DUMPSARENA

Service Provider Professional (JNCIP-SP)

Juniper JN0-664

Version Demo

Total Demo Questions: 13

Total Premium Questions: 139

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, OSPF	12
Topic 2, IS-IS	20
Topic 3, BGP	32
Topic 4, MPLS	7
Topic 5, Layer 2 VPNs	20
Topic 6, Layer 3 VPNs	24
Topic 7, Multicast	11
Topic 8, Class of Service	13
Total	139

QUESTION NO: 1

Which two statements are correct about IS-IS interfaces? (Choose two.)

A. If a broadcast interface is in both L1 and L2, one combined hello message is sent for both levels.

B. If a point-to-point interface is in both L1 and L2, separate hello messages are sent for each level.

C. If a point-to-point interface is in both L1 and L2, one combined hello message is sent for both levels.

D. If a broadcast interface is in both L1 and L2, separate hello messages are sent for each level

ANSWER: C D

Explanation:

IS-IS uses different hello behavior depending on the circuit type. On a point-to-point circuit enabled for both Level 1 and Level 2, a single point-to-point IS-IS hello (IIH) is used. The point-to-point IIH includes a circuit-type field that identifies whether the sender supports Level 1, Level 2, or both levels; consequently, one combined hello can advertise dual-level operation on that adjacency. This behavior is defined by the point-to-point IS-IS extensions in RFC 5303. On a broadcast (LAN) circuit, however, Level 1 and Level 2 use distinct LAN IIH packet formats and independently perform LAN adjacency functions, including designated IS election at each enabled level. Therefore, an interface participating in both levels sends separate hello messages for Level 1 and Level 2. This distinction is important when validating adjacencies and interpreting IS-IS packet captures: a dual-level point-to-point link shows a combined point-to-point hello capability, while a dual-level broadcast LAN has per-level hello traffic. See RFC 5303, Three-Way Handshake for IS-IS Point-to-Point Adjacencies and RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments.

QUESTION NO: 2

You are configuring anycast RP for load balancing and redundancy in your PIM-SM domain. You want to share active sources between RPs.

In this scenario, what are two solutions that will accomplish this task? (Choose two.)

A. Configure MSDP on each RP router.

B. Configure anycast PIM with the rp-set statement on each RP router.

C. Configure anycast PIM with the rp-set statement on each source DR router.

D. Configure MSDP on each source DR router.

ANSWER: A B

Explanation:

“Configure MSDP on each RP router.” is required because an anycast rendezvous point deployment uses multiple physical RP routers that advertise or are configured with the same RP address. A source’s PIM Register message can reach one of those physical RPs, but the other RPs must also learn that the source is active so that receivers whose joins reach a different RP can discover and join the source tree. MSDP provides this synchronization by exchanging Source-Active messages between the RP routers. Each RP therefore maintains awareness of multicast sources registered at its anycast peers.

“Configure anycast PIM with the rp-set statement on each RP router.” establishes the anycast-RP peer set on the routers performing the RP role. In Junos, the anycast PIM configuration identifies the local RP address and the other members of the RP set, allowing the physical RPs to operate as one logical anycast RP while retaining distinct unicast addresses for inter-RP communication. Used together, anycast PIM supplies RP redundancy and load distribution, while MSDP distributes active-source information among those RPs. Juniper documents this model and the associated configuration under Anycast RP Overview and MSDP Overview.

QUESTION NO: 3

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

In this scenario, which solution do you propose?

A. circuit cross-connect

B. BGP Layer 2 VPN

C. LDP Layer 2 circuit

D. translational cross-connect

ANSWER: C

Explanation:

LDP Layer 2 circuit is the appropriate service because it provides a point-to-point Layer 2 VPN over an MPLS provider network while using the Label Distribution Protocol to signal the pseudowire. The provider edge routers use targeted LDP sessions and Layer 2 circuit label bindings to establish the service; BGP is not required for VPN signaling or route distribution. Customer Ethernet frames are encapsulated with a pseudowire label and transported across the MPLS core through an existing transport LSP. This separation between the access circuit, pseudowire, and MPLS transport path lets the provider carry multiple independent Layer 2 circuits over the same core LSP, which supports efficient scaling. Because the service presents transparent Layer 2 connectivity between customer endpoints, the customer retains control of its private-network routing design: its CE routers or switches can run the desired IP routing protocols without provider participation in those customer routes. Juniper documents Layer 2 circuits as point-to-point Layer 2 VPN services and supports LDP-signaled pseudowires for this purpose. The LDP pseudowire signaling behavior is standardized in RFC 4447; see also Juniper's Layer 2 VPN overview.

