CCNA (Cisco Certified Network Associate) Anki Practice Test (Sample)

Study Guide

BY EXAMZIFY

Everything you need from our exam experts!

Table of Contents

Introduction

Preparing for a certification exam can feel overwhelming, but with the right tools, it becomes an opportunity to build confidence, sharpen your skills, and move one step closer to your goals. At Examzify, we believe that effective exam preparation isn't just about memorization, it's about understanding the material, identifying knowledge gaps, and building the test-taking strategies that lead to success.

This guide was designed to help you do exactly that.

Whether you're preparing for a licensing exam, professional certification, or entry-level qualification, this book offers structured practice to reinforce key concepts. You'll find a wide range of multiple-choice questions, each followed by clear explanations to help you understand not just the right answer, but why it's correct.

The content in this guide is based on real-world exam objectives and aligned with the types of questions and topics commonly found on official tests. It's ideal for learners who want to:

• Practice answering questions under realistic conditions,
• Improve accuracy and speed,
• Review explanations to strengthen weak areas, and
• Approach the exam with greater confidence.

We recommend using this book not as a stand-alone study tool, but alongside other resources like flashcards, textbooks, or hands-on training. For best results, we recommend working through each question, reflecting on the explanation provided, and revisiting the topics that challenge you most.

Remember: successful test preparation isn't about getting every question right the first time, it's about learning from your mistakes and improving over time. Stay focused, trust the process, and know that every page you turn brings you closer to success.

Let's begin.

How to Use This Guide

This guide is designed to help you study more effectively and approach your exam with confidence. Whether you're reviewing for the first time or doing a final refresh, here's how to get the most out of your Examzify study guide:

1. Start with a Diagnostic Review

Skim through the questions to get a sense of what you know and what you need to focus on. Your goal is to identify knowledge gaps early.

2. Study in Short, Focused Sessions

Break your study time into manageable blocks (e.g. 30 – 45 minutes). Review a handful of questions, reflect on the explanations.

3. Learn from the Explanations

After answering a question, always read the explanation, even if you got it right. It reinforces key points, corrects misunderstandings, and teaches subtle distinctions between similar answers.

4. Track Your Progress

Use bookmarks or notes (if reading digitally) to mark difficult questions. Revisit these regularly and track improvements over time.

5. Simulate the Real Exam

Once you're comfortable, try taking a full set of questions without pausing. Set a timer and simulate test-day conditions to build confidence and time management skills.

6. Repeat and Review

Don't just study once, repetition builds retention. Re-attempt questions after a few days and revisit explanations to reinforce learning. Pair this guide with other Examzify tools like flashcards, and digital practice tests to strengthen your preparation across formats.

There's no single right way to study, but consistent, thoughtful effort always wins. Use this guide flexibly, adapt the tips above to fit your pace and learning style. You've got this!

Questions

1. **Can VTP clients add, modify, or delete VLANs?**
   A. Can
   B. Can't
   C. Can only add
   D. Can only modify

2. **What is the maximum number of hosts per Class C subnet?**
   A. 512
   B. 256
   C. 254
   D. 128

3. **Which protocol has less overhead compared to TCP?**
   A. ICMP
   B. UDP
   C. FTP
   D. SFTP

4. **When the administrative mode of SW1 G0/0 is dynamic auto and SW2 G0/0 is set to trunk, what is the operational mode of SW1 G0/0?**
   A. Trunk
   B. Dynamic Access
   C. Static Access
   D. Disabled

5. **What is the relationship between RSTP port cost and data traffic?**
   A. Higher port costs prioritize traffic
   B. Port costs determine the best path for traffic
   C. Port costs have no effect on traffic
   D. Lower port costs create more traffic

6. **What is the IPv6 multicast address that corresponds to all nodes?**

   A. FF02::2

   B. FF02::1

   C. FF00::1

   D. FF01::1

7. **What is the primary benefit of using EtherChannel?**

   A. Improved bandwidth and redundancy

   B. Reduced network complexity

   C. Enhanced packet security

   D. Simplified network management

8. **In which scenario would the 'Fragment Offset' field be used?**

   A. To specify packet urgency

   B. To indicate packet precedence

   C. To identify the position of a fragment

   D. To mark the end of transmission

9. **Which layer of the OSI model is responsible for error detection and correction relating to the physical layer?**

   A. Layer 1 - Physical

   B. Layer 3 - Network

   C. Layer 5 - Session

   D. Layer 2 - Data Link

10. **Which type of ARP message is sent without a request?**

    A. Request ARP

    B. Gratuitous ARP

    C. Response ARP

    D. Broadcast ARP

Answers

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

Explanations

1. Can VTP clients add, modify, or delete VLANs?

 A. Can

 B. Can't

 C. Can only add

 D. Can only modify

VTP (VLAN Trunking Protocol) clients cannot add, modify, or delete VLANs. The primary role of a VTP client is to receive VTP advertisements from a VTP server and update its VLAN configuration based on those advertisements. While VTP servers are authorized to make changes to VLAN configurations in the domain, VTP clients simply adopt the VLAN information provided by the server. This behavior ensures that VLAN management is centralized, reducing the risk of inconsistencies that can arise when multiple devices are allowed to alter VLAN configurations independently. The VTP client relies on the server for its VLAN database, which means it does not possess the capability to create or modify VLANs on its own. This distinction is crucial for maintaining network stability and uniformity across switches in a VTP domain.

