

HP.HPE7-S01.v2026-10-10.q28

Exam Code:	HPE7-S01
Exam Name:	Advanced HPE Compute Architect Solutions Written Exam
Certification Provider:	HP
Free Question Number:	28
Version:	v2026-10-10
# of views:	117
# of Questions views:	280
https://www.freegas.com/qa/HP/HPE7-S01/HP.HPE7-S01.v2026-10-10.q28.html	

NEW QUESTION: 1

A customer runs latency-sensitive network functions in virtual machines and reports that the hypervisor virtual switch adds unacceptable overhead.

What should you recommend?

- A. Enable SR-IOV and assign virtual functions directly to the virtual machines
- B. Increase the number of vCPUs assigned to each virtual machine
- C. Enable jumbo frames on the virtual switch uplinks
- D. Move the virtual machines to a dedicated port group with NIC teaming

Answer: (SHOW ANSWER)

SR-IOV allows a physical adapter to present virtual functions that are passed directly to guest VMs, bypassing the hypervisor networking stack and reducing latency and CPU overhead. Note that SR-IOV can restrict features such as live migration in some configurations.

NEW QUESTION: 2

A customer wants a mirrored, hardware-protected boot volume that consumes no drive bays and is isolated from the data drives.

Which option should you recommend?

- A. Two SFF SATA SSDs configured as RAID 1 on the MR controller
- B. HPE NS204i-u NVMe Hot Plug Boot Optimized Storage Device
- C. A single internal USB boot device with a backup image
- D. Software RAID with the Intel VROC key

Answer: (SHOW ANSWER)

The NS204i series provides two M.2 NVMe drives in a dedicated boot device with hardware RAID

1 managed independently of the data storage controller. It does not consume front drive bays and keeps the OS boot volume separate from data.

NEW QUESTION: 3

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 (LEAVE A REPLY)

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.

NEW QUESTION: 4

A customer has an HPE Private Cloud AI solution, running HPE AI Essentials 1.11. The customer has identified a NIM for an LLM that they want to use. The customer has imported NeMo microservices on the AI Essentials platform. The customer plans to use NeMo Customizer's LoRA to fine-tune the LLM for their purposes.

What should you recommend that the customer do to deploy the fine-tuned model on HPE Private Cloud AI with HPE Machine Learning Inference Software?

- A. Set up a custom OpenLLM registry in which they add the fine-tuned weights.
- B. Store the fine-tuned weights in the local S3 object store.
- C. Use environment variables when they deploy the model to point to the NeMo entity store running on AI Essentials.
- D. Submit a support request to add the fine-tuned model to the NIMs hosted on HPE Private Cloud AI.

Answer: C (LEAVE A REPLY)

LoRA fine-tuning creates adapter weights, not a new NIM image. NVIDIA NIM supports serving LoRA adapters by setting environment variables such as NIM_PEFT_SOURCE to the NeMo Entity Store/Data Store, allowing the NIM to pull the fine-tuned adapters for inference.

NEW QUESTION: 5

An architect needs to apply a tested, mutually validated set of firmware and drivers to a ProLiant Gen11 fleet.

What should the architect use?

- A. HPE Service Pack for ProLiant (SPP) applied with Smart Update Manager (SUM)
- B. Individual component packages downloaded from the HPE Support Center
- C. HPE Intelligent Provisioning firmware update wizard

D. Redfish firmware inventory export

Answer: A (LEAVE A REPLY)

The SPP is a validated bundle of firmware, drivers, and system software tested together for a given server generation. SUM is the deployment engine that analyzes the target and installs only applicable components in the correct order.

NEW QUESTION: 6

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

- A.** A BMC ASIC enables secure boot and 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 (LEAVE A REPLY)

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>

NEW QUESTION: 7

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 (LEAVE A REPLY)

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>

NEW QUESTION: 8

A customer needs to replace their ESXi servers, so you propose HPE ProLiant DL servers. The customer requires TPM attestation.

What should you recommend?