QUESTION NO: 4

Click the Exhibit button.

PE-1 and PE-2 are configured with LDP-signaled pseudowires to provide connectivity between CE-1 and CE-2. You notice no connectivity exists between CE-1 and CE-2.

Referring to the exhibit, which two statements describe potential causes for this fault? (Choose two.)

A. The VC IDs are mismatched.

B. There is no LSP configured from PE-1 to PE-2.

C. Interface ge-0/0/0 on PE-1 is down.

D. There is no LSP configured from PE-2 to PE-1.

ANSWER: A D

Explanation:

LDP-signaled Layer 2 pseudowires require both provider-edge routers to advertise and bind the same virtual-circuit identifier for the pseudowire. Therefore, *The VC IDs are mismatched.* is a valid cause: the two endpoints will not associate their LDP label mappings as the same Layer 2 circuit, preventing the pseudowire from becoming operational. The virtual-circuit ID is configured under the relevant `protocols l2circuit` neighbor and interface hierarchy and must agree at both ends.

There is no LSP configured from PE-2 to PE-1. is also a potential cause. A pseudowire needs MPLS transport connectivity between its endpoint PEs. In particular, PE-2 must have a usable MPLS label-switched path toward PE-1 so that it can send labeled pseudowire traffic across the provider network. The transport may be provided by LDP or RSVP, depending on the network design, but an operational path in that direction is essential for bidirectional CE-to-CE service. Juniper's Layer 2 Circuit documentation describes the VC-ID requirement and the dependence of L2 circuits on MPLS transport: Juniper Layer 2 circuit overview. Related configuration details are available in the Juniper Layer 2 circuits documentation.

QUESTION NO: 5

Which two statements are correct about RSVP-TE fast reroute using link protection? (Choose two.)

A. The point of local repair switches traffic to a backup path after a protected link fails.

B. The backup path can bypass the failed link without waiting for end-to-end LSP reconvergence.

C. Link protection always bypasses the adjacent downstream router.

D. The ingress router must immediately signal a new primary LSP before traffic can be repaired.

ANSWER: A B

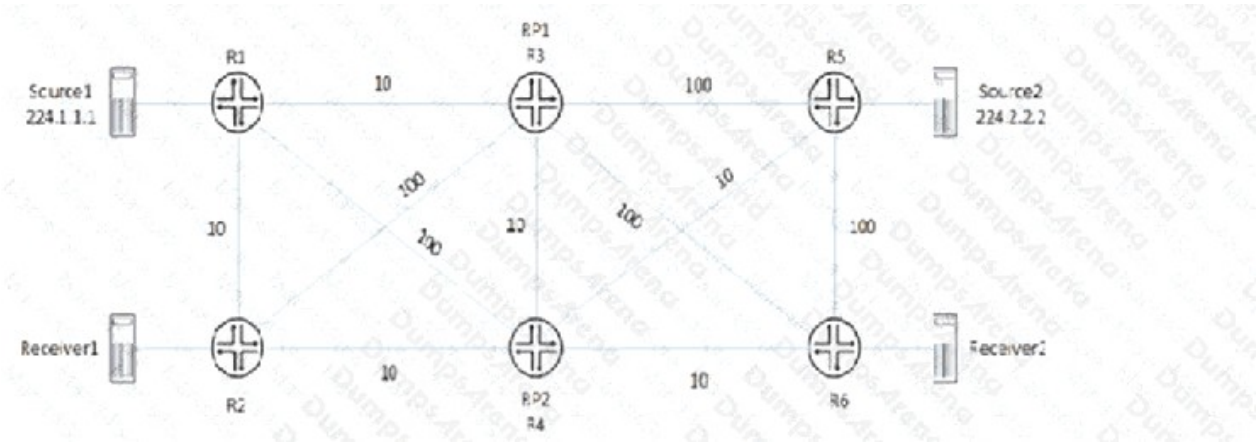
Explanation:

With link protection, the point of local repair establishes a backup path that bypasses the protected link. When that link fails, the point of local repair can rapidly switch traffic to the backup without waiting for ingress-to-egress LSP signaling and convergence. This local repair behavior is intended to minimize traffic loss following a protected failure.

