

Boost up Your Certification Score

HP

HPE7-S02

Advanced HPE Compute Integrator Solutions



For More Information – Visit link below:

<https://www.examsboost.com/>

Product Version

- ✓ **Up to Date products, reliable and verified.**
- ✓ **Questions and Answers in PDF Format.**

Latest Version: 6.0

Question: 1

For an HPE HPC environment using the Lustre file system, which component of the HPE Cray ClusterStor E1000 architecture is responsible for storing file metadata and managing the namespace?

- A. Object Storage Targets (OST)
- B. InfiniBand Leaf Switches
- C. Metadata Servers (MDS) and Metadata Targets (MDT)
- D. Lustre Management Server (LMS)

Answer: C

Question: 2

In an enterprise machine learning pipeline built on HPE AI Essentials, what is the specific role of the integrated "Feast" component?

- A. It acts as the primary distributed computing framework for scaling Python workloads.
- B. It provides a production-ready feature store for unified storage and serving of ML features for both batch and real-time inference.
- C. It tracks experiment parameters and provides a centralized model registry.
- D. It performs hyperparameter tuning using the Katib operator.

Answer: B

Question: 3

Which component of HPE AI Essentials provides an interactive SQL interface, allowing data scientists to query "Data Lakes" and "Data Warehouses" directly from their notebooks?

- A. Istio
- B. NVIDIA NIM
- C. KubeFlow
- D. EzPresto

Answer: D

Question: 4

A Cloud Administrator has installed the vSphere CSI driver on a Tanzu cluster. They are now creating a StorageClass for developers. What is the correct value for the provisioner field in the manifest to link this StorageClass to the vSphere CSI driver?

- A. csi.vsphere.vmware.com
- B. kubernetes.io/vsan
- C. hpe.com/storage-csi
- D. nvidia.com/gpu-storage

Answer: A

Question: 5

When configuring a vSAN Storage Policy for use with VMware Tanzu on an HPE compute cluster, which "Failure Tolerance Method" (FTM) is recommended for all-flash configurations to achieve maximum space efficiency while maintaining redundancy?

- A. RAID-1 (Mirroring)
- B. RAID-5/6 (Erasure Coding)
- C. RAID-0 (Performance)
- D. Distributed Fault Tolerance

Answer: B

Question: 6

A data scientist wants to use a pre-trained Large Language Model (LLM) but needs to ensure it only answers questions based on the company's internal private documents. Which architectural pattern, natively supported and documented within HPE AI Essentials, should be implemented?

- A. Retrieval-Augmented Generation (RAG).
- B. Training a new model from scratch.
- C. Zero-shot prompting on a public API.
- D. Moving all documents to a public S3 bucket.

Answer: A

Question: 7

While troubleshooting a cluster managed by HPE Performance Cluster Manager (HPCM), an administrator finds a node in the "Unknown" state. Which HPCM command should be used to check the console output of that node remotely to see where it hung during boot?

- A. `cpower --on`
- B. `ssh root@[node_ip]`
- C. `cm-nodes -v`
- D. `conman [node_name]`

Answer: D

Question: 8

How does HPE Private Cloud AI utilize NVIDIA NIM (NVIDIA Inference Microservices) to accelerate time-to-value for generative AI applications?

- A. NIM replaces the need for a hypervisor by running inference natively on the system BIOS.
- B. NIM provides pre-built, optimized, and containerized AI models that are ready to be served via industry-standard APIs.
- C. NIM dynamically overclocks the GPU memory bus based on real-time inference latency.
- D. NIM automatically generates synthetic training datasets for fine-tuning LLMs.

Answer: B

Question: 9

What is a primary hardware benefit of utilizing NVIDIA BlueField-3 DPUs in an HPE compute node designed for AI?

- A. It provides integrated liquid cooling for the CPU.
- B. It increases the amount of HBM3 memory available to the GPU.
- C. It offloads networking, security, and storage tasks from the host CPU, preserving more CPU cycles for application logic.
- D. It replaces the need for a physical boot drive (M.2) in the server.

Answer: C

Question: 10

During a "Day 2" audit of an HPE Private Cloud AI cluster, the administrator finds that several nodes have "Drifted" from their assigned OneView Server Profile.

Which OneView feature can be used to automatically bring these nodes back into compliance?

- A. Profile Re-apply
- B. Firmware Force-Flash
- C. iLO Auto-Correction
- D. Profile "Remediate" within the OneView dashboard

Answer: D

Thank You for Trying Our Product

For More Information – **Visit link below:**

<https://www.examsboost.com/>

15 USD Discount Coupon Code:

G74JA8UF

FEATURES

- ✓ **90 Days Free Updates**
- ✓ **Money Back Pass Guarantee**
- ✓ **Instant Download or Email Attachment**
- ✓ **24/7 Live Chat Support**
- ✓ **PDF file could be used at any Platform**
- ✓ **50,000 Happy Customer**



Visit us at: <https://www.examsboost.com/test/hpe7-s02>