show ospf3 neighbor
ID             Interface  State  Pri    Dead
172.16.1.1     ge-0/0/0.0 Full   128    39
  Neighbor-address fe80::20c:29ff:fef9:7f7b
  Area 0.0.0.0
172.16.1.1     ge-0/0/0.0 Full   128    37
  Neighbor-address fe80::20c:29ff:fef9:7f7b
  Area 0.0.0.1
172.16.1.1     ge-0/0/1.0 Full   128    37
  Neighbor-address fe80::20c:29ff:fef9:7f85
  Area 0.0.0.1
```



Referring to the exhibit, which OSPFv3 configuration is implemented on router R1?

- A. set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0
- B. set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 set protocols ospf3 area 0.0.0.1 virtual-link neighbor-id 172.16.1.2
- C. set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 secondary
- D. set protocols ospf3 area 0.0.0.0 interface ge-0/0/0.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/1.0 set protocols ospf3 area 0.0.0.1 interface ge-0/0/0.0 interface-type p2p

**Correct Answer:** D



Section: (none)

Explanation

Explanation/Reference:

#### QUESTION 17

```
[edit]
user@router# show protocols bgp group ISP
export ISP;
neighbor 172.16.20.20 {
    peer-as 65310;
}
neighbor 192.168.51.200 {
    peer-as 65441;
}
```

```
[edit]
user@router# show policy-options
policy-statement ISP {
    term ROUTES {
        from {
            protocol bgp;
            as-path LOCAL-ROUTES;
        }
        then accept;
    }
}
as-path LOCAL-ROUTES "(.*)";
```



Your network is connected to two different ISPs and you notice that they are using your network for transit traffic.

In this scenario, which two configuration statements will solve this problem? (Choose two.)

- A. set policy-options policy-statement ISP term REST then reject
- B. set policy-options policy-statement ISP term ROUTES then reject
- C. set policy-options as-path LOCAL-ROUTES "()"
- D. set policy-options as-path LOCAL-ROUTES "(65310|65441)+"

**Correct Answer:** AC


Section: (none)

Explanation

Explanation/Reference:

#### QUESTION 18





A customer wants to reduce LSP flooding in their IS-IS network.

Which parameter should you change to accomplish this task?

- A. `[edit protocols isis] user@router# set spf-options rapid-runs 5`
- B. `[edit protocols isis interface <interface-name>] user@router# set csnp-interval 65535`
- C. `[edit protocols isis interface <interface-name>] user@router# set lsp-interval 1000`
- D. `[edit protocols isis interface <interface-name>] user@router# set mesh-group <mesh-group-number>`



**Correct Answer: B**


**Section: (none)**

**Explanation**

**Explanation/Reference:**

**QUESTION 19**





A network designer wants to ensure that traffic from R2 destined for 2001:db9:ffff:ff00::/62 always traverses the R2-R1 link if that link is available.

Referring to the exhibit, which configuration change will satisfy this requirement? A.

```

user@R1# show protocols isis
export leak-v6;

user@R1# show policy-options
policy-statement leak-v6 {
  term DC-routes {
    from {
      protocol isis;
      level 2;
      route-filter 2001:db9:ffff:ff00::/62 orlonger;
    }
    to level 1;
    then accept;
  }
}
  
```

```

user@R2# show protocols isis
export leak-v6;

user@R2# show policy-options
policy-statement leak-v6 {
  term DC-routes {
    from {
      protocol isis;
      level 2;
      route-filter 2001:db9:ffff:ff00::/62 orlonger;
    }
    to level 1;
    then accept;
  }
}
  
```

B.



```
user@R1# show protocols isis
import leak-v6;

user@R1# show policy-options
policy-statement leak-v6 {
  term DC-routes {
    from {
      protocol isis;
      level 1;
      route-filter 2001:db9:ffff:ff00::/62 orlonger;
    }
    to level 2;
    then accept;
  }
}
```

C. }

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 20**



```
[edit]
user@R2# run show isis adjacency
Interface  System L   State      Hold (secs) SNPA
ge-0/0/0.0  R3    1   Up         6 0:50:56:93:54:4b
ge-0/0/0.0  R3    2   Up         7 0:50:56:93:54:4b
ge-0/0/1.0  R4    2   Up         7 0:50:56:93:54:4b
```

```
[edit]
user@R2# show
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 172.16.2.2/30;
      }
      family iso;
    }
  }
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.1.2/30;
      }
      family iso;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 22.22.22.22/32;
      }
      family iso;
        address 49.0001.0022.2222.0022.00;
      }
    }
  }
}
protocols {
  isis {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
    interface lo0.0 {
      level 1 disable;
    }
  }
}
```

```

    family inet {
      address 10.1.1.1/30;
    }
    family iso;
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 10.1.1.5/30;
    }
    family iso;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 44.44.44.44/32;
    }
    family iso;
      address 49.0004.0044.4444.0044.00;
    }
  }
}
}
protocols {
  isis {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
    interface lo0.0 {
      level 1 disable;
    }
  }
}
```



R2 has IS-IS adjacencies with R3 and R4. You want to ensure that R2 has both a level 1 and level 2 adjacency to both R3 and R4, but R2 only has one adjacency with R4.

Referring to the exhibit, which configuration change will solve this issue?

- A. Change the IS-IS area on R4 to match R2.
- B. Remove the level 1 disable configuration from R4.
- C. Remove the level 1 disable configuration from R2.
- D. Change the IS-IS area on R2 to match R4.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 21

You are asked to configure a series of interface policers and firewall filters, which include policers, on the same device. You must ensure that the two configuration methods do not conflict. What are two considerations when performing this task? (Choose two.)

- A. On inbound traffic, firewall filters are applied before interface policers.
- B. On outbound traffic, interface policers are applied before firewall filters.
- C. On outbound traffic, firewall filters are applied before interface policers.
- D. On inbound traffic, interface policers are applied before firewall filters.

**Correct Answer:** CD

**Section:** (none)

**Explanation**

**Explanation/Reference:**



#### QUESTION 22



