

COMPTIA CLOUD+ (CVO-002) PRACTICE TEST (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

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

<https://comptiacloudpluscv0002.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. For a cloud-based system with an availability SLA of 99.99%, which unavailability time would be considered acceptable during a 30 day test window?**
 - A. 1 minute
 - B. 4 minutes
 - C. 10 minutes
 - D. 60 minutes

- 2. In migrating three on premise servers to virtual machines, which minimum CPU resources on the second host would best support the new VMs?**
 - A. 1,600 MHz for the application server and 2,300 MHz for each web server
 - B. 3,300 MHz for the application server and 4,600 MHz for each web server
 - C. 3,300 MHz for the application server and 5,250 MHz for each web server
 - D. 5,250 MHz for the application server and 3,100 MHz for each web server

- 3. Which high-availability solution involves multiple servers that each service requests concurrently and can take over the load of a failed member?**
 - A. Active-passive
 - B. Active-active
 - C. Passive-passive
 - D. Passive-active

- 4. Which cloud service model would be recommended to a company for hardware capacity to host a production database application?**
 - A. CaaS
 - B. XaaS
 - C. IaaS
 - D. PaaS

- 5. Which statement about Type 1 hypervisors is true?**
 - A. They run directly on hardware without a host OS.
 - B. They require a host operating system.
 - C. They are always open source.
 - D. They cannot support multiple virtual machines.

- 6. Which environmental constraint is most likely to disqualify a cross-border datacenter site for a defense project?**
- A. Downtime impact
 - B. Legal restrictions
 - C. Peak time frames
 - D. Bandwidth
- 7. In a cloud-based log aggregation setup, the log sizes have increased by 50% daily and the system stops publishing logs at noon. What is the most likely cause?**
- A. The logging directory does not have sufficient storage space.
 - B. The syslog service is not running on the servers.
 - C. The data limit has been exceeded at the SaaS provider.
 - D. There is a cloud service provider outage.
- 8. Which protocol is commonly used to gather performance data from network devices and assist with monitoring?**
- A. SNMP
 - B. SSH
 - C. FTP
 - D. DNS
- 9. A CIO notes that a new 20TB CRM application increases power usage. Which solution provides a performance benefit and reduces power consumption?**
- A. SSD
 - B. SAS
 - C. SATA
 - D. FC
- 10. Which storage type can take advantage of Universal Naming Convention addressable storage?**
- A. SAN
 - B. NAS
 - C. DAS
 - D. SATA

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. For a cloud-based system with an availability SLA of 99.99%, which unavailability time would be considered acceptable during a 30 day test window?

- A. 1 minute
- B. 4 minutes**
- C. 10 minutes
- D. 60 minutes

An availability SLA of 99.99% means only a tiny amount of downtime is allowed over the test window. To see how much, convert the 30 days into minutes: $30 \text{ days} \times 24 \text{ hours} \times 60 \text{ minutes} = 43,200 \text{ minutes}$. The allowed downtime is 0.01% of that, which is $0.0001 \times 43,200 = 4.32 \text{ minutes}$. So, during a 30-day window, about 4.3 minutes of unavailability are permissible. Among the options, 4 minutes fits within this limit (it's under 4.32 minutes), while 10 minutes and 60 minutes exceed it. Although 1 minute would also be acceptable, the question aims for the maximum downtime still within the SLA, which is about 4 minutes.

2. In migrating three on premise servers to virtual machines, which minimum CPU resources on the second host would best support the new VMs?

- A. 1,600 MHz for the application server and 2,300 MHz for each web server**
- B. 3,300 MHz for the application server and 4,600 MHz for each web server
- C. 3,300 MHz for the application server and 5,250 MHz for each web server
- D. 5,250 MHz for the application server and 3,100 MHz for each web server

When migrating to virtual machines, you size the CPU resources for each VM to meet its workload while ensuring the host can supply the total demand without bottlenecks. For three VMs—a application server and two web servers—you sum the minimum CPU needs of all VMs to determine the host's required capacity. The allocation that assigns 1,600 MHz to the application server and 2,300 MHz to each web server provides a total of 6,200 MHz across the host. This is the smallest total that still meets each VM's needs, making efficient use of the host's CPU resources. The other options require more total CPU for the VMs, which would waste capacity or reduce headroom for spikes. In practice, you'd also account for hypervisor overhead and potential overcommit, but as a minimum-size solution, this distribution is the most appropriate.

3. Which high-availability solution involves multiple servers that each service requests concurrently and can take over the load of a failed member?

- A. Active-passive
- B. Active-active**
- C. Passive-passive
- D. Passive-active

Active-active high availability means several servers are online at the same time and all handle requests concurrently, with a load balancer distributing traffic among them. If one server fails, the others are already processing work and can take over its load without waiting for a failover, keeping the service available and responsive. This setup boosts throughput and resilience because it doesn't rely on a single active node. The trade-off is that you must keep data and state synchronized across all nodes to avoid inconsistencies, which adds complexity.

4. Which cloud service model would be recommended to a company for hardware capacity to host a production database application?

- A. CaaS
- B. XaaS
- C. IaaS
- D. PaaS**

A production database in the cloud benefits most from a model that provides scalable resources while handling the underlying management for you. PaaS offers a platform with built-in services for running applications and databases, and it includes managed database services, automated backups, patching, high availability, and easy scaling. This means you get reliable hardware-backed capacity and the ability to grow resources as demand increases without having to manage the operating system, middleware, or database maintenance tasks yourself. If you truly needed full control over the OS and database configuration, IaaS would be an option, but for hosting a production database with scalable capacity and reduced operational overhead, a platform-managed approach is typically the best fit.