- A. Using the OneView for vCenter plug-in to deploy the ESXi OS images to the servers
- B. Replacing the servers' default TPM certificates with a certificate signed by a private CA
- C. Purchasing servers with the trusted supply chain option
- D. Using the servers' UEFI Secure Boot mode

Answer: D (LEAVE A REPLY)

ESXi TPM attestation requires TPM 2.0 plus UEFI Secure Boot so vCenter can verify the host boot chain integrity. Trusted Supply Chain improves hardware provenance, but it is not the specific requirement for TPM attestation.

Reference: [https://techdocs.broadcom.com/us/en/vmware-](https://techdocs.broadcom.com/us/en/vmware-cis/vsphere/vsphere/7-0/vsphere-security/securing-esxi-hosts/securing-esxi-hosts-with-trusted-platform-module.html)

[cis/vsphere/vsphere/7-0/vsphere-security/securing-esxi-hosts/securing-esxi-hosts-with-trusted-platform-module.html](https://techdocs.broadcom.com/us/en/vmware-cis/vsphere/vsphere/7-0/vsphere-security/securing-esxi-hosts/securing-esxi-hosts-with-trusted-platform-module.html)

NEW QUESTION: 9

An HPE Private Cloud AI customer needs help planning how to use Airflow within that solution.

The customer currently has several GitHub repos in which different teams store their DAGs.

What should you advise AI administrators to do?

- A. Configure the Airflow framework to reference each existing GitHub repo using individual secrets for the credentials.
- B. Replicate the built-in Airflow framework to the namespace of each user with an existing GitHub repo.
- C. Create a recurring Spark pipeline to copy the DAGs from each GitHub to the Shared Data Volume from which Airflow pulls the DAGs.
- D. Set up one Github repo, to which they will refer the Airflow framework, with the existing GitHub repos added as submodules.

Answer: (SHOW ANSWER)

HPE AI Essentials Airflow is configured to pull DAGs from a specified Git repository, branch, and subdirectory. For multiple team DAG repositories, the practical approach is to point Airflow to one parent repo and include the other repos as Git submodules.

Reference:

[https://support.hpe.com/hpesc/public/docDisplay?](https://support.hpe.com/hpesc/public/docDisplay?docId=a00aie111hen_us&docLocale=en_US&page=Airflow%2FAirflow-dags-git-repos.html)

[docId=a00aie111hen_us&docLocale=en_US&page=Airflow%2FAirflow-dags-git-repos.html](https://support.hpe.com/hpesc/public/docDisplay?docId=a00aie111hen_us&docLocale=en_US&page=Airflow%2FAirflow-dags-git-repos.html)

NEW QUESTION: 10

A customer complains that firmware and BIOS settings drift between identical Synergy compute modules over time.

What should you recommend to enforce consistency?

- A. Apply server profiles created individually per compute module
- B. Use server profile templates and monitor for profile compliance
- C. Schedule quarterly SPP updates from the Composer
- D. Enable iLO Federation groups across all frames

Answer: B (LEAVE A REPLY)

Server profile templates define the desired configuration (firmware baseline, BIOS settings, connections, boot order). HPE OneView flags profiles that deviate from the template as non-compliant and allows remediation, which prevents configuration drift.

NEW QUESTION: 11

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 (LEAVE A REPLY)

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.

NEW QUESTION: 12

A customer must encrypt OS volumes using BitLocker with hardware-backed key protection on new ProLiant Gen11 servers.

What must be present?

- A. A TPM 2.0 module enabled in UEFI mode
- B. A Smart Array controller with a Secure Encryption license
- C. Self-encrypting drives with a local key manager
- D. iLO Advanced Premium Security Edition

Answer: A (LEAVE A REPLY)

BitLocker uses TPM 2.0 to seal the volume encryption key to platform measurements.

Controller-based Secure Encryption and SEDs protect data at rest at the drive level but do not provide the TPM key protector BitLocker requires.

NEW QUESTION: 13

A customer wants to deploy an AI-powered chatbot using HPE Private Cloud AI but needs help putting together all the components.

What should you recommend the customer use?