```
user@router> show bgp neighbor 192.168.100.2
Peer: 192.168.100.2+50862 AS 65512 Local: 192.168.100.1+179 AS 65512
  Group: INT          Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal      State: Established      Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Refresh>
  Options: <GracefulShutdownRcv>
  Local Address: 192.168.100.1 Holdtime: 90 Preference: 170
  Graceful Shutdown Receiver local-preference: 0
  Number of flaps: 0
  Peer ID: 192.168.100.2    Local ID: 192.168.100.1    Active Holdtime: 90
  Keepalive Interval: 30    Group index: 0    Peer index: 0    SNMP index: 3
  I/O Session Thread: bgpio-0    State: Enabled
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast inet-vpn-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer supports 4 byte AS extension (peer-as 65512)
  Peer does not support Addpath
  NLRI(s) enabled for color nexthop resolution: inet-unicast
...
```



Referring to the exhibit, the local BGP router is receiving IPv4 routes from the BGP neighbor, but it is not receiving L3 VPN routes from the BGP neighbor.

Which two actions should you take to solve this problem? (Choose two.)

- A. Configure the family `inet-vpn unicast` statement on the BGP neighbor.
- B. Configure the family `inet unicast` statement on the local BGP router.
- C. Configure the family `inet-vpn unicast` statement on the local BGP router.
- D. Configure the family `inet unicast` statement on the BGP neighbor.

**Correct Answer:** AC


**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 23**





Your network connects two segments of your customer's network as shown in the exhibit. They need to exchange routes between Segment 1 and Segment 2 but both segments use the same AS number.

Which two steps will accomplish this task? (Choose two.)

- A. Configure the `routing-options autonomous-system loops 1` parameter on routers R1 and R4.
- B. Configure the BGP group with the `advertise-peer-as` parameter on routers R2 and R3.
- C. Configure the `routing-options autonomous-system loops 1` parameter on routers R2 and R3.
- D. Configure the BGP group with the `as-override` parameter on routers R1 and R4.

**Correct Answer:** AB

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 24** You have a mixed vendor EVPN environment and you need to ensure VXLAN interoperability between all devices.

In this scenario, which statement is correct?

- A. You should only use pure Type 2 routes.
- B. You should only use pure Type 5 routes.
- C. You should only use Type 2 and Type 5 routes.
- D. You should only use Type 6 and Type 2 routes.

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 25**

The link between CE1 and PE1 has a history of flapping. To avoid the impact that flapping causes to the network, you decide to use route damping.

Which statement is correct in this scenario?

- A. Dampening is enabled on interfaces.
- B. Dampened routes decay at a sliding rate known as half-life.
- C. Routes become dampened when the configured `max-suppress` value is reached.
- D. Dampened routes become active when their figure of merit drops below the reuse value.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**



**QUESTION 26** You are deploying a new EVPN service for your customers.

You must build the service based on the following requirements:

- both Layer 2 and Layer 3 functionality must be supported;
- your customers must be able to support multiple VLANs in the same EVPN instance (EVI).

In this scenario, which two types of routing instances should be configured? (Choose two.)

- A. virtual switch
- B. virtual router
- C. VRF
- D. EVPN

**Correct Answer:** CD

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 27** Which two statements about IS-IS are correct? (Choose two.)

- A. Level 1 intermediate systems exchange routing information with Level 1 intermediate systems on other IS-IS areas.
- B. An IS-IS router sets the attached bit in the PDUs it sends to a Level 1 area to indicate that it is a backbone router.
- C. A Level 1 router can only form adjacencies with other Level 1 routers.
- D. Level 2 routers can form adjacencies with either Level 1 or Level 2 routers.

**Correct Answer:** BC

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 28** Which two statements are correct about Opaque LSAs in OSPF? (Choose two.)

- A. Type 10 LSAs are used for MPLS traffic-engineering and have area scope.
- B. Type 11 LSAs are used for MPLS traffic-engineering and have area scope.
- C. Type 11 LSAs are used for MPLS label exchange and have link-local scope
- D. Type 9 LSAs are used for graceful-restart and have link-local scope

**Correct Answer:** AD

**Section:** (none)

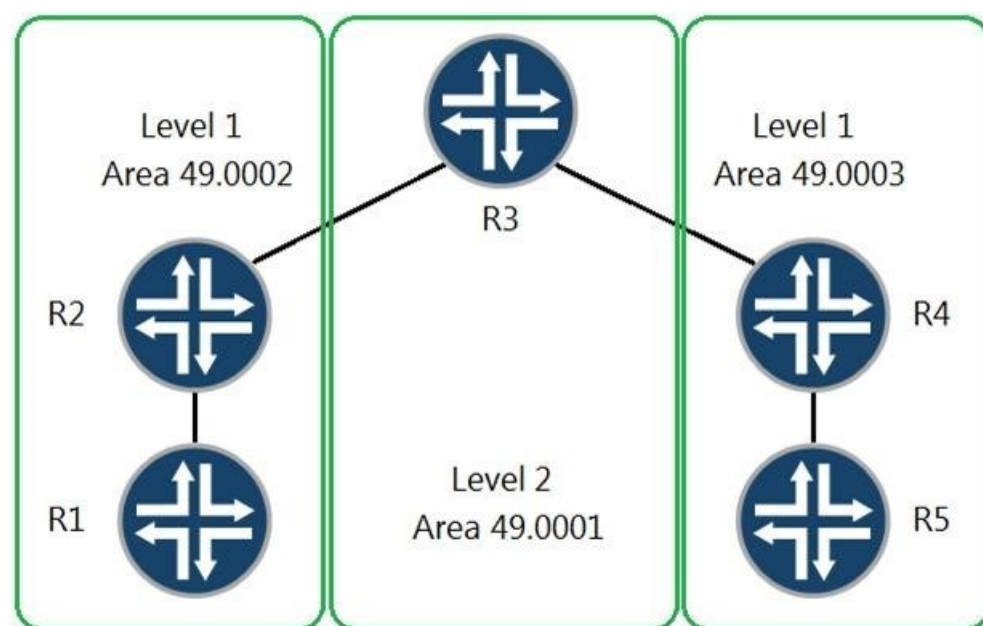
**Explanation**

**Explanation/Reference:**

**QUESTION 29**







All adjacencies have been formed, no extra options have been configured, and no policies have been written.

Referring to the exhibit, which two statements are correct? (Choose two.)

- A. R2 will create a default route and send it as a TLV to R1
- B. R1 cannot reach R5
- C. R1 can reach R5
- D. R1 will create its own default route that points to R2

**Correct Answer:** AC

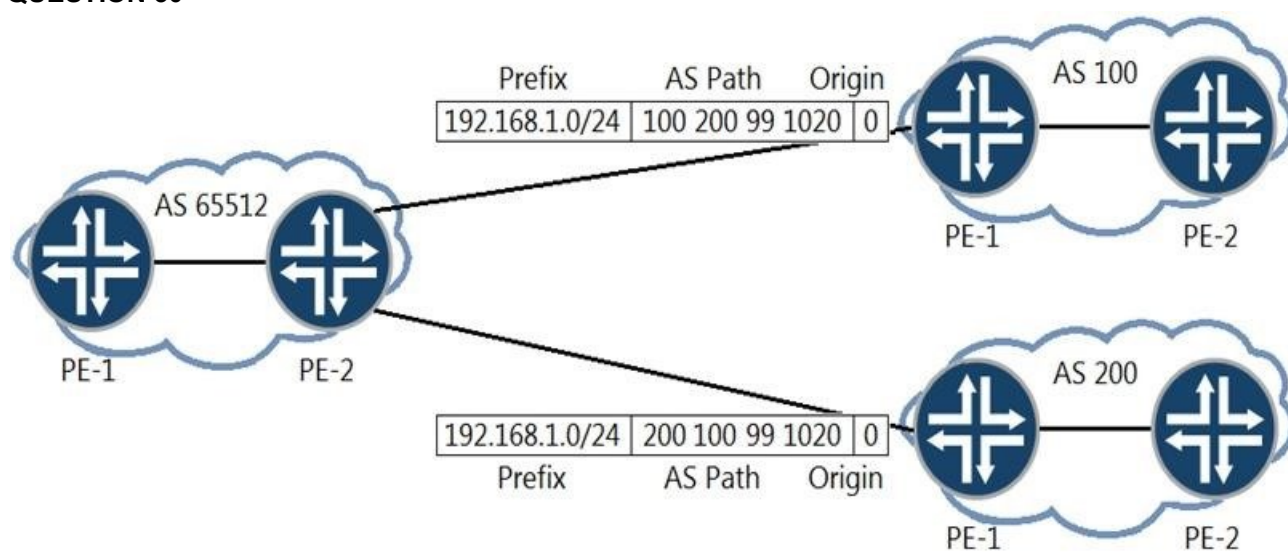
**Section:** (none)

**Explanation**

**Explanation/Reference:**



### QUESTION 30



You are the administrator of AS 65512. You are learning the 192.168.1.0/24 prefix from both AS 100 and AS 200. You want traffic destined to the 192.168.1.0/24 prefix to exit your AS towards AS 200.

How would you accomplish this task?



- A. Configure an import routing policy on PE-2 to set a higher MED on the path learned from AS 100.
- B. Configure an import routing policy on PE-2 to modify the origin attribute on the path learned from AS 100.
- C. Configure an import routing policy on PE-2 to set a higher local preference value on the path learned from AS 200.
- D. Configure an import routing policy on PE-2 to append the AS path attribute on the path learned from AS 100.

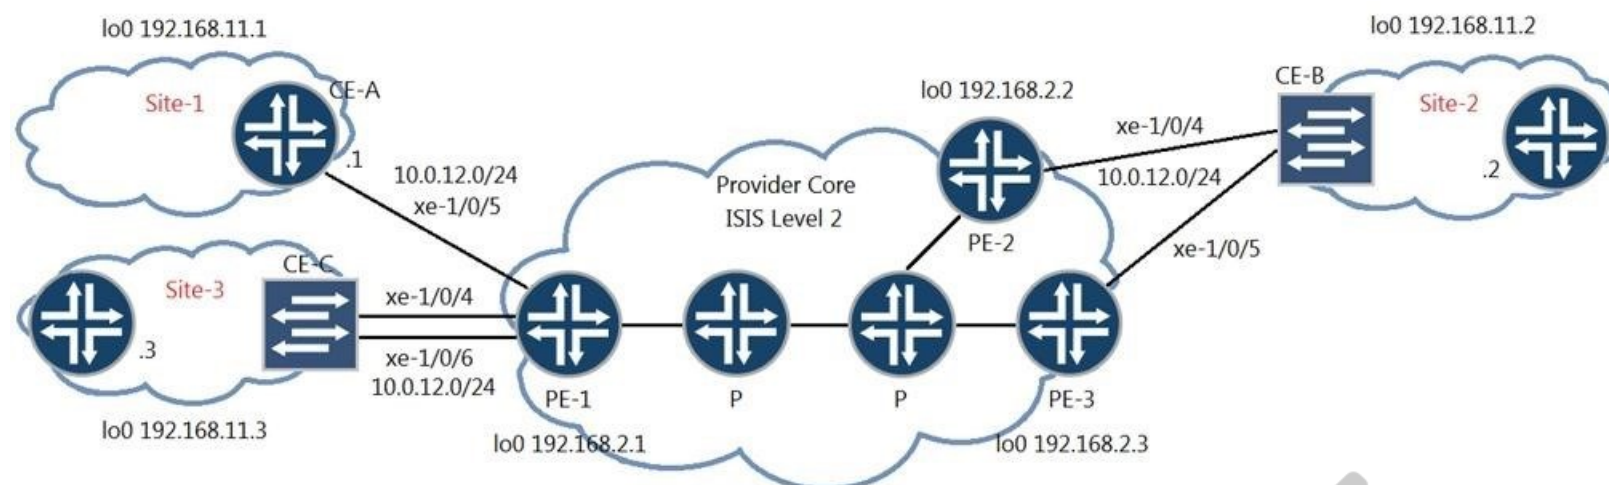
**Correct Answer: C**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

### QUESTION 31



You have the LDP signaled VPLS topology as shown in the exhibit. CE-B at Site-2 is multihomed to both PE-2 and PE-3.

In this scenario, where would you configure loop prevention?

- A. PE-1
- B. CE-B
- C. PE-3
- D. PE-2

**Correct Answer: A**


**Section: (none)**

**Explanation**

**Explanation/Reference:**

### QUESTION 32





