PassQueen
a.k.a dumpsqueen.com

Juniper Networks Certified Internet Professional SP (JNCIP-SP)

Juniper JN0-660

Version Demo

Total Demo Questions: 26

Total Premium Questions: 261

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, OSPF	34
Topic 2, IS-IS	26
Topic 3, BGP	41
Topic 4, MPLS	38
Topic 5, Layer 2 VPNs	29
Topic 6, Layer 3 VPNs	26
Topic 7, Multicast	30
Topic 8, Mix Questions	37
Total	261

QUESTION NO: 1

Which two statements are true about Bidirectional Forwarding Detection (BFD)? (Choose two.)

A. BFD detects forwarding-path failures using periodic control packets.

B. BFD can cause an associated routing-protocol adjacency to be torn down rapidly.

C. BFD replaces the need to configure an IGP in an MPLS core.

D. BFD advertises IP reachability information between neighbors.

ANSWER: A B

Explanation:

BFD provides rapid detection of a forwarding-path failure between two systems. It uses periodic control packets sent at a negotiated interval and declares the session down when the configured detection time expires without receiving expected packets.

Routing protocols can use BFD to react more quickly to a failure than they would by relying only on their native hello and hold timers. For example, BFD can be associated with BGP, OSPF, and IS-IS adjacencies on supported Junos platforms. When BFD reports a failure, the associated protocol can promptly tear down the affected adjacency and begin reconvergence.

See the Juniper BFD overview and RFC 5880.

QUESTION NO: 2

Click the Exhibit button.

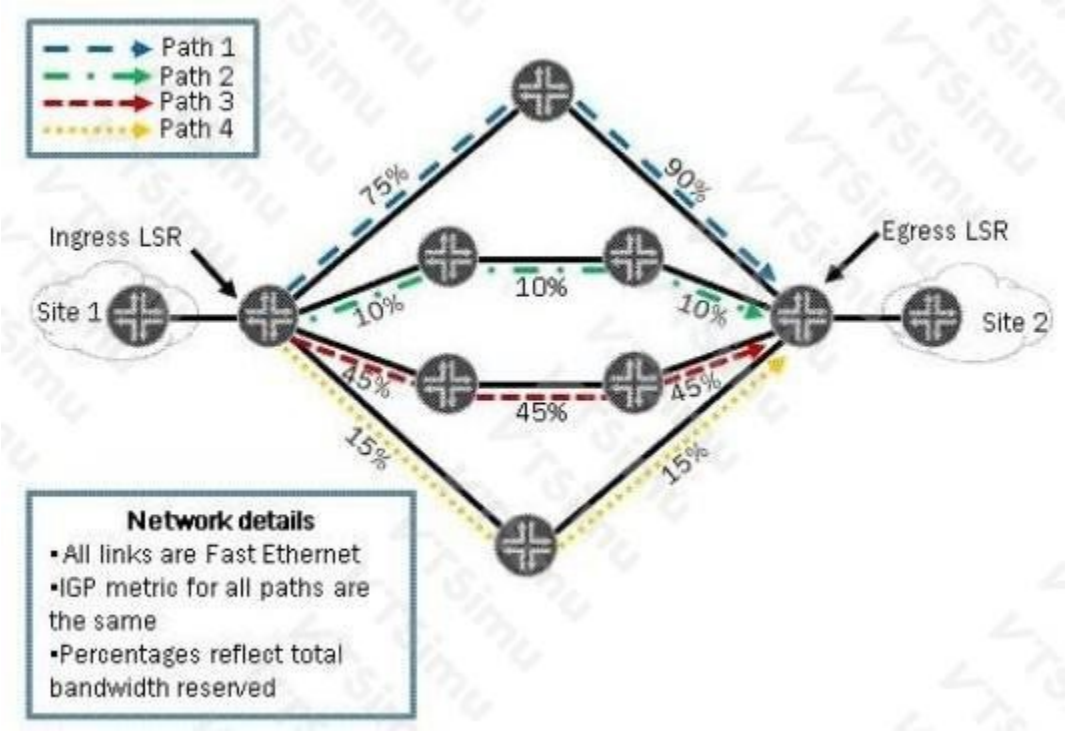

You have an MPLS network and you have configured least-fill as your CSPF tiebreaker.

Using the information in the exhibit, which path will be used to signal a new LSP requiring 12 Mbps?

A. Path 1

B. Path 2

C. Path 3

D. Path 4

ANSWER: D

Explanation:

Path 4 is selected. Constrained Shortest Path First first eliminates any candidate route that cannot satisfy the new LSP's 12-Mbps bandwidth requirement. It then evaluates the remaining feasible paths according to the configured path-selection behavior. When equal-cost candidate paths require a CSPF tie-break, the *least-fill* setting favors the path with the lowest existing bandwidth utilization. Equivalently, it spreads RSVP-TE reservations across available network resources rather than concentrating additional reservations on a link or path that is already more heavily filled.

Based on the bandwidth allocations shown in the exhibit, Path 4 is the eligible path with the least fill for the requested reservation. Selecting it preserves the more-used resources for traffic that may have fewer viable alternatives and supports better load distribution for subsequent RSVP-signaled LSPs. The reservation is therefore signaled along Path 4, assuming all links on that path continue to advertise sufficient reservable bandwidth at signaling time. Junos documentation describes CSPF as calculating an RSVP LSP route subject to constraints such as bandwidth and documents CSPF path-selection behavior in the MPLS traffic-engineering context: Juniper CSPF documentation and Juniper RSVP overview.

QUESTION NO: 3

Click the Exhibit button.

```
user@router# show routing-options multicast
scope 1 {
  prefix 224.0.1.39/32;
  interface fe-0/0/0.0;
}
```

Referring to the exhibit, which statement is correct?

A. Only multicasts packets (224.0.1.39) are allowed on the input and output direction.

B. Auto-RP discovery messages are filtered in the input and output direction.

C. Rendezvous point announcements are filtered in the output direction.

D. This filter does not work because the input or output parameter is missing.

ANSWER: C

Explanation:

Rendezvous point announcements are filtered in the output direction. The filter shown in the exhibit matches the well-known multicast destination address used for Auto-RP RP-Announce messages, 224.0.1.39, and the filter is applied in the output direction. As a result, RP announcement packets matching that address are prevented from being transmitted through the interface on which the filter is applied. This is useful when controlling the propagation scope of multicast control-plane information, ensuring that RP announcement traffic does not leave a designated multicast domain or administrative boundary.

The address assignment is significant: IANA identifies 224.0.1.39 as the Cisco-RP-Announce multicast address. Therefore, matching this destination specifically targets rendezvous point announcement traffic rather than ordinary multicast data flows. Junos firewall filters can be applied to traffic in a specified direction on an interface, allowing an operator to enforce this type of multicast-control filtering at the interface boundary. See the IANA IPv4 Multicast Address Space Registry for the address assignment and Juniper's firewall filter overview for direction-aware packet filtering behavior.

QUESTION NO: 4

Click the Exhibit button.

