DUMPSARENA

Enterprise Routing and Switching Professional (JNCIP-ENT)

Juniper JN0-649

Version Demo

Total Demo Questions: 8

Total Premium Questions: 89

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

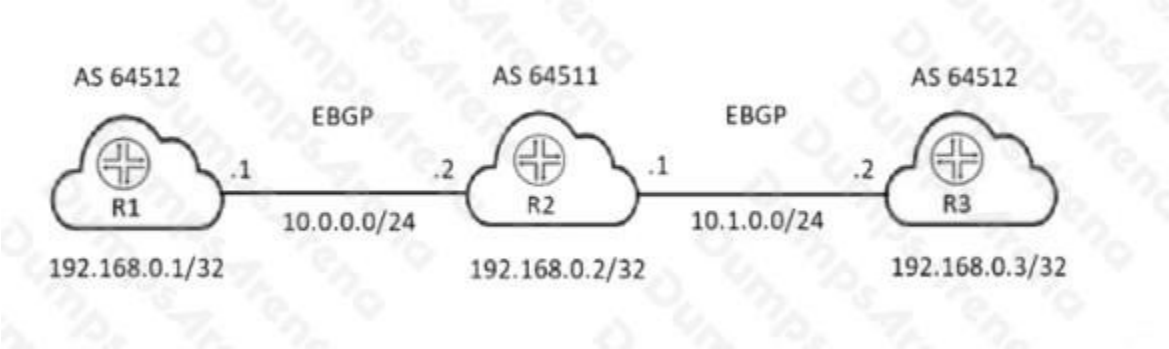
Topic Break Down

Topic	No. of Questions
Topic 1, Layer 2 switching	26
Topic 2, Layer 2 security	10
Topic 3, Protocol-independent routing	1
Topic 4, OSPF	12
Topic 5, IS-IS	4
Topic 6, BGP	21
Topic 7, High availability	1
Topic 8, Multicast	8
Topic 9, Mix Questions	6
Total	89

QUESTION NO: 1

You are asked to establish full connectivity between all devices in the BGP network.

Referring to the exhibit, which two configuration changes will allow BGP route advertisements? (Choose two.)



A. On R2, include the loops 2 statement at the [edit protocols bgp family inet unicast] hierarchy.

B. On R1 and R3, include the loops 2 statement at the [edit protocols bgp family inet unicast] hierarchy.

C. On R1 and R3, include the advertise-peer-as statement at the [edit protocols bgp group external] hierarchy.

D. On R2, include the advertise-peer-as statement at the [edit protocols bgp group external] hierarchy.

ANSWER: B D

Explanation:

On R1 and R3, configuring `loops 2` under the IPv4 unicast BGP family permits each router to accept advertisements whose AS paths contain its own local AS number, up to the configured occurrence count. This is required when routes traverse R2 and return toward the other router that uses the same autonomous system. The setting allows the needed paths to be installed and propagated while retaining a bounded AS-path loop check. Junos documents the behavior and hierarchy of this feature in the loops statement reference.

On R2, configuring `advertise-peer-as` in the external BGP group permits R2 to advertise a route learned from one peer to another peer in that same peer AS. This is the necessary advertisement behavior for carrying reachability between R1 and R3 through R2 when both edge routers share an AS. With this control enabled on R2 and the permitted local-AS occurrences configured on R1 and R3, the BGP speakers can exchange the required routes for full connectivity. Juniper describes this behavior in the advertise-peer-as statement reference.

QUESTION NO: 2

You are deploying IS-IS on a broadcast Ethernet segment. All routers on the segment use the same IS-IS interface priority.

Which value is used to elect the Designated Intermediate System?

A. The highest system ID

B. The lowest system ID

C. The highest router ID

D. The lowest interface address

ANSWER: A

Explanation:

The router with the highest system ID is elected as the Designated Intermediate System (DIS) when interface priorities are equal. IS-IS first compares the configured interface priorities. If those values tie, the highest system ID is selected.

The DIS performs functions on the broadcast network that include generating the pseudonode LSP used to represent the multiaccess segment in the link-state database. Unlike OSPF DR operation, IS-IS does not elect a backup DIS. The DIS election behavior is defined in RFC 1195 and is described in Juniper's IS-IS overview.

