

HP

HPE7-S01


Advanced HPE Compute Architect

For More Information – Visit link below:

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

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/hpe7-s01>

Latest Version: 6.0

Question: 1

An architect is detailing the networking for an HPE Private Cloud AI solution. Which hardware feature of the HPE Slingshot interconnect ensures that high-priority AI training traffic is not delayed by standard management or background storage traffic?

- A. Advanced Quality of Service (QoS) with hardware-based traffic class prioritization.
- B. Standard Link Aggregation Control Protocol (LACP) using multiple physical ports.
- C. Software-defined bandwidth throttling configured within the HPE OneView portal.
- D. Virtual Local Area Network (VLAN) tagging performed by the host operating system.

Answer: A

Question: 2

You are designing an AI solution that must interoperate with a third-party Red Hat OpenShift cluster. Which HPE component ensures that GPUs can be dynamically allocated to OpenShift pods?

- A. NVIDIA NVLink Bridge Software.
- B. HPE OneView for VMware vCenter.
- C. NVIDIA GPU Operator for Kubernetes.
- D. HPE iLO Amplifier Pack.

Answer: C

Question: 3

When presenting a design for an AI Inference solution, the customer asks for a "Cost-Effective" option for small models. Which HPE server and GPU combination should you propose?

- A. HPE Cray EX2500 with 8x NVIDIA H100 GPUs.
- B. HPE ProLiant DL320 Gen11 with 1x NVIDIA L4 GPU.
- C. HPE ProLiant DL380a Gen11 with 4x NVIDIA H100 GPUs.
- D. HPE Synergy 480 Gen11 with 2x NVIDIA A100 GPUs.

Answer: B

Question: 4

During the design phase, you find that the customer's facility water temperature is higher than standard. Which component of the HPE Cray EX2500 should you verify for compatibility?

- A. The Secondary Loop of the Cooling Distribution Unit (CDU).
- B. The internal fans of the Management Network switches.
- C. The local SSD endurance rating of the boot drives.
- D. The wattage of the redundant Titanium power supplies.

Answer: D

Question: 5

What is a critical architectural requirement for the storage layer in an AI solution powered by HPE Cray ClusterStor E1000?

- A. Support for long-term tape archiving.
- B. Connectivity via 10GbE copper BASE-T.
- C. Deduplication and compression to minimize flash wear.
- D. High-throughput parallel file system access (Lustre) to saturate GPU compute.

Answer: B

Question: 6

In a validated design for a medium-sized AI inference solution, why would you recommend HPE Alletra Storage MP over HPE MSA 2060?

- A. Because the MSA 2060 does not support any form of data encryption.
- B. Because Alletra MP provides the NVMe performance and 100% availability required for AI.
- C. Because Alletra MP is the only storage that can be physically mounted in a 2U rack.
- D. Because the MSA 2060 requires a specialized liquid-cooled cabinet to function.

Answer: A

Question: 7

A customer asks for a "Validated Solution" for Generative AI. Which HPE resource provides the pre-tested configuration and deployment guide?

- A. HPE Reference Architecture (RA).
- B. HPE QuickSpecs for individual servers.
- C. NVIDIA YouTube channel tutorials.
- D. HPE InfoSight General Dashboard.

Answer: A

Question: 8

To meet the 5PB high-bandwidth storage requirement for genomic datasets, which storage solution should be integrated into the design?

- A. HPE Alletra 9000 with All-NVMe Flash.
- B. HPE MSA 2062 with Tiered Storage modules.
- C. HPE StoreOnce 5650 with Catalyst software.
- D. HPE Cray ClusterStor E1000 using Lustre.

Answer: D

Question: 9

When demonstrating NVIDIA Base Command Manager, which feature shows the health and performance of the entire AI cluster in a single view?

- A. The Jupyter Notebook experiment log.
- B. The HPE Alletra MP storage console.
- C. The Cluster Management GUI Dashboard.
- D. The NVIDIA NIM API documentation.

Answer: C

Question: 10

Which software component in the HPE AI stack is used to manage and deploy "NIMs" (NVIDIA Inference Microservices)?

- A. NVIDIA AI Enterprise
- B. HPE InfoSight for Servers
- C. HPE OneView for HPE Synergy
- D. NVIDIA BlueField Manager

Answer: A

Thank You for Trying Our Product
Special 16 USD Discount Coupon: NSZUBG3X
Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>