

HPE ARUBA NETWORKING CERTIFIED PROFESSIONAL - CAMPUS ACCESS (ACP- CA) PRACTICE TEST (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

**VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE**

<https://hpeacpca.examzify.com>



Copyright © 2026 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 accurate, complete, and timely information about this product from reliable sources.

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	5
Answers	8
Explanations	10
Next Steps	15

SAMPLE

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

SAMPLE

- 1. Which command configures OSPF as passive on VLAN interfaces 20, 30, and 40?**
 - A. vlan 20,30,40 ospf passive
 - B. interface vlan 20,30,40 ip ospf passive
 - C. router ospf 1 area 0 passive-interface vlan 20,30,40
 - D. router ospf 1 area 0 redistribute local
- 2. How is Multicast Transmission Optimization implemented in an Aruba wireless network?**
 - A. The optimal rate is based on the highest broadcast rate across all associated clients
 - B. When enabled the minimum default rate for multicast traffic is set to 12 Mbps for 5 GHz
 - C. The optimal rate is based on the lowest broadcast rate across all associated clients
 - D. The optimal rate is based on the lowest unicast rate across all associated clients
- 3. Which AP models are mentioned for the remote office deployment?**
 - A. AP-515 and AP-575S
 - B. AP-300 and AP-400
 - C. AP-205 and AP-470
 - D. AP-700 and AP-800
- 4. Which default AOS-CX user role provides View switch configuration and REST API PUT/POST access?**
 - A. Administrators
 - B. auditors
 - C. sysops
 - D. helpdesk
- 5. Which of the following is a typical step in configuring DHCP Snooping to prevent rogue servers, as shown in example configurations?**
 - A. Enable DHCP Snooping with Option 82 enabled
 - B. Enable DHCP Snooping with Option 82 disabled
 - C. Disable DHCP Snooping entirely
 - D. Enable DHCP Snooping and use static bindings only

6. A maximum aggregate ISL bandwidth of 200G is supported by which stacking technology?
- A. VSF
 - B. VSX
 - C. Both VSF and VSX
 - D. Neither VSF nor VSX
7. In the Aruba CX 6300 port-access example, which interface is configured for the port-access commands?
- A. interface 1/1/1
 - B. interface 1/2/1
 - C. interface 2/1/1
 - D. interface 1/1/2
8. In the WPA 4-Way Handshake, which action occurs after the PMK is proven?
- A. Distributes an encrypted GTK to the client
 - B. Exchanges messages for generating PTK
 - C. Sets first initialization vector (IV)
 - D. Proves knowledge of the PMK
9. Which VLANs are allowed on the trunk in the correct port-access configuration for the APs?
- A. 100,200,300
 - B. 100,200
 - C. 200,300
 - D. 100
10. Which MSTP-related setting must be identically configured with MSTP to load-balance VLANs on Access-1 switches?
- A. Spanning-tree bpdu-guard setting
 - B. Spanning-tree instance vlan mapping
 - C. spanning-tree Cist mapping
 - D. spanning-tree root-guard setting

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which command configures OSPF as passive on VLAN interfaces 20, 30, and 40?

- A. `vlan 20,30,40 ospf passive`
- B. interface vlan 20,30,40 ip ospf passive**
- C. `router ospf 1 area 0 passive-interface vlan 20,30,40`
- D. `router ospf 1 area 0 redistribute local`

OSPF passive interfaces are used to prevent OSPF from forming neighbors and from sending Hellos on those interfaces. When you mark an interface as passive, it still carries routing information learned from other active interfaces, but it won't actively participate in OSPF neighbor discovery on that link. To apply this to VLAN interfaces 20, 30, and 40, you configure the passive setting on each interface so that the SVIs for those VLANs don't become OSPF neighbors. The approach shown—configuring `ip ospf passive` on the VLAN interfaces themselves (often written as `interface vlan 20,30,40` followed by `ip ospf passive`)—directly targets those SVIs. This is why it's the best fit: it precisely disables OSPF adjacencies on the specified VLAN interfaces. Other patterns either use syntax that isn't valid on many devices for listing multiple interfaces or apply passive settings at the router process level, which is not the per-interface control you want. Redistributing local routes has nothing to do with making interfaces passive.

