



HP HPE5-H51

Delta - Advanced HPE Compute Solutions

Thank You for Downloading HPE5-H51 Updated Exam Questions

<https://certs4sale.com/hp/hpe5-h51-pdf-exam-dumps>

Question: 1

You are configuring an HPE system to handle advanced AI workloads. Which of the following combinations of technologies is the most suitable for GPU-intensive tasks in an on-premises infrastructure?

- A. Intel Xeon CPUs and NVIDIA A100 GPUs
- B. ARM processors and HPE accelerators
- C. AMD EPYC CPUs and NVIDIA T4 GPUs
- D. Intel Xeon CPUs and AMD Radeon GPUs

Answer: A

Explanation:

The combination of Intel Xeon CPUs and NVIDIA A100 GPUs is ideal for advanced AI workloads because of their high performance and compatibility for handling GPU-intensive tasks in an on-premises infrastructure.

Question: 2

You are tasked with optimizing a hybrid cloud environment for a client's data analytics workload. Which workload characteristic would most likely influence the need for as-a-service infrastructure over on-premises?

- A. High latency tolerance
- B. Predictable resource consumption
- C. Rapid scalability
- D. High upfront hardware costs

Answer: C

Explanation:

The need for rapid scalability is often a key factor for choosing as-a-service infrastructure because it allows the client to dynamically increase or decrease resources based on demand without large upfront hardware costs.

Question: 3

Which of the following is a primary benefit of using HPE memory-driven computing for big data workloads?

- A. Higher power efficiency
- B. Reduced GPU reliance
- C. Increased memory bandwidth and reduced latency
- D. Reduced need for network switches

Answer: C

Explanation:

HPE memory-driven computing is optimized for big data workloads due to its increased memory bandwidth and reduced latency, which significantly enhances performance for large-scale data

processing.

Question: 4

In the context of an AI workload, which CPU characteristic is most important when choosing a processor?

- A. Single-thread performance
- B. Number of cores
- C. Cache size
- D. Integrated graphics capability

Answer: B

Explanation:

For AI workloads, the number of cores is a critical factor as these tasks often benefit from parallel processing, where having more cores can improve the system's ability to handle large, complex computations simultaneously.

Question: 5

You are designing a high-performance computing (HPC) solution for a client running molecular simulations. Which memory technology should be prioritized to ensure low latency and high throughput?

- A. DDR4 SDRAM
- B. DDR3 SDRAM
- C. GDDR6
- D. HBM (High Bandwidth Memory)

Answer: D

Explanation:

For HPC solutions such as molecular simulations, High Bandwidth Memory (HBM) provides low latency and high throughput, making it the preferred memory technology for these intensive computing tasks.

Question: 6

You are deploying an HPE Synergy system to a data center with both compute and storage modules. Set up the logical interconnect to ensure optimal performance for high-bandwidth workloads. Include steps for configuring redundancy in the interconnects.

Answer:

To configure the logical interconnect for an HPE Synergy system:

1. Access HPE OneView and navigate to the Logical Interconnect Groups (LIG) section.
2. Create or modify a logical interconnect group.
3. Assign uplink sets for network connectivity, ensuring each set supports high bandwidth.
4. Enable link aggregation to increase throughput.
5. Configure redundancy by assigning active-passive or active-active failover across multiple interconnect modules.
6. Deploy the logical interconnect group to the available frames.

Question: 7

A client requires an on-premises compute platform to support a mission-critical ERP system. High availability and memory capacity are key concerns. Which HPE platform would you position?

- A. HPE Synergy
- B. HPE ProLiant
- C. HPE Apollo
- D. HPE Superdome Flex

Answer: D

Explanation:

HPE Superdome Flex is designed for mission-critical workloads with a focus on high availability and large memory capacity, making it ideal for running an ERP system.

Question: 8

You are positioning the HPE Synergy platform for a customer with a mixed workload environment that requires the ability to provision resources quickly. What feature of Synergy would be most beneficial for this requirement?

