

Nutanix Certified Professional Multicloud Infrastructure (NCP-MCI) Practice Exam (Sample)

Study Guide



Everything you need from our exam experts!

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	16

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. What are immutable copies of a vDisk called when read by multiple VMs?**
 - A. A. Disk Clones
 - B. B. Golden Images
 - C. C. Volume Groups
 - D. D. Shadow Clones
- 2. What is a common reason for data recovery failure in a Nutanix cluster?**
 - A. Inadequate network bandwidth
 - B. Improper node configuration
 - C. Insufficient replica count
 - D. Outdated software versions
- 3. What is the preferred method for handling VM communication in an AHV environment?**
 - A. Using a dedicated management network.
 - B. Implementing storage-only networks.
 - C. Utilizing VLAN tagging for enhanced isolation.
 - D. Employing a flat network configuration.
- 4. What function does the Nutanix Controller Virtual Machine (CVM) serve?**
 - A. To manage user interfaces
 - B. To oversee application performance
 - C. To manage storage services for infrastructure
 - D. To act as a backup solution
- 5. What are the first two steps to create a trunked interface on a VM in AOS 5.15x? (Choose two)**
 - A. Use acli
 - B. Log in over PE web UI
 - C. SSH to CVM
 - D. Update VM dialog

- 6. What is the primary use of Nutanix Cluster Check (NCC)?**
- A. To manage network traffic
 - B. To store backup data
 - C. To monitor and validate the health of a Nutanix cluster
 - D. To configure virtual machines
- 7. What is the function of High Availability in a Nutanix cluster?**
- A. Ensuring seamless data migration
 - B. Migrating workloads to different data centers
 - C. Maintaining system uptime during failures
 - D. Distributing storage evenly across nodes
- 8. How can an administrator analyze high cluster latency data with the least number of steps?**
- A. Modify Data Density in the dashboard.
 - B. Click on the chart in the widget to expand the data elements.
 - C. Create a new Analysis chart based on cluster name.
 - D. Access the cluster quick access widget for more details.
- 9. What is the effect on the replicas of the VMs if a cluster with RF2 loses two drives on different nodes?**
- A. Some VM data may be lost
 - B. No VMs lose data if the node has two or more SSDs
 - C. Some VMs may reboot and gain access to data
 - D. No VMs lose data because of RF2
- 10. What should an administrator consider to improve performance if SSD tiers are insufficient for workloads' IOPS demands?**
- A. Increasing the number of nodes in the cluster
 - B. Adding more RAM to the hosts
 - C. Upgrading the SSDs to higher capacity
 - D. Increasing the SSD tier size

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What are immutable copies of a vDisk called when read by multiple VMs?

- A. A. Disk Clones
- B. B. Golden Images
- C. C. Volume Groups
- D. D. Shadow Clones**

The term "Shadow Clones" refers to immutable copies of a virtual disk (vDisk) that can be read simultaneously by multiple virtual machines (VMs). This allows for efficient management of resources and ensures that the original vDisk remains unchanged while providing a consistent read experience. Shadow Clones enable various VMs to access a common, unaltered version of the vDisk without affecting performance or the integrity of the source disk. This technology can be particularly useful in scenarios where you want to deploy multiple identical environments or when conducting testing and development, as it minimizes storage consumption and optimizes resource utilization. The other options would not accurately describe this concept: Disk Clones typically refer to copies of a disk that may not remain immutable; Golden Images are often used to create standard templates for VMs but don't imply an immutable reference across multiple instances; Volume Groups pertain more to aggregation of disks rather than individual disk states for multi-VM reads.

2. What is a common reason for data recovery failure in a Nutanix cluster?

- A. Inadequate network bandwidth
- B. Improper node configuration
- C. Insufficient replica count**
- D. Outdated software versions

A common reason for data recovery failure in a Nutanix cluster is insufficient replica count. In a distributed environment like Nutanix, data is typically duplicated across multiple nodes to ensure redundancy and availability. If the replica count is lower than the required threshold for data recovery, there may not be enough copies of data available to successfully restore it in the event of a failure. This is particularly critical in scenarios where multiple nodes or disks fail, as the system relies on having sufficient replicas to reconstruct the data accurately. The chosen answer aligns with best practices in data management, where it is essential to maintain an adequate number of replicas to ensure that data can be reliably recovered. Without this redundancy, the risk of data loss becomes significantly higher, leading to recovery failures when trying to access or restore information. While inadequate network bandwidth, improper node configuration, and outdated software versions can certainly impact overall cluster performance or availability, they do not directly influence the redundancy and availability of data like the replica count does. Adequate replicas directly determine the system's ability to recover from instances of data loss, making it a critical facet of data recovery strategies in Nutanix environments.

