

# JNCIS - Enterprise Routing and Switching (JNCIS-ENT) Practice Exam (Sample)

## Study Guide



**Everything you need from our exam experts!**

**Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.**

**ALL RIGHTS RESERVED.**

**No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.**

**Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.**

**SAMPLE**

## **Questions**

SAMPLE

- 1. What is the default TTL value in a VRRP advertisement?**
  - A. 0 TTL**
  - B. 1 TTL**
  - C. 64 TTL**
  - D. 255 TTL**
- 2. To enable nonstop active routing (NSR), which two commands must be used?**
  - A. Set routing-options nonstop-routing**
  - B. Set routing-options graceful-restart**
  - C. Set chassis redundancy graceful-switchover**
  - D. Set routing-options graceful-switchover**
- 3. When a switch receives a broadcast frame, how does it handle it?**
  - A. It sends the frame to all switch ports, regardless of VLAN.**
  - B. It only forwards it to the designated VLAN ports.**
  - C. It queues the frame until the destination port becomes available.**
  - D. It drops the frame if no MAC address is learned.**
- 4. What determines which switch will become the designated bridge on the same LAN segment?**
  - A. The port cost for the associated interface**
  - B. The bridge ID priority**
  - C. The total root cost of the switch**
  - D. The priority of the associated interface**
- 5. Which statement describes the default Junos OS behavior for OSPF?**
  - A. External LSAs are advertised in a stub area.**
  - B. An ABR does not announce a default route into a stub area.**
  - C. Stub area internal routers generate a default route.**
  - D. Only totally stubby areas need a default route.**

- 6. You want to allow RIP routes to be redistributed by an ASBR into the connected OSPF area. Which OSPF area type would be used to satisfy the requirement?**
- A. Transit area**
  - B. Stub area**
  - C. Totally stubby area**
  - D. Not-so-stubby area**
- 7. An ABR is connected to an NSSA, which configuration disables the exporting of Type 7 LSAs into the connected NSSA?**
- A. No-type7**
  - B. NSSA**
  - C. Area-range**
  - D. No-nssa-abr**
- 8. Which type of OSPF area must not contain any ASBRs and must not receive Extern LSAs?**
- A. Not-so-stubby area**
  - B. Totally stubby area**
  - C. Stub area**
  - D. Backbone area**
- 9. Which option shows the configuration syntax for a routed VLAN interface on an EX Series switch?**
- A. interfaces { vlan { unit 10 { family inet { address 1.1.1.1/24; } } } } vlans { VLAN10 { vlan-id 10; routing-interface vlan.10; } }**
  - B. interfaces { vlan { unit 10 { family inet { address 1.1.1.1/24; } } } } vlans { VLAN10 {vlan-id 10; l3-interface vlan.10; } }**
  - C. vlans { VLAN10 { vlan-id 10; family inet { address 1.1.1.1/24; } } }**
  - D. interfaces { vlan { unit 10 { family inet { address 1.1.1.1/24; } } } } vlans { VLAN10 { vlan-id 10; } }**

- 10. What causes an increase in dropped packets for a GRE tunnel when the DF bit is set?**
- A. The GRE tunnel has a lower MTU than the physical interface.**
  - B. GRE tunnels do not support fragmentation.**
  - C. GRE tunnels do not support the DF bit.**
  - D. The GRE tunnel has a higher MTU than the physical interface.**

## **Answers**

SAMPLE

1. D
2. A
3. A
4. B
5. B
6. D
7. D
8. C
9. B
10. A

SAMPLE



## **Explanations**

SAMPLE

**1. What is the default TTL value in a VRRP advertisement?**

- A. 0 TTL**
- B. 1 TTL**
- C. 64 TTL**
- D. 255 TTL**

In VRRP (Virtual Router Redundancy Protocol), the default Time to Live (TTL) value for advertisement messages is set to 255. This indicates that the VRRP advertisements can traverse multiple networks before being discarded. TTL is a mechanism used to prevent packets from circulating indefinitely in case of routing loops. By having a TTL value of 255, the VRRP advertisements can be received by devices on the same local segment or by those upstream, as they can effectively reach multiple hops without being dropped. Choosing this default value allows network devices to participate in routing redundancy efficiently, ensuring quick failover should the primary router fail. This is especially important in scenarios where multiple routers are configured to manage the same virtual IP address, and timely communication between them is crucial for maintaining network availability and reliability.

