

적중율 좋은 JN0-683 인기 자격증 인증 시험 자료 공부 문제



Itexamdump JN0-683 최신 PDF 버전 시험 문제집을 무료로 Google Drive에서 다운로드하세요:
https://drive.google.com/open?id=11_zplzbLRO8uDOg6Xoh10-_tDO3INV0E

만약 Itexamdump를 선택하였다면 여러분은 받은 성공한 것입니다. 여러분은 아주 빠르게 안전하게 또 쉽게 Juniper JN0-683 인증 시험 자격증을 취득하실 수 있습니다. 우리 Itexamdump에서 제공되는 모든 덤프들은 모두 100% 보장도를 자랑하며 그리고 우리는 일년 무료 업데이트를 제공합니다.

Juniper JN0-683 시험요강:

주제	소개
주제 1	<ul style="list-style-type: none"> 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.
주제 2	<ul style="list-style-type: none"> Data Center Interconnect: For Data Center Engineers, this part focuses on interconnecting data centers, covering Layer 2 and Layer 3 stretching, stitching fabrics together, and using EVPN-signaled VXLAN for seamless communication between data centers.
주제 3	<ul style="list-style-type: none"> 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.

>> JN0-683 인기 자격증 인증 시험 자료 <<

JN0-683 최신 시험 덤프 자료, JN0-683 시험 대비 최신 덤프 공부 자료

Juniper JN0-683 인증 시험은 전문적인 관련 지식을 테스트하는 인증 시험입니다. Itexamdump는 여러분이 Juniper JN0-683 인증 시험을 통과할 수 있도록 도와주는 사이트입니다. 많은 분들이 많은 시간과 돈을 들여 혹은 여러 학원 등을 다니면서 Juniper JN0-683 인증 시험 패스에 노력을 다합니다. 하지만 우리 Itexamdump에서는 20시간 좌우만 투자하면 무조건 Juniper JN0-683 시험을 패스할 수 있도록 도와드립니다.

최신 JNCIP-DC JN0-683 무료 샘플 문제 (Q26-Q31):

질문 # 26

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. ECN marks packets based on WRED settings.
- B. Only the source and destination devices need ECN enabled.

- C. ECN is negotiated only among the switches that make up the IP fabric for each queue.
- D. The switch experiencing the congestion notifies the source device.

정답: A

설명:

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.

DECN 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 (e.g., with RoCE) to mark packets when queue thresholds are exceeded.

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.

* No Negotiation for ECN: The same guide clarifies that ECN operates independently on each switch, with no negotiation between devices, unlike PFC, which uses DCBX for negotiation.

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.

질문 # 27

Click the Exhibit button.

```

user@Leaf-1> show configuration switch-options
service-id 1;
route-distinguisher 192.168.100.51:1;
vrf-target target:65000:55;
user@Leaf-2> show configuration switch-options
vtep-source-interface lo0.0;
route-distinguisher 192.168.100.51:2;
vrf-target target:65000:54;

```

Connections between hosts connected to Leaf 1 and Leaf-2 are not working correctly.

Referring to the exhibit, which two configuration changes are required to solve the problem? (Choose two.)

- A. Configure the set switch-options route-distinguisher 192.168.100.51:2 parameter on Leaf-1.
- B. Configure the set switch-options service-id 1 parameter on Leaf-2.
- C. Configure the set switch-options vrf-target target:65000:55 parameter on Leaf-2.
- D. Configure the set switch-options vtep-source-interface lo0.0 parameter on Leaf-1.

정답: C,D

질문 # 28

As part of the onboarding process for new switches being added to your data centers, your company uses Juniper Networks' ZTP process. As part of the ZTP process, a script is executed by the devices being onboarded.

Which statement is correct in this scenario?

- A. The Junos ZTP process supports Python, SLAX, and Perl.
- B. The Junos ZTP process supports Jscript, Ansible, and Perl.
- C. The Junos ZTP process supports Shell, Python, and SLAX.
- D. The Junos ZTP process supports Shell, Jscript, and Ansible.

정답: C

설명:

Juniper Networks' ZTP (Zero Touch Provisioning) process automates the deployment of new devices by allowing them to fetch and execute scripts for configuration and setup as they are powered on and connected to the network.

Supported Scripting Languages:

The Junos OS supports several scripting languages that can be used during the ZTP process:

Shell scripts are often used for general automation tasks. Python is a widely supported language in Junos, offering powerful scripting capabilities for automating network tasks.

SLAX (Service Logic Execution Environment) is a scripting language specific to Junos, designed to automate configuration tasks and simplify network operations.

질문 # 29

You manage an IP fabric with an EVPN-VXLAN overlay. You have multiple tenants separated using multiple unique VRF instances. You want to determine the routing information that belongs in each routing instance's routing table.

In this scenario, which property is used for this purpose?

- A. the route distinguisher value
- B. the VRF target community
- C. the routing instance type
- D. the VRF table label

정답: A

설명:

* Understanding VRF and Routing Instances:

* In an EVPN-VXLAN overlay network, multiple tenants are separated using unique VRF (Virtual Routing and Forwarding) instances. Each VRF instance maintains its own routing table, allowing for isolated routing domains within the same network infrastructure.

* Role of Route Distinguisher:

* Route Distinguisher (RD):The RD is a unique identifier used in MPLS and EVPN environments to distinguish routes belonging to different VRFs. The RD is prepended to the IP address in the route advertisement, ensuring that routes from different tenants remain unique even if they use the same IP address range.

* Correct Property:

* D. the route distinguisher value:This is the correct answer because the RD is crucial in determining which routing information belongs to which VRF instance. It ensures that each VRF's routing table only contains relevant routes, maintaining isolation between tenants.

Data Center References:

* The RD is a key element in MPLS and EVPN-based multi-tenant environments, ensuring proper routing segregation and isolation for different VRFs within the data center fabric.

질문 # 30

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

정답: A,D

설명:

* 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.

질문 # 31

.....

경쟁율이 심한 IT시대에Juniper JN0-683인증시험을 패스하여 자격증을 취득함으로 IT업계 관련 직종에 종사하고자 하는 분들에게는 아주 큰 가산점이 될수 있고 자신만의 위치를 보장할수 있으며 더욱이는 한층 업된 삶을 누릴수 있을수도 있습니다. Juniper JN0-683 덤프로 Juniper JN0-683 시험에서 실패하면 덤프비용을 보상해드리기에 안심하고 시험준비하셔야 합니다.

JN0-683최신 시험덤프자료 : <https://www.itexamdump.com/JN0-683.html>

- JN0-683최신 덤프문제모음집 圖 JN0-683최신버전덤프 □ JN0-683높은 통과율 인기 덤프자료 ④ 오픈 웹 사이트 > kr.fast2test.com □ 검색 { JN0-683 } 무료 다운로드 JN0-683높은 통과율 인기 덤프자료
- 최근 인기시험 JN0-683인기자격증 인증시험자료 대비자료 □ > JN0-683 <를 무료로 다운로드하려면 { www.itdumpskr.com } 웹사이트를 입력하세요 JN0-683퍼펙트 덤프 최신 데모문제
- JN0-683인기자격증 인증시험자료 시험 최신버전 덤프자료 샘플문제 □ 시험 자료를 무료로 다운로드하려