- A. HPE Smart Memory
- B. Embedded iLO management
- C. Synergy Composer
- D. HPE OneView

Answer: C

Explanation:

The Synergy Composer is a key feature of HPE Synergy that allows for the rapid provisioning of resources by automating and managing the allocation of compute, storage, and networking resources within a composable infrastructure.

Question: 8

You are tasked with configuring HPE GreenLake Central for a client who wants to implement a hybrid cloud strategy. Provide steps to ensure seamless management between on-premises resources and cloud resources.

Answer:

Steps to configure HPE GreenLake Central for hybrid cloud:

1. Log in to HPE GreenLake Central and connect both on-premises and cloud resources.
2. Set up real-time monitoring for resource usage.
3. Define scaling policies to automatically adjust resource allocation between on-prem and cloud environments.
4. Use OneView integration to manage the on-premises infrastructure.
5. Set up billing reports to track consumption across both environments.

Question: 9

A client is looking for a platform to support both on-premises and cloud-based workloads in a pay-per-use model. Which HPE platform should you recommend?

- A. HPE GreenLake
- B. HPE Apollo 6500
- C. HPE ProLiant DL360
- D. HPE Synergy

Answer: A

Explanation:

HPE GreenLake offers a pay-per-use model that supports both on-premises and cloud-based workloads, allowing customers to dynamically scale resources based on demand.

Question: 10

You are deploying an HPE Apollo 4000 system for big data analytics. Which features of this platform should you highlight to the customer as beneficial for handling large datasets?

- A. Integrated GPUs
- B. High storage density
- C. High throughput networking
- D. NVMe SSD support

Answer: B, C, D

Explanation:

The HPE Apollo 4000 is optimized for big data analytics with its high storage density, high throughput networking, and support for NVMe SSDs to handle the large datasets and fast I/O operations typical of big data workloads.

Question: 11

Which of the following benefits does HPE GreenLake provide for customers seeking a cloud-like experience on-premises? (Select TWO)

- A. Upfront hardware purchase
- B. Pay-per-use pricing
- C. Fixed pricing model
- D. Infrastructure as a service (IaaS)

Answer: B, D

Explanation:

HPE GreenLake offers a pay-per-use pricing model and delivers Infrastructure as a Service (IaaS) on-premises, giving customers the flexibility of a cloud-like experience without upfront hardware purchases.

Question: 12

You are responsible for setting up HPE Apollo 4000 for big data analytics with a focus on maximizing storage throughput. Configure the storage system to support high-speed data processing and optimize I/O performance.

Answer:

Steps to configure HPE Apollo 4000 for big data analytics:

1. Install and configure NVMe SSDs in the Apollo 4000 system for high-speed storage.
2. Set up RAID 0 or RAID 10 for enhanced throughput and performance.
3. Enable parallel I/O operations to optimize data access for analytics workloads.
4. Configure HPE OneView to monitor storage performance and I/O operations.
5. Test the system by running a big data analytics job and analyzing storage throughput.

Question: 13

You are working with a client to gather information about their existing environment before designing a solution. Which tools or methods could help you analyze the current environment?

- A. Business impact analysis
- B. Network monitoring tools
- C. Interviews with IT staff
- D. HPE OneView

Answer: C, D

Explanation:

HPE OneView can provide detailed infrastructure insights, and interviews with IT staff help gather specific operational challenges and needs. Network monitoring and business impact analysis are also useful but depend on the focus of the solution.

Question: 14

Which feature of HPE OneView provides the capability to manage and monitor multiple HPE platforms from a single interface?

- A. Unified API
- B. Hyperconverged management
- C. Infrastructure templates
- D. Single-pane-of-glass management

Answer: D

Explanation:

HPE OneView offers single-pane-of-glass management, enabling centralized management of multiple platforms, including compute, storage, and networking resources, from one interface.