QUESTION NO: 3

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

```
user@switch> show poe interface
Interface    Admin       Oper     Max        Priority       Power          Class
             status      status   power                     consumption
 ge-0/0/0    Enabled     OFF      15.4W      Low            0.0W           not-applicable
 ge-0/0/1    Enabled     OFF      15.4W      Low            0.0W           not-applicable
 ge-0/0/2    Enabled     OFF      15.4W      Low            0.0W           not-applicable
 ge-0/0/3    Enabled     OFF      15.4W      Low            0.0W           not-applicable
 ge-0/0/4    Enabled     OFF      15.4W      Low            0.0W           not-applicable
 ge-0/0/5    Enabled     OFF      15.4W      Low            0.0W           not-applicable
 ge-0/0/6    Enabled     OFF      15.4W      Low            0.0W           not-applicable
 ge-0/0/7    Enabled     OFF      15.4W      Low            0.0W           not-applicable
 ge-0/0/8    Enabled     OFF      15.4W      Low            0.0W           not-applicable
 ge-0/0/9    Enabled     OFF      15.4W      Low            0.0W           not-applicable
ge-0/0/10    Enabled      ON      25.4W(L)   Low            11.0W           4
ge-0/0/11    Enabled      ON      25.4W(L)   High           11.4W           4
  (L) LLDP-negotiated value on the port.
user@switch> show poe controller
Controller  Maximum   Power         Guard    Management   Status        Lldp
index       power     consumption   band                               Priority
   0        100.00W   22.40W         10W     Class        AT_MODE       Disabled
```

A. The maximum wattage that this switch can allocate to attached Ethernet devices is 100 watts.

B. If the total power consumption exceeds 90 watts, the ge-0/0/11 interface will continue to receive power.

C. PoE is not enabled on the ge-0/0/0 interface.

D. The ge-0/0/10 interface supports PoE+.

ANSWER: A C

Explanation:

The exhibit's PoE operational information identifies 100 watts as the power available for allocation to powered devices. This value is the switch's available PoE power budget: it represents the total amount of power that can be supplied across eligible Ethernet interfaces, subject to the configured power-management behavior and the power requirements of connected devices. Therefore, the statement that the switch can allocate a maximum of 100 watts to attached Ethernet devices correctly reflects the displayed budget.

The interface-level PoE state shown for `ge-0/0/0` is disabled. A disabled PoE state means that the port is not enabled to provide inline power to a connected powered device, regardless of whether the physical interface might otherwise be capable of carrying Ethernet traffic. Junos provides both system-wide PoE information and per-interface PoE administrative and operational status, allowing an operator to verify exactly this condition. See Juniper's show poe interface command reference and its Power over Ethernet documentation for the meaning of PoE budgets and interface status.

QUESTION NO: 4

You are asked to troubleshoot voice quality issues on your newly implement VoIP network. You notice that the voice packets are being dropped. You have verified that the packets are correctly marked for expedited forwarding queue.