2. How is Multicast Transmission Optimization implemented in an Aruba wireless network?

- A. The optimal rate is based on the highest broadcast rate across all associated clients
- B. When enabled the minimum default rate for multicast traffic is set to 12 Mbps for 5 GHz
- C. The optimal rate is based on the lowest broadcast rate across all associated clients
- D. The optimal rate is based on the lowest unicast rate across all associated clients**

Multicast Transmission Optimization is about making multicast frames reliably reach every client by choosing a rate that all listeners can handle. In Aruba, this is done by basing the multicast rate on the lowest unicast data rate among all associated clients. Since multicast is delivered to everyone without per-client acknowledgments, the AP keeps the multicast rate at or below what the slowest client can support. This ensures that even the client with the weakest link can decode the multicast frames, reducing the chance of frame loss for the group and improving overall efficiency for group traffic. If you tried to use a higher rate tied to faster clients, the slowest devices would miss frames, hurting everyone. Dynamic adjustments occur as client conditions change, keeping multicast performance aligned with real-time capabilities.

3. Which AP models are mentioned for the remote office deployment?

- A. AP-515 and AP-575S**
- B. AP-300 and AP-400
- C. AP-205 and AP-470
- D. AP-700 and AP-800

When planning a remote office deployment, you look for compact, cost-effective access points that deliver solid Wi-Fi coverage for small spaces and fit easily into centralized management. The models AP-515 and AP-575S are designed for this remote/branch-office use case, offering the right balance of form factor, performance, and management fit for a single-site or small office environment. They provide sufficient throughput and client support for typical remote office needs without the extra scale or features meant for larger campuses or specialized outdoor deployments. The other model groups tend to target different scenarios: older or lower-end generations that may not meet current remote-office performance expectations, or devices built for larger venues, high density, or outdoor use. That's why the pair listed here—AP-515 and AP-575S—best matches remote office deployment needs.

4. Which default AOS-CX user role provides View switch configuration and REST API PUT/POST access?

- A. Administrators
- B. auditors
- C. sysops**
- D. helpdesk

In AOS-CX, user roles define what you can view and what you can change. Being able to view the switch configuration is a read-side permission, while using REST API PUT/POST is a write operation that applies changes to the device. The sysops role is designed for operators who need to inspect configurations and perform configuration changes, including through the REST API. That combination—viewing the switch configuration and executing REST API calls that modify the configuration—fits sysops best. Administrators have broader access beyond that, auditors are typically read-only and cannot perform configuration changes, and the helpdesk role is usually more limited and focused on basic support tasks rather than direct configuration edits via REST.

5. Which of the following is a typical step in configuring DHCP Snooping to prevent rogue servers, as shown in example configurations?

- A. Enable DHCP Snooping with Option 82 enabled
- B. Enable DHCP Snooping with Option 82 disabled**
- C. Disable DHCP Snooping entirely
- D. Enable DHCP Snooping and use static bindings only

DHCP Snooping protects endpoints by watching DHCP traffic and constructing a binding table, so only legitimate DHCP server responses are allowed back to clients on trusted ports. The Relay Agent Information option (Option 82) carries switch-port details to the DHCP server, which can support more precise policy and address assignment. In many Aruba example configurations aimed at preventing rogue servers, the typical step shown is enabling DHCP Snooping while keeping Option 82 disabled. This provides the essential protection without adding the complexity of relay-agent data management, ensuring rogue servers on untrusted ports can't respond to clients while the switch filters communications based on the binding table. Enabling Option 82 is an additional capability that isn't necessary for this baseline protection, and other options either remove protection entirely or impose a stricter, less flexible setup.