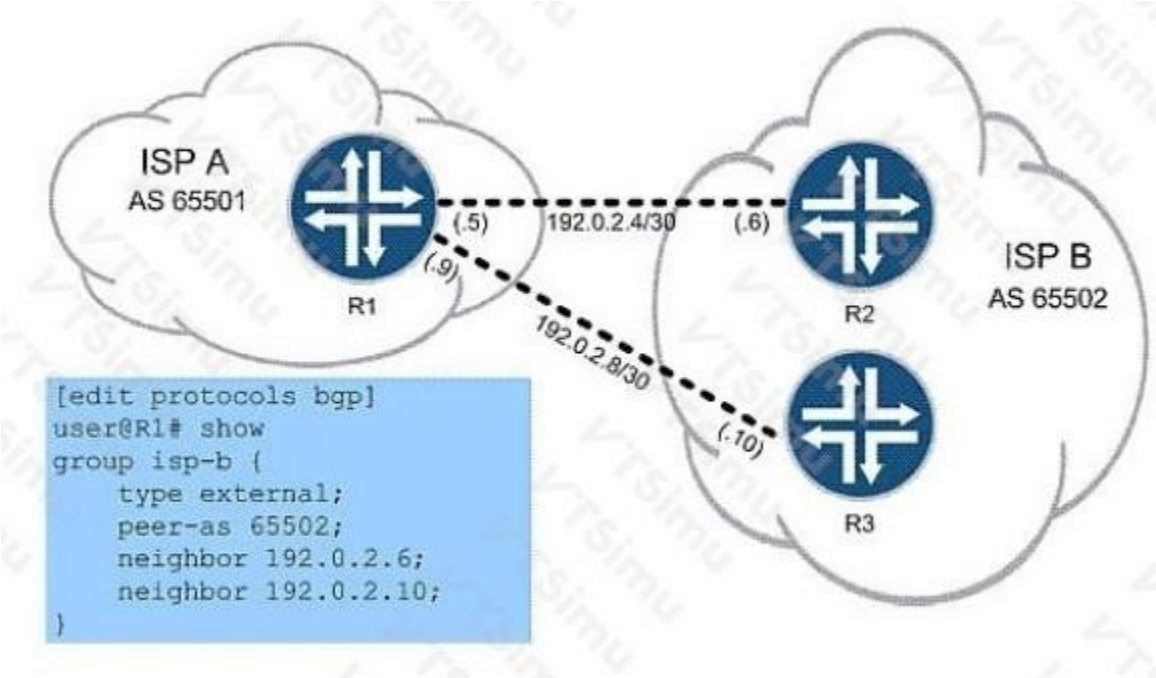

You work for ISP A, as shown in the exhibit, and must configure R1 to use load balancing across both available links for all routes to ISP B's network. You start by configuring this policy:

policy-statement load-balance {

then {

load-balance per-packet;

}

}

Which two commands do you use to finish the configuration? (Choose two.)

A. set protocols bgp group isp-b multihop

B. set routing-options forwarding-table export load-balance

C. set routing-options forwarding-table import load-balance

D. set protocols bgp group isp-b multipath

ANSWER: B D

Explanation:

`set protocols bgp group isp-b multipath` is required because BGP normally selects only one best path for a destination, even when it receives equivalent routes through separate external peerings. Enabling multipath at the BGP group level allows Junos to retain multiple eligible equal-cost BGP paths and install them for forwarding, provided the paths meet BGP multipath eligibility requirements. This gives R1 both links toward ISP B for the relevant destinations.

`set routing-options forwarding-table export load-balance` applies the previously defined forwarding-table export policy. The policy action `load-balance per-packet` instructs the Packet Forwarding Engine to distribute packets across the installed equal-cost next hops on a per-packet basis. Applying it as a forwarding-table *export* policy is the correct attachment point because the policy controls the forwarding-table treatment of routes after route selection, rather than importing routes into the routing table. Together, BGP multipath supplies the multiple usable BGP next hops and the forwarding-table export policy enables the requested per-packet load balancing behavior. Juniper documents BGP multipath

configuration at BGP Multipath Load Balancing and forwarding-table load-balancing policy use at Configuring Per-Packet Load Balancing.

QUESTION NO: 5

Refer to the exhibit.

```
user@router# show routing-instances
VRF {
    instance-type vrf;
    interface fe-1/2/3.0;
    route-distinguisher 65535:1;
    vrf-target target:65535:1;
}

user@router# run show route advertising-protocol bgp 4.4.4.4 extensive
VRF.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
* 172.16.1.0/30 (1 entry, 1 announced)
BGP group RR type Internal
    Route Distinguisher: 65535:1
    BGP label allocation failure: Need a nexthop address on LAN
    Nexthop: Not advertised
    Flags: Nexthop Change
    Localpref: 100
    AS path: I
    Communities: target:65535:1
```

What are two ways to resolve the error BGP label allocation failure: Need a next hop address on LAN? (Choose two.)

A. Configure a /32 static route with the next hop as itself to any remote address on the network assigned to the CE-facing interface.

B. Resolve the MPLS issue between the PE routers.

C. Configure a static ARP entry on fe-l/#3.0.

D. Configure vrf-table-label in the routing-instance.

ANSWER: A D

Explanation:

Configure a /32 static route with the next hop as itself to any remote address on the network assigned to the CE-facing interface. This creates the required host route and makes the remote next-hop address resolvable through the local LAN-

facing interface. Consequently, the routing system has a valid next-hop context when it needs to allocate and advertise a label for the relevant VPN route. A static route is configured under the routing-options hierarchy and can use a next hop that establishes this necessary recursive resolution behavior.

Configure vrf-table-label in the routing-instance. This causes Junos to allocate an MPLS label for the VRF routing table rather than requiring individual labeled forwarding behavior to depend on the unavailable LAN next-hop condition. Packets arriving with that table label are directed into the appropriate VRF for an IP lookup, which is particularly useful for VPN forwarding designs involving CE-facing interfaces. Juniper documents `vrf-table-label` as assigning a label to the VRF routing table, and documents static-route configuration and next-hop behavior in the routing protocol reference. See Juniper vrf-table-label statement reference and Juniper static routing documentation.

QUESTION NO: 6

What are three reasons an OSPF neighbor ship would be stuck in ExStart? (Choose three.)

A. The LSA database exchange is not yet completed.

B. There is an MTU mismatch between the OSPF routers.

C. There is an interface-type mismatch between the OSPF routers.

D. There is a unicast communication problem between the OSPF routers.

E. Both OSPF routers are using the same router ID.

ANSWER: B D E

Explanation:

An MTU mismatch between the OSPF routers can prevent Database Description packet negotiation from progressing. During this stage, a router includes its interface MTU in Database Description packets; if the receiving router cannot accept the advertised value, the adjacency commonly remains in ExStart. A unicast communication problem between the OSPF routers also directly affects this state. Although OSPF discovery and Hello traffic can succeed using multicast on a shared network, Database Description packets used for master/slave negotiation are sent to the neighbor's unicast address. Filtering, forwarding, or reachability problems affecting that traffic can therefore prevent the transition to Exchange. Both OSPF routers using the same router ID also prevents stable master/slave selection, because router ID is used during the negotiation process to determine which neighbor is master. These conditions interrupt the Database Description negotiation that defines ExStart, before the routers can reliably begin exchanging the contents of their link-state databases. RFC 2328 describes neighbor-state progression, Database Description exchanges, and router-ID-based master selection in OSPF Version 2. Juniper's operational OSPF guidance is available in the OSPF troubleshooting documentation.

QUESTION NO: 7

Click the Exhibit button.

```
user@router> show route table inet.0 protocol ospf
inet.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
224.0.0.5/32       *[OSPF/10] 1w2d 21:43:55, metric 1
                      MultiRecv

user@router> show ospf database
    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   67.176.255.1     67.176.255.1     0x80000176  1138  0x22 0x3f8e  48
Router   67.176.255.2     67.176.255.2     0x80000151  1131  0x22 0xa349  48
Router  *67.176.255.3     67.176.255.3     0x80000152  1137  0x22 0xa733  72
Router   67.176.255.4     67.176.255.4     0x80000167  1138  0x22 0x9937  48
Network  67.176.10.2      67.176.255.1     0x80000007  1138  0x22 0x7ae2  40
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   10.128.0.0       67.176.255.2     0x8000013a  2118  0x22 0xda22  36
Extern   10.128.55.32     67.176.255.2     0x8000013a  1624  0x22 0xe1f   36
Extern   10.128.128.0     67.176.255.2     0x8000013a   628  0x22 0xd5a6  36
Extern   10.128.128.32    67.176.255.2     0x8000013a   127  0x22 0xe7fb  36

user@router> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
67.176.10.5      ge-0/0/8.0             Full      67.176.255.4     128    37
67.176.10.2      ge-0/0/8.0             Full      67.176.255.1     128    39
67.176.10.3      ge-0/0/8.0             ExStart   67.176.255.2       1    36

user@router> show ospf interface ge-0/0/8.0 extensive
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/8.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            3
  Type: P2MP, Address: 67.176.10.4, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 2
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

Based on the exhibit, why is the router not installing routes in the OSPF database into the routing table?

A. The neighbor is stuck in the ExStart state

B. The routes are going to a different table than inet.0

C. There is an interface type mismatch

D. There is an MTU mismatch

ANSWER: C

Explanation:

There is an interface type mismatch. OSPF interface type determines how the protocol models the attached network and how it forms and maintains adjacencies. For example, a broadcast interface uses designated-router behavior and network LSAs, whereas a point-to-point interface uses a different link model. The OSPF peers on a link must use compatible network types so that the received link-state information can be correctly associated with a usable next hop and incorporated into SPF results. When the interface types do not match, the router can retain or display LSDB information while being unable to produce valid, installable OSPF routes for the affected topology. Consequently, the prefixes visible in the OSPF database do not appear in the routing table as active OSPF routes. The remedy is to verify the OSPF interface configuration on both ends of the link and configure matching interface types appropriate for the underlying medium, such as point-to-point for a direct router-to-router link. Juniper documents OSPF interface types and their operational behavior in its OSPF interface configuration reference and OSPF interface types documentation.

QUESTION NO: 8

You want to copy only OSPF prefixes from site.inet.0 to company.inet.0.