Referring to the exhibit, what must you configure to solve the problem?

```
[edit]
user@R1# show class-of-service
classifiers {
    dscp voip {
        import default;
    }
}
interfaces {
    ge-1/0/0 {
        unit 0 {
            classifiers {
                dscp voip;
            }
        }
    }
}
user@R1> show interfaces ge-1/0/0 extensive
Physical interface: ge-1/0/0, Enabled, Physical link is Up
  Interface index: 154, SNMP ifIndex: 527, Generation: 157
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error: None, Loop Detect PDU Error:
None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control:
Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Pad to minimum frame size: Disabled
  Media type: Copper
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000

  Auto-negotiation: Enabled, Remote fault: Online
  Pad to minimum frame size: Disabled
  Media type: Copper
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Schedulers     : 0
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state: unsuppressed
  Current address: 4c:96:14:93:9a:95, Hardware address: 4c:96:14:93:9a:95
  Last flapped   : 2022-05-16 11:44:33 PDT (21:23:22 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              894761                 0 bps
   Output bytes  :              681004               240 bps
   Input  packets:               13083                 0 pps
   Output packets:               11321                 0 pps
   IPv6 transit statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Input errors:

    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0, L2
mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0,
MTU errors: 0,
    Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0                          430544                 8126              456123
    1                            4294                 1654                2817
    2                               0                    0                   0
    3                           11194                11194                   0
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
  Ethernet FEC statistics              Errors
    FEC Corrected Errors                   0
```

```
  FEC Uncorrected Errors                0
  FEC Corrected Errors Rate             0
  FEC Uncorrected Errors Rate           0
MAC statistics:                      Receive         Transmit
  Total octets                        947941           752356
  Total packets                        13084            11320
  Unicast packets                         92               93
  Broadcast packets                       37               34
  Multicast packets                    12955            11193
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
  Total errors                             0                0
Filter statistics:
  Input packet count                   13083
  Input packet rejects                     0
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                              11320
  Output packet pad count                              0
  Output packet error count                            0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
```

```
  Fragment frames                      0
  VLAN tagged frames                   0
  Code violations                      0
  Total errors                         0                0
Filter statistics:
  Input packet count                13083
  Input packet rejects                 0
  Input DA rejects                     0
  Input SA rejects                     0
  Output packet count                          11320
  Output packet pad count                          0
  Output packet error count                        0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
      Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault: OK
  Local resolution:
      Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0 (0x00)
CoS information:
  Direction : Output
  CoS transmit queue              Bandwidth              Buffer Priority   Limit
                             %           bps      %        usec
  0 best-effort             95     950000000     95           0      low    none
  3 network-control          5      50000000      5           0      low    none
Interface transmit statistics: Disabled
```

A. You must configure a multifield classifier to put the VoIP traffic in the correct queue.

B. You must configure a rewrite rule to ensure that the traffic is scheduled properly in the device.

C. You must configure a scheduler to allocate bandwidth to the expedited forwarding queue.

D. You must configure a policer to ensure that the queue is not being starved.

ANSWER: C

Explanation:

You must configure a scheduler to allocate bandwidth to the expedited forwarding queue. In Junos Class of Service, classification identifies traffic and assigns it to a forwarding class, while a scheduler defines how the associated output queue is serviced. Marking the VoIP packets for the expedited-forwarding forwarding class correctly identifies the traffic, but it does not by itself guarantee the queue sufficient transmission resources or queue treatment. A scheduler supplies parameters such as transmit rate, priority, buffer size, and drop-profile behavior. The scheduler must then be associated with the expedited-forwarding class through a scheduler map applied to the relevant output interface. For latency-sensitive voice, the expedited-forwarding queue is normally given strict-high priority and an appropriate bandwidth allocation, with queue-buffer settings sized for the expected traffic characteristics. This allows the egress interface to service voice traffic predictably and

prevents drops caused by an unsuitable default scheduling or buffering treatment. Juniper's CoS documentation describes schedulers as the mechanism that controls bandwidth allocation and queue servicing, and scheduler maps as the binding between forwarding classes and schedulers. See Juniper's scheduler overview and Juniper's scheduler-map overview.

QUESTION NO: 5

Your enterprise network uses routing instances to support multitenancy. Your Junos devices use BGP to peer to multiple BGP devices. You must ensure that load balancing is achieved within the routing instance.

Which two statements would accomplish this task? (Choose two.)

A. Configure the multipath option at the [edit protocols bgp group neighbor] hierarchy.

B. Configure the multipath option at the [edit protocols bgp group] hierarchy.

C. Configure a load-balance per-packet policy and apply it at the [edit routing-options forwarding-table] hierarchy.

D. Configure the multipath option at the [edit routing-instances routing-options] hierarchy.

ANSWER: B D

Explanation:

Configuring the multipath option at the [edit protocols bgp group] hierarchy enables BGP multipath for peers in the specified BGP group. When BGP installs multiple eligible paths with equal cost, the forwarding plane can use those equal-cost next hops for load sharing. In a routing-instance context, the same BGP group-level configuration is applied beneath that instance's BGP protocol hierarchy, keeping route selection and forwarding behavior associated with the tenant's separate routing table.

Configuring the multipath option at the [edit routing-instances <instance-name> routing-options] hierarchy enables equal-cost multipath processing specifically for the named routing instance. This is important in a multitenant design because routing instances maintain independent routing tables and forwarding contexts. Enabling multipath in the instance ensures that qualifying equal-cost BGP routes learned for that tenant can be retained and used for load balancing without relying on the default routing instance's behavior. Junos documents BGP multipath configuration at protocol and group scope, as well as multipath behavior under routing options. See the Junos BGP multipath statement reference and the Junos routing-options multipath reference.

