

Juniper JN0-650

Enterprise Routing and Switching, Professional

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.




<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/jn0-650>

Latest Version: 8.3

Question: 1

Exhibit.

The exhibit is a screenshot of a Junos OS terminal window titled "Exhibit". It shows the output of the command "show dot1x interface ge-0/0/10.0 detail". The output lists various configuration parameters for the interface, including role, administrative state, supplicant mode, number of retries, quiet period, transmit period, MAC Radius status, reauthentication status, and timeouts. The current supplicant mode is "Single" and the number of connected supplicants is 2.

```
user@switch> show dot1x interface ge-0/0/10.0 detail
ge-0/0/10.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60seconds
  Transmit period: 30seconds
  Mac Radius: Disabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600seconds
  Supplicant timeout: 30seconds
  Server timeout: 30seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: 24
  Number of connected supplicants: 2
```

You want to limit port access to only one device at a time.

Referring to the exhibit, which configuration change will accomplish this task?

- A. Enable MAC RADIUS restrict.
- B. Change the supplicant mode to multiple.
- C. Change the supplicant mode to single-secure.
- D. Change the maximum EAPOL request to 1.

Answer: C

Explanation:

In Junos OS, the supplicant-mode configuration under protocols dot1x determines how the switch handles multiple MAC addresses on a single physical port. According to the exhibit, the current mode is set to Single, and the Number of connected supplicants is 2. This indicates that the port is currently allowing multiple devices, which contradicts the goal of limiting access to only one device at a time.

Here is the breakdown of why Option C is the correct solution based on Juniper's standard behavior:

Supplicant Mode: Single (Current State): In this mode, the first device to authenticate opens the port for all subsequent devices. As long as the first device remains authenticated, other devices

can send traffic through the port without individual authentication. This is why the exhibit shows 2 connected supplicants despite the mode being "Single."

Supplicant Mode: Single-Secure (The Solution): This mode strictly limits the port to only one MAC address. Once a device successfully authenticates via 802.1X, the switch drops any traffic coming from any other MAC address on that port. If the authenticated device logs off or the session times out, the port becomes available for a new device, but never more than one simultaneously. * Supplicant Mode: Multiple (Option B): This mode allows multiple supplicants to authenticate individually. Each MAC address must go through its own authentication process. This would allow more than one device, which is the opposite of the user's requirement. MAC RADIUS Restrict (Option A): This feature is used to force MAC-based authentication and does not inherently limit the number of devices to one in the same way that changing the supplicant mode does.

Maximum EAPOL requests (Option D): This parameter defines how many times the switch will send an EAP-Request/Identity frame to a supplicant before giving up. Changing this to 1 does not restrict the number of devices allowed on the port; it only changes the retry logic for a single authentication attempt.

Configuration Example for Junos OS 24.4: To implement this change, you would use the following command: `set protocols dot1x edit interface ge-0/0/10.0 supplicant-mode single-secure`

Question: 2

Your OSPF network consists of a mix of 1GbE and 10GbE interfaces. By default, all interfaces have the same cost in your OSPF network. You are asked to ensure that the 10GbE interfaces are more preferred when available

In this scenario, which two statements would accomplish this behavior? (Choose two.)

- A. You should define the reference bandwidth as 10G. which will assign the 1GbE interfaces a higher cost
- B. You should manually assign the interface metric for each 10GbE interface to be higher than the 1GbE interfaces in your OSPF network.
- C. You should define the reference bandwidth as 1G. which will assign the 1GbE interfaces a higher cost.
- D. You should manually assign the interface metric for each 1GbE interface to be higher than the 10GbE interfaces in your OSPF network.

Answer: A, D

Explanation:

OSPF determines the best path to a destination by calculating the metric (cost) of each link. By default, Junos OS uses a reference bandwidth of 100 Mbps to calculate this cost using the formula:

$$Cost = \frac{\text{Reference Bandwidth}}{\text{Interface Bandwidth}}$$

When the reference bandwidth is left at the default 100 Mbps, any interface with a speed of 100 Mbps or higher (including 1 GbE and 10 GbE) is assigned a cost of 1 because the minimum OSPF cost is 1. This results in equal-cost paths, preventing the router from preferring the faster 10 GbE link.

To ensure 10 GbE interfaces are preferred, you must create a cost differential:

Option A (Reference Bandwidth): By increasing the reference bandwidth to 10G (or higher), the calculation changes. For a 10 GbE link, the cost becomes $\$10,000 / 10,000 = 1\$$. For a 1 GbE link, the cost becomes $\$10,000 / 1,000 = 10\$$. Since OSPF prefers the path with the lowest cumulative cost, the 10 GbE link is now preferred.

Option D (Manual Metric): You can manually override the automatic cost calculation by assigning a higher metric specifically to the 1 GbE interfaces. If a 1 GbE interface is manually set to a cost of 50 and the 10 GbE interface remains at 1 (or is set to a lower value), the router will prioritize the 10 GbE path.

Option B is incorrect because a higher metric makes a path less preferred. Option C is incorrect because a 1G reference bandwidth would still result in both 1 GbE and 10 GbE interfaces having a cost of 1.