**2. To enable nonstop active routing (NSR), which two commands must be used?**

- A. Set routing-options nonstop-routing**
- B. Set routing-options graceful-restart**
- C. Set chassis redundancy graceful-switchover**
- D. Set routing-options graceful-switchover**

To enable nonstop active routing (NSR) in a Junos environment, it is crucial to understand the functions that NSR serves in maintaining routing protocols during a control plane failover. One of the key requirements for enabling NSR is setting the routing options specifically for nonstop routing. Using the command that sets `routing-options nonstop-routing` is essential as it directly enables NSR. This option ensures that the router maintains the routing information in the routing table even if a routing protocol is in the process of restarting. With this configuration, the router can seamlessly continue to forward packets without disruption. In conjunction with this command, it's common to utilize another command related to graceful restart to ensure compatibility and to enhance the overall stability of routing protocols during a switchover. However, the exact combination that is necessary for NSR activation specifically hinges on enabling nonstop routing through the aforementioned command. The interplay of these commands in NSR is critical, as enabling nonstop routing is the foundational step for achieving the high availability that NSR provides, allowing for minimal interruption of service.

**3. When a switch receives a broadcast frame, how does it handle it?**

**A. It sends the frame to all switch ports, regardless of VLAN.**

**B. It only forwards it to the designated VLAN ports.**

**C. It queues the frame until the destination port becomes available.**

**D. It drops the frame if no MAC address is learned.**

When a switch receives a broadcast frame, it operates under the principle that broadcast frames are intended to reach all devices within the same network segment or VLAN. Therefore, the switch sends the broadcast frame out through all of its active ports, ensuring it reaches every device that is part of the segment, regardless of VLAN boundaries. This behavior is fundamental to how Ethernet and Layer 2 switching operates. In a typical scenario, a broadcast frame is meant for all devices within a local broadcast domain, which is defined by a VLAN. By forwarding the broadcast out all switch ports, the switch enables devices that need to receive the broadcast message to do so, thereby allowing for processes such as ARP requests and other necessary communication to occur within the local network. The other choices presented do not align with how broadcast frames are handled by switches. In particular, forwarding only to designated VLAN ports does not accommodate all devices that could be part of the same broadcast domain if multiple VLANs are in use. Moreover, queuing the frame is not a standard action for broadcasts; they are typically transmitted immediately. Lastly, dropping the frame if no MAC address is known contradicts the basic principle of broadcasting, which is aimed at reaching all devices, regardless of prior knowledge about them. Thus,

**4. What determines which switch will become the designated bridge on the same LAN segment?**

**A. The port cost for the associated interface**

**B. The bridge ID priority**

**C. The total root cost of the switch**

**D. The priority of the associated interface**

The designated bridge on a LAN segment is chosen based on the bridge ID priority, which is a combination of the bridge priority and the MAC address of the switch. The bridge ID essentially identifies the switch in the Spanning Tree Protocol (STP) process. When determining which switch will become the designated bridge, the switch with the lowest bridge ID is chosen. The bridge ID is composed of two parts: the bridge priority (a configurable value) and the switch's MAC address. The switch with the highest priority (lowest numerical value) is preferred; if there is a tie, the switch with the lowest MAC address is selected. This prioritization ensures a stable and predictable network topology by allowing the network administrator to influence which switch takes on the designated role, facilitating better control over network resources and traffic flow. The designated bridge is crucial because it forwards frames to and from the root bridge across the LAN segment, maintaining the Spanning Tree topology and preventing loops. Other options may refer to different aspects of STP, such as port costs and root costs, but they do not directly determine the designated bridge on a segment. The choice of the designated bridge specifically hinges on the bridge ID priority as the primary criterion.