```
[edit routing-options]
user@router# show
interface-routes {
    rib-group inet site-to-company;
}
rib-groups {
    site-to-company {
        import-rib [ site.inet.0 company.inet.0 ];
        import-policy site-to-company;
    }
}
```

What is the proper configuration syntax you must use to accomplish this task?

A.

```
[edit]
user@router# show policy-options policy-statement site-to-company
term ospf {
    from protocol ospf;
    then accept;
}
term default {
    then reject;
}
[edit routing-options]
user@router# show
interface-routes {
    rib-group inet site-to-company;
}
rib-groups {
    site-to-company {
        export-rib [ site.inet.0 company.inet.0 ];
        import-policy site-to-company;
    }
}
```

B.

```
[edit]
user@router# show policy-options policy-statement site-to-company
term ospf {
    from protocol ospf;
    then accept;
}
term default {
    then reject;
}
[edit routing-options]
user@router# show
interface-routes {
    rib-group inet site-to-company;
}
rib-groups {
    site-to-company {
        import-rib [ company.inet.0 site.inet.0 ];
rib-groups {
    site-to-company {
        import-rib [ company.inet.0 site.inet.0 ];
        import-policy site-to-company;
    }
```

C.

```
[edit]
user@router# show policy-options policy-statement site-to-company
term ospf {
    from protocol ospf;
    then accept;
}
term default {
    then reject;
}
[edit routing-options]
user@router# show
rib-groups {
    site-to-company {
        import-rib [ site.inet.0 company.inet.0 ];
        import-policy site-to-company;
    }
}
```

D.

```
[edit]
user@router# show policy-options policy-statement site-to-company
term ospf {
    from protocol ospf;
    then accept;
}
term default {
    then reject;
}
```

A. Option A

B. Option B

C. policy-options { policy-statement ospf-only { term ospf { from protocol ospf; then accept; } } } routing-options { rib-groups { site-to-company { import-rib [site.inet.0 company.inet.0]; import-policy ospf-only; } } } protocols { ospf { rib-group site-to-company; } }

D. Option D

ANSWER: C

Explanation:

The configuration that defines a RIB group with `site.inet.0` as the primary import RIB and `company.inet.0` as the secondary import RIB, applies an OSPF-only import policy, and attaches that RIB group to OSPF is correct. In Junos, the first routing table in an `import-rib` list is the primary table into which the routing protocol installs its routes. Subsequent tables are secondary tables that receive copies of the eligible routes. Therefore, placing `site.inet.0` first preserves it as the OSPF route destination, while `company.inet.0` receives the selected copies.

The RIB group `import-policy` evaluates routes before they are installed in secondary RIBs. A policy term matching `from protocol ospf` and accepting those routes ensures that only OSPF-learned prefixes are copied. Associating the RIB group under the OSPF protocol hierarchy causes this behavior to apply to routes learned by that OSPF instance. Juniper documents RIB groups, including primary and secondary import-RIB behavior, at RIB groups CLI reference. Policy matching based on route protocol is documented in the Junos routing policy match conditions documentation.

QUESTION NO: 9

Click the Exhibit button.

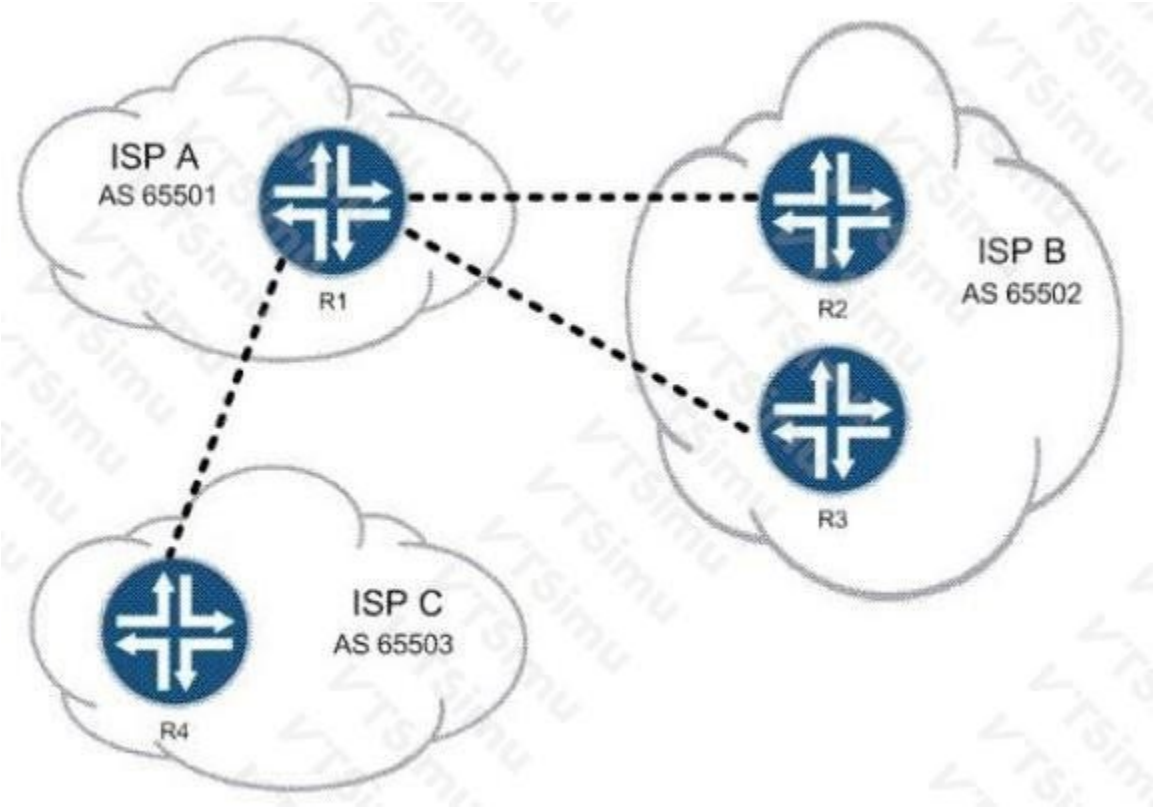

Your employer is ISP

A. Use policy to filter routes on AS number.

B. Use the well-known no-export community.

C. Use the MED to prefer the proper routes.

D. Use communities to identify and filter routes.

ANSWER: A D

Explanation:

Use policy to filter routes on AS number and use communities to identify and filter routes are both valid ways to enforce the no-transit requirement while preserving customer connectivity. An import or export routing policy can match BGP AS paths that identify routes learned from ISP B or ISP C. The policy can then prevent routes learned from one upstream from being advertised to the other upstream. Customer-originated routes can still be advertised to both providers, and customer prefixes can still receive routes from both providers, allowing the ISP's customers to reach destinations behind each external ISP without making the local autonomous system a transit path between them.

BGP communities provide a scalable policy-control mechanism for the same requirement. On import, the ISP can tag routes according to their source, such as routes received from ISP B or ISP C. Export policies can then match those tags and suppress advertisement toward the opposite upstream while allowing the appropriate advertisements toward customers. This avoids relying solely on a specific AS-path structure and is especially useful when policy needs to be applied consistently across multiple border routers or external peers. Junos routing policy supports matching AS-path information directly, and it supports attaching and matching BGP communities for route-control decisions. See Juniper's AS-path routing-policy documentation and BGP community policy documentation.

QUESTION NO: 10

What are two elements that must be configured on an interface to enable RSVP signaling for MPLS LSPs? (Choose two.)

A. Add the relevant interfaces under the [edit protocols rsvp] hierarchy.

B. Add the family mpls statement to the relevant interfaces under the [edit interfaces] hierarchy.

C. Add the relevant interfaces under the [edit protocols ldp] hierarchy.

D. Add the family iso statement to the relevant interfaces under the [edit interfaces] hierarchy.

ANSWER: A B

Explanation:

The interface must have the **family mpls** protocol family configured so that it can accept and forward labeled MPLS traffic. RSVP-signaled MPLS LSPs require MPLS packet processing on every participating transit interface.

The interface must also be enabled under the **[edit protocols rsvp]** hierarchy. This activates RSVP on the interface, allowing it to exchange RSVP Path and Resv messages and to maintain RSVP reservations. In a typical MPLS traffic-engineering design, the same interfaces also participate in an IGP with traffic-engineering extensions enabled. See the Junos RSVP configuration documentation and the Junos MPLS configuration documentation.

QUESTION NO: 11

Click the Exhibit button.

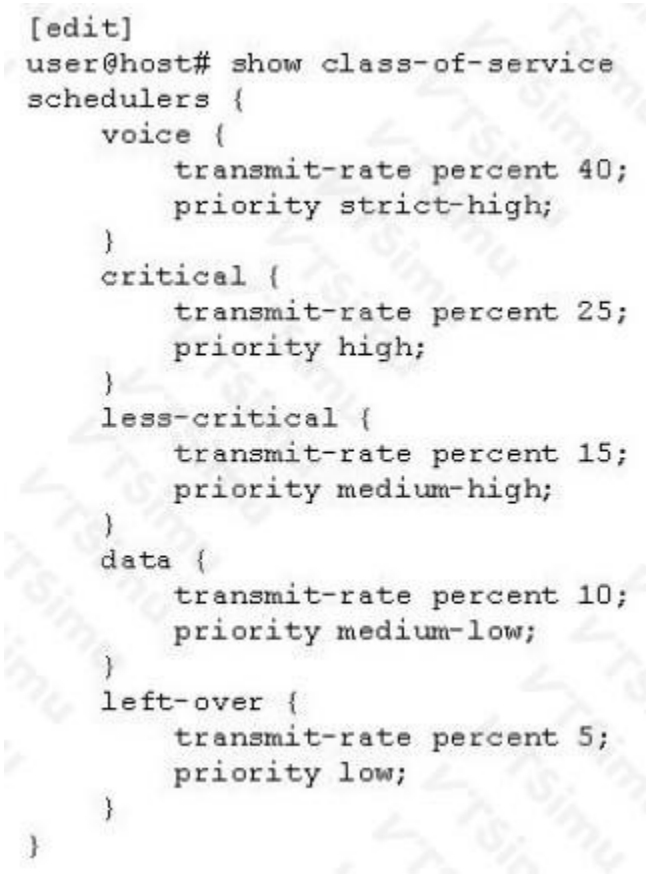