QUESTION NO: 6

Your organization has recently acquired another company. You must carry all of the company's existing VLANs across the corporate backbone to the existing branch locations without changing addressing and with minimal configuration.

Which technology will accomplish this task?

A. Q-in-Q all-in-one bundling

B. PVLAN isolated VLAN

C. MVRP registration normal

D. EVPN-VXLAN anycast gateway

ANSWER: A

Explanation:

Q-in-Q all-in-one bundling is the appropriate technology because it transparently transports multiple existing customer VLANs across a provider or corporate Ethernet backbone. The ingress edge device adds an outer service VLAN tag to frames while retaining each original customer VLAN tag. Consequently, the acquired company's VLAN identifiers and Layer 3 addressing remain unchanged end to end, allowing its existing Layer 2 segmentation to extend to the branch locations without requiring individual backbone VLAN configuration for every carried VLAN.

In Junos terminology, all-in-one bundling maps all customer VLANs received on a customer-facing interface into one service VLAN for transport. This is particularly suitable when a large or unknown set of customer VLANs must be carried with minimal operational overhead. At the remote edge, the outer service tag is removed and the original tagged Ethernet frames are delivered, preserving the acquired network's VLAN structure. Juniper documents Q-in-Q tunneling as a method for carrying customer VLAN traffic through a service-provider network and identifies all-in-one bundling as the model that bundles all customer VLANs into a single service VLAN. See Juniper Q-in-Q Tunneling documentation and Juniper Q-in-Q Tunneling on ELS switches.

QUESTION NO: 7

You are examining DR and BDR election behavior on an OSPF broadcast network.

Which two statements are correct? (Choose two.)

A. A router with an interface priority of 0 is ineligible to become the DR or BDR.

B. The highest router ID wins when eligible routers have equal interface priorities.

C. The router with the lowest router ID is preferred as the DR.

D. A newly connected router with a higher priority immediately replaces the existing DR.

ANSWER: A B

Explanation:

An OSPF interface configured with a router priority of **0** is ineligible to become either the designated router or backup designated router. This setting is useful when a router should form adjacencies on a multiaccess segment but must not assume a designated-router role.

When eligible routers have the same OSPF interface priority, the router with the numerically highest router ID is preferred during the DR or BDR election. OSPF uses the priority first and the router ID as the tie breaker among otherwise equally eligible candidates.

The DR election is nonpreemptive. A newly available router with a higher priority does not replace an existing DR until an election is triggered, such as after the current DR fails. See RFC 2328: OSPF Version 2.

QUESTION NO: 8

You recently committed a change to a router to reject OSPF routes sourced from area 10. However, you are still seeing area 10 routes in the routing table.

Referring to the exhibit, which statement is correct?

```
[edit policy-options]
policy-statement advertise-ospf-routes {
  term find-ospf {
    from {
      protocol ospf;
    }
    then {
      accept;
    }
  }
  term reject-area-10 {
    from {
      protocol ospf;
      area 10;
    }
    then {
      reject;
    }
  }
}
```

A. The OSPF protocol is first matched by find-ospf and accepted.

B. The routes only timeout after 24 hours.

C. The routes remain in the table until the device is rebooted.

D. The routes remain in the table until the routing daemon is restarted.

ANSWER: A

Explanation:

The OSPF protocol is first matched by find-ospf and accepted. Junos routing policies evaluate terms in their configured order, from the first term through subsequent terms. A term that contains an accept action is terminating: once a route matches that term, Junos accepts the route and stops evaluating the policy for that route. In the displayed policy, the find-ospf term matches routes based on the OSPF protocol and then accepts them. Consequently, an OSPF route from area 10 is accepted by that earlier term before Junos reaches the later term intended to reject area 10 routes. To enforce the desired filtering behavior, the area-specific rejection term must be evaluated before the broader OSPF accept term, or the broad match must be narrowed so it does not preempt the area 10 match. This behavior follows the standard ordered, first-match processing model for Junos routing policy terms and terminating actions. See Juniper's routing policy evaluation documentation and its routing policy term configuration overview.