```

user@R2# show interfaces
ge-0/0/0 {
  unit 0 {
    family iso;
    family inet6 {
      address 2001:db8::2/64;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 172.16.2.2/24;
    }
    family iso;
    family inet6 {
      address 2001:db8:1::2/64;
    }
  }
}
user@R2# show protocols isis
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface lo0.0 {
  passive;
}
topologies ipv6-unicast;

```

```

user@R1# show interfaces
ge-0/0/0 {
  unit 0 {
    family iso;
    family inet6 {
      address 2001:db8::1/64;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family iso;
    family inet6 {
      address 2001:db8:1::1/64;
    }
  }
}
user@R1# show protocols isis
interface ge-0/0/0.0;
interface ge-0/0/1.0;
interface lo0.0 {
  passive;
}
topologies ipv6-unicast;

```



A network administrator is migrating from IPv4 to IPv6 and one of the IS-IS adjacencies is not coming up between R1 and R2.

Which action will solve the problem?

- A. Remove `topologies ipv6-unicast` from `protocols isis` on R2.
- B. Configure `topologies ipv4-unicast` from `protocols isis` on R2.
- C. Remove `topologies ipv6-unicast` from `protocols isis` on R1.
- D. Configure an IPv4 address on interface `ge-0/0/1` on R1.

**Correct Answer:** D


**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 33**