Link protection does not necessarily protect against failure of the next-hop router itself. Node protection requires a backup path that avoids both the protected link and the adjacent downstream node. RSVP-TE fast reroute is defined in RFC 4090; Juniper also describes the feature in its RSVP fast reroute documentation.

QUESTION NO: 6

Exhibit



Referring to the exhibit, PIM-SM is configured on all routers, and Anycast-RP with Anycast-PIM is used for the discovery mechanism on RP1 and RP2. The interface metric values are shown for the OSPF area.

In this scenario, which two statements are correct about which RP is used? (Choose two.)

A. Source2 will use RP2 and Receiver2 will use RP2 for group 224.2.2.2.

B. Source2 will use RP1 and Receiver2 will use RP1 for group 224.2.2.2.

C. Source1 will use RP1 and Receiver1 will use RP1 for group 224.1.1.1.

D. Source1 will use RP1 and Receiver1 will use RP2 for group 224.1 1 1

ANSWER: A C

Explanation:

With Anycast-RP, RP1 and RP2 advertise the same anycast RP address. Each router therefore selects its active RP based on the normal unicast route to that shared address; in this topology, OSPF path metrics determine which physical RP is

reached. The source-side PIM router encapsulates Register messages toward the closest reachable Anycast-RP, while receiver-side routers build (*,G) shared-tree state toward the Anycast-RP using their own unicast best path. Anycast-PIM supplies the coordination between the individual RP routers so that multicast-source information learned at either RP is available to the other RP. Consequently, Source2 and Receiver2 select the RP reached through RP2 for group 224.2.2.2, and Source1 and Receiver1 select the RP reached through RP1 for group 224.1.1.1. The selection is determined independently at the relevant source and receiver routers according to the exhibited OSPF metrics, rather than by a requirement that all multicast groups globally use one physical RP. Juniper documents that Anycast-RP uses a shared RP address and that Anycast-PIM is used to synchronize source information among RPs. See Juniper Anycast RP documentation and Juniper PIM-SM overview.

QUESTION NO: 7

You are configuring a BGP signaled Layer 2 VPN across your MPLS enabled core network. Your PE-2 device connects to two sites within the s VPN

In this scenario, which statement is correct?

A. By default on PE-2, the site's local ID is automatically assigned a value of 0 and must be configured to match the total number of attached sites.

B. You must create a unique Layer 2 VPN routing instance for each site on the PE-2 device.

C. You must use separate physical interfaces to connect PE-2 to each site.

D. By default on PE-2, the remote site IDs are automatically assigned based on the order that you add the interfaces to the site configuration.

ANSWER: D

Explanation:

“By default on PE-2, the remote site IDs are automatically assigned based on the order that you add the interfaces to the site configuration.” is correct. In a BGP-signaled Layer 2 VPN, BGP distributes the Layer 2 VPN endpoint and label-block information needed for PE devices to establish the appropriate pseudowires. A site identifier distinguishes a site within the Layer 2 VPN and is used with the advertised label block to determine the label and pseudowire association for remote connectivity. When the PE has multiple locally attached sites and the relevant site identifiers are not explicitly configured, Junos allocates the identifiers in interface-configuration order: the first configured site is assigned the first available identifier, followed by the next site. This provides deterministic local provisioning while allowing BGP to advertise the information that remote PE devices require to build the full mesh of Layer 2 VPN connectivity. The behavior is part of the BGP-based VPLS/L2VPN control-plane model defined for carrying VPN membership and label-block information in BGP. See Juniper’s BGP-signaled Layer 2 VPN overview and RFC 4761 for the BGP Layer 2 VPN signaling model.

