

Exam JN0-683 Cost & Passing JN0-683 Score Feedback



P.S. Free 2025 Juniper JN0-683 dumps are available on Google Drive shared by Test4Cram: <https://drive.google.com/open?id=1Zz686mwVUXcNc196j8to-n74XQ9LKLl2>

Customers always attach great importance to the quality of JN0-683 exam torrent. We can guarantee that our study materials deserve your trustee. We have built good reputation in the market now. After about ten years' development, we have owned a perfect quality control system. All JN0-683 exam prep has been inspected strictly before we sell to our customers. Generally, they are very satisfied with our JN0-683 Exam Torrent. Also, some people will write good review guidance for reference. Maybe it is useful for your preparation of the JN0-683 exam. In addition, you also can think carefully which kind of study materials suit you best. If someone leaves their phone number or email address in the comments area, you can contact them directly to get some useful suggestions.

With the consistent reform in education, our JN0-683 test question also change with the newest education regulation. We have strong confidence in offering the first-class JN0-683 study prep to our customers. So what you have learned is fully conforming to the latest test syllabus. Also, our specialists can predicate the JN0-683 exam precisely. Firstly, our company has summed up much experience after so many years' accumulation. The model test is very important. You are advised to master all knowledge of the model test. Therefore, we sincerely wish you can attempt to our JN0-683 Test Question. Practice and diligence make perfect. Every one looks forward to becoming an excellent person. You will become the lucky guys after passing the JN0-683 exam.

>> Exam JN0-683 Cost <<

Passing JN0-683 Score Feedback, Reliable JN0-683 Test Notes

In the past few years, our JN0-683 study materials have helped countless candidates pass the JN0-683 exam. After having a related certification, some of them encountered better opportunities for development, some went to great companies, and some became professionals in the field. JN0-683 Study Materials have stood the test of time and market and received countless praises. Through the good reputation of word of mouth, more and more people choose to use JN0-683 study torrent to prepare for the JN0-683 exam, which makes us very gratified.

Juniper JN0-683 Exam Syllabus Topics:

Topic	Details
Topic 1	• VXLAN: This part requires knowledge of VXLAN, particularly how the control plane manages communication between devices, while the data plane handles traffic flow. Demonstrate knowledge of how to configure, Monitor, or Troubleshoot VXLAN.
Topic 2	• EVPN-VXLAN Signaling: This section assesses an understanding of Ethernet VPN (EVPN) concepts, including route types, multicast handling, and Multiprotocol BGP (MBGP). It also covers EVPN architectures like CRB and ERB, MAC learning, and symmetric routing.

- Data Center Deployment and Management: This section assesses the expertise of data center networking professionals like architects and engineers, focusing on key deployment concepts. Topics include Zero-touch provisioning (ZTP), which automates device setup in data centers without manual input.

Juniper Data Center, Professional (JNCIP-DC) Sample Questions (Q66-Q71):

NEW QUESTION # 66

You are designing an IP fabric for a large data center, and you are concerned about growth and scalability. Which two actions would you take to address these concerns? (Choose two.)

- A. Use EX4300 Series devices as the spine devices.
- B. Use OFX5700 Series devices as the super spines.
- C. Design a three-stage Clos IP fabric.
- D. Design a five-stage Clos IP fabric.

Answer: B,D

NEW QUESTION # 67

You are deploying an IP fabric with an oversubscription ratio of 3:1. In this scenario, which two statements are correct? (Choose two.)

- A. The oversubscription ratio remains the same when you remove leaf devices.
- B. The oversubscription ratio decreases when you add leaf devices.
- C. The oversubscription ratio remains the same when you add leaf devices.
- D. The oversubscription ratio increases when you remove leaf devices.

Answer: C,D

Explanation:

* Understanding Oversubscription Ratio in IP Fabrics:

* The oversubscription ratio in an IP fabric typically refers to the ratio of the available bandwidth at the edge of the network (leaves) to the available bandwidth at the core or spine. A 3:1 oversubscription ratio means that for every 3 units of bandwidth at the leaves, there is 1 unit of bandwidth at the spine.