```
[edit routing-instances]
user@PE-1# show
CE-A-1 {
  instance-type vrf;
  interface ge-0/0/9.0;
  route-distinguisher 10.222.222.4:1;
  vrf-target target:65511:101;
  routing-options {
    static {
      route 192.168.100.0/24
    }
  }
  next-hop 192.168.0.2;
}
CE-A-2 {
  instance-type vrf;
  interface ge-0/0/8.0;
  route-distinguisher 10.222.222.4:3;
  vrf-target target:65511:101;
  routing-options {
    static {
      route 192.168.101.0/24
    }
  }
  next-hop 192.168.1.2;
}
```

Referring to the exhibit, there is a Layer 3 VPN setup that connects sites CE-A-1, CE-A-2, and CE-A-3 together. Host-1 can communicate with Host-3, but Host-1 cannot communicate with Host-2.

What must you do to solve the problem?

- A. Change the route distinguisher in both routing instances to the same value.
- B. Use the `next-table` configuration statement for static routes in the corresponding routing instances.
- C. Use BGP instead of static routing between the CE and PE devices.
- D. Use the `auto-export` command in both routing instances.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 34**



```
user@host> show pim join 234.100.0.1 extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 234.100.0.1
  Source: 192.168.100.2
  Flags: sparse, spt
  Active upstream interface: ge-1/0/0.0
  Active upstream neighbor: 192.168.101.2
  MoFRR Backup upstream interface: ge-1/0/1.0
  MoFRR Backup upstream neighbor: 192.168.102.2
  Upstream state: Join to Source, No Prune to RP
  Keepalive timeout: 300
  Uptime: 00:00:15
  Downstream neighbors:
    Interface: ge-1/2/0.0
      192.168.103.2 State: Join Flags: S Timeout: Infinity
      Uptime: 00:00:15 Time since last Join: 00:00:15
  Number of downstream interfaces: 1
```

Which three statements are true about the `show pim join` output shown in the exhibit? (Choose three.)

- A. This is a source-specific multicast stream.
- B. The multicast receiver is still using the RP to receive the stream.
- C. The multicast stream has been configured with a backup path to allow for fast reroute.
- D. The multicast stream does not have an RP.
- E. The shortest path to the source is through the RP.



**Correct Answer:** CDE

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 35** Why do interprovider option B VPNs scale better than interprovider option A VPNs?

- A. The ASBRs in interprovider option B VPNs do not need per-VPN VRF tables.
- B. The ASBRs in interprovider option A VPNs do not need per-VPN VRF tables.
- C. The ASBRs in interprovider option A VPNs only carry internal routes.
- D. The ASBRs in interprovider option B VPNs only carry internal routes.

**Correct Answer:** A


**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 36**





```
users@R1> show ospf3 database inter-area-prefix detail
```

```
OSPF3 database, Area 0.0.0.0
Type      ID          Adv Rtr      Seq          Age    Cksum    Len
InterArPfx 0.0.0.11      172.16.1.1  0x80000001   4      0xaa9a   36
Prefix 2001:db9:ffff:ff00::/64
Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.12      172.16.1.1  0x80000001   4      0x8c6e   44
Prefix 2001:db9:ffff:ff00::1/128
Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.13      172.16.1.1  0x80000001   4      0xa899   36
Prefix 2001:db9:ffff:ff01::/64
Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.14      172.16.1.1  0x80000001   4      0x8a6d   44
Prefix 2001:db9:ffff:ff01::1/128
Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.15      172.16.1.1  0x80000001   4      0xa698   36
Prefix 2001:db9:ffff:ff02::/64
Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.16      172.16.1.1  0x80000001   4      0x886c   44
Prefix 2001:db9:ffff:ff02::1/128
Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.17      172.16.1.1  0x80000001   4      0xa497   36
Prefix 2001:db9:ffff:ff03::/64
Prefix-options 0x0, Metric 0
InterArPfx 0.0.0.18      172.16.1.1  0x80000001   4      0x866b   44
Prefix 2001:db9:ffff:ff03::1/128
Prefix-options 0x0, Metric 0
```



Referring to the exhibit, which command would reduce the size of the OSPF database and corresponding routes? A.

```
user@R1# show protocols ospf3
area 0.0.0.1 {
    area-range 2001:db9:ffff:ff00::/62;
}

user@R1# show policy-options policy-statement summary-2001
term 10 {
    from {
        route-filter 2001:db9:ffff:ff00::/62 prefix-length-range /64-/128;
    }
    then accept;
}

user@R1# show protocols ospf3
area 0.0.0.0 {
    inter-area-prefix-import summary-2001;
}
```

B.



```

user@R1# show policy-options policy-statement summary-2001
term 10 {
  from {
    route-filter 2001:db9:ffff:ff00::/62 prefix-length-range /64-/128;
  }
  then accept;
}
user@R1# show protocols ospf3
area 0.0.0.1 {
  inter-area-prefix-export summary-2001;
}

user@R1# show protocols ospf3
area 0.0.0.1 {
  stub no-summaries;
}

```

C.

D.


**Correct Answer: A**

**Section: (none)**

**Explanation**

**Explanation/Reference:**

**QUESTION 37**



You are providing carrier-of-carrier VPN services for AS 100 and AS 200. You want to distribute MPLS labels between your PE routers and the AS 100 and AS 200 CE routers.

What must be enabled to accomplish this task?



- A. Use BGP with the `inet-vpn` address family enabled.
- B. Use BGP with the `labeled-unicast` address family enabled.
- C. Use RSVP with the `lsp-set` parameter enabled.
- D. Use RSVP with the `tunnel-services` parameter enabled.

**Correct Answer:** A

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 38** Which two statements regarding Ethernet segments (ES) are correct? (Choose two.) A. The Type-4 EVPN route will be used to elect the designated forwarder for the ES.

- B. The Type-3 EVPN route will be used for the aliasing function to load-balance to the ES.
- C. The Type-1 EVPN route will indicate if the ES is all-active or single-active.
- D. The Type-2 EVPN route will indicate if there is a designated forwarder on the ES.


**Correct Answer:** AC

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 39**



You have a Layer 3 VPN established between PE-1 and PE-2 to allow communication between CE-1 and CE-2.

Referring to the exhibit, which statement is correct?

- A. You will need an OSPF import policy on PE-1 to receive the BGP routes, learned from PE-2, through the Layer 3 VPN.
- B. You will need a BGP export policy on PE-1 to redistribute the routes, learned from CE-1, through the Layer 3 VPN.
- C. You will need a VRF import policy on PE-1 to advertise the BGP routes, learned from CE-1, through the Layer 3 VPN.
- D. You will need an OSPF export policy on PE-2 to redistribute the BGP routes, learned from PE-1, through the Layer 3 VPN.

**Correct Answer:** D

**Section:** (none)

**Explanation**



Explanation/Reference:

QUESTION 40