QUESTION NO: 8

You are using a Layer 3 VPN to connect two customer sites. The VPN routes for the customer networks appear as hidden in the bgp. 13vpn. o routing table on the PE routers.

What is causing this problem?

A. The routes use overlapping IP addresses.

B. There is not an established MPLS LSP between the two PE routers.

C. There is a routing loop in the service provider backbone.

D. Route targets are not configured.

ANSWER: B

Explanation:

There is not an established MPLS LSP between the two PE routers. In a BGP Layer 3 VPN, a PE router advertises VPN-IPv4 routes with a BGP next hop that normally identifies the remote PE router. To install an eligible VPN route for forwarding, Junos must be able to resolve that BGP next hop through the MPLS transport infrastructure. Typically, an IGP supplies reachability to the remote PE loopback address and RSVP or LDP supplies an MPLS label-switched path, with the resulting transport next-hop information available for resolution.

When no usable LSP exists to the remote PE, the VPN route cannot obtain a resolved MPLS forwarding next hop. Junos retains the received route in `bgp.l3vpn.0`, but marks it hidden because it is not usable for route selection and installation into the applicable VRF routing table. Restoring MPLS transport reachability between the PE routers allows recursive next-hop resolution, permits the VPN route to become active, and enables the PE to impose the transport and VPN labels required to carry customer traffic across the provider backbone.

Juniper describes BGP Layer 3 VPN route distribution and MPLS transport requirements in its BGP Layer 3 VPN documentation and explains route next-hop resolution in the Junos route resolution documentation.

QUESTION NO: 9

When using OSPFv3 for an IPv4 environment, which statement is correct?

A. OSPFv3 only supports IPv4.

B. OSPFv3 supports both IPv6 and IPv4, but not in the same routing instance.

C. OSPFv3 is not backward compatible with IPv4

D. OSPFv3 supports IPv4 only on interfaces with family inet6 defined

ANSWER: D

Explanation:

OSPFv3 supports carrying IPv4 routing information through its IPv4 address-family extensions, but the protocol itself still uses IPv6 for its control-plane transport. In particular, OSPFv3 neighbor discovery, adjacency establishment, and protocol packet exchange use IPv6 link-local addresses. Consequently, on Junos, an interface participating in OSPFv3 for an IPv4 environment must have IPv6 enabled by configuring `family inet6`. The interface can simultaneously have its IPv4 addressing under `family inet`, while OSPFv3 uses the IPv6 capability to communicate with adjacent OSPFv3 routers and advertises IPv4 reachability for the IPv4 address family. This design follows the OSPFv3 address-family mechanism: IPv4 prefixes can be carried by OSPFv3, but the underlying OSPFv3 protocol transport remains IPv6-based. See the Juniper OSPFv3 overview and RFC 5838, which defines support for additional address families in OSPFv3.

QUESTION NO: 10

Exhibit

Referring to the exhibit, you must provide Internet access for VPN-A using CE-1 as the hub C

E.

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

A. You must use RIB groups to leak routes between the inet.0 and vpn-a.inet.0 tables.

B. RIB groups are not needed to leak routes between the inet.0 and vpn-a.inet.0 tables.

C. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> GW-1.