* Impact of Adding or Removing Leaf Devices:

* Removing Leaf Devices: When you remove leaf devices, the amount of total edge bandwidth decreases while the bandwidth in the spine remains constant. This causes the oversubscription ratio to increase because there is now less total bandwidth to distribute across the same amount of spine bandwidth.

* Adding Leaf Devices: Conversely, when you add leaf devices, the total edge bandwidth increases. Since the spine bandwidth remains the same, the oversubscription ratio would remain the same if the additional leaves consume their share of the available bandwidth proportionally.

Conclusion:

* Option C: Correct - Removing leaf devices increases the oversubscription ratio.

* Option D: Correct - Adding leaf devices typically maintains the oversubscription ratio assuming uniform bandwidth distribution.

NEW QUESTION # 68

Which statement is correct about a collapsed fabric EVPN-VXLAN architecture?

- A. Using Virtual Chassis at the leaf layer increases resiliency.
- B. Fully meshed back-to-back links are needed between the spine devices.
- C. Border gateway functions occur on border leaf devices.
- D. It supports multiple vendors in the fabric as long as all the spine devices are Juniper devices deployed with L2 VTEPs

Answer: C

Explanation:

* Collapsed Fabric Architecture:

* A collapsed fabric refers to a simplified architecture where the spine and leaf roles are combined, often reducing the number of

devices and links required.

- * In this architecture, the spine typically handles core switching, while leaf switches handle both access and distribution roles.

- * Understanding Border Gateway Functionality:

- * Border gateway functions include connecting the data center to external networks or other data centers.

- * In a collapsed fabric, these functions are usually handled at the leaf level, particularly on border leaf devices that manage the ingress and egress of traffic to and from the data center fabric.

- * Correct Statement:

- * D. Border gateway functions occur on border leaf devices: This is accurate in collapsed fabric architectures, where the border leaf devices take on the role of managing external connections and handling routes to other data centers or the internet.

Data Center References:

- * The collapsed fabric model is advantageous in smaller deployments or scenarios where simplicity and cost-effectiveness are prioritized. It reduces complexity by consolidating functions into fewer devices, and the border leaf handles the critical task of interfacing with external networks.

In conclusion, border gateway functions are effectively managed at the leaf layer in collapsed fabric architectures, ensuring that the data center can communicate with external networks seamlessly.

NEW QUESTION # 69

Exhibit.



You want to enable the border leaf device to send Type 5 routes of local networks to the border leaf device in another data center. What must be changed to the configuration shown in the exhibit to satisfy this requirement?

- A. Add encapsulation vxlan to the evpn hierarchy.
- **B. Move vrf-target target: 65000:1 to the evpn hierarchy.**
- C. Add a VLAN configuration with an 13-interface to the tenant1 routing instance.
- D. Change: 5001 in the route-distinguisher to : 10010.

Answer: B

Explanation:

In this scenario, you want the border leaf device to advertise Type 5 EVPN routes to another border leaf in a different data center. Type 5 routes in EVPN are used to advertise IP prefixes, which means that for proper route advertisement, you need to configure the correct settings within the evpn hierarchy.

Step-by-Step Analysis:

- * Understanding EVPN Type 5 Routes:

- * EVPN Type 5 routes are used to advertise IP prefixes across EVPN instances, which allow different data centers or networks to exchange routing information effectively.

- * VRF Target Setting:

* The vrf-target configuration is crucial because it defines the export and import policies for the VRF within the EVPN instance. For EVPN Type 5 routes to be advertised to other border leaf devices, the vrf-target needs to be correctly configured under the evpn hierarchy, not just within the routing instance.

Command to solve this:

move vrf-target target:65000:1 to evpn

* Other Options:

* Option B: Adding a VLAN configuration would not address the requirement to advertise Type 5 routes.