```
[edit]
user@R4# run show route hidden extensive

inet.0: 7 destinations, 7 routes (5 active, 0 holddown, 1 hidden)
11.11.11.0/24 (1 entry, 0 announced)
  BGP Preference: 170/-101
    Next hop type: Unusable, Next hop index: 0
    Address: 0xbc4dbb4
    Next-hop reference count: 2
    State: <Hidden Int Ext>
    Peer AS: 65002
    Age: 18
    Validation State: unverified
    Task: BGP_65002_65002.22.22.22
    AS path: 65001 I
    Communities: no-export no-advertise
    Accepted
    Localpref: 100
    Router ID: 22.22.22.22
    Indirect next hops: 1
      Protocol next hop: 172.16.1.1
      Indirect next hop: 0x0 - INH Session ID: 0x0
```

```
[edit protocols bgp]
user@R2# show
group 65001 {
  neighbor 172.16.1.1 {
    export no-advertise;
    peer-as 65001;
  }
}
group 65002 {
  type internal;
  local-address 22.22.22.22;
  neighbor 44.44.44.44 {
    export no-advertise;
  }
}
import no-export;
export nhs;
local-as 65002;
```

```
[edit]
user@R2# show policy-options
policy-statement no-advertise {
  term 1 {
    then {
      community add no-advertise;
    }
  }
}
policy-statement no-export {
  term 1 {
    then community add no-export;
  }
}
policy-statement nhs {
  term 1 {
    then {
      next-hop self;
    }
  }
}
community no-advertise members no-advertise;
community no-export members no-export;
```





R2 is receiving a route from an EBGp neighbor and is advertising the route to R4.

Referring to the exhibit, which configuration on R2 will solve the issue with the route on R4?

- A. Move the `no-advertise` export policy from group 65002 to a global BGP policy.
- B. Move the `nhs` policy from a global BGP export policy to an export policy under group 65002.
- C. Move the `no-export` policy from a global BGP import policy to an import policy under group 65001.
- D. Move the `no-advertise` export policy from group 65001 to a global BGP policy.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 41

```
[edit routing-instances]
user@PE-1# show
vpn-a {
    instance-type vrf;
    interface ge-1/1/4.0;
    route-distinguisher 192.168.1.1:1;
    vrf-target target:65111:101;
    protocols {
        bgp {
            group my-ext-group {
                type external;
                peer-as 65601;
                neighbor 10.0.10.2;
            }
        }
    }
}
```



You have an established Layer 3 VPN between two PE devices. You are asked to only send certain routes from PE-1 over the VPN to the remote site while maintaining all the routes on the PE-1 device. You created a policy that matches the specific routes and then tags these routes with the appropriate target community values.

In this scenario, which configuration changes must be made to satisfy the requirement?

- A. Configure the export parameter and apply the policy to the `my-ext-group` BGP group configuration.
- B. Configure the `vrf-export` parameter and apply the policy under the `edit routing-instancesvpn-a` hierarchy.
- C. Configure a RIB group and apply the policy as an import policy to routes distributed into the `bgp.13vpn.0` routing table.
- D. Configure the import parameter and apply the policy to the `my-ext-group` BGP group configuration.

**Correct Answer:** B


**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 42





Referring to the exhibit, which three statements about route 10.222/16 are correct when using the default BGP advertisement rules? (Choose three)

- A. R1 will prepend AS10 when advertising 10.222/16 to R2.
- B. R2 will prepend AS10 when advertising 10.222/16 to R7.
- C. R2 will advertise 10.222/16 to R4 with itself as the next hop.
- D. R1 will advertise 10.222/16 to R2 with itself as the next hop.
- E. R7 will advertise 10.222/16 to R9 with itself as the next hop.




**Correct Answer:** BCE

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 43



Traffic is being sent from R2 to R3. The link between R2 and R3 fails. Referring to the exhibit, which statement is correct?



- A. Traffic will automatically reroute using the shortest path, which is R2 to R1 to R3.
- B. Traffic will automatically reroute distributed between all available paths.
- C. Manual intervention is required for traffic to be rerouted.
- D. Traffic will automatically reroute using R2 to R4 to R1 to R3.


**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 44



R1 assigns incoming voice traffic to the ef forwarding class. All other traffic is assigned to the best-effort forwarding class. You have configured a CoS re-write rule on R1 to include the correct CoS bit values in packets sent towards R2. You want R2 to classify traffic using the CoS markings created by R1.

Which two configuration steps are necessary to accomplish this task? (Choose two.)

- A. Assign the behavior aggregate classifier to the ge-0/0/1.0 interface on R2.
- B. Assign the CoS re-write rule to the ge-0/0/1.0 interface on R2.
- C. Configure a CoS re-write rule on R2 and assign matching CoS values.
- D. Configure a behavior aggregate classifier on R2.




**Correct Answer:** AB

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 45



All networks shown in the exhibit contain more than one BGP speaker. You operate ISP A and must ensure that Customer Y sends their traffic to you over the directly connected link. Customer Y is not to be used for transit into your network.

What would you do to accomplish this task?

- A. Advertise routes to Customer X with the custom defined 0:0 community.
- B. Advertise routes to Customer X with the well-known no-advertise community.
- C. Advertise routes to Customer Y with the custom defined 65535:65535 community.



D. Advertise routes to Customer Y with the well-known no-export community.


**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 46



Referring to the exhibit, you need to implement VPLS between CE-A1 and CE-A2. You must ensure that no loops are created due to the multihoming of the connection from CE-A2 to PE2 and PE3. Based on the type of VPLS, which two solutions will satisfy this requirement? (Choose two.)

- A. In an LDP VPLS, configure a primary and backup neighbor.
- B. In an LDP VPLS, configure multihoming and local preference on PE-2 and PE-3.
- C. In a BGP VPLS, configure multihoming and local preference on PE-2 and PE-3.
- D. In a BGP VPLS, configure a primary and backup neighbor.




**Correct Answer:** AC

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 47



You have a Layer 3 VPN established between PE-1 and PE-2 to allow communication between CE-1 and CE-2. You want to establish communication between CE-1 and CE-2.



Referring to the exhibit, which statement is correct?

- A. You will need a BGP export policy on PE-1 to redistribute the OSPF routes, learned from PE-2, to the CE1 BGP neighbor.
- B. You will need a VRF import policy on PE-2 to advertise the OSPF routes, learned from CE-2, through the Layer 3 VPN.
- C. You will need a VRF export policy on PE-2 to redistribute the OSPF routes, learned from CE-2, through the Layer 3 VPN.
- D. You will need a VRF import policy on PE-1 to receive the OSPF routes, learned from PE-2, through the Layer 3 VPN.


**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 48



Refer ring to the exhibit, which two statements are correct for a route advertised by R1 towards R4?  
(Choose two.)

- A. The BGP next hop is set to 14.1.1.1 by R2.
- B. The AS path is set to 150 by R2.
- C. The BGP next hop is set to 192.168.2.2 by R2.
- D. The AS path is set to null by R2.

**Correct Answer:** AD

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 49