2. What is the maximum number of hosts per Class C subnet?

 A. 512

 B. 256

 C. 254

 D. 128

In a Class C subnet, the default subnet mask is 255.255.255.0, which can also be represented as /24 in CIDR notation. This mask designates the first 24 bits of the IP address for the network portion and leaves the remaining 8 bits for the host portion. To calculate the number of usable hosts in a Class C subnet, we use the formula $2^n - 2$, where n is the number of bits available for host addresses. In this case, n is 8, since there are 8 bits available in a Class C subnet. Thus, the calculation becomes: $2^8 - 2 = 256 - 2 = 254$ The subtraction of 2 accounts for the network address and the broadcast address, which cannot be assigned to hosts. Therefore, the maximum number of usable hosts in a Class C subnet is 254.

3. Which protocol has less overhead compared to TCP?

A. ICMP

B. UDP

C. FTP

D. SFTP

UDP, or User Datagram Protocol, is a connectionless protocol that is part of the Internet Protocol suite. Unlike TCP (Transmission Control Protocol), which establishes a connection before data transmission and includes mechanisms for ensuring reliable delivery and ordering of packets, UDP does not require such extensive overhead.   UDP simply encapsulates data into packets called datagrams and sends them to the destination without establishing a connection or ensuring that the packets arrive in order. This lack of session management and acknowledgments means that UDP introduces significantly less overhead than TCP. As a result, UDP is often used in applications where speed is critical and where occasional data loss is acceptable, such as video streaming, online gaming, or voice over IP (VoIP). ICMP, while it also has lower overhead than TCP, primarily serves diagnostic and control purposes rather than general application data transport. FTP (File Transfer Protocol) and SFTP (SSH File Transfer Protocol) are both built on top of TCP and thus inherit its overhead characteristics, making them less suitable for scenarios prioritizing low-latency communication.

4. When the administrative mode of SW1 G0/0 is dynamic auto and SW2 G0/0 is set to trunk, what is the operational mode of SW1 G0/0?

A. Trunk

B. Dynamic Access

C. Static Access

D. Disabled

When examining the scenario where SW1's G0/0 port is set to dynamic auto and SW2's G0/0 port is configured as trunk, it's important to understand how the Dynamic Trunking Protocol (DTP) operates. In this case, SW2's port, being set to trunk, actively negotiates to establish a trunk connection with connected devices. The dynamic auto mode on SW1's G0/0 allows that port to respond to DTP negotiations. Since the port on SW2 is configured as trunk, it will send out DTP frames to negotiate the connection type. The interaction between these two settings leads to SW1's G0/0 port transitioning to trunk mode as it responds favorably to SW2's request for trunking. Thus, SW1 G0/0, in this situation, will successfully establish a trunk link with SW2 and operate in trunk mode. This is appropriate because the device is able to open up a trunk link through DTP negotiation established by the port on SW2 that is explicitly requesting a trunk configuration.

5. What is the relationship between RSTP port cost and data traffic?

A. Higher port costs prioritize traffic

B. Port costs determine the best path for traffic

C. Port costs have no effect on traffic

D. Lower port costs create more traffic

Port costs in Rapid Spanning Tree Protocol (RSTP) play a critical role in determining the best path for data traffic within a network. The protocol uses port costs to evaluate different potential paths that data might take towards its destination. Each interface or port on a switch is assigned a cost based on factors like bandwidth; lower costs are assigned to higher-bandwidth interfaces. When RSTP analyzes routes to determine the active path for forwarding traffic, it selects the path with the lowest cumulative port cost. Therefore, if multiple paths are available, traffic will traverse the path that has the least cost, thereby facilitating efficient data transmission. This mechanism is vital for optimizing network performance, minimizing latency, and ensuring that bandwidth is used effectively on high-capacity links. The other choices do not accurately convey the role of port costs in RSTP. For instance, higher port costs do not prioritize traffic; rather they make those paths less desirable for data transmission. Saying port costs have no effect on traffic contradicts the fundamental workings of RSTP. Moreover, lower port costs do not create more traffic; instead, they provide a more efficient route for existing traffic, enabling more effective data flow across the network.

6. What is the IPv6 multicast address that corresponds to all nodes?

A. FF02::2

B. FF02::1

C. FF00::1

D. FF01::1

The IPv6 multicast address that corresponds to all nodes in a local link is FF02::1. This address is specifically designed to reach all IPv6 nodes on the local segment of the network. When a packet is sent to this address, every IPv6-enabled device on the local link will receive the packet, which is essential for certain types of network communications where broadcasting to all devices is required. This address operates within the link-local scope, meaning it is only relevant within the confines of the local network segment. Other options listed do not serve this purpose: FF02::2 is designated for all routers on the local link, not for all nodes. FF00::1 is not a valid multicast address because multicast addresses must start with FF and are defined with certain format requirements. FF01::1 is intended for all nodes but within a specific scope of the network, seeking connectivity at the node-local level rather than the wider link-local scope. Thus, FF02::1 stands out as the correct choice for communicating with all applicable nodes on the local link in an IPv6 network.