* Option C: Adding VXLAN encapsulation may be necessary for other scenarios but does not directly address the Type 5 route advertisement.

* Option D: Changing the route-distinguisher will differentiate routes but does not impact the advertisement of Type 5 routes to other data centers.

By moving the vrf-target to the evpn hierarchy, you enable the proper route advertisement, ensuring that the Type 5 routes for local networks are shared with other data center border leaf devices. This is aligned with best practices for multi-data center EVPN implementations, which emphasize the correct placement of routing policies within the EVPN configuration.

NEW QUESTION # 70

You are asked to set up an IP fabric that supports AI or ML workloads. You have chosen to use lossless Ethernet in this scenario, which statement is correct about congestion management?

- A. Only the source and destination devices need ECN enabled.
- B. The switch experiencing the congestion notifies the source device.
- C. ECN marks packets based on WRED settings.
- D. ECN is negotiated only among the switches that make up the IP fabric for each queue.

Answer: C

Explanation:

Step 1: Understand the Context of Lossless Ethernet and Congestion Management

* Lossless Ethernet in IP Fabrics: AI/ML workloads often require high throughput and low latency, with minimal packet loss.

Lossless Ethernet is achieved using mechanisms like Priority Flow Control (PFC), which pauses traffic on specific priority queues to prevent drops during congestion. This is common in data center IP fabrics supporting RoCE (RDMA over Converged Ethernet), a protocol often used for AI/ML workloads.

* Congestion Management: In a lossless Ethernet environment, congestion management ensures that the network can handle bursts of traffic without dropping packets. Two key mechanisms are relevant here:

* Priority Flow Control (PFC): Pauses traffic on a specific queue to prevent buffer overflow.

* Explicit Congestion Notification (ECN): Marks packets to signal congestion, allowing end devices to adjust their transmission rates (e.g., by reducing the rate of RDMA traffic).

* AI/ML Workloads: These workloads often use RDMA (e.g., RoCEv2), which relies on ECN to manage congestion and PFC to ensure no packet loss. ECN is critical for notifying the source device of congestion so it can throttle its transmission rate.

Step 2: Evaluate Each Statement

A: The switch experiencing the congestion notifies the source device.

* In a lossless Ethernet environment using ECN (common with RoCEv2 for AI/ML workloads), when a switch experiences congestion, it marks packets with an ECN flag (specifically, the ECN-Echo bit in the IP header). These marked packets are forwarded to the destination device.

* The destination device, upon receiving ECN-marked packets, sends a congestion notification back to the source device (e.g., via a CNP - Congestion Notification Packet in RoCEv2). The source device then reduces its transmission rate to alleviate congestion.

* How this works in Junos: On Juniper switches (e.g., QFX series), you can configure ECN by setting thresholds on queues. When the queue depth exceeds the threshold, the switch marks packets with ECN. For example:

text

Copy

```
class-of-service {
  congestion-notification-profile ecn-profile {
    queue 3 {
      ecn threshold 1000; # Mark packets when queue depth exceeds 1000 packets
    }
  }
}
```

* Analysis: The switch itself does not directly notify the source device. Instead, the switch marks packets, and the destination device notifies the source. This statement is misleading because it implies direct notification from the switch to the source, which is not how ECN works in this context.

* This statement is false.

B: Only the source and destination devices need ECN enabled.

* ECN requires support at multiple levels:

* Source and Destination Devices: The end devices (e.g., servers running AI/ML workloads) must support ECN. For example, in RoCEv2, the NICs on the source and destination must be ECN-capable to interpret ECN markings and respond to congestion (e.g., by sending CNPs).

* Switches in the IP Fabric: The switches must also support ECN to mark packets during congestion. In an IP fabric, all switches along the path need to be ECN-capable to ensure consistent congestion management. If any switch in the path does not support ECN, it might drop packets instead of marking them, breaking the lossless behavior.

* Junos Context: On Juniper devices, ECN is enabled per queue in the class-of-service (CoS) configuration, as shown above. All switches in the fabric should have ECN enabled for the relevant queues to ensure end-to-end congestion management.