```
[edit]
user@host# show class-of-service
schedulers {
    voice {
        transmit-rate percent 40;
        priority strict-high;
    }
    critical {
        transmit-rate percent 25;
        priority high;
    }
    less-critical {
        transmit-rate percent 15;
        priority medium-high;
    }
    data {
        transmit-rate percent 10;
        priority medium-low;
    }
    left-over {
        transmit-rate percent 5;
        priority low;
    }
}
```

On your MX Series router, traffic using the critical scheduler is out of profile. All other data is currently in profile. Referring to the exhibit, which statement is correct?

A. The critical queue is serviced before the less-critical queue.

B. The critical queue is serviced after the left-over queue.

C. The critical queue is serviced before the data queue.

D. The critical queue is serviced before the voice queue.

ANSWER: B

Explanation:

The critical queue is serviced after the left-over queue. In Junos CoS scheduling, a scheduler's configured transmit rate represents its committed, or in-profile, service allocation. As long as traffic in another scheduler remains in profile, that scheduler is entitled to its configured bandwidth before bandwidth is distributed to traffic that has exceeded its profile. Because the critical scheduler is explicitly stated to be out of profile while the other traffic is in profile, critical traffic is competing only for excess capacity after the in-profile scheduler commitments have been honored. The left-over scheduler is among the traffic that remains in profile in the scenario, so it is serviced before the out-of-profile critical scheduler. Scheduler priority governs distribution in the relevant contention stage, but it does not allow an out-of-profile scheduler to displace bandwidth that is still required to satisfy in-profile scheduler allocations. This behavior lets Junos preserve guaranteed service commitments while still permitting an over-profile queue to use capacity when it becomes available. Juniper's scheduler overview and scheduler configuration documentation describe the relationship between transmit rates, excess bandwidth, and scheduler service: Junos CoS schedulers overview and Configuring CoS schedulers.

QUESTION NO: 12

Click the Exhibit button.

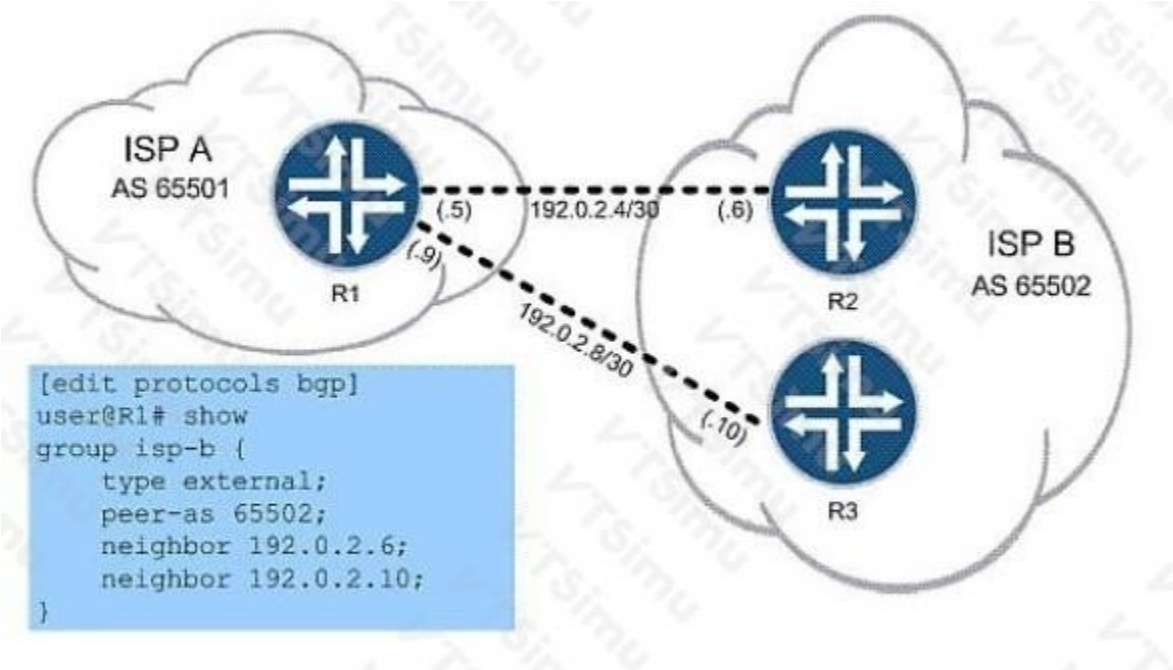

Referring to the exhibit, you work for ISP A and are asked to configure R1 to forward traffic for all routes across both available links, to both routers in ISP B's network. Which three configuration commands do you use? (Choose three.)

A. set protocols bgp group isp-b multihop

B. set policy-options policy-statement load-balance then load-balance per-packet

C. set routing-options forwarding-table import load-balance

D. set protocols bgp group isp-b multipath

E. set routing-options forwarding-table export load-balance

ANSWER: B D E

Explanation:

The correct configuration combines BGP multipath route installation with a forwarding-table policy that instructs the Packet Forwarding Engine to use the equal-cost next hops on a per-packet basis. `set protocols bgp group isp-b multipath` permits BGP to retain and install multiple equivalent paths learned from the ISP B peers rather than selecting only one BGP best path. This provides R1 with eligible next hops for both external links. The policy statement containing `then load-balance per-packet` defines the desired forwarding behavior: packets can be distributed across the installed equal-cost next hops. Applying that policy with `set routing-options forwarding-table export load-balance` exports the policy to forwarding-table processing, allowing the forwarding plane to program and use both paths. Together, these commands ensure that all eligible destinations learned through the two ISP B routers can be forwarded using both available links. Juniper documents BGP multipath behavior and configuration in its BGP load balancing documentation, and documents forwarding-table export policies in the forwarding-table policy documentation.

QUESTION NO: 13

Which two statements are true about the PIM designated router (DR) on a multiaccess network? (Choose two.)

A. A DR is elected separately on each PIM-enabled multiaccess interface.

B. The DR sends PIM Register messages for directly connected sources.

C. The DR is always the multicast rendezvous point.

D. The DR is elected using the lowest interface IP address.

ANSWER: A B

Explanation:

The PIM DR is elected separately for each PIM-enabled multiaccess interface. The election is based on the highest PIM DR priority, with the highest IP address used as the tie-breaker when priorities are equal. The elected router performs functions that require one router to act on behalf of the shared LAN.

For a directly connected active source in a PIM sparse-mode deployment, the DR encapsulates the source traffic in PIM Register messages and sends the messages to the rendezvous point. This lets the RP learn about the source and begin constructing multicast forwarding state. See RFC 7761 for PIM-SM DR election and Register procedures.

QUESTION NO: 14

You are asked to set up a route reflection cluster for a group of IBGP peers that are fully meshed.

You want to only reflect routes that arrive outside the cluster to the IBGP peers.

Which route reflector configuration accomplishes this task?

A.

