



Vendor: HP

Exam Code: HPE7-S01

Exam Name: Advanced HPE Compute Architect Solutions
Written Exam

Version: DEMO

QUESTION 1

A customer plans to use Ray with HPE Private Cloud AI and wants to know more about the framework's integration with the solution's GPUs.

What should you explain?

- A. The solution's Kuberay cluster has a GPU worker, and AI administrators can define the number of GPUs for that worker.
- B. The solution's Kuberay cluster supports GPUs by default and can autoscale up to the maximum number of GPUs on a single AI-optimized node.
- C. The solution's Kuberay cluster supports GPUs by default and can autoscale up to the maximum number of GPUs in the solution.
- D. The solution's default Kuberay cluster does not support GPU workers, but the customer can deploy an additional Kuberay cluster that does.

Answer: A

Explanation:

Ray GPU support creates a separate GPU worker pod/group for GPU-accelerated Ray jobs, and administrators can configure its GPU resources; this GPU worker is not autoscaled.

Reference: <https://docs.ezmeral.hpe.com/unified-analytics/15/Ray/ray-gpu-support.html>

QUESTION 2

Which statement correctly describes hardware security capabilities for HPE Cray XD systems?

- A. A BMC ASIC enables secure boot any root of trust capabilities similar to those in HPE iLO.
- B. The HPE Cray XD systems include a sidecar node dedicated to security functions, such as monitoring firmware and OS for compromise.
- C. The HPE Cray XD systems use the same iLO ASIC and firmware as the HPE ProLiant DL line.
- D. BMC software integrates with an iLO ASIC to provide hardware root of trust capabilities.

Answer: A

Explanation:

HPE Cray XD systems use an HPE Cray XD BMC that provides hardware root of trust, recovery, chain of trust, and secure boot capabilities. They do not use the same iLO ASIC/firmware as HPE ProLiant DL servers.

Reference: <https://www.hpe.com/us/en/collaterals/collateral.a50004301enw.html>

QUESTION 3

Which statement correctly describes how you select HPE servers to act as nodes in a VMware vSAN ESA cluster?

- A. You can select vSAN ReadyNode configurations from the Broadcom Compatibility Guide or build your own configurations based on VMware guidelines.
- B. You can build your own configuration based on VMware guidelines or select predefined HPE Smart Templates for vSAN.
- C. You must select vSAN-compatible storage controllers from the Broadcom Compatibility Guide, but you make any choice for other hardware components.
- D. You must choose an appropriate HPE vSAN ReadyNode based on the Broadcom Compatibility Guide.

Answer: D

Explanation:

For vSAN ESA, HPE directs customers to use certified vSAN ESA ReadyNode configurations listed in the Broadcom Compatibility Guide; BYO/custom configurations are associated with vSAN OSA, not ESA.

Reference: <https://www.hpe.com/us/en/collaterals/collateral.a00067741enw.html>

QUESTION 4

You are configuring a proposal for HPE ProLiant servers to act as vSAN nodes in a new ESA cluster.

What is one piece of information needed to determine the nodes' drive capacity requirements?

- A. The RAID modes supported by the nodes' storage controllers
- B. The nodes' iSCSI capabilities
- C. The number of node failures to tolerate
- D. The desired ratio of cache versus capacity drives

Answer: C

Explanation:

vSAN ESA capacity sizing depends on the storage policy, especially Failures to Tolerate (FTT), because FTT determines the RAID/data placement overhead and therefore the required usable drive capacity.

Reference: <https://blogs.vmware.com/cloud-foundation/2022/11/02/capacity-overheads-for-the-esa-in-vsan-8/>

QUESTION 5

You want to demonstrate HPE Private Cloud AI running a complex AI application, such as a video search app, to a customer.

What is one good resource for putting together this demonstration more quickly?

- A. HPE GitHub repo for Airflow demos
- B. The tutorials folder in the HPE AI Essentials shared directory
- C. NVIDIA CUDA catalog
- D. NVIDIA AI blueprints

Answer: D

Explanation:

NVIDIA AI Blueprints are predefined, customizable reference workflows/sample AI applications designed to accelerate building and demonstrating complex AI use cases. HPE AI Essentials specifically describes Blueprints as workflows for creating and deploying generative AI applications.