**5. Which statement describes the default Junos OS behavior for OSPF?**

- A. External LSAs are advertised in a stub area.
- B. An ABR does not announce a default route into a stub area.**
- C. Stub area internal routers generate a default route.
- D. Only totally stubby areas need a default route.

The statement reflecting the default behavior of the Junos OS for OSPF is that an ABR (Area Border Router) does not announce a default route into a stub area. In OSPF, stub areas are designed to suppress external route advertisements, which include routes to external networks (NSSA) and inter-area routes. The primary purpose of a stub area is to reduce the amount of routing information exchanged and to maintain a simpler routing table within the area. In a stub area, routers will have a simplified view of the network, and the default route (0.0.0.0) can be configured to be injected by the ABR. However, the default behavior without special configuration is for the ABR to not provide a default route into the stub area, maintaining the integrity and simplicity of routing within that area. This is in accordance with the OSPF standards governing stub area functionality. The other options suggest behaviors that conflict with the principles of OSPF's stub area configurations. For instance, external LSAs being advertised contradicts the purpose of a stub area, which aims to suppress such advertisements. Stub area internal routers do not naturally generate a default route unless configured specifically for this purpose, which is not the default behavior in

**6. You want to allow RIP routes to be redistributed by an ASBR into the connected OSPF area. Which OSPF area type would be used to satisfy the requirement?**

- A. Transit area
- B. Stub area
- C. Totally stubby area
- D. Not-so-stubby area**

To enable redistribution of RIP routes into an OSPF area through an Autonomous System Boundary Router (ASBR), a Not-So-Stubby Area (NSSA) is the appropriate choice. NSSA allows for route redistribution from an external protocol like RIP into the OSPF domain while maintaining OSPF's hierarchical design. In detail, NSSA is a variation of a stub area that allows external routes to be imported, but it's limited to only those external routes and not the entire OSPF routing table. It ensures that internal OSPF routes remain unaffected while still allowing selected external routing information to be utilized. This makes it possible for the ASBR to advertise RIP routes into the OSPF network without compromising the area's stub characteristics. In contrast, stub areas and totally stubby areas will block external routing information and limit the types of routes that can be advertised, making them unsuitable for the requirement to redistribute RIP routes. A transit area is designed to facilitate the passage of OSPF packets between different OSPF areas and doesn't inherently support the redistribution of routes from other protocols. Thus, selecting a Not-So-Stubby Area is essential to achieve the desired integration of RIP into the OSPF routing environment properly.

**7. An ABR is connected to an NSSA, which configuration disables the exporting of Type 7 LSAs into the connected NSSA?**

- A. No-type7**
- B. NSSA**
- C. Area-range**
- D. No-nssa-abr**

In an OSPF (Open Shortest Path First) network, an ABR (Area Border Router) connects different OSPF areas. When an ABR connects to a Not-So-Stubby Area (NSSA), it can receive Type 7 LSAs (Link-State Advertisements) but may need to control how these LSAs are handled to maintain the proper OSPF routing environment. The correct configuration that disables the exporting of Type 7 LSAs into the connected NSSA is known as "no-nssa-abr." This command effectively prevents the ABR from translating any Type 7 LSAs received from the NSSA into Type 5 external LSAs, which would then be flooded into other areas. By doing this, the ABR ensures that the characteristics of the NSSA are preserved and that external routing information does not inadvertently propagate into the NSSA, which could lead to routing inefficiencies or even loops. In contrast, the other options do not specifically achieve this outcome. The "no-type7" command is not a valid OSPF command. The "NSSA" option doesn't specify a configuration command that would prevent Type 7 LSAs from being exported. The "area-range" command is used to