```
[edit protocols bgp]
user@router# show
group int-peers {
    type internal;
    local-address 172.16.1.1;
    cluster 172.16.1.1;
    advertise-external;
    neighbor 172.16.2.2;
    neighbor 172.16.3.3;
}
```

C.

```
[edit protocols bgp]
user@router# show
group int-peers {
    type internal;
    local-address 172.16.1.1;
    cluster 172.16.1.1;
    no-client-reflect;
    neighbor 172.16.2.2;
    neighbor 172.16.3.3;
}
```

B.

```
[edit protocols bgp]
user@router# show
group int-peers {
    type internal;
    local-address 172.16.1.1;
    cluster 172.16.1.1;
    export next-hop-self;
    neighbor 172.16.2.2;
    neighbor 172.16.3.3;
}
```

D.

```
[edit protocols bgp]
user@router# show
group int-peers {
    type internal;
    local-address 172.16.1.1;
    cluster 172.16.1.1;
    damping;
    neighbor 172.16.2.2;
    neighbor 172.16.3.3;
}
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: A

Explanation:

The depicted configuration in *Option A* is correct because it uses the BGP `no-client-reflect` behavior for the route-reflector cluster. A route reflector normally reflects routes learned from one route-reflector client to its other clients, as well as advertising appropriate routes between clients and nonclients. Enabling `no-client-reflect` changes that client-to-client behavior: routes received from a client within the cluster are not reflected to other clients. The reflector still advertises routes

learned from nonclient peers—peers outside the route-reflector client cluster—to its clients. This precisely provides the requested policy of reflecting only routes that arrive from outside the cluster to the IBGP peers, while retaining the route reflector as the distribution point for externally learned routes. The setting is useful where the client IBGP peers are already fully meshed and therefore do not need the route reflector to redistribute client-originated paths among themselves. Junos documents `no-client-reflect` as preventing advertisements of routes received from a client to other clients; this preserves reflection of routes received from nonclients toward clients. See the Junos no-client-reflect statement reference and Juniper's BGP route reflection documentation.

QUESTION NO: 15

Which two statements are true about OSPF virtual links? (Choose two.)

A. A virtual link is configured between two ABRs.

B. A virtual link can use a stub area as its transit area.

C. A virtual link provides logical connectivity to Area 0.

D. A virtual link replaces the need for Area 0 in an OSPF domain.

ANSWER: A C

Explanation:

An OSPF virtual link is configured between two Area Border Routers. It provides a logical connection through a nonbackbone transit area to restore connectivity to Area 0 when a physical backbone connection is not available.

The transit area for a virtual link cannot be a stub area. A virtual link must carry the OSPF routing information required for backbone operation, while stub areas restrict the flooding of external routing information. Virtual links should generally be avoided in new designs when a physical or tunneled backbone connection can be provided, but they remain available for specific topology requirements.

See the Juniper OSPF area configuration documentation and RFC 2328.

QUESTION NO: 16

Click the Exhibit button.

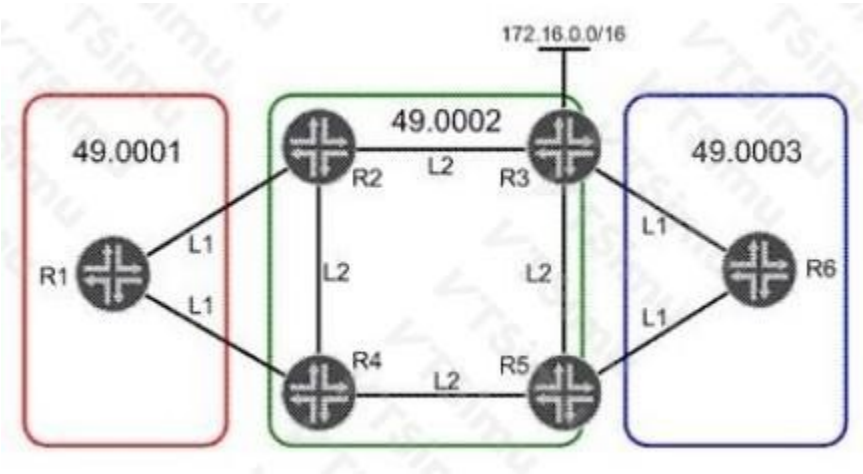

In the exhibit, network 172.16.0.0/16 is redistributed into IS-IS in Area 49.0002. R1 must use R2 to access 172.16.0.0/16. All other traffic leaving Area 49.0001 must use R4.Which three steps will accomplish this task? (Choose three.)

A. Configure R1 to ignore the attached bit.

B. Disable the attached bit on R4 in Area 49.0001.

C. Enable an L2 adjacency on the link between R1 and R2.

D. Leak network 172.16.0.0/16 into L1 on R2.

E. Redistribute a static default route into L1 on R4.

ANSWER: A D E

Explanation:

Configure R1 to ignore the attached bit, leak network 172.16.0.0/16 into L1 on R2, and redistribute a static default route into L1 on R4. The attached-bit mechanism normally causes an L1 router to use an L1/L2 router as its default exit from the area. In this topology, R1 must not use that automatically derived default behavior, because doing so can select R2 as the exit for traffic other than the specific redistributed prefix. Ignoring the attached bit on R1 removes this implicit default-route decision.

Leaking 172.16.0.0/16 from R2's Level 2 database into its Level 1 database advertises a specific, reachable L1 route to R1. Longest-prefix matching then ensures that traffic for 172.16.0.0/16 follows R2. R4 should instead originate an explicit default route into Level 1 by redistributing its static default route. This supplies R1 with a deliberate L1 default route for all destinations that do not match the leaked 172.16.0.0/16 route, thereby directing all other interarea-bound traffic through R4. Junos IS-IS policy can control route export by level, including exporting routes into Level 1. See Juniper's IS-IS routing policy documentation and IS-IS overview.

QUESTION NO: 17

You are setting up MPLS RSVP LSPs between R1 and R6 through your core network. You must ensure that the R1 has redundant ERO paths. You must also ensure that both paths are signaled and ready for traffic.

Which action will accomplish these requirements?

A. Create two primary paths.

B. Create a primary path and create a secondary path.

C. Create a primary path and create a secondary path with the standby parameter.

D. Create a primary path and create a secondary path with the active parameter.

ANSWER: C

Explanation:

Create a primary path and create a secondary path with the standby parameter. In Junos, an RSVP label-switched path can have explicitly routed primary and secondary paths, allowing each path to use its own configured explicit route object. The primary path carries traffic during normal operation. Configuring the secondary path as `standby` causes RSVP to signal that backup path and reserve its resources while the primary remains operational. Therefore, the alternate path is already established and can take over promptly if the primary path becomes unavailable, rather than requiring RSVP signaling only after a failure occurs.

This design provides both requested properties: redundant ERO-based path selection and a pre-signaled backup that is ready for traffic. The behavior is particularly useful where fast restoration and resource availability are required, because the ingress router has an established alternate LSP available when it must switch away from the active primary route. Juniper documents RSVP LSP path configuration and the use of primary, secondary, and standby paths in its MPLS configuration guidance: Juniper: Configuring MPLS Paths and Juniper: RSVP Overview.

QUESTION NO: 18

What must you do to use the Layer 2 interworking interface to connect a BGP Layer 2 VPN with an LDP Layer 2 circuit? (Choose two.)