```

user@router# run show class-of-service rewrite-rule name
traffic-class
rewrite rule: traffic-class, code point type: exp, index:
58866
Forwarding class      Loss Priority      Code Point
best-effort           low               000
best-effort           high             001
expedited-forwarding  low              111
expedited-forwarding  high             011
assured-forwarding    low              100
assured-forwarding    high             101
network-control       low              110
network-control       high             111

```

Your router should be configured with a rewrite rule which alters the default behavior of expedited forwarding as shown in the exhibit.

In this scenario, which configuration is correct?

```

[edit class-of-service]
user@router# show
rewrite-rules {
  exp traffic-class {
    import best-effort;
    import assured-forwarding;
    import network-control;
    forwarding-class expedited-forwarding {
      loss-priority low code-point 111;
    }
  }
}

[edit class-of-service]
user@router# show
rewrite-rules {
  exp traffic-class {
    import rewrite-rule best-effort;
    import rewrite-rule expedited-forwarding;
    import rewrite-rule assured-forwarding;
    import rewrite-rule network-control;
    forwarding-class expedited-forwarding {
      loss-priority low code-point 111;
    }
  }
}

```



A.



B.

```
[edit class-of-service]
user@router# show
rewrite-rules {
    exp traffic-class {
        import best-effort;
        import assured-forwarding;
        import expedited-forwarding;
        import network-control;
    }
}

[edit class-of-service]
user@router# show
rewrite-rules {
    exp traffic-class {
        import default;
        forwarding-class expedited-forwarding {
            loss-priority low code-point 111;
        }
    }
}
```

C.



D.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 50**



```
[edit routing-instances vpn-x]
user@router# show
instance-type l2vpn;
interface ge-1/0/1.513;
interface ge-1/0/1.512;
route-distinguisher 192.168.1.2:1;
vrf-import import-vpn-x;
vrf-export export-vpn-x;
protocols {
  l2vpn {
    encapsulation-type ethernet-vlan;
    site ce-a {
      site-identifier 2;
      interface ge-1/0/1.512;
      interface ge-1/0/1.513;;
    }
  }
}
```

You have the Layer 2 VPN configuration shown in the exhibit. You are asked to determine the remote site ID for ge-1/0/1.512.

In this scenario, what is the remote site ID?

- A. 4
- B. 5
- C. 1
- D. 3

**Correct Answer: C**

**Section: (none)**


**Explanation**

**Explanation/Reference:**



#### QUESTION 51

```
user@R2# show protocols isis
level 1 disable;
interface ge-0/0/0.0;
interface ge-0/0/1.0 {
  level 2 metric 300;
}
```



```
user@R1# show protocols isis
level 1 disable;
interface ge-0/0/0.0;

user@R3# show protocols isis
level 1 disable;
interface ge-0/0/1.0;
```

**AREA 49.0002**



Referring to the exhibit, what will the IS-IS cost be for R3 to reach R1?

- A. 301
- B. 73
- C. 20
- D. 310

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 52**





```

user@router> show route protocol bgp hidden extensive

inet.0: 66 destinations, 66 routers (66 active, 0 holddown, 0 hidden)

CES.inet.0: 11 destinations, 11 routes (3 active, 0 holddown, 1 hidden)
10.1.1.0/24 (1 entry, 0 announced)
  BGP      Preference: 170/-101
            Route Distinguisher: 65512:1
            Next hop type: Unusable, Next hop index: 0
            Address: 0xc7412d0
            Next-hop reference count: 16
            State: <Secondary Hidden Int Ext ProtectionCand>
            Local AS: 65512 Peer AS: 65512
            Age: 1:53
            Validation State: unverified
            Task: BGP_65512.192.168.100.1
            AS path: I
            Communities: target:65512:100
            Import Accepted
            VPN Label:17
            Localpref: 100
            Router ID: 192.168.100.1
            Primary Routing Table: bgp.13vpn.0
            Indirect next hops: 1
              Protocol next hop: 192.168.100.1
              Label operation: Push 17
              Label TTL action: prop-ttl
              Load balance label: Label 17: None;
              Indirect next hop: 0x0 - INH Session ID: 0x0

...

65512:1:10.1.1.0/24 (1 entry, 0 announced)
  -BGP     Preference: 170/-101
            Route Distinguisher: 65512:1
            Next hop type: Unusable, Next hop index: 0
            Address: 0xc7412d0
            Next-hop reference count: 16
            State: <Hidden Int Ext Changed ProtectionPath ProtectionCand>
            Local AS: 65512 Peer AS: 65512
            Age: 1:53
            Validation State: unverified
            Task: BGP_65512.192.168.100.1
            AS path: I
            Communities: target:65512:100
            Import Accepted
            VPN Label: 17
            Localpref: 100
            Router ID: 192.168.100.1
            Secondary Tables: CE5.inet.0
            Indirect next hops: 1
              Protocol next hop: 192.168.100.1
              Label operation: Push 17
              Label TTL action: prop-ttl
              Load balance label: Label 17: None;
              Indirect next hop: 0x0 - INH Session ID: 0x0

```



Referring to the exhibit, a Layer 3 VPN is configured, however, the routes are being hidden.

What is the problem?

- A. The BGP peer is not reachable through the IGP.
- B. An active MPLS tunnel does not exist between the peers.



- C. A route distinguisher mismatch exists between the peers.
- D. A VRF target community mismatch exists between the peers.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 53** You are creating an LDP-signaled Layer 2 circuit between two sites. Site1 and Site2 use different VLAN IDs to connect to your PE devices.

In this scenario, which encapsulation type must be used on the logical interfaces?

- A. vlan-ccc
- B. vlan-vpls
- C. vlan-bridge
- D. vlan-tcc

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 54** What are two reasons an IBGP learned route would be hidden? (Choose two.)

- A. The route is suppressed by damping.
- B. The route has a next hop of the local routing device.
- C. The route is rejected by an export policy.
- D. The route has an empty AS path.

**Correct Answer:** AB

**Section:** (none)

**Explanation**

**Explanation/Reference:**

**QUESTION 55**

	AS-Path	MED	Local Preference	Origin
ISP-A	100 200 1	50	100	I
ISP-B	3000 1500	50	100	E
ISP-C	5000 4000	50	100	?
ISP-D	1000 7000	50	100	I

You are receiving the same 200.0.0.0/24 BGP route from four different ISPs.

Referring to the exhibit, which ISP's route would be selected as active?

- A. ISP-B
- B. ISP-A
- C. ISP-C



D. ISP-D

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 56