3. What is the preferred method for handling VM communication in an AHV environment?

- A. Using a dedicated management network.
- B. Implementing storage-only networks.
- C. Utilizing VLAN tagging for enhanced isolation.**
- D. Employing a flat network configuration.

In an AHV environment, utilizing VLAN tagging for enhanced isolation is the preferred method for handling VM communication. VLAN (Virtual Local Area Network) tagging allows for the segmentation of network traffic, which enhances security and performance by isolating network traffic according to specific logical groupings. This becomes particularly important in a virtualized environment like Nutanix AHV, where multiple virtual machines may need to communicate without interfering with each other's network traffic. By applying VLAN tagging, network administrators can ensure that communication between certain VMs is restricted and controlled, reducing the risk of unauthorized access and helping to streamline network management. In contrast, a dedicated management network may focus on administrative tasks rather than isolating VM traffic, and implementing storage-only networks is primarily concerned with storage access rather than general VM communication. A flat network configuration lacks the necessary segmentation, which can lead to security vulnerabilities and network congestion as all traffic is treated equally, resulting in a less efficient and less secure network environment. Therefore, leveraging VLAN tagging provides a robust framework for maintaining effective and secure VM communications within an AHV infrastructure.

4. What function does the Nutanix Controller Virtual Machine (CVM) serve?

- A. To manage user interfaces
- B. To oversee application performance
- C. To manage storage services for infrastructure**
- D. To act as a backup solution

The Nutanix Controller Virtual Machine (CVM) plays a crucial role in the management of storage services within the Nutanix infrastructure. It is responsible for handling tasks related to storage operations, including data storage, replication, deduplication, and compression. The CVM ensures that virtual machines (VMs) have access to the storage they need by managing the underlying distributed storage fabric that Nutanix is built upon. In the context of a hyper-converged infrastructure, the CVM acts as a key component that integrates compute and storage resources, making scale-out architectures efficient and effective. By managing storage services, the CVM allows for seamless data handling and provides the necessary services to ensure data availability and reliability across the platform. The other options focus on functions that are not the primary role of the CVM. For instance, managing user interfaces pertains more to the management console or admin interfaces rather than storage management. Oversight of application performance typically involves other monitoring tools or systems, and while backup solutions are vital, they are not the direct responsibility of the CVM, which concentrates specifically on storage management tasks.

5. What are the first two steps to create a trunked interface on a VM in AOS 5.15x? (Choose two)

- A. Use acli**
- B. Log in over PE web UI
- C. SSH to CVM
- D. Update VM dialog

To create a trunked interface on a virtual machine (VM) in AOS 5.15x, the first two critical steps involve utilizing both the acli (Acropolis Command Line Interface) and the PE web UI (Prism Element web User Interface). Understanding the roles of these tools in the process is key. Utilizing acli is essential because it allows for more nuanced and scriptable commands to manage networking features and configurations directly. It provides capabilities to set up trunking, enabling the VM to handle multiple VLANs, which is vital for environments that require network segmentation or more complex configurations. This command-line interface is powerful for tasks like deploying new resources or changing configurations on the fly. Accessing the PE web UI is another critical step, as it allows for a more visual approach to managing and configuring VMs. The web interface provides a user-friendly way to navigate through the various settings and options that are available, making it easier for users who prefer graphical interfaces or are less familiar with command-line operations. The other options, such as SSH to the CVM (Controller VM) and updating the VM dialog, are not among the initial steps for trunk interface creation. While SSH may be used for various management tasks, it is not the

6. What is the primary use of Nutanix Cluster Check (NCC)?

- A. To manage network traffic
- B. To store backup data
- C. To monitor and validate the health of a Nutanix cluster**
- D. To configure virtual machines

Nutanix Cluster Check (NCC) serves a critical role in maintaining the integrity and performance of Nutanix clusters. Its primary function is to monitor and validate the health of a Nutanix cluster, ensuring that the cluster operates optimally and without issues. NCC performs various health checks that look for potential problems within the cluster, such as node failures, network connectivity issues, or storage configuration discrepancies. By conducting these checks, NCC helps administrators proactively identify and address issues before they can lead to downtime or degraded performance. Options related to managing network traffic, storing backup data, or configuring virtual machines do not align with the specific purpose of NCC. These activities fall under different tools and functions within the Nutanix architecture and ecosystem, which focus on tasks that are separate from the real-time health monitoring provided by NCC.