- A. An NVIDIA AI blueprint
- B. A Ray Serve deployment
- C. An Airflow DAG
- D. A Kubeflow pipeline

Answer: A (LEAVE A REPLY)

NVIDIA AI Blueprints provide prebuilt reference workflows for assembling complex AI applications such as chatbots, including models, NIMs, retrieval components, and deployment guidance. Ray Serve, Airflow, and Kubeflow can run pieces of a workflow, but they do not provide the end-to-end application blueprint.

NEW QUESTION: 14

What is a benefit of the storage component of HPE Private Cloud AI (small, medium, and large configurations)?

- A. It is built from the all-flash local drives on the AI-optimized nodes to provide the lowest latency.
- B. It is built on GreenLake for File Storage, which supports both NFS and Lustre for mixed AI workloads.
- C. It is built on GreenLake for File Storage, which supports low latency communications.
- D. It is built from hybrid local drives on the AI-optimized nodes to balance performance and cost-effectiveness.

Answer: (SHOW ANSWER)

Small, medium, and large HPE Private Cloud AI configurations use GreenLake for File with Object Storage / Alletra X10000 File and Object Storage, not local node drives. This provides shared, high-performance AI storage for the platform.

Reference: <https://www.hpe.com/us/en/collaterals/collateral.a50009216enw.html>

NEW QUESTION: 15

An HPE Synergy customer wants compute modules to boot a bare-metal OS image with no OS installed on local disks, so that hardware can be repurposed within minutes.

Which component provides this capability?

- A. HPE Synergy Composer
- B. HPE Image Streamer
- C. HPE Intelligent Provisioning
- D. HPE OneView Global Dashboard

Answer: B (LEAVE A REPLY)

HPE Synergy Image Streamer provides stateless, image-based OS deployment. Golden images and deployment plans are applied through server profiles, allowing compute modules to boot from the Image Streamer appliance instead of local storage.

NEW QUESTION: 16

A customer requires assurance that server firmware cannot be compromised by malicious code injected below the OS level.

Which HPE ProLiant feature addresses this requirement?

- A. UEFI Secure Boot with Microsoft-signed keys only
- B. Silicon Root of Trust anchored in the iLO ASIC
- C. Trusted Platform Module (TPM) 2.0 measured boot
- D. Intelligent Provisioning firmware validation

Answer: B (LEAVE A REPLY)

Silicon Root of Trust embeds an immutable fingerprint in the iLO silicon during manufacturing.

Every firmware component is cryptographically validated against that fingerprint before it is allowed to execute, and compromised firmware triggers automatic recovery. Secure Boot and TPM operate higher in the boot chain and depend on firmware already being trustworthy.

Valid HPE7-S01 Dumps shared by PrepPdf.com for Helping Passing HPE7-S01 Exam! PrepPdf.com now offer the **newest HPE7-S01 exam dumps**, the PrepPdf.com HPE7-S01 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com HPE7-S01 dumps with Test Engine here:

<https://www.preppdf.com/HP/HPE7-S01-prepaway-exam-dumps.html> (86 Q&As Dumps,

40%OFF Special Discount: Exam-Tests)

NEW QUESTION: 17

Which factors should you consider when determining the GPU memory requirements for inferencing an AI model?

- A. The number of parameters in the model and the floating point precision
- B. The memory available on the host and the number of convolutional layers in the model
- C. The memory available on the host and the number of parameters in the model
- D. The floating point precision and the number of layers in the model

Answer: A (LEAVE A REPLY)

GPU memory for inferencing is primarily driven by the model size, which depends on the number of parameters and the precision used to store them, such as FP32, FP16, BF16, or INT8.

NEW QUESTION: 18

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

Answer: A (LEAVE A REPLY)

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>

NEW QUESTION: 19

You are proposing 32 HPE Cray XD670 systems with HPE Performance Cluster Manager (HPCM) to a customer.

What is one benefit of HPCM that you should emphasize?

- A.** HPCM makes it much easier to deploy and patch the systems' operating systems.
- B.** HPCM provides GPU-compatible scheduling software for HPC jobs.
- C.** HPCM integrates with HPE Compute Ops Management to provide a single interface for monitoring a system's health.
- D.** HPCM enables the systems to support secure boot and hardware root of trust.

