

HCIA Cloud Computing Practice Test (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. In which scenario is virtualization applicable?**
 - A. Only for high-load computing environments**
 - B. If a company has low host resource utilization**
 - C. When all employees use highly different applications**
 - D. For scientific institutions with generic operating systems**
- 2. What is the main purpose of the KVM module in the context of virtualization?**
 - A. Creating Vcpu and allocating virtual memory.**
 - B. Simulating I/O operations.**
 - C. Managing virtual networks.**
 - D. Handling storage provisioning.**
- 3. In Huawei FusionCompute, what happens when a host in the cluster fails?**
 - A. The virtual machine automatically migrates to other hosts in the cluster.**
 - B. The virtual machine fails and needs to be restored manually.**
 - C. All virtual machines are paused until the host is restored.**
 - D. The virtual machines continue to operate on the failed host.**
- 4. How can computing virtualization be classified?**
 - A. It can be divided into type I and type II virtualization.**
 - B. Full virtualization and paravirtualization are also classifications of computing virtualization.**
 - C. Both the above statements are true.**
 - D. None of the above statements are true.**
- 5. What is the best description of the relationship between big data and cloud computing?**
 - A. Cloud computing is a collection of data that cannot be captured, managed, and processed with conventional software tools within a certain time frame.**
 - B. Big data is a pay-as-you-go model that provides usable, portable, and on-demand network access.**
 - C. Cloud computing cannot function as the underlying computing resource for big data processing.**
 - D. Big data can serve as a type of cloud computing service for users to choose according to their needs.**

- 6. How can templates for virtual machines be created?**
- A. By downloading them from the internet**
 - B. By converting or cloning existing virtual machines**
 - C. Using a third-party application**
 - D. Through manual code entry**
- 7. When deploying virtual machines, what is a major advantage of using templates?**
- A. It allows the user to operate without any additional software**
 - B. It ensures unique configurations for each virtual machine**
 - C. It reduces the time required for setup and configuration**
 - D. It increases the complexity of maintenance tasks**
- 8. Which of the following actions is NOT part of snapshot management?**
- A. Copying a snapshot**
 - B. Modifying a snapshot**
 - C. Creating a snapshot**
 - D. Deleting a snapshot**
- 9. What is the classification of CPU virtualization technology?**
- A. Full virtualization, paravirtualization, software virtualization, hardware-assisted virtualization**
 - B. Software virtualization only**
 - C. Hardware-assisted virtualization only**
 - D. Only full and paravirtualization**
- 10. Which of the following descriptions regarding memory reuse in Huawei Fusion Compute is incorrect?**
- A. The administrator can enable or disable the memory reuse function at any time.**
 - B. After the memory reuse function is turned off, the internal sum of all virtual machines is less than or equal to the total available physical memory.**
 - C. Using memory reuse can degrade virtual machine performance.**
 - D. Turning off memory reuse requires that the memory reuse rate of all hosts in the cluster is less than or equal to 100%.**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. In which scenario is virtualization applicable?

- A. Only for high-load computing environments
- B. If a company has low host resource utilization**
- C. When all employees use highly different applications
- D. For scientific institutions with generic operating systems

Virtualization is particularly applicable in scenarios where a company has low host resource utilization. This is because virtualization allows multiple virtual machines (VMs) to run on a single physical server, which can lead to more efficient use of resources. When a company's physical servers are underutilized, virtualization enhances capacity utilization by enabling diverse workloads to share the same hardware. This leads to cost savings, simpler management, and the ability to quickly provision new environments without the need for additional physical hardware. In this context, when host resources are not being fully utilized, implementing virtualization allows for a better return on investment and improved operational efficiency, making it a suitable approach for organizations looking to optimize their infrastructure. Since virtualization helps maximize the usage of existing resources, it is an advantageous solution for businesses facing low utilization rates.

2. What is the main purpose of the KVM module in the context of virtualization?

- A. Creating Vcpu and allocating virtual memory.**
- B. Simulating I/O operations.
- C. Managing virtual networks.
- D. Handling storage provisioning.

The primary function of the KVM (Kernel-based Virtual Machine) module in virtualization is to create virtual CPUs (vCPUs) and allocate virtual memory. KVM transforms the Linux kernel into a hypervisor, allowing it to manage multiple virtual machines (VMs). Each VM operates as an independent entity with its own operating system, and KVM enables the allocation of virtual hardware resources, which include vCPUs and memory, to these VMs. This capability is essential for delivering the computational power needed to run numerous isolated environments on a single physical server. Other options address important aspects of virtualization but do not highlight the core purpose of KVM itself. Simulating I/O operations, managing virtual networks, and handling storage provisioning are essential services in a virtualized environment, but they are not the primary function of KVM. Instead, these tasks are usually handled by other components within the virtualization infrastructure or through additional management layers that work alongside KVM.