* Analysis: This statement is incorrect because it's not just the source and destination devices that need ECN enabled—switches in the fabric must also support ECN for it to work effectively across the network.

* This statement is false.

C: ECN marks packets based on WRED settings.

* WRED (Weighted Random Early Detection): WRED is a congestion avoidance mechanism that drops packets probabilistically before a queue becomes full, based on thresholds. It's commonly used in non-lossless environments to manage congestion by dropping packets early.

* ECN with WRED: In a lossless Ethernet environment, ECN can work with WRED-like settings, but instead of dropping packets, it marks them with an ECN flag. In Junos, ECN is configured with thresholds that determine when to mark packets, similar to how WRED uses thresholds for dropping packets. For example:

```
class-of-service {
  congestion-notification-profile ecn-profile {
    queue 3 {
      ecn threshold 1000; # Mark packets when queue depth exceeds 1000 packets
    }
  }
}
```

* How ECN Works in Junos: The ECN threshold acts like a WRED profile, but instead of dropping packets, the switch sets the ECN bit in the IP header when the queue depth exceeds the threshold. This is a key mechanism for congestion management in lossless Ethernet for AI/ML workloads.

* Analysis: This statement is correct. ECN in Junos uses settings similar to WRED (i.e., thresholds) to determine when to mark packets, but marking replaces dropping in a lossless environment.

* This statement is true.

D: ECN is negotiated only among the switches that make up the IP fabric for each queue.

* ECN Negotiation: ECN is not a negotiated protocol between switches. ECN operates at the IP layer, where switches mark packets based on congestion, and end devices (source and destination) interpret those markings. There's no negotiation process between switches for ECN.

* Comparison with PFC: This statement might be confusing ECN with PFC, which does involve negotiation. PFC uses LLDP (Link Layer Discovery Protocol) or DCBX (Data Center Bridging Exchange) to negotiate lossless behavior between switches and endpoints for specific priority queues.

* Junos Context: In Junos, ECN is a unilateral configuration on each switch. Each switch independently decides to mark packets based on its own queue thresholds, and there's no negotiation between switches for ECN.

* Analysis: This statement is incorrect because ECN does not involve negotiation between switches. It's a marking mechanism that operates independently on each device.

* This statement is false.

Step 3: Identify the Correct Statement

From the analysis:

* A is false: The switch does not directly notify the source device; the destination does.

* B is false: ECN must be enabled on switches in the fabric, not just the source and destination.

* C is true: ECN marks packets based on thresholds, similar to WRED settings.

* D is false: ECN is not negotiated between switches.

The question asks for the correct statement about congestion management, and C is the only true statement.

However, the question asks for two statements, which suggests there might be a discrepancy in the question framing, as only one statement is correct based on standard Juniper and lossless Ethernet behavior. In such cases, I'll assume the intent is to identify the single correct statement about congestion management, as

"choose two" might be a formatting error in this context.

Step 4: Provide Official Juniper Documentation Reference

Since I don't have direct access to Juniper's proprietary documents, I'll reference standard Junos documentation practices, such as those found in the Junos OS Class of Service Configuration Guide from Juniper's TechLibrary:

* ECN in Lossless Ethernet: The Junos OS CoS Configuration Guide explains that ECN is used in lossless Ethernet environments

The configuration uses a threshold-based mechanism, similar to WRED, but marks packets instead of dropping them. This is documented under the section for congestion notification profiles.

This aligns with the JNCIP-DC exam objectives, which include understanding congestion management mechanisms like ECN and PFC in data center IP fabrics, especially for AI/ML workloads.

• • • • •

Passing JN0-683 Score Feedback: https://www.test4cram.com/JN0-683_real-exam-dumps.html

- BTW, DOWNLOAD part of Test4CramJN0-683 dumps from Cloud Storage: <https://drive.google.com/open?id=1Zz686mwVUXcNc196j8to-n74XO9LKLi2>