A. You must configure a policy importing the bgp.l2vpn.O routing table into the inet.0 routing table.

B. You must enable the l2iw protocol.

C. You must have a Tunnel Services PIC installed to used the L2 interworking interface.

D. You must include the iw0 statement under the [edit interfaces] hierarchy.

ANSWER: B D

Explanation:

To connect a BGP-signaled Layer 2 VPN to an LDP-signaled Layer 2 circuit through Junos Layer 2 interworking, you must enable the `l2iw` protocol and configure the logical Layer 2 interworking interface, `iw0`. The `protocols l2iw` hierarchy activates the Layer 2 interworking feature and is where the interworking configuration associates the BGP Layer 2 VPN side with the LDP Layer 2 circuit side. This feature provides the required adaptation between the distinct control-plane signaling methods while forwarding Layer 2 traffic across the interconnected services. The `iw0` interface is the dedicated interworking interface and must be declared under the `interfaces` hierarchy before it can be referenced by the Layer 2 interworking configuration. In practical configuration terms, creating the service requires both the protocol feature to be enabled and the `iw0` interface to exist. Juniper's Layer 2 VPN documentation describes Layer 2 interworking as the mechanism for connecting BGP Layer 2 VPNs and LDP Layer 2 circuits; see the Junos Layer 2 VPN documentation. The relevant configuration hierarchy and syntax are also covered in the Junos OS CLI Reference.

QUESTION NO: 19

You want to use IS-IS on a GRE interface where the underlying Layer 3 MTU is 1500. Which statement is correct?

A. IS-IS can be used because every IS-IS interface must be capable of transmitting packets at least as large as 1476 bytes, and the GRE header is 24 bytes.

B. IS-IS cannot be used because the IS-IS hello is not allowed to be fragmented and has the DF bit set.

C. IS-IS can be used, but the networking device directly attached to the circuit must be capable of fragmentation.

D. IS-IS cannot be used, but the router can enable a GRE key that serves the same function as IS-IS.

ANSWER: C

Explanation:

IS-IS can be used, but the networking device directly attached to the circuit must be capable of fragmentation. IS-IS requires implementations to support PDUs of at least 1492 bytes. A GRE tunnel transported across an IPv4 network with a 1500-byte Layer 3 MTU adds encapsulation overhead: a minimum 20-byte outer IPv4 header plus a 4-byte GRE header. Consequently, the largest unfragmented GRE payload is 1476 bytes, which is smaller than the IS-IS minimum PDU size requirement. An IS-IS PDU that reaches the required size can still be sent through the GRE tunnel when the directly connected device can fragment the outer IP packet for transmission across the 1500-byte underlay. The receiving endpoint reassembles the outer IP fragments before GRE decapsulation, allowing the original IS-IS PDU to be delivered intact. This depends on normal IPv4 fragmentation behavior being available on the GRE transport path; it is not an IS-IS PDU-level fragmentation mechanism. RFC 1195 specifies the IS-IS PDU size capability requirement, while Juniper's GRE documentation describes GRE as IP encapsulation carried over an underlying IP network. See RFC 1195 and Juniper GRE configuration documentation.

QUESTION NO: 20

Click the Exhibit button.

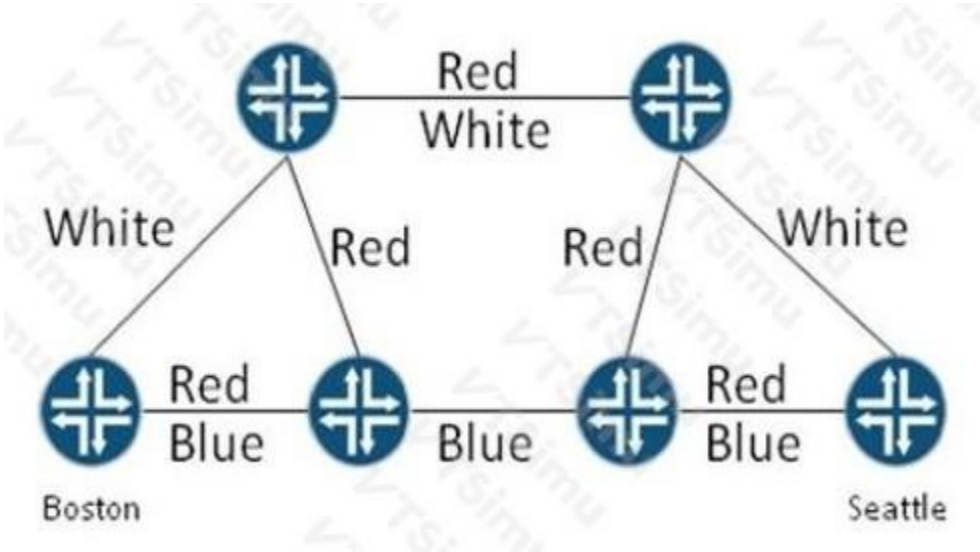

On the network shown in the exhibit, a network administrator is attempting to bring up an LSP between Boston and Seattle using administrative groups.

Which two of the following LSP configurations allow the LSP to establish? (Choose two.)

A.
```
[edit protocols mpls label-switched-path Boston-to-Seattle]
user@Boston# show
to 192.168.10.100;
admin-group {
     include-any White;
     exclude Red;
}
```

B.
```
[edit protocols mpls label-switched-path Boston-to-Seattle]
user@Boston# show
to 192.168.10.100;
admin-group include-all [ Red White Blue ];
```

C.
```
[edit protocols mpls label-switched-path Boston-to-Seattle]
user@Boston# show
to 192.168.10.100;
admin-group {
     include-any [ Red Blue ];
     include-all Blue;
}
```