D. Internet traffic from Site 2 takes the path of PE-2 -> PE-1 -> CE-1 -> PE-1 -> GW-1.

ANSWER: B D

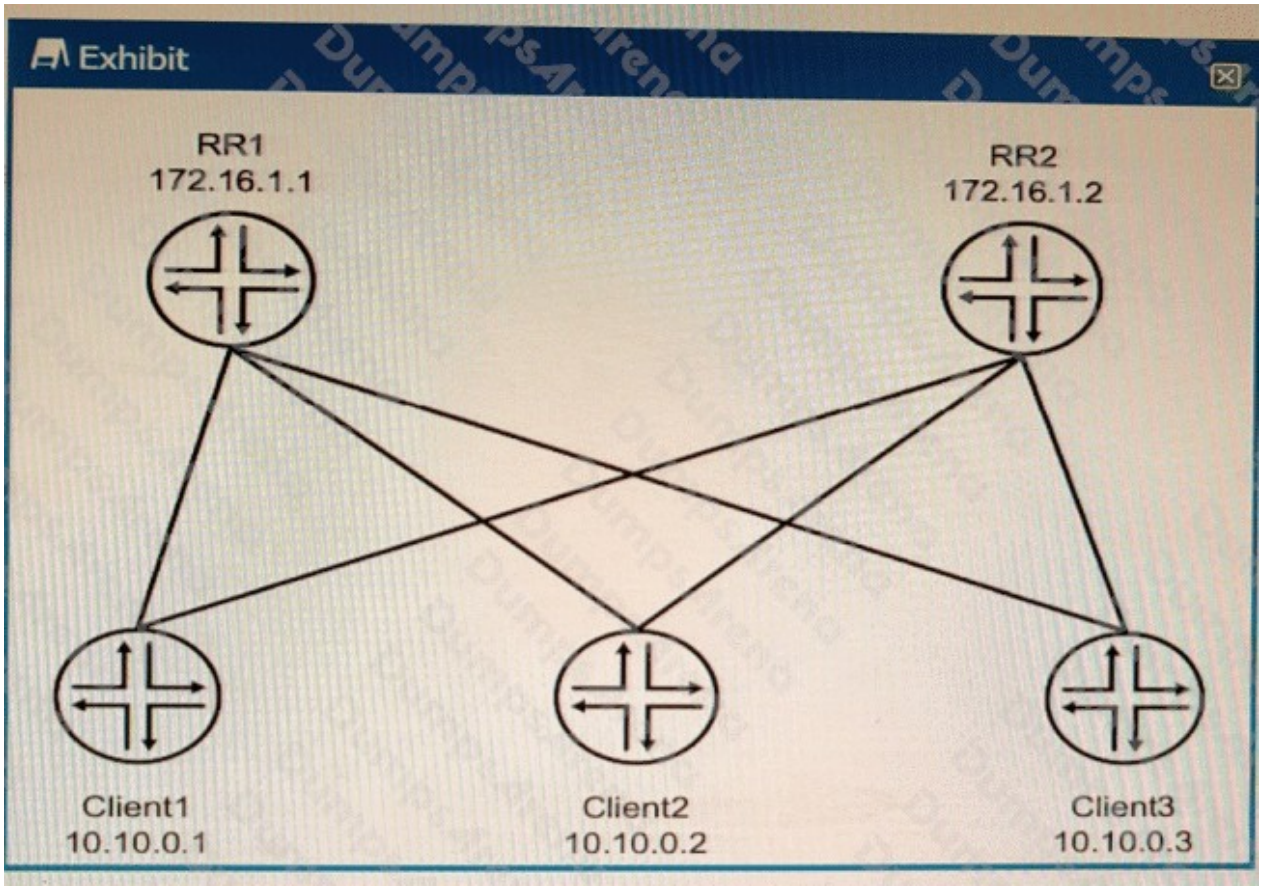
Explanation:

RIB groups are not mandatory for exchanging reachability between the global `inet.0` table and a VPN routing table. Junos supports policy-controlled route leaking through routing-instance import and export mechanisms, allowing the required Internet reachability to be made available without making a RIB group a design requirement. This allows the provider to retain controlled separation between VPN and Internet routing while selectively sharing only the necessary routes. Juniper documents routing-instance route leaking and the use of import policy under the `instance-import` statement at Junos instance-import documentation.

With CE-1 serving as the centralized Internet-access hub for VPN-A, traffic sourced at Site 2 is first transported across the MPLS VPN from PE-2 to PE-1. PE-1 forwards the traffic to CE-1, because CE-1 is the designated hub that provides the Internet exit function. CE-1 then sends the traffic back toward PE-1 on its Internet-facing path, after which PE-1 forwards it to GW-1. Therefore, the forwarding sequence is PE-2 to PE-1 to CE-1 to PE-1 to GW-1. This is the expected hairpin forwarding behavior for a centralized-services VPN design. See Juniper's Layer 3 VPN overview for the VPN forwarding model.