Question: 15

You are tasked with deploying HPE Apollo 4000 for big data analytics. Configure the system to optimize storage performance and ensure high throughput for large data processing tasks.

Answer:

Steps to configure HPE Apollo 4000 for big data analytics:

1. Install high-performance NVMe SSDs to ensure fast data access.
2. Configure RAID 0 or RAID 10 for optimal balance between performance and redundancy.
3. Enable parallel I/O operations for efficient data processing.
4. Set up HPE OneView to monitor storage performance and identify bottlenecks.
5. Test the system by running a sample big data workload and analyzing throughput metrics.

Question: 16

Which critical step should be taken to verify the installed HPE Synergy solution matches the performance requirements outlined in the Customer Intent Documentation (CID)?

- A. Perform stress tests on the system's memory and CPU
- B. Review the installation logs for any errors
- C. Conduct user acceptance testing
- D. Schedule a follow-up performance review after one year

Answer: A

Explanation:

Performing stress tests on key components such as memory and CPU ensures that the HPE Synergy solution meets the performance requirements outlined in the CID.

Question: 17

A customer using HPE GreenLake has requested a solution to monitor their hybrid environment and provide recommendations for firmware and OS updates. Which tool should be deployed to offer AI-driven recommendations and automated monitoring?

- A. HPE OneView
- B. HPE InfoSight
- C. HPE iLO Amplifier Pack
- D. HPE Central Management Console

Answer: B

Explanation:

HPE InfoSight is an AI-driven tool that offers automated monitoring and provides recommendations for firmware and OS updates, helping ensure the customer's environment remains optimized.

Question: 17

A customer has deployed HPE GreenLake and requires automated monitoring of OS and application updates. Which tool should be used to manage updates and ensure continuous performance monitoring?

- A. HPE iLO
- B. HPE GreenLake Central
- C. HPE OneView
- D. HPE InfoSight

Answer: D

Explanation:

HPE InfoSight is the best tool for automated monitoring of OS and application updates, offering insights and recommendations to ensure performance remains optimal throughout the lifecycle of the solution.

Question: 18

You are troubleshooting a high memory utilization issue in an HPE Synergy system running data analytics workloads. Which step should be taken first?

- A. Increase memory capacity
- B. Enable memory compression
- C. Increase CPU core allocation
- D. Analyze application memory usage patterns

Answer: D

Explanation:

Before making hardware changes, it's essential to analyze application memory usage patterns to identify inefficiencies or memory leaks that may be consuming unnecessary resources.

Question: 19

A customer's HPE Apollo 6500 system is experiencing slow GPU processing during AI model training. Develop a strategy to optimize GPU utilization and reduce processing times.

Answer:

Steps to optimize GPU utilization:

1. Adjust the batch size of the AI model to maximize GPU performance.
2. Monitor GPU temperature and performance to identify any thermal throttling.
3. Enable NVLink (if available) to improve GPU-to-GPU communication.
4. Run test training sessions post-optimization to validate the improvement in processing times.

Question: 20

A customer using HPE Apollo 6000 reports high CPU utilization during data processing workloads. To optimize the system's performance, which test should be performed first to identify the bottleneck?

- A. Test memory bandwidth
- B. Run synthetic disk I/O tests
- C. Monitor CPU core utilization per workload
- D. Check network throughput

Answer: C

Explanation:

Monitoring CPU core utilization per workload helps identify if the workload is properly distributed across the CPU cores, ensuring that no single core is overloaded, which can degrade performance.

THANK YOU FOR DOWNLOADING HPE5-H51 UPDATED EXAM QUESTIONS

Note: Thanks For Trying The Demo Of Our HPE5-H51 Exam Product

Visit Our Site to Purchase the Full Set of Actual HPE5-H51 Exam Questions With Answers.

Money Back Guarantee



Click The Link Below

<https://certs4sale.com/hp/hpe5-h51-pdf-exam-dumps>