**8. Which type of OSPF area must not contain any ASBRs and must not receive Extern LSAs?**

- A. Not-so-stubby area**
- B. Totally stubby area**
- C. Stub area**
- D. Backbone area**

The choice for the type of OSPF area that must not contain any Autonomous System Boundary Routers (ASBRs) and must not receive External Link-State Advertisements (Extern LSAs) is a stub area. A stub area is designed to reduce the amount of routing information that routers need to maintain. By disallowing ASBRs, stub areas ensure that no external routes are advertised, which simplifies the routing table and minimizes the overhead associated with external route advertisements. Routers within a stub area will only exchange intra-area routes and inter-area routes, which makes the OSPF routing protocol more efficient within that area. By preventing Extern LSAs in the stub area, OSPF optimizes resource usage. This is especially beneficial in scenarios where external routes are not needed locally and helps maintain the integrity and simplicity of the routing topology within that area. The other types of areas, such as not-so-stubby areas or totally stubby areas, allow for some level of external routing information, with totally stubby areas being a more restricted version allowing only a default route. Backbone areas can contain ASBRs and receive external LSAs as they are the core of the OSPF network, connecting all other area types.

9. Which option shows the configuration syntax for a routed VLAN interface on an EX Series switch?

- A. `interfaces { vlan { unit 10 { family inet { address 1.1.1.1/24; } } } } vlans { VLAN10 { vlan-id 10; routing-interface vlan.10; } }`
- B. interfaces { vlan { unit 10 { family inet { address 1.1.1.1/24; } } } } vlans { VLAN10 {vlan-id 10; l3-interface vlan.10; } }**
- C. `vlans { VLAN10 { vlan-id 10; family inet { address 1.1.1.1/24; } } }`
- D. `interfaces { vlan { unit 10 { family inet { address 1.1.1.1/24; } } } } vlans { VLAN10 { vlan-id 10; } }`

The configuration syntax for a routed VLAN interface on an EX Series switch is accurately represented in the choice indicating that the VLAN is associated through the l3-interface statement. This is essential as it specifies that the VLAN is intended for Layer 3 routing, thus enabling IP addressing and routing capabilities. In this case, the presence of the l3-interface command directly links the VLAN itself to the specific IP address configured under that VLAN. This means the switch will treat this VLAN interface as a routed interface, allowing for inter-VLAN routing and ensuring that traffic can be appropriately managed and forwarded between VLANs. The inclusion of the family inet section under the unit also confirms that the interface is configured to support IPv4 addressing, which is critical for routing. Overall, using the l3-interface designation ensures that the switch understands that it should establish a routing context for the VLAN, facilitating communication with other networks or devices. Other options lack certain critical elements or employ incorrect terminology that does not accurately represent how routed VLANs need to be configured on EX series switches, resulting in improper configurations for intended functionalities.

10. What causes an increase in dropped packets for a GRE tunnel when the DF bit is set?

- A. The GRE tunnel has a lower MTU than the physical interface.**
- B. GRE tunnels do not support fragmentation.
- C. GRE tunnels do not support the DF bit.
- D. The GRE tunnel has a higher MTU than the physical interface.

When the Don't Fragment (DF) bit is set in a GRE tunnel, it indicates that the packet should not be fragmented during transmission. If the GRE tunnel has a lower Maximum Transmission Unit (MTU) than the physical interface, this can lead to packets being larger than the allowable size for the GRE tunnel. When packets that are too large attempt to traverse the tunnel without the ability to fragment, they are dropped because the network devices cannot break them down into smaller packets for transmission. Consequently, this results in an increase in dropped packets. The DF bit prevents fragmentation, so if the tunnel's MTU is lower than the size of the packet being sent, the packet will not be forwarded, leading to loss. This situation emphasizes the importance of ensuring that the MTU settings for GRE tunnels are correctly configured to prevent packet loss. In contrast, the other choices do not accurately address the relationship between MTU settings and the DF bit in GRE tunnels, leading to a misunderstanding of how packets are handled in such configurations.