5. Which statement about Type 1 hypervisors is true?

- A. They run directly on hardware without a host OS.**
- B. They require a host operating system.
- C. They are always open source.
- D. They cannot support multiple virtual machines.

Bare-metal virtualization means the hypervisor sits directly on the physical hardware, with no general-purpose host operating system in between. This arrangement lets the hypervisor control access to CPU, memory, storage, and networking, providing strong isolation and efficient resource management for multiple virtual machines. That's why the statement is correct: a Type 1 hypervisor runs directly on hardware and does not require a host OS. In practice, this is what distinguishes it from Type 2 hypervisors, which run on top of a host operating system and are more common for desktop testing rather than data-center deployments. It's also not true that Type 1 hypervisors must be open source—there are both open-source and proprietary options. And yes, Type 1 hypervisors are designed to support many virtual machines, not just one.

6. Which environmental constraint is most likely to disqualify a cross-border datacenter site for a defense project?

- A. Downtime impact
- B. Legal restrictions**
- C. Peak time frames
- D. Bandwidth

Regulatory compliance around cross-border data handling is the main factor here. For a defense project, where sensitive information and access are tightly controlled, legal restrictions determine whether a site can even be used. Data residency requirements, export controls, and restrictions on who can access or transfer defense data across borders can make a location unusable regardless of other capabilities. If a jurisdiction imposes rules that prevent storing or processing the data in that country or requires special authorization that isn't feasible, the site is effectively disqualified. Downtime impact, peak time frames, and bandwidth are important for performance and resilience, but they can often be mitigated with design choices, redundancy, capacity planning, and robust connectivity. They don't inherently bar a site in the way legal and regulatory restrictions do.

7. In a cloud-based log aggregation setup, the log sizes have increased by 50% daily and the system stops publishing logs at noon. What is the most likely cause?

A. The logging directory does not have sufficient storage space.

B. The syslog service is not running on the servers.

C. The data limit has been exceeded at the SaaS provider.

D. There is a cloud service provider outage.

When log volumes keep growing and publishing stops at a specific time, the most likely cause is that the logging directory used for staging or buffering logs has run out of storage. The system continues to generate logs (hence the 50% daily increase) but has no space to write new entries or to queue them for sending to the central system. Once the disk fills up, the logging agent can't write or flush logs, causing publishing to halt around the time the space is exhausted (noon in this case). To fix, you'd typically free up space, expand the disk, or implement longer-term retention or rotation so that new logs can continue to be written and forwarded. Why the other possibilities fit less well: If the syslog service weren't running, you'd expect a drop in log generation at the source rather than a continued increase in local log sizes. If the data limit at the SaaS provider were the issue, you'd see quota-related errors from the cloud service rather than a gradual daily growth followed by a clean stop, and the behavior would depend on provider policies rather than the local storage state. A cloud provider outage would more likely cause a broad service disruption rather than a symptoms pattern tied to growing local log volumes and a timed halt.

8. Which protocol is commonly used to gather performance data from network devices and assist with monitoring?

A. SNMP

B. SSH

C. FTP

D. DNS

The main idea is using a protocol designed for management and monitoring to collect performance data from devices. Simple Network Management Protocol is built for this purpose: a central manager polls devices for metrics stored in their MIBs, such as interface utilization, CPU and memory usage, disk space, and uptime. Devices run an SNMP agent to respond with the requested values, and SNMP can also send traps to alert when a metric crosses a threshold. This makes SNMP a natural fit for monitoring tools like Nagios, Zabbix, or PRTG, because it provides standardized, scalable access to a wide range of performance data. Other options are intended for different tasks—SSH is for secure remote command execution, FTP for file transfers, and DNS for name resolution—not for collecting standardized monitoring metrics.

9. A CIO notes that a new 20TB CRM application increases power usage. Which solution provides a performance benefit and reduces power consumption?

- A. SSD**
- B. SAS
- C. SATA
- D. FC

Using solid-state storage delivers both higher performance and lower power per operation. Without moving parts, SSDs have much faster access times and far higher IOPS than spinning disks, so the CRM application can read and process data much more quickly. That boosts throughput and reduces latency under load. At the same time, SSDs typically draw less power per I/O and generate less heat than HDD-based drives, helping reduce overall power consumption even as performance improves. Among the options, the solid-state drive option directly provides these combined benefits, whereas the other choices refer to traditional spinning-disk interfaces or technologies that don't inherently deliver the same power-efficient performance improvements.

10. Which storage type can take advantage of Universal Naming Convention addressable storage?

- A. SAN
- B. NAS**
- C. DAS
- D. SATA

Universal Naming Convention is used to address resources on a network by a path, such as \\server\share, which is how you access file-level shares over the network. This makes network-attached storage the best fit because NAS devices present folders as shares over the network and clients connect using UNC paths. In contrast, SAN provides block storage accessed via LUNs over Fibre Channel or iSCSI and doesn't use UNC paths to access files. Direct-attached storage is not network-accessible in the same way, and SATA is simply a drive interface, not a networked storage addressing method.

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://comptiacloudpluscsv0002.examzify.com>

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

SAMPLE