6. A maximum aggregate ISL bandwidth of 200G is supported by which stacking technology?

- A. VSF**
- B. VSX
- C. Both VSF and VSX
- D. Neither VSF nor VSX

Stacking determines how much traffic can flow between switches in the same stack. VSF (Virtual Switching Framework) creates a single logical switch across the stacked units and uses the inter-switch links in a way that supports a higher combined bandwidth. In the Aruba ACP-CA context, the maximum aggregate inter-switch link (ISL) bandwidth achievable with VSF is 200 Gbps, so that technology best fits the 200G requirement. VSX (Virtual Switching Extension) uses a different stacking approach with a different bandwidth profile and does not reach this 200 Gbps aggregate limit in the same scenarios, so it isn't the correct fit here.

7. In the Aruba CX 6300 port-access example, which interface is configured for the port-access commands?

- A. interface 1/1/1**
- B. interface 1/2/1
- C. interface 2/1/1
- D. interface 1/1/2

Port-access configuration is applied to the interface level, not globally. In the Aruba CX 6300 example, the port-access commands are placed under the configuration for a specific interface—the one connected to the endpoint being secured. That interface is 1/1/1, so the sample shows all port-access settings inside the context for that port. Since port-access controls how a single port behaves (802.1X, MAB, VLAN assignment, etc.), the example focuses on configuring that particular port rather than applying the same settings to multiple ports. The other ports would require their own port-access blocks if they were part of the scenario, but they aren't shown in this example.

8. In the WPA 4-Way Handshake, which action occurs after the PMK is proven?

- A. Distributes an encrypted GTK to the client
- B. Exchanges messages for generating PTK**
- C. Sets first initialization vector (IV)
- D. Proves knowledge of the PMK

In the WPA 4-Way Handshake, once both sides have verified they know the PMK, the next step is to derive and install the per-session keys by exchanging messages to generate the Pairwise Transient Key (PTK). This PTK is created using the PMK together with fresh nonces (ANonce and SNonce) and is what ultimately protects the unicast data between the client and the AP. After the PTK is established, the GTK for multicast/broadcast traffic is distributed and installed in the subsequent messages. The IVs and other encryption parameters are part of the ongoing encrypted data flow, not a separate handshake action.

9. Which VLANs are allowed on the trunk in the correct port-access configuration for the APs?

- A. 100,200,300
- B. 100,200
- C. 200,300**
- D. 100

In port-access for APs, the switch port acts like a trunk that carries both a native management VLAN (untagged) and additional VLANs that tag wireless traffic. The VLANs used to carry SSID traffic—those that map to the wireless networks—must be allowed as tagged traffic on the trunk. Choosing the VLANs 200 and 300 reflects that those are the VLANs assigned to the APs for wireless data (the tagged traffic). VLAN 100 is the native management VLAN, which remains untagged on the trunk and isn't listed among the allowed tagged VLANs. If you tried to allow 100 as a tagged VLAN, it would disrupt the native untagged management traffic and cause APs to misbehave. So, the trunk should permit the VLANs used for user data (200 and 300), while the management VLAN (100) stays native.

10. Which MSTP-related setting must be identically configured with MSTP to load-balance VLANs on Access-1 switches?

- A. Spanning-tree bpdu-guard setting
- B. Spanning-tree instance vlan mapping**
- C. spanning-tree Cist mapping
- D. spanning-tree root-guard setting

In MSTP, traffic for each VLAN is handled by a specific MST instance. To evenly distribute load across multiple links between switches, the mapping of VLANs to MST instances must be the same on all switches in the region. This ensures that the same VLANs take the same relative paths and opportunities for alternative routes, allowing multiple spanning-tree instances to route traffic concurrently rather than funneling all VLANs down a single link. The setting that controls this is the spanning-tree instance to VLAN mapping, which explicitly assigns which VLANs belong to which MST instances. If this mapping differs between Access-1 switches, the network can't achieve consistent, balanced paths for VLANs, defeating the purpose of MSTP-based load balancing. The other options—bpdu-guard, root-guard, or a CIST mapping—relate to protecting the topology or to the internal spanning tree structure, not to how VLANs are distributed across MST instances for load balancing.

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://hpeacpca.examzify.com>

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

SAMPLE