Question: 3

You have two multicast receivers connected to the same VLAN. You notice that the switch that they are connected to is forwarding multicast traffic out of all the ports in the same VLAN, instead of just the two ports for the connected multicast receivers
In this scenario, what would you configure to optimize multicast forwarding?

- A. promiscuous mode
- B. spanning tree
- C. IGMP snooping
- D. IRB interface

Answer: C

Explanation:

In a standard Layer 2 switch environment without specific multicast optimizations, multicast traffic is treated similarly to broadcast traffic and is flooded out of all ports within a VLAN (except the port where it was received). This is inefficient as it consumes unnecessary bandwidth on ports where no receivers are present.

To optimize this, IGMP Snooping is the standard solution.

How it works: The switch "snoops" or monitors IGMP (Internet Group Management Protocol) messages between hosts and routers.

Result: When a receiver sends an IGMP Membership Report (Join) for a specific multicast group, the switch records which port that request came from. It then builds a Layer 2 multicast forwarding table so that traffic for that group is only forwarded to the specific ports that have active listeners.

Junos OS 24.4 Configuration: On modern Juniper switches, you enable this globally or per VLAN using the `set protocols igmp-snooping vlan <vlan-name>` command.

Option A (Promiscuous mode) is used for monitoring or specific security configurations (like Private VLANs) and would not limit multicast flooding. Option B (Spanning Tree) manages loop prevention and does not differentiate between traffic types for forwarding optimization. Option D (IRB interface) is used for Layer 3 routing between VLANs but does not inherently provide Layer 2 multicast optimization within a single VLAN.

Question: 4

You are implementing an EVPN-VXLAN edge-routed bridging design using Layer 3 gateway operations. In this scenario, which statement is correct?

- A. Each distribution switch has unique IP addresses for IRB interfaces; routing protocols run on IRB interfaces
- B. Distribution switches share the same anycast IP addresses for IRB interfaces; routing occurs at the distribution layer.
- C. Only core switches have IRB interfaces; all Layer 3 routing happens in the core.
- D. IRB interfaces are disabled; all routing happens through external routers only.

Answer: B

Explanation:

In an EVPN-VXLAN Edge-Routed Bridging (ERB) architecture, also known as a collapsed fabric, the Layer 3 default gateway functionality is moved from the core/spine layer down to the edge (the leaf or distribution layer).

Anycast Gateways: To support seamless host mobility and redundancy, multiple distribution/leaf switches are configured with the same anycast IP address and MAC address on their IRB interfaces for a given VLAN. This allows a host to move between different switches without needing to update its default gateway configuration or ARP cache.

Distributed Routing: Routing occurs locally at the edge (distribution layer). Traffic destined for a different subnet is routed by the first switch it hits (the ingress leaf), rather than being backhauled to a central core router.

Symmetric vs. Asymmetric IRB: Junos OS 24.4 supports both models, but the ERB design typically utilizes symmetric routing for better scalability, where each leaf only needs to know the routes for its locally connected VNIs and uses a transit VNI for inter-subnet communication.

Option A is incorrect because while unique IPs can be used (Method 1 in some docs), the defining characteristic of an efficient ERB design is the use of shared Anycast IPs for the gateway. Option C describes a Centrally-Routed Bridging (CRB) design, not ERB.

Question: 5

Which statement about LLDP and LLDP-MED operations on EX Series devices is correct?

- A. LLDP only operates on interfaces configured for Layer 2
- B. EX Series devices flood LLDP frames across a Layer 2 domain to calculate a network topology
- C. EX Series devices support LLDP-MED power negotiation, enabling dynamic allocation of PoE power based on endpoint device needs.
- D. LLDP-MED focuses on discovering network connectivity devices like routers and switches.

Answer: C

Explanation:

Junos OS 24.4 on EX Series switches provides robust support for LLDP (Link Layer Discovery Protocol) and its extension, LLDP-MED (Media Endpoint Discovery).

LLDP-MED Power Negotiation: This feature allows a switch (Power Sourcing Equipment or PSE) and a connected device (Powered Device or PD), such as an IP phone or access point, to negotiate power requirements beyond the standard IEEE 802.3af/at classes. The switch can dynamically allocate the exact amount of power the device needs (in 0.1W increments), which optimizes the power budget of the switch.

LLDP Scope: LLDP is a Link Layer protocol (Layer 2), but it is not restricted to Layer 2 interfaces; it can also operate on Layer 3 interfaces to advertise system identity and capabilities. This makes Option A incorrect.

Link-Local Protocol: LLDP frames use a specific multicast MAC address (01:80:c2:00:00:0e) that is not flooded or forwarded by switches. They are strictly link-local between two directly connected neighbors. This makes Option B incorrect.

Endpoint Focus: LLDP-MED is specifically designed for Media Endpoint Devices (like VoIP phones), providing TLVs for network policy (VLAN/QoS), location identification, and inventory management. Standard LLDP is used for discovering network connectivity devices. This makes Option D incorrect.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>