

Advanced HPE Compute Integrator Solutions Written Exam

HP HPE7-S02

Version Demo

Total Demo Questions: 5

Total Premium Questions: 55

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Plan and Design HPE Compute Solutions	18
Topic 2, Install, Configure, and Set Up HPE Compute Solutions	20
Topic 3, Troubleshoot and Repair HPE Compute Solutions	2
Topic 4, Manage, Monitor, and Maintain HPE Compute Solutions	14
Topic 5, Mix Questions	1
Total	55

QUESTION NO: 1

A customer is running SUSE Rancher on HPE ProLiant DL servers and HPE Alletra arrays.

How should you install the HPE CSI Driver for Kubernetes?

- A. Download the latest Helm chart provided by HPE for the driver. Install that chart.
- B. Obtain the latest Ansible playbook for the driver from the HPE GitHub site. Run the playbook.
- C. Use SUSE Rancher Apps & Marketplace (the official Partner repository) to find and install the HPE CSI Driver.
- D. Install an HPE Data Ops Manager plug-in for SUSE Rancher. Use the plug-in to deploy the CSI driver.

ANSWER: C

Explanation:

Use SUSE Rancher Apps & Marketplace (the official Partner repository) to find and install the HPE CSI Driver. This is the supported Rancher-native deployment approach because Rancher manages Kubernetes applications as Helm chart releases through its Apps & Marketplace interface. An administrator can add or select the HPE Storage partner chart repository, locate the HPE CSI Driver chart, supply the required storage-system and Kubernetes settings, and deploy it to the intended cluster and namespace. This provides a consistent lifecycle workflow for installation, configuration review, version management, and later upgrades from the Rancher management interface.

The HPE CSI Driver supplies Container Storage Interface integration between Kubernetes workloads and supported HPE storage platforms, including the provisioning and attachment of persistent volumes. In a Rancher-managed environment, deploying the published HPE package through the Rancher catalog aligns the CSI installation with the platform's application-management model rather than treating it as an unrelated manual deployment. HPE's CSI documentation describes Helm-based deployment and the driver's Kubernetes integration, while Rancher documents its Apps & Marketplace workflow for installing and managing Helm charts. See the [HPE CSI Driver documentation](#) and the [Rancher Helm Charts documentation](#).

QUESTION NO: 2

What option does HPE Private Cloud AI provide for organizations who want to retain Prometheus metrics over the long term?

- A. Enabling metric collection on the HPE Private Cloud AI workspace from HPE GreenLake cloud
- B. Installing Grafana on the pod running the HPE AI Essentials UI
- C. Exporting alerts from the HPE AI Essentials UI
- D. Registering an OTel endpoint

ANSWER: D

Explanation:

Registering an OTel endpoint is the correct choice because it lets HPE Private Cloud AI send observability telemetry, including Prometheus metrics, to an external OpenTelemetry-compatible endpoint. This provides a path to use an organization's existing observability platform or metrics backend for durable storage, governance, querying, and retention policies that extend beyond the local platform's operational metrics window. The registered endpoint acts as the integration point through which telemetry can be exported for centralized monitoring and long-term historical analysis. This approach is particularly useful where teams need to correlate AI infrastructure and workload metrics with enterprise-wide application, platform, or security telemetry, while applying their own retention and compliance controls. HPE Private Cloud AI is designed to integrate AI operations with enterprise infrastructure and management capabilities, as described on the [HPE Private Cloud AI overview](#). OpenTelemetry is an open observability framework that standardizes the collection and export of telemetry to compatible backends; its metrics model and export architecture are documented in the [OpenTelemetry metrics documentation](#).

QUESTION NO: 3

A customer is designing an HPE Slingshot fabric for a much larger number of compute-node NICs than its current installation. The customer accepts that each switch will devote a greater proportion of ports to fabric connectivity in exchange for increased overall scale.

Which topology class best fits this requirement?

- A. Class 1
- B. Class 2
- C. Class 3
- D. A topology class selected according to the installed operating system

ANSWER: C

Explanation:

A Class 3 topology is intended to scale to a larger overall endpoint population. It uses more switch groups and higher-level inter-switch links, so more switch ports are allocated to fabric connectivity rather than endpoint-facing Level 0 connections.

A Class 1 topology provides higher compute-node NIC density per switch because it reserves fewer ports for the additional fabric layers. That characteristic is beneficial for smaller designs, but it does not meet the customer's primary requirement for greater total fabric scale.

QUESTION NO: 4

What is one action that an HPE Private Cloud AI Administrator can perform from the Alerts view in HPE AI Essentials?

- A. Restart the pod that generated the alert
- B. Configure alert forwarding rules
- C. View logs from the pod associated with the alert
- D. Correlate the alert with other alerts from the HPE GreenLake cloud workspace

ANSWER: C

Explanation:

View logs from the pod associated with the alert is correct because the Alerts view in HPE AI Essentials is an operational investigation entry point. An administrator can select an alert and access log information for the Kubernetes pod associated with the reported condition. This makes it possible to move directly from a symptom to workload-level diagnostic evidence, including application messages, container startup events, service failures, and resource-related errors.

Reviewing the affected pod's logs supports the standard HPE Private Cloud AI troubleshooting workflow: identify the alert, determine the impacted workload, inspect the available diagnostic details, and use that evidence to decide on an appropriate corrective action. The alert therefore provides context and navigation into the relevant workload telemetry rather than merely presenting a notification. This capability helps administrators investigate conditions without making an immediate change to the running environment, which is especially valuable when preserving evidence of an intermittent or application-specific failure.

HPE product documentation and learning materials describe AI Essentials as providing platform observability and alert-driven troubleshooting for managed AI workloads. See the [HPE Private Cloud AI documentation portal](#) and the [HPE Private Cloud AI solution overview](#) for product documentation access and operational context.

QUESTION NO: 5

An organization has an HPE Cray Storage Systems C500 solution. IT is interested in scripting management of these systems.

Which API can they use?

- A. HPE GreenLake
- B. Redfish
- C. HPE OneView
- D. Slurm

ANSWER: B

Explanation:

Redfish is the appropriate API for programmatic and scripted management of an HPE Cray Storage Systems C500 solution. Redfish is a standards-based, RESTful management interface that exposes resources through HTTPS and JSON payloads. This approach allows administrators to use common automation tools and scripts to query system inventory and health, retrieve status and event information, and perform supported management actions without requiring an interactive graphical interface. The API's resource-oriented model is particularly useful in high-performance computing environments, where repeatable automation is important for operating storage infrastructure consistently at scale.

Redfish is defined by the Distributed Management Task Force as an open industry standard for managing servers, storage, networking, and converged infrastructure. Its use of standard web mechanisms also makes it suitable for integration with configuration-management platforms, orchestration workflows, and custom operational tooling. HPE documents Redfish as a management interface within its infrastructure-management ecosystem, enabling API-driven administration where supported by the platform. For the Redfish standard and its data model, see the [DMTF Redfish standard page](#). HPE's Redfish developer resources are available through the [HPE iLO Developer portal](#).