Answer: A (LEAVE A REPLY)

HPE Performance Cluster Manager provides bare-metal provisioning, image management, software updates, monitoring, and cluster health management for large HPC clusters.

Scheduling is handled by tools such as Slurm, not HPCM.

Reference: <https://www.hpe.com/us/en/collaterals/collateral.a00044858enw.html>

NEW QUESTION: 20

A customer runs latency-sensitive network functions in virtual machines and reports that the hypervisor virtual switch adds unacceptable overhead.

What should you recommend?

- A.** Move the virtual machines to a dedicated port group with NIC teaming
- B.** Enable jumbo frames on the virtual switch uplinks
- C.** Enable SR-IOV and assign virtual functions directly to the virtual machines
- D.** Increase the number of vCPUs assigned to each virtual machine

Answer: C (LEAVE A REPLY)

NEW QUESTION: 21

A customer must decommission ProLiant Gen11 servers and needs to sanitize all internal drives, NVRAM, and iLO settings before the servers leave the data center.

What should you recommend?

- A. Run Intelligent Provisioning and perform a factory reset
- B. Use iLO One-button secure erase
- C. Delete all logical drives in the Smart Array configuration utility
- D. Reflash the system ROM with the latest SPP

Answer: (SHOW ANSWER)

One-button secure erase in iLO performs a NIST SP 800-88 compliant sanitization of supported drives, non-volatile memory, and iLO/BIOS settings, and produces an erase report for audit.

Deleting logical drives or reflashing firmware does not remove user data.

NEW QUESTION: 22

What is one benefit of the NVIDIA AI Enterprise license provided in HPE Private Cloud AI?

- A. Drivers to permit executing training pipelines on GPUs outside of containers
- B. Autoscaling of models across one or more GPUs based on current demand
- C. Entitlement to use NVIDIA components in blueprints for accelerating AI application development
- D. NVIDIA Data Center GPU Manager (DCGM) available from HPE Private Cloud AI platform

Answer: (SHOW ANSWER)

The NVIDIA AI Enterprise license in HPE Private Cloud AI provides entitlement to NVIDIA enterprise AI software components used in AI workflows/blueprints, helping customers build and deploy AI applications faster.

Reference: <https://docs.nvidia.com/ai-enterprise/software/latest/index.html>

<https://www.nvidia.com/en-us/data-center/products/ai-enterprise/>

NEW QUESTION: 23

A customer with 30 rack-mount ProLiant servers wants to view aggregated health status and perform group power operations without deploying any management software.

What should you recommend?

- A. iLO Federation
- B. HPE OneView Standard
- C. HPE Compute Ops Management Enhanced
- D. Agentless Management Service (AMS)

Answer: (SHOW ANSWER)

iLO Federation uses multicast peer discovery to let any iLO in a group display aggregated health, license, and firmware information and perform group power and configuration actions, with no external management server required.

NEW QUESTION: 24

A customer reports lower-than-expected memory bandwidth on a two-socket ProLiant Gen11 server populated with 12 DIMMs per socket in an unbalanced pattern.

What should you recommend?

- A. Enable Memory Interleaving in the UEFI workload profile
- B. Populate DIMMs so that all memory channels per processor are equally populated
- C. Replace RDIMMs with LRDIMMs of the same capacity
- D. Set the memory operating speed manually to the maximum supported value

Answer: ([SHOW ANSWER](#))

DDR5 platforms achieve maximum bandwidth only when every memory channel of a processor is populated identically. Unbalanced population disables interleaving across channels and reduces effective bandwidth regardless of BIOS settings.

NEW QUESTION: 25

You are setting up a demonstration of HPE Private Cloud AI. You will show how easy it is to deploy a NIM with HPE Machine Learning Inference Software. As part of the demonstration, you worked with a developer to create a simple web application that sends queries to this NIM.

What is one important thing to keep in mind about your plan?

- A. You must establish a private NGC registry and collect its credentials to prepare for defining the source in HPE Machine Learning Inference Software.
- B. You must set up a storageclass that references HPE GreenLake for File Storage in advance to ensure that the deployment proceeds correctly.
- C. You will need to set up the application to include an API token, associated with the deployment, in the headers of requests.
- D. You must develop an API that understands the application's queries and incorporates those in the NIM as a custom image.