D.
```
[edit protocols mpls label-switched-path Boston-to-Seattle]
user@Boston# show
to 192.168.10.100;
admin-group {
     include-any Red;
     include-all Blue;
}
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: C D

Explanation:

Administrative groups, also called link colors, are RSVP-TE/MPLS traffic-engineering constraints that control which links are eligible during constrained shortest-path first computation. A link can be assigned one or more administrative-group bits, and an LSP can request that specific group bits be included, excluded, or avoided. For the Boston-to-Seattle LSP to establish, the configured constraint must leave at least one continuous eligible path from the ingress to the egress. The two marked configurations do this: their administrative-group matching criteria permit a complete end-to-end route through the topology while honoring the requested group behavior. Junos evaluates the administrative-group constraints against the link attributes advertised for the TE topology, then signals the resulting RSVP LSP over the selected path. This capability is useful for deliberately selecting links with a desired transport characteristic or separating LSPs onto distinct classes of infrastructure. Juniper documents administrative groups as TE link attributes that can be used with LSP include/exclude constraints; the relevant configuration and operational behavior are described in the Juniper administrative groups documentation. RSVP-TE path-selection constraints are also defined by the RSVP-TE specification (RFC 3209).

QUESTION NO: 21

Click the Exhibit button.

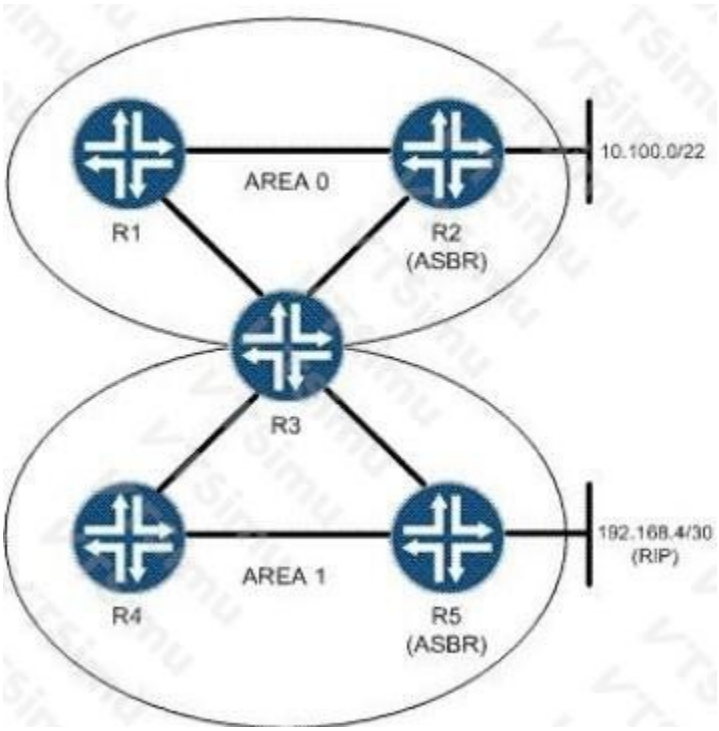

You are asked to configure an OSPF network based on the topology shown in the exhibit. You must keep the link-state database in Area 1 as small as possible. What will accomplish this?

A. Area 0 should be configured as a stub area so that it will not announce routes into Area 1.

B. Area 1 should be configured as an NSSA to limit the size of the link-state database.

C. Area 1 should be configured as a stub area with no-summaries to limit the size of the link-state database.

D. Area 0 should be configured with a virtual link to R4 to limit the size of the Area 1 link-state database.

ANSWER: C

Explanation:

Area 1 should be configured as a stub area with no-summaries to limit the size of the link-state database. In Junos, configuring an area as a stub area with the `no-summaries` setting creates a totally stub area behavior. The ABR does not flood interarea summary LSAs into that area and instead injects a default route for destinations outside the area. This means routers within Area 1 retain the intra-area link-state information needed to calculate paths within their own area, while avoiding individual topology and reachability advertisements for networks in other OSPF areas. Stub-area operation also prevents external LSAs from being flooded into the area. Together, suppression of both interarea summaries and external reachability information minimizes the LSDB that Area 1 routers must maintain, while preserving a route toward the ABR for nonlocal destinations. Juniper documents the OSPF stub-area configuration, including the `no-summaries` option, at Juniper OSPF stub statement documentation. The underlying OSPF area and summary-LSA behavior is defined in RFC 2328.

QUESTION NO: 22

The network design team has decided to activate multicast in the network. Auto-RP has been selected as the RP mechanism. Which PIM operational mode must be enabled in this network?

A. sparse mode

B. sparse-dense mode

C. dense mode

D. source specific multicast

ANSWER: B

Explanation:

sparse-dense mode is required for classic Auto-RP operation because Auto-RP distributes rendezvous-point information using multicast discovery and announcement packets. Candidate RPs send RP-announce messages to the well-known multicast group 224.0.1.39, and mapping agents advertise the selected RP-to-group mappings to 224.0.1.40. These groups are in the 224.0.1.x control range and the Auto-RP messages must be flooded through the multicast domain so routers can learn RP mappings before they have group-specific sparse-mode state.

PIM sparse-dense mode supplies this required behavior: groups for which an RP mapping is known can operate using sparse-mode, shared-tree behavior, while groups without an RP mapping can use dense-mode flooding behavior. Consequently, the Auto-RP control groups remain reachable during RP discovery, allowing routers to receive and install the mappings needed for normal sparse multicast operation. This was the traditional deployment model for Auto-RP; networks using pure sparse mode commonly require an Auto-RP listener mechanism or an alternative RP-discovery method. Juniper's PIM documentation describes PIM operating modes and RP-based multicast operation: Juniper PIM Protocol Overview. Cisco's Auto-RP configuration documentation also identifies sparse-dense mode as the standard mode for Auto-RP operation: Cisco Auto-RP Overview.

QUESTION NO: 23

Click the Exhibit button

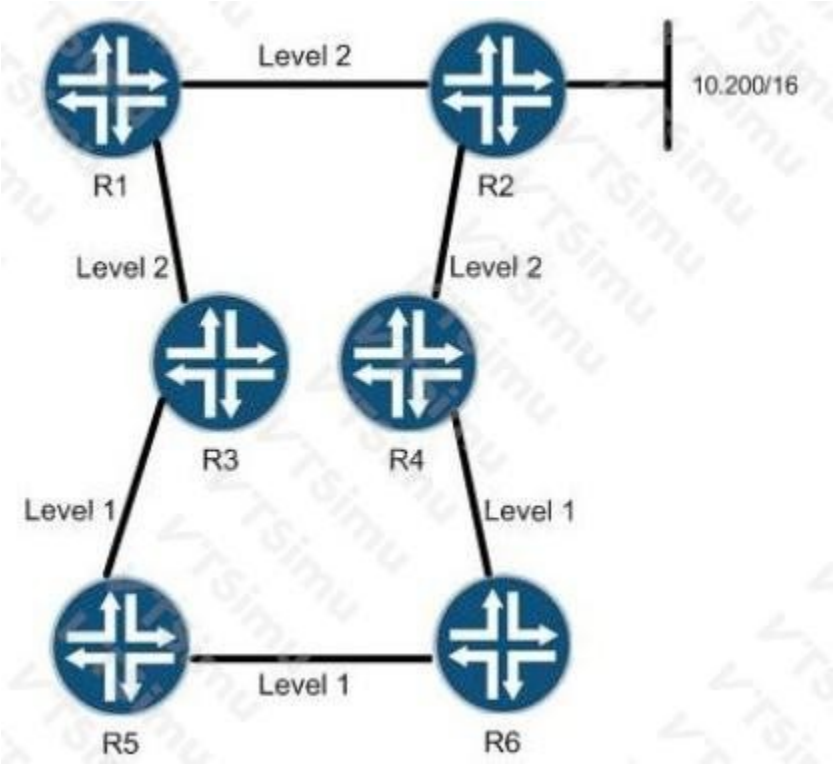

The 10.200/16 network is announced as an IS-IS route by R2 to its IS-IS neighbors. R3 and R4 are configured with an IS-IS export policy, which announces this route to R5 and R6.

Which statement is true?

A. When viewed on R5 the 10.200/16 route will be marked down.

B. When viewed on R5 the 10.200/16 route will be marked up.

C. The 10.200/16 route will not be visible on R5.

D. The 10.200/16 route will be marked with the overload bit.

ANSWER: A

Explanation:

When viewed on R5 the 10.200/16 route will be marked down. The export performed by R3 and R4 is a Level 2-to-Level 1 route leak. IS-IS uses the up/down bit in the IP reachability information to identify prefixes that have been leaked downward from the Level 2 area into a Level 1 area. Consequently, R5 receives the 10.200/16 prefix with the down bit set.

This marking is a loop-prevention mechanism. A Level 1/Level 2 router can recognize that the prefix originated from Level 2 and was subsequently advertised into Level 1. It therefore avoids advertising that same leaked reachability back upward into Level 2, where it could create an interlevel routing loop or unnecessary duplicate reachability information. The bit affects the route advertisement semantics; it does not mean that the prefix is unusable for forwarding by R5. R5 can install and use the route while retaining the down-bit attribute associated with the leaked IS-IS advertisement.

The IS-IS IP reachability encoding and up/down-bit behavior are defined in RFC 1195, Section 3.10.1. Juniper's IS-IS configuration and operational documentation is available in the Junos IS-IS documentation.

QUESTION NO: 24

You have established an MSDP peering between multicast domains for your AS and a neighboring AS. You are concerned about unnecessary flooding or looping of source active messages between MSDP peers. Which MSDP mechanism is used to prevent this problem?

A. The receiving MSDP peer performs a multicast RPF check to ensure that the SA messages are forwarded away from only the originating multicast source.

B. The receiving MSDP peer sends register-stop messages towards the originating RP that forwarded the SA messages.

C. The receiving MSDP peer performs a multicast peer-RPF check to ensure that SA messages are accepted only from the path toward the originating RP.

D. The receiving MSDP peer sends prune messages to all upstream neighbors to avoid receiving additional SA messages.

ANSWER: C

Explanation:

The receiving MSDP peer performs a multicast peer-RPF check to ensure that SA messages are accepted only from the path toward the originating RP. This is the MSDP peer reverse-path-forwarding mechanism, commonly called peer-RPF. Each Source-Active message identifies the originating rendezvous point, and a receiving MSDP router verifies that the message arrived from the MSDP peer that is the expected next hop when routing toward that originating RP. If the SA message arrives from a different peer, the router treats it as having arrived on an invalid reverse path and does not accept it for further propagation.

This check is specifically designed for interdomain MSDP topologies, where multiple peering paths can otherwise cause the same SA information to circulate repeatedly. By tying acceptance of an SA message to the unicast route toward the originating RP, peer-RPF establishes a loop-free direction for SA-message distribution and limits needless replication. It also adapts to routing changes because the expected MSDP peer is determined from the current route to the originating RP. The behavior and purpose of MSDP peer-RPF are defined in RFC 3618, Multicast Source Discovery Protocol; Juniper's multicast documentation is available through the Juniper Documentation portal.

QUESTION NO: 25

Click the Exhibit button.

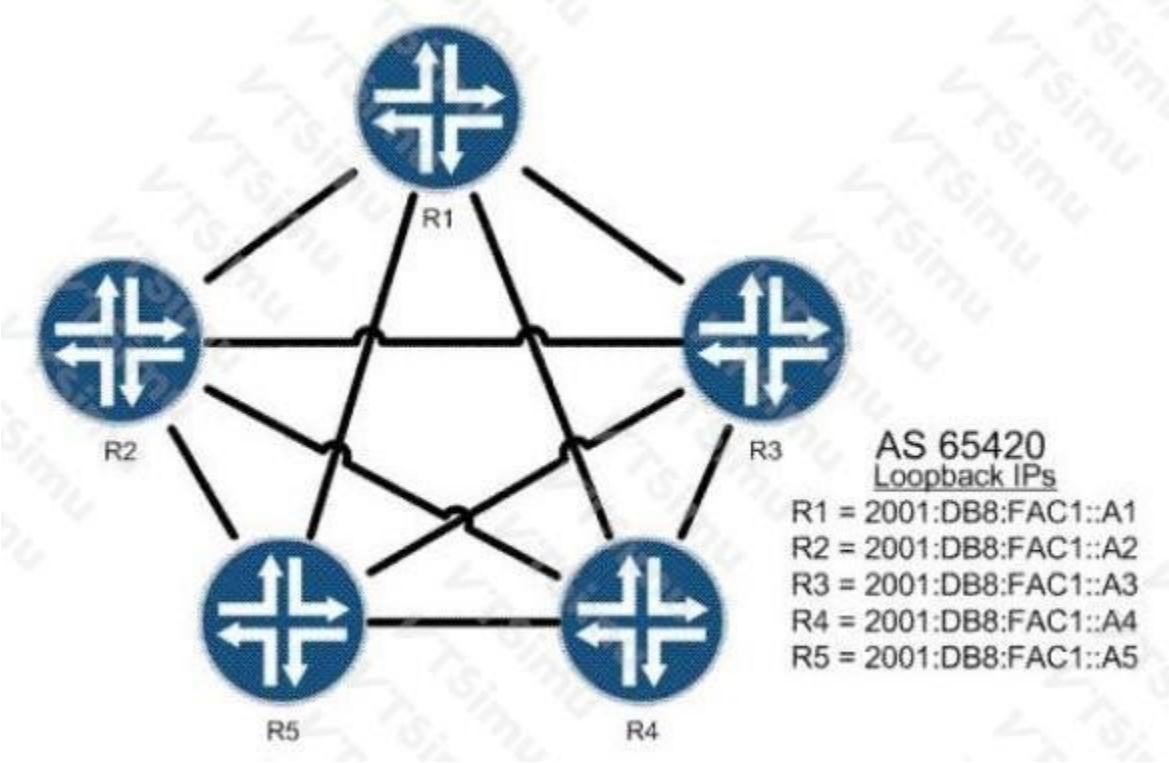

All routers shown in the exhibit are BGP neighbors. R1 must be a route reflector, and R2 through R5 must be clients. R2 should only receive one copy of all routes sent from R4.

Which configuration is valid?

A.
```
[edit protocols bgp]
root@R3# show
group AS65420 {
    type reflector;
    local-address 2001:db8:fac1::a3;
    neighbor 2001:db8:fac1::a1;
    neighbor 2001:db8:fac1::a2;
    neighbor 2001:db8:fac1::a4;
    neighbor 2001:db8:fac1::a5;
}
```

B.
```
[edit protocols bgp]
root@R1# show
group AS65420 {
    type internal;
    local-address 2001:db8:fac1::a1;
    cluster 10.1.1.1;
    no-client-reflect;
    neighbor 2001:db8:fac1::a2;
    neighbor 2001:db8:fac1::a3;
    neighbor 2001:db8:fac1::a4;
    neighbor 2001:db8:fac1::a5;
}
```

C.
```
[edit protocols bgp]
root@R3# show
group AS65420 {
    type internal;
    local-address 2001:db8:fac1::a3;
    cluster 10.1.1.1;
    no-client-reflect;
    neighbor 2001:db8:fac1::a1;
    neighbor 2001:db8:fac1::a2;
    neighbor 2001:db8:fac1::a4;
    neighbor 2001:db8:fac1::a5;
}
```

D.
```
[edit protocols bgp]
root@R1# show
group AS65420 {
    type internal;
    local-address 2001:db8:fac1::a1;
    neighbor 2001:db8:fac1::a2;
    neighbor 2001:db8:fac1::a3;
    neighbor 2001:db8:fac1::a4;
    neighbor 2001:db8:fac1::a5;
    neighbor 2001:db8:fac1::a3;
    neighbor 2001:db8:fac1::a4;
    neighbor 2001:db8:fac1::a5;
}
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