7. What is the function of High Availability in a Nutanix cluster?

- A. Ensuring seamless data migration
- B. Migrating workloads to different data centers
- C. Maintaining system uptime during failures**
- D. Distributing storage evenly across nodes

High Availability in a Nutanix cluster serves the critical function of maintaining system uptime during failures. This means that in the event of a hardware or software fault, the cluster is designed to automatically recover and continue providing services with minimal interruption. By leveraging features such as data redundancy, fault tolerance, and automated failover, Nutanix ensures that even if one or more components fail, the overall system remains operational, allowing users to maintain access to applications and data without significant downtime. In contrast to the other choices, seamless data migration and workload migration to different data centers address different operational concerns that are not directly related to High Availability. While distributing storage evenly across nodes is important for optimal performance and efficiency, it does not specifically pertain to the primary goal of ensuring continuous operation during failures. High Availability focuses on resilience and uninterrupted service, making it an essential aspect of Nutanix cluster architecture.

8. How can an administrator analyze high cluster latency data with the least number of steps?

- A. Modify Data Density in the dashboard.
- B. Click on the chart in the widget to expand the data elements.**
- C. Create a new Analysis chart based on cluster name.
- D. Access the cluster quick access widget for more details.

The correct approach is to click on the chart in the widget to expand the data elements. This action is the most efficient way to analyze high cluster latency data because it allows the administrator to quickly visualize and examine specific metrics and trends related to latency without the need for additional configuration or setup. When an administrator interacts directly with the chart in the dashboard, they can gain immediate insights into the data points, which helps identify patterns or spikes in latency. This method is straightforward and minimizes the steps involved in accessing critical information. Other options require more time or additional actions. Modifying data density in the dashboard affects how the data is visually represented but does not provide a direct and detailed investigation into the latency specifics. Creating a new analysis chart based on cluster name involves additional steps, which may not be necessary if the data is already present in the existing widget. Accessing the cluster quick access widget could lead to more details, but it might not be as focused or directly tied to the latency data at hand as simply expanding the chart.

9. What is the effect on the replicas of the VMs if a cluster with RF2 loses two drives on different nodes?

- A. Some VM data may be lost**
- B. No VMs lose data if the node has two or more SSDs
- C. Some VMs may reboot and gain access to data
- D. No VMs lose data because of RF2

In a cluster configured with a replication factor of 2 (RF2), each piece of data is stored in two different locations across the cluster. This means that for any given virtual machine (VM), there are two replicas available that can be used to recover data in the event of failure. If a cluster with RF2 loses two drives located on different nodes, the replicas of the VMs can be affected. Since RF2 relies on having two copies of the data, losing two drives presents a risk to the integrity of the VMs' data. Depending on which specific drives failed, it is possible that both replicas of a particular piece of data could be lost. This would lead to the scenario where some VM data may be irretrievably lost, as there wouldn't be a remaining copy available to recover the information. The other options do not account for the implications of losing two drives on different nodes. For instance, stating that no VMs lose data presumes that at least one replica remains intact, which may not be the case if the specific drives that failed held both replicas of the same data. Additionally, the idea that some VMs may reboot and gain access to data misinterprets the impact of drive failures; while some VMs might reboot

10. What should an administrator consider to improve performance if SSD tiers are insufficient for workloads' IOPS demands?

- A. Increasing the number of nodes in the cluster
- B. Adding more RAM to the hosts
- C. Upgrading the SSDs to higher capacity
- D. Increasing the SSD tier size**

When addressing the performance of workloads that exceed the IOPS capabilities of existing SSD tiers, increasing the SSD tier size is a viable solution. This approach effectively provides additional storage space within the existing SSD infrastructure, allowing for a larger area from which the system can draw IOPS. By expanding the capacity of the SSD tier, it can accommodate more data and enable enhanced performance for demanding applications, which often require faster data access and retrieval times. Furthermore, increasing the SSD tier size can help balance the load more efficiently across the drives, thereby improving the overall throughput and reducing latency. This becomes crucial in environments where workloads are I/O intensive and where the existing tier is unable to keep up with the transaction rates or read/write demands. While other options, such as adding more nodes or increasing RAM, may also contribute to performance enhancements, they do not directly address the specific issue of insufficient SSD tier capacity in relation to IOPS demands. Upgrading to higher-capacity SSDs might seem similar, but it does not inherently increase the available performance unless those SSDs also come with higher IOPS specifications.

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://nutanixcertifiedprofessionalmulticloudinfrastructure-ncpmci.examzify.com>

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

SAMPLE