3. In Huawei FusionCompute, what happens when a host in the cluster fails?

- A. The virtual machine automatically migrates to other hosts in the cluster.
- B. The virtual machine fails and needs to be restored manually.**
- C. All virtual machines are paused until the host is restored.
- D. The virtual machines continue to operate on the failed host.

In Huawei FusionCompute, when a host in the cluster fails, the expected behavior involves the management of virtual machines (VMs) to ensure service continuity. The correct answer reflects that the virtual machine fails and needs to be restored manually. When a host fails, the underlying hypervisor is unable to maintain the operation of the VMs hosted on that failed node. Consequently, the VMs go offline and cannot continue to operate as there is no available host to manage their execution. After the issue with the host is identified, the failed virtual machines must be restored manually, usually by restarting them on a functioning host within the cluster. This emphasizes the importance of planning for failover and restoration processes in cloud environments to minimize downtime. In contrast, options mentioning automatic migration, pausing all VMs, or continuing operations on a failed host inaccurately portray the expected behavior. Automatic migration occurs in environments with high availability features, which may not apply in every scenario. Pausing VMs or continuing operations without a functioning host would contradict the inherent limitations of the physical infrastructure and the management protocols established in FusionCompute. Hence, the understanding that VMs will require manual intervention after a host failure is key to effective cloud infrastructure management.

4. How can computing virtualization be classified?

- A. It can be divided into type I and type II virtualization.
- B. Full virtualization and paravirtualization are also classifications of computing virtualization.
- C. Both the above statements are true.**
- D. None of the above statements are true.

Computing virtualization can indeed be classified in multiple ways, and both type I and type II virtualization, as well as full virtualization and paravirtualization, represent valid methods of classification. Type I virtualization, also known as bare-metal virtualization, occurs directly on the host's hardware without an intermediate operating system, allowing for greater efficiency and resource management. This type is commonly used in data centers and enterprise environments. Type II virtualization, on the other hand, runs on top of an existing operating system and is often used for desktop virtualization, providing ease of use and convenience for end-users. In addition to this primary classification, virtualization can also be categorized into full virtualization and paravirtualization. Full virtualization creates an entirely virtual machine that mimics the hardware of the host system, allowing guest operating systems to run unmodified. Paravirtualization requires modifications to the guest OS to communicate with the hypervisor, optimizing performance by reducing overhead. Thus, since both types of classifications—type I/type II and full virtualization/paravirtualization—are accurate and significant, the correct answer that encompasses both statements is indeed that both classifications are valid.

5. What is the best description of the relationship between big data and cloud computing?

- A. Cloud computing is a collection of data that cannot be captured, managed, and processed with conventional software tools within a certain time frame.**
- B. Big data is a pay-as-you-go model that provides usable, portable, and on-demand network access.**
- C. Cloud computing cannot function as the underlying computing resource for big data processing.**
- D. Big data can serve as a type of cloud computing service for users to choose according to their needs.**

The best description of the relationship between big data and cloud computing is that big data can serve as a type of cloud computing service for users to choose according to their needs. This highlights the versatility and scalability that cloud computing offers in handling vast amounts of data often associated with big data applications. Cloud platforms provide infrastructure and services that can easily accommodate the storage, processing, and analysis needed for big data use cases. By leveraging cloud computing, organizations can access the necessary resources on demand, enabling them to analyze large datasets without the need for significant upfront investments in hardware and software. This flexibility allows users to scale resources up or down based on their requirements, making it an ideal solution for managing the dynamic nature of big data workloads. In contrast, the other options present limitations or definitions that do not accurately describe the symbiotic relationship between the two. Option A describes cloud computing incorrectly as simply a collection of data, rather than emphasizing its capabilities. Option B mischaracterizes big data as a financial model rather than a data type or analytics challenge. Option C incorrectly suggests that cloud computing cannot support big data processing, while in actuality, many cloud services are built specifically to enable big data analytics.

6. How can templates for virtual machines be created?

- A. By downloading them from the internet**
- B. By converting or cloning existing virtual machines**
- C. Using a third-party application**
- D. Through manual code entry**

Creating templates for virtual machines is commonly achieved by converting or cloning existing virtual machines. This method allows you to take an already configured and operational virtual machine and create a duplicate that retains the operating system, installed applications, settings, and configurations. This is highly beneficial as it streamlines the deployment process for new virtual machines, ensuring consistency across deployments. When using cloned or converted virtual machines as templates, administrators can quickly deploy multiple instances with the same configuration, which enhances efficiency and helps maintain standardization within the infrastructure. Cloning is especially useful in large environments where numerous virtual machine instances need to be maintained, ensuring they share the same setup and applications. The other methods, although feasible in some contexts, do not serve as the primary or most efficient way to create templates. For instance, downloading templates from the internet may lead to security concerns or compatibility issues. Using a third-party application could introduce unnecessary complexity or might not provide the required level of control. Manual code entry is not practical or efficient for creating virtual machine templates due to the intricate configurations required. Therefore, converting or cloning existing virtual machines stands out as the most effective approach for creating templates.