The valid configuration is the one that establishes R1 as the route reflector and explicitly designates R2, R3, R4, and R5 as route-reflector clients under their respective BGP neighbor statements. A route reflector reflects routes learned from one client to its other clients, removing the need for a full IBGP mesh among those client routers. For routes originated or advertised by R4, R1 applies the route-reflection attributes needed to prevent reflection loops, including ORIGINATOR_ID and CLUSTER_LIST processing. This lets R2 receive the reflected route from R1 as the intended single advertisement through the route-reflector design. Junos implements this relationship by using the `route-reflector-client` statement in the appropriate BGP group or neighbor configuration on R1; the client designation is configured on the route reflector, not on the client routers. The configuration also must maintain a consistent BGP autonomous system for the IBGP sessions and use the correct peer addresses shown in the topology. Juniper documents route-reflector client configuration and route reflection behavior at route-reflector-client Statement and in the BGP Route Reflectors documentation.

QUESTION NO: 26

Click the Exhibit button.

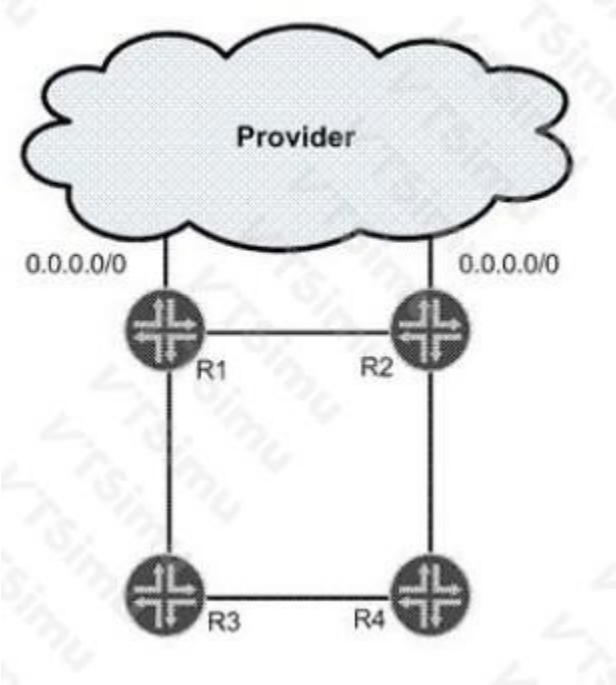

In the exhibit, R1 and R2 have a static default route configured that points toward the provider. Both routers redistribute the default route into OSPF. R2 is the preferred gateway to reach the provider. R1 is the backup gateway. All link metrics are equal. Which two steps ensure that traffic to the provider flows through R2 while the network is working properly? (Choose two.)

A. Redistribute the default route as a Type 1 external route on router R1 and a Type 2 external route on router R2.

B. Redistribute the default route as a Type 2 external route on router R1 and a Type 1 external route on router R2.

C. Modify the preference value of the default route on router R1 so that it is less preferred than OSPF external routes.

D. Modify the preference value of the default route on router R2 so that it is more preferred than OSPF external routes.

ANSWER: B C

Explanation:

Redistributing the default route as a Type 2 external route on router R1 and a Type 1 external route on router R2 ensures that other OSPF routers select the route advertised by R2 when both gateways are available. In OSPF route selection, an external Type 1 route is preferred over an external Type 2 route. A Type 1 external metric also includes the internal cost to reach the advertising ASBR, whereas a Type 2 metric is evaluated differently and is intended for cases where the external cost dominates. Thus, advertising R2's default as Type 1 establishes the desired primary exit, while R1's Type 2 advertisement remains available for backup use after R2's route is withdrawn.

Router R1 itself initially prefers its directly configured static default route because Junos assigns static routes a default route preference of 5, substantially better than the default preference of 150 for OSPF external routes. Raising the preference of R1's static default route to a value less preferred than OSPF external routes causes R1 to install and use the OSPF external default learned through R2 during normal operation. If R2 becomes unavailable, R1 can again use its own static default route as the backup path. Juniper documents external-route behavior in its OSPF external routes documentation and default route preferences in its route preference documentation.