```
user@R1> show configuration protocols evpn
encapsulation vxlan;
default-gateway no-gateway-community;
extended-vni-list all;
```

```
user@R1> show configuration switch-options
vtep-source-interface lo0.0;
route-distinguisher 192.168.101.2:65101;
vrf-import EVPN-IMPORT;
vrf-target {
    target:1:100;
    auto;
}
```

```
user@R2> show configuration protocols evpn
vni-options {
    vni 22030 {
        vrf-target target:65101:22030;
    }
}
encapsulation vxlan;
default-gateway no-gateway-community;
extended-vni-list all;
```

```
user@R2> show configuration switch-options
vtep-source-interface lo0.0;
route-distinguisher 192.168.101.2:65101;
vrf-target {
    target:1:100;
    auto;
```



You are using EVPN to provide Layer 2 stretched VLANs between two sites. You notice that the MAC addresses in either site are not showing up on the remote site. Referring to the exhibit, what are two ways to solve this problem? (Choose two.)

- A. On R1, issue the `set switch-options vrf-target target:65101:22030` command.
- B. On R2, issue the `delete protocols evpn vni-options vni 22030` command.
- C. On R2, issue the `set switch-options vrf-target target:65101:22030` command.
- D. On R1, issue the `set protocols evpn vni-options vni 22030 vrf-target target :65101:22030` command.


**Correct Answer:** AC

**Section:** (none)



**Explanation**  
**Explanation/Reference:**

#### QUESTION 57



Referring to the exhibit, you want to reduce the CPU processing load on PE 2 by preventing the receipt of routes belonging to Customer B.

In this scenario, which layer 3 VPN scaling mechanism provides this functionality?

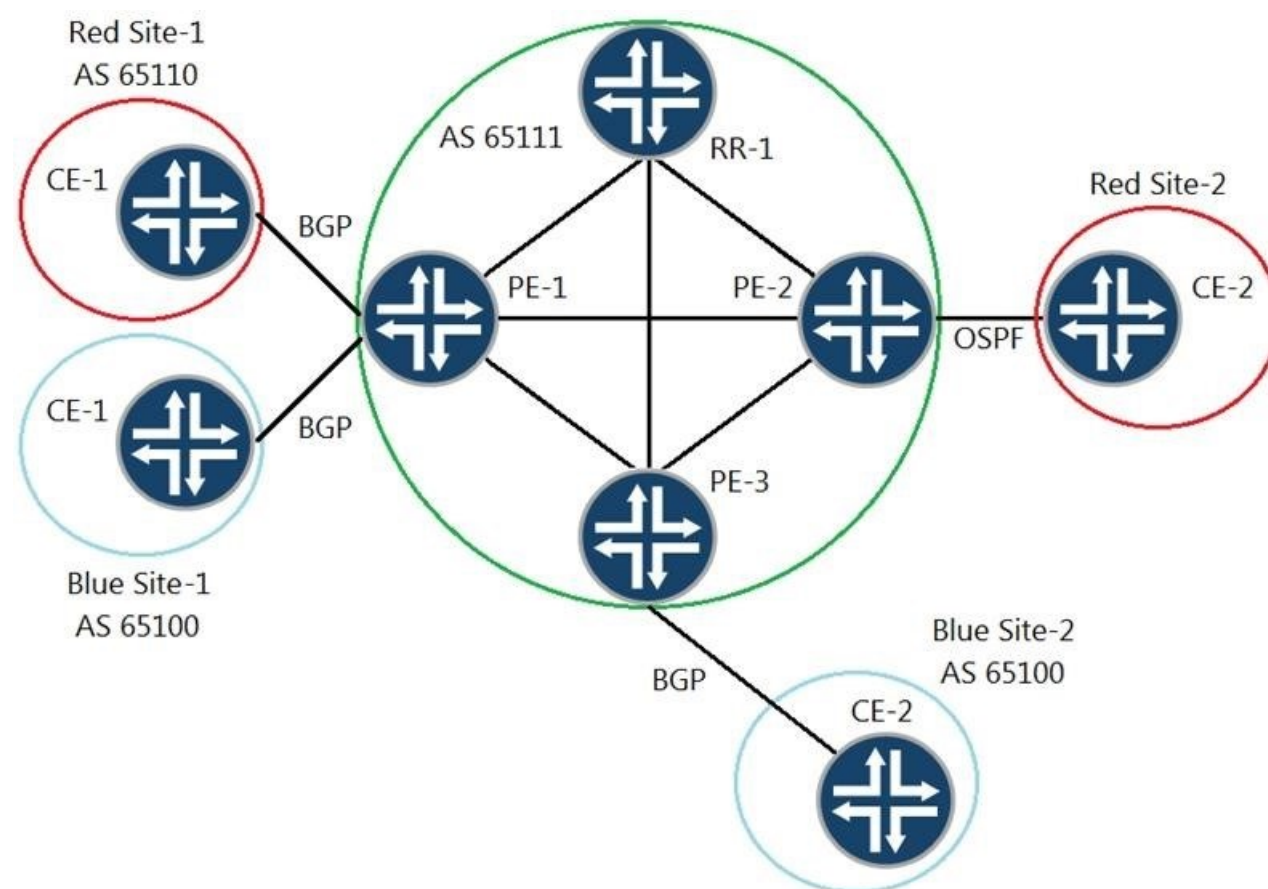
- A. route refresh
- B. route origin
- C. route reflection
- D. route target filtering

**Correct Answer:** D  
**Section:** (none)  
**Explanation**

**Explanation/Reference:**

#### QUESTION 58





You have a Layer 3 VPN established between PE-1 and PE-2 as well as between PE-1 and PE-3. You are using a route reflector (RR-1) to distribute VPN routes to your IBGP peers. You are asked to ensure that only relevant routes are sent from RR-1 to each of the PE routers.

Referring to the exhibit, which statement is correct?

- A. You should use VRF export policies on RR-1 to control which routes are sent to each PE router.
- B. You should use route target filtering only on RR-1 to control which routes are sent to each PE router.
- C. You should use firewall filtering on RR-1 and all the PE devices to control which routes are sent to each PE router.
- D. You should use route target filtering on RR-1 and all the PE devices to control which routes are sent to each PE router.

**Correct Answer:** B

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 59

You are considering different MPLS VPN connectivity options for a new customer deployment. Your customer requires shared LSPs, Layer 2 connectivity, and auto-provisioning.

Which type of VPN satisfies the requirements?

- A. BGP Layer 3 VPNs
- B. circuit cross-connects
- C. LDP Layer 2 circuits
- D. BGP Layer 2 VPNs

**Correct Answer:** D

**Section:** (none)

**Explanation**



**Explanation/Reference:****QUESTION 60**

You are establishing a Layer 3 VPN between two PE devices. Currently you have a single internal IPv4 BGP peering between the PE devices. You must ensure that the IPv4 and IPv6 routes from both CE devices are exchanged between these sites.

Which two statements are correct in this scenario? (Choose two.)

- A. You must enable IPv6 tunneling on the LSPs between the PE devices.
- B. You must establish an IPv6 BGP peering between the two PEs.
- C. You must enable the `inet6-vpn` NLRI on both PE devices.
- D. You must enable the `inet-vpn` NLRI on both PE devices.

**Correct Answer:** CD

**Section:** (none)

**Explanation**

**Explanation/Reference:****QUESTION 61**

```
[edit routing-instances]
user@PE-1# show
CE-1 {
  protocols {
    bgp {
      group CE-1 {
        type external;
        peer-as 65555;
        neighbor 10.1.1.100;
      }
    }
  }
  instance-type vrf;
  interface ge-0/0/2.0;
  route-distinguisher 65512:1;
  vrf-target target:65512:100;
}
CE-2 {
  protocols {
    bgp {
      group CE-2 {
        type external;
        peer-as 63333;
        neighbor 10.1.2.100;
      }
    }
  }
  instance-type vrf;
  interface ge-0/0/3.0;
  route-distinguisher 65512:2;
  vrf-target target:65512:100;
}
```





Two CE devices (CE-1 and CE-2) belong to the same customer and connect into a single PE device (PE-1). However, the CE devices cannot communicate with each other. You want to allow the CE devices to communicate with each other.

Referring to the exhibit, which action would solve the problem?

- A. Configure both routing instances with the `set routing-options autonomous-system loops 3` statement.
- B. Configure both routing instances with the `as-override` statement within the BGP protocol.
- C. Configure both routing instances with the `vrf-table-label` statement.
- D. Configure both routing instances with the `set routing-options auto-export` statement.


**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 62



```

user@R2> show isis database extensive level 2
Header: LSP ID: R1.00-00, Length: 457 bytes
  Allocated length: 491 bytes, Router ID: 10.254.0.1
  Remaining lifetime: 1130 secs, Level: 2, Interface: 73
  Estimated free bytes: 0, Actual free bytes: 34
  Aging timer expires in: 1130 secs
  Protocols: IP, IPv6
  
```



```

Packet: LSP ID: R1.00-00, Length: 457 bytes, Lifetime : 1196 secs
  Checksum: 0xef18, Sequence: 0x1d, Attributes: 0x7 <L1 L2 Overload>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0
  
```

```

TLVs:
  Area address: 49.0002 (3)
  LSP Buffer Size: 1492
  Speaks: IP
  Speaks: IPV6
  IP router id: 10.254.0.1
  IP address: 10.254.0.1
  IPv6 TE Router ID: 2001:db8::1
  Hostname: R1
  IS neighbor: R1.02, Internal, Metric: default 10
  IS neighbor: R1.03, Internal, Metric: default 10
  Extended IS Reachability TLV, Type: 22, Length: 90
  IS extended neighbor: R1.02, Metric: default 10 SubTLV len: 34
    IP address: 172.16.1.1
    IPv6 address: 2001:db8::1
    Local interface index: 73, Remote interface index: 0
  Router Capability: Router ID 10.254.0.1, Flags: 0x00
    IPv6 TE Router Id: 2001:db8::1
  No queued transmissions
  
```

A network administrator is investigating why traffic from R2 is not being forwarded to R1.



Referring to the `show isis database` command output shown in the exhibit, what is causing this problem on the network?

- A. R1 and R2 are in different IS-IS areas.
- B. The preferred interface between R1 and R2 is experiencing errors.
- C. R1 is configured to drop all incoming traffic.
- D. R2 is ignoring specific LSPs from R1 in its SPF calculations.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 63

```
[edit protocols bgp]
user@R1# show
group INT {
    type internal;
    local-address 192.168.100.1;
    family inet {
        unicast;
    }
    family inet6 {
        unicast;
    }
    neighbor 192.168.100.2;
}
```



```
[edit protocols bgp]
user@R2# show
group INT {
    type internal;
    local-address 192.168.100.2;
    export nhs;
    neighbor 192.168.100.1;
}
```

Referring to the exhibit, which statement is true?

- A. The BGP session between R1 and R2 will establish correctly and only the inet6 unicast NLRI will pass routing information.
- B. The BGP session between R1 and R2 will fail to establish correctly due to an NLRI mismatch.
- C. The BGP session between R1 and R2 will establish correctly and the inet unicast and the inet6 unicast NLRIs will pass routing information.
- D. The BGP session between R1 and R2 will establish correctly and only the inet unicast NLRI will pass routing information.

**Correct Answer:** D

**Section:** (none)

**Explanation**

**Explanation/Reference:**

#### QUESTION 64



Referring to the exhibit, which two statements are correct? (Choose two.)


- A. Term 0 is missing a terminating action that allows BGP routes to be accepted.
- B. Term 0 is missing a route-filter that specifies the allowed routes.
- C. The reject at the end of the policy is preventing the routes from being accepted.
- D. You cannot have terminating actions outside of terms.

**Section: (none)**

### Explanation

**Explanation/Reference:**

### QUESTION 65



Referring to the exhibit, which path would traffic passing through R1 take to get to R6?



- A.  $R_1 \rightarrow R_2 \rightarrow R_4 \rightarrow R_6$
- B.  $R_1 \rightarrow R_2 \rightarrow R_3 \rightarrow R_6$
- C.  $R_1 \rightarrow R_3 \rightarrow R_5 \rightarrow R_6$
- D.  $R_1 \rightarrow R_3 \rightarrow R_4 \rightarrow R_6$

**Correct Answer:** C

**Section:** (none)

**Explanation**

**Explanation/Reference:**

