Vendor: HP

Exam Code: HPE7-S02

Exam Name: Advanced HPE Compute Integrator Solutions Written Exam

Version: DEMO

QUESTION 1
You are proposing HPE Private Cloud AI to a financial institution and want to demo the solution running a complex AI-enabled security application. What can help you to accelerate the development of the demo?

A. Finding sample Kubeflow pipelines in the HPE AI Essentials shared volume
B. Looking for an appropriate NVIDIA blueprint in the NGC catalog
C. Using tools within NVIDIA AI Enterprise Base Command Manager
D. Searching for tutorials in the HPE AI Essentials documentation

Answer: B
Explanation:
NVIDIA Blueprints provide sample/reference AI applications that accelerate development of complex demos and enterprise AI workflows; they are available through NVIDIA AI Enterprise/NGC resources.

QUESTION 2
You need to connect several existing HPE Cray XD servers to a new HPE Slingshot network featuring HPE Slingshot ToR switches. The servers have existing 200 Gbps Ethernet NICs. What is required?

A. Replacing the NICs with InfiniBand NDR NICs
B. Checking whether the NICs are qualified for connecting to the HPE Slingshot switch
C. Replacing the NICs with HPE Slingshot NICs
D. Installing an LNet router between the servers and HPE Slingshot switches

Answer: B
Explanation:
HPE Slingshot can work with select qualified 100/200 GbE Ethernet adapters, not only HPE Slingshot NICs. Therefore, the existing 200 GbE NICs do not automatically need replacement, but they must be verified as qualified for the Slingshot ToR switch.

QUESTION 3
What is one difference between a Class 1 and a Class 3 HPE Slingshot topology?

A. The Class 