Reference:

https://support.hpe.com/hpesc/public/docDisplay?docId=sd00007592en_us&docLocale=en_US&page=GUID-7899509F-3811-4AB7-9528-DCD0A948F4FE.html

QUESTION 6

Which statement correctly describes services for HPE ProLiant servers using direct liquid cooling

(DLC)?

- A. HPE Direct Liquid Cool Startup services are required for all racks of servers using DLC.
- B. HPE Direct Liquid Cool Startup services must be included for the first rack of servers that a customer deploys in a data center, but not for additional racks.
- C. HPE provides DLC consulting services and encourages partners to offer their own installation services.
- D. HPE provides optional DLC installation services, which partners can supplement with their own.

Answer: D

Explanation:

HPE ProLiant DLC services are not mandatory for every rack or only the first rack. HPE positions installation/startup services as optional lifecycle service choices, and HPE Authorized Channel Partners can be involved in service delivery.

QUESTION 7

What is one reason to suggest a VMware vSAN OSA deployment rather than an ESA deployment?

- A. The customer wants to use features such as compression, which are not supported with ESA.
- B. The customer wants to create a vSAN cluster that provides storage for other clusters but does not host VMs.
- C. The customer wants a cost-effective solution that incorporates HDDs.
- D. The customer prioritizes obtaining the highest write performance.

Answer: C

Explanation:

vSAN OSA supports hybrid configurations with SSD cache and HDD capacity drives, while ESA is optimized for modern all-flash/NVMe hardware and is the higher-performance architecture.

Reference: <https://techdocs.broadcom.com/us/en/vmware-cis/vsan/vsan/8-0/vsan-administration/building-a-virtual-san-cluster.html>

QUESTION 8

You are demonstrating HPE Private Cloud AI and HPE Machine Learning Inference Software to a customer. The customer says that the organization has a custom model, stored in an OpenLLM hub, which they would like to deploy with HPE Machine Learning Inference Software.

What should you explain?

- A. HPE Machine Learning Inference Software supports pulling this model from the hub and deploying it.
- B. While HPE Machine Learning Inference Software does not support deploying custom models, Ray Serve, also provided by HPE Private Cloud AI, does.
- C. While HPE Machine Learning Inference Software does not support deploying custom models, KServe, also provided by HPE Private Cloud AI, does.
- D. HPE Machine Learning Inference Software can deploy this model if the customer packages it with Bento Archive and saves it to a PVC in HPE Private Cloud AI.

Answer: A

Explanation:

MLIS supports OpenLLM registries, which allow it to pull OpenLLM-compatible models directly

from the Hugging Face/OpenLLM hub and use them to create deployments.

Reference:

https://support.hpe.com/hpesc/public/docDisplay?docId=a00aie112hen_us&docLocale=en_US&page=MLIS%2Fregistries%2Fopenllm.html

QUESTION 9

Company TXXX is working on autonomous vehicle simulation and needs a system optimized for high-throughput training of generative AI models. The customer already uses HPE iLO to manage other servers and would like to continue with that option.

Which HPE server model aligns with these goals?

- A. HPE Cray XD670
- B. HPE ProLiant DL580 Gen12
- C. HPE ProLiant Compute XD685
- D. HPE Cray XD2000

Answer: C

Explanation:

The HPE ProLiant Compute XD685 is optimized for high-throughput AI/model training workloads and uses HPE iLO management, matching the customer's need to continue managing systems through iLO.

QUESTION 10

A customer needs an upgrade for their ESXi hosts, and you are proposing HPE ProLiant DL servers. When you recommend using HPE custom ESXi images with these servers, the customer says that they have their own tools that they include on all images.

What should you explain?

- A. The customer can rely on the tools and drivers included in the HPE custom image to replace the functions of the tools the customer is currently using.
- B. The customer can create their own custom image with their tools and with the HPE bundles available in the vibsdepot.
- C. The customer should use an HPE OneView for vCenter OS Build Plan to combine their custom ESXi image with the HPE custom ESXi image.
- D. The customer must use the custom image for ESXi to load successfully on the servers. However, the customer can add their tools to the OS manually after the server boots.

Answer: B

Explanation:

HPE provides custom ESXi images, but customers can also use VMware Image Builder to create their own ESXi image, including their tools plus the required HPE drivers and bundles from HPE vibsdepot.

Reference: <https://vibsdepot.hpe.com>

<https://www.hpe.com/us/en/collaterals/collateral.a50010842enw.html>

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



10% Discount Coupon Code: ASTR14