Answer: ([SHOW ANSWER](#))

HPE AI Essentials inference endpoints use access tokens for remote/client access, so the web application must pass the deployment's API token in the request headers when querying the NIM endpoint.

Reference:

[https://support.hpe.com/hpesc/public/docDisplay?](https://support.hpe.com/hpesc/public/docDisplay?docId=a00aie18hen_us&docLocale=en_US&page=pcai%2FGenAI%2Finfer-model-endpoint.html)

[docId=a00aie18hen_us&docLocale=en_US&page=pcai%2FGenAI%2Finfer-model-endpoint.html&](https://support.hpe.com/hpesc/public/docDisplay?docId=a00aie18hen_us&docLocale=en_US&page=pcai%2FGenAI%2Finfer-model-endpoint.html)

NEW QUESTION: 26

A customer wants to use the Spark framework included in HPE Private Cloud AI to ingress and process data stored in an external MinIO object store. The customer wants to avoid distributing the store's access and secret keys to AI Users.

What should you advise AI administrators to do?

- A.** Integrate the MinIO object store with the SSO and authorization framework provided by HPE Private Cloud AI's Keycloak.
- B.** Create a custom spark image that is pre-integrated with the MinIO object store.
- C.** Set up an HPE AI Essentials data source to the MinIO object store.
- D.** Replicate the data from the MinIO object store to a NAS store compatible with the HPE AI Essentials Data Volumes format.

Answer: ([SHOW ANSWER](#))

HPE AI Essentials lets administrators connect external S3 object stores such as MinIO as data sources. Users can then access them from clients such as Spark without being given the access key or secret key.

Reference:

[https://support.hpe.com/hpesc/public/docDisplay?](https://support.hpe.com/hpesc/public/docDisplay?docId=a00aie18hen_us&docLocale=en_US&page=ManageClusters%2Fconnect-object-stores.html)

[docId=a00aie18hen_us&docLocale=en_US&page=ManageClusters%2Fconnect-object-stores.html](https://support.hpe.com/hpesc/public/docDisplay?docId=a00aie18hen_us&docLocale=en_US&page=ManageClusters%2Fconnect-object-stores.html)

NEW QUESTION: 27

A customer wants to understand more about how HPE AI Essentials secures communications between pods running on HPE AI Essentials.

Which correctly describes how these pods authenticate each other when establishing TLS communications?

- A.** Using SVIDs obtained from the built-in SPIRE server
- B.** Using certificates issued by the HPE GreenLake identity provider
- C.** Using certificates signed by the built-in Istio gateway
- D.** Using tokens obtained from the built-in Keycloak

Answer: ([SHOW ANSWER](#))

HPE AI Essentials uses SPIRE/SPIFFE identities for workload authentication. Pods obtain SVIDs from the built-in SPIRE server, which are then used to establish authenticated TLS communication between workloads.

NEW QUESTION: 28

A government customer requires servers whose components are traceable and assembled under controlled conditions to reduce supply chain risk.

What should you recommend?

- A.** HPE Trusted Supply Chain servers
- B.** ProLiant servers with the high-security iLO security state enabled
- C.** Servers ordered with HPE Factory Express Level 4 integration
- D.** Standard ProLiant servers with TPM 2.0 modules installed

Answer: **A** ([LEAVE A REPLY](#))

HPE Trusted Supply Chain servers are built in controlled facilities by vetted personnel, ship with security settings enabled and verified firmware, and provide component traceability.

Factory Express provides custom integration but does not address supply chain provenance.

Valid HPE7-S01 Dumps shared by PrepPdf.com for Helping Passing HPE7-S01 Exam! PrepPdf.com now offer the **newest HPE7-S01 exam dumps**, the PrepPdf.com HPE7-S01 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com HPE7-S01 dumps with Test Engine here:
<https://www.preppdf.com/HP/HPE7-S01-prepaway-exam-dumps.html> (86 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)