7. When deploying virtual machines, what is a major advantage of using templates?

- A. It allows the user to operate without any additional software**
- B. It ensures unique configurations for each virtual machine**
- C. It reduces the time required for setup and configuration**
- D. It increases the complexity of maintenance tasks**

Using templates when deploying virtual machines provides a significant advantage in reducing the time required for setup and configuration. Templates streamline the process of creating virtual machines by pre-defining key settings, applications, and configurations. This means that rather than manually installing and configuring each new virtual machine from scratch, users can quickly deploy new instances based on a predefined template that already includes all the necessary elements. This automated approach not only speeds up the deployment process but also promotes consistency across virtual machines. As a result, users can efficiently scale their environments while minimizing the risk of configuration errors that could arise from manual setup processes. Furthermore, leveraging templates can improve overall resource management in a cloud environment, allowing organizations to deploy resources rapidly as demand fluctuates, ensuring they can respond effectively to changing needs.

8. Which of the following actions is NOT part of snapshot management?

- A. Copying a snapshot**
- B. Modifying a snapshot**
- C. Creating a snapshot**
- D. Deleting a snapshot**

Snapshot management refers to the processes and actions taken to handle snapshots of virtual machines or storage systems, which are essentially point-in-time copies of data. Creating a snapshot involves capturing the current state of a virtual machine or data. Deleting a snapshot removes an existing snapshot, and modifying a snapshot typically involves changing its properties, although direct modification may not always occur since snapshots are primarily immutable once created. Copying a snapshot, however, is not typically considered a standard action within snapshot management itself. While you can create additional snapshots based on existing ones or utilize the data within a snapshot in different ways, the act of copying a snapshot is more about data management rather than snapshot management per se. This distinguishes it from the core actions associated with managing snapshots, which revolve around their lifecycle—creation, deletion, and modification. In summary, the correct response identifies an action that doesn't align directly with the traditional framework of managing snapshots, which is focused on the utilization and maintenance of snapshots rather than replicating them.

9. What is the classification of CPU virtualization technology?

- A. Full virtualization, paravirtualization, software virtualization, hardware-assisted virtualization**
- B. Software virtualization only
- C. Hardware-assisted virtualization only
- D. Only full and paravirtualization

The classification of CPU virtualization technology includes multiple modes such as full virtualization, paravirtualization, software virtualization, and hardware-assisted virtualization. Each of these types represents a different approach to how virtualization can be implemented. Full virtualization allows the guest operating system to run unmodified because it relies on the hypervisor to intercept and manage all the calls made by the guest to the underlying hardware. This means that the virtualization layer fully emulates the hardware, allowing any OS to operate as if it were on a physical machine. Paravirtualization, on the other hand, requires the guest operating system to be modified to be aware of the hypervisor. This approach can lead to better performance as the guest OS can make direct calls to the hypervisor, reducing overhead. Software virtualization refers to virtualization achieved primarily through software methods, which may not provide the same level of performance as hardware-based solutions, but it can be easier to implement on systems with less hardware support for virtualization. Hardware-assisted virtualization leverages specific processor features (such as Intel VT-x or AMD-V) that provide direct support for virtualization in the hardware itself. This can improve performance and simplify the hypervisor's task. Thus, the choice that encompasses the comprehensive classifications and different methods of CPU virtualization technology is

10. Which of the following descriptions regarding memory reuse in Huawei Fusion Compute is incorrect?

- A. The administrator can enable or disable the memory reuse function at any time.
- B. After the memory reuse function is turned off, the internal sum of all virtual machines is less than or equal to the total available physical memory.
- C. Using memory reuse can degrade virtual machine performance.**
- D. Turning off memory reuse requires that the memory reuse rate of all hosts in the cluster is less than or equal to 100%.

Using memory reuse in Huawei Fusion Compute is designed to optimize resource allocation, particularly when handling virtual machines. The correct choice reflects a common concern regarding performance implications. Memory reuse allows multiple virtual machines to share the same physical memory resources, which can lead to increased efficiency. However, this shared approach might lead to a situation where the performance of virtual machines can be impacted negatively, especially if there are excessive memory allocations leading to contention among virtual machines for the same memory resources. By enabling memory reuse, administrators can maximize the use of available physical memory, but this approach can introduce overhead as the system manages how the memory is accessed and shared. Performance issues may arise, particularly under heavy load or when running memory-intensive applications, where the separation and retrieval of memory from various virtual machines can lead to latency and reduced efficiency. Therefore, the assertion that using memory reuse can degrade virtual machine performance accurately reflects the reality of balancing resource optimization against individual VM performance, making it the correct answer concerning the question of incorrect descriptions regarding memory reuse.