7. What is the primary benefit of using EtherChannel?

A. Improved bandwidth and redundancy

B. Reduced network complexity

C. Enhanced packet security

D. Simplified network management

Using EtherChannel primarily improves bandwidth and redundancy within a network. By aggregating multiple physical links into a single logical link, EtherChannel allows for increased throughput since the combined bandwidth of the individual links is utilized. This means that if one link fails or is taken offline for maintenance, the remaining links continue to operate, ensuring that network connectivity remains intact. This redundancy enhances the reliability of the network, making it less susceptible to single points of failure. The increased capacity comes from the fact that multiple Ethernet links can handle more data than a single link alone, which effectively allows for load balancing across the aggregated links. This feature is crucial in high-traffic scenarios where bandwidth demands are significant. While the other options may provide certain benefits, they do not encapsulate the primary driving force behind the implementation of EtherChannel as effectively as the combination of improved bandwidth and redundancy. Reducing network complexity, enhancing packet security, and simplifying network management are benefits that may stem from proper configuration and use of EtherChannel, but they are not its main focus.

8. In which scenario would the 'Fragment Offset' field be used?

A. To specify packet urgency

B. To indicate packet precedence

C. To identify the position of a fragment

D. To mark the end of transmission

The 'Fragment Offset' field is used to identify the position of a fragment within an original packet when that packet is fragmented during transmission. Fragmentation occurs when a packet is too large to be transmitted over a network with a smaller maximum transmission unit (MTU). When a packet is fragmented, each fragment must contain information that helps reassemble the packet back to its original form at the destination. The Fragment Offset provides the necessary details about where a particular fragment fits in relation to the entire packet. It specifies the distance (in 8-byte blocks) of the fragment's data from the beginning of the original unfragmented packet. This crucial information allows the recipient to correctly order and combine all fragments to recreate the complete data packet accurately. In contrast, other options such as specifying packet urgency, indicating precedence, or marking the end of transmission do not involve the Fragment Offset. These functions are handled by different fields or protocols within the packet header and serve different purposes in network communication.

9. Which layer of the OSI model is responsible for error detection and correction relating to the physical layer?

A. Layer 1 - Physical

B. Layer 3 - Network

C. Layer 5 - Session

D. Layer 2 - Data Link

The Data Link layer, which is Layer 2 of the OSI model, plays a critical role in ensuring reliable communication across the physical medium by providing error detection and correction. This layer is responsible for data framing, which includes adding headers and trailers to packets to form frames. Among its key functions is the ability to detect errors that may occur while data is being transmitted over the physical layer. It does this by using techniques such as checksums or cyclic redundancy checks (CRC). When errors are detected, the Data Link layer can initiate mechanisms for error correction, which may include retransmission of the corrupted frames. This error handling capability is vital for maintaining data integrity and reliable communication, particularly in a noisy environment where transmission errors are likely to happen. In contrast, the responsibilities of the Physical layer (Layer 1) focus on the actual transmission of raw bits over the physical medium and do not include error checking or correction. The Network layer (Layer 3) manages routing and forwarding of data packets but does not focus on error detection at the link level, while the Session layer (Layer 5) is more concerned with managing sessions and dialogues between applications rather than data transmission reliability.

10. Which type of ARP message is sent without a request?

A. Request ARP

B. Gratuitous ARP

C. Response ARP

D. Broadcast ARP

Gratuitous ARP is a type of ARP message that is sent without a prior request. This message is typically generated by a host to inform the local network about its own IP-to-MAC address mapping. When a device sends a gratuitous ARP, it broadcasts its own MAC address to the network while claiming an IP address that it has taken ownership of or reconfigured. This serves several purposes: it helps other devices on the network update their ARP tables with the correct MAC address for a specific IP address, which is particularly useful in scenarios such as IP address conflicts, or when a device moves to a new network segment. Unlike a traditional ARP request, which is sent to inquire about the MAC address associated with a specified IP address, a gratuitous ARP does not wait for a response and actively announces the sender's own address information. This behavior is why it distinctly stands out from other types of ARP messages.

Next Steps

Congratulations on reaching the final section of this guide. You've taken a meaningful step toward passing your certification exam and advancing your career.

As you continue preparing, remember that consistent practice, review, and self-reflection are key to success. Make time to revisit difficult topics, simulate exam conditions, and track your progress along the way.

If you need help, have suggestions, or want to share feedback, we'd love to hear from you. Reach out to our team at hello@examzify.com.

Or visit your dedicated course page for more study tools and resources:

https://ccnaanki.examzify.com

We wish you the very best on your exam journey. You've got this!