QUESTION NO: 11

Exhibit



The environment is using BGP All devices are in the same AS with reachability redundancy Referring to the exhibit, which statement is correct?

A. RR1 is peered to Client2 and RR2

B. RR2 is in an OpenConfirm State until RR1 becomes unreachable.

C. Client1 is peered to Client2 and Client3.

D. Peering is dynamically discovered between all devices.

ANSWER: A

Explanation:

RR1 is peered to Client2 and RR2 is correct. The topology shown uses route reflection within one autonomous system to provide scalable internal BGP connectivity while retaining redundant reachability. RR1 has an internal BGP client relationship with Client2, allowing RR1 to reflect eligible routes to and from that client. RR1 also has an internal BGP peering with RR2. This route-reflector-to-route-reflector session is important because it shares routing information between the redundant route reflectors, helping each reflector provide a consistent control-plane view and preserve reachability if client connectivity or a reflector path changes.

Route reflection is an internal BGP mechanism that removes the requirement for every internal BGP speaker to maintain a full mesh. A route reflector distinguishes client and nonclient peer relationships and applies the route-reflection rules when advertising routes learned through internal BGP. In a redundant design, clients can peer with more than one route reflector, and the reflectors can also peer with each other, as depicted here. The session relationships indicated in the exhibit therefore directly establish that RR1 peers with Client2 and RR2. Juniper's route-reflection overview describes these client and nonclient relationships and their use in reducing internal BGP mesh requirements: Juniper BGP Route Reflectors documentation. The underlying protocol behavior is specified in RFC 4456.

QUESTION NO: 12

Which origin code is preferred by BGP?

A. Internal

B. External

C. Incomplete

D. Null

ANSWER: A

Explanation:

Internal is correct because it represents the IGP origin value, which has the most preferred (lowest) origin-code value in BGP path selection. BGP uses the origin attribute to indicate how a route entered BGP. A route originating through an IGP-based mechanism carries the IGP/internal origin, while routes learned through the legacy EGP mechanism use the EGP/external origin and routes redistributed into BGP use the incomplete origin. When BGP reaches the origin-attribute comparison in its best-path process, it selects the path with the lowest origin value: IGP/internal is preferred before EGP/external, which is preferred before incomplete. This comparison is made only after earlier, more influential BGP attributes, such as local preference and AS-path length, have not distinguished the candidate paths. The origin attribute is therefore a tie-breaking criterion, but its ordering is deterministic and important when otherwise comparable routes exist. Juniper documents the BGP path-selection process, including preference for the lowest origin code, in its BGP Path Selection Process documentation. The origin attribute and its values are also standardized in RFC 4271, Section 5.1.1.

QUESTION NO: 13

Which statement is correct about IS-IS when it performs the Dijkstra algorithm?

A. The local router moves its own local tuples into the candidate database

B. When a new neighbor ID in the tree database matches a router ID in the LSDB, the neighbor ID is moved to the candidate database

C. Tuples with the lowest cost are moved from the tree database to the LSDB.

D. The algorithm will stop processing once the tree database is empty.

ANSWER: A

Explanation:

The local router moves its own local tuples into the candidate database is correct. During IS-IS shortest-path-first processing, the router begins its calculation from itself as the root of the topology. Its local reachability information is used to initialize the candidate set, which represents nodes or prefixes that have been discovered but whose least-cost path has not yet been finalized. The SPF process then selects the candidate with the lowest accumulated metric, installs that result into the shortest-path tree, and examines the selected node's advertised links to discover or update additional candidates. This iterative process derives the shortest paths from the local router to all reachable destinations based on the link-state information learned through IS-IS link-state packets. Initializing the candidate database from local tuples is therefore necessary: it establishes the local router as the calculation origin and supplies the first directly connected topology entries from which the SPF tree can expand. IS-IS is a link-state protocol, so each router independently performs this calculation against its topology database after relevant topology changes. Juniper's IS-IS documentation describes the protocol's link-state operation and SPF route calculation behavior; the underlying IS-IS algorithm and database model are standardized in ISO-derived IS-IS specifications and extended for IP by RFC 1195. See Juniper IS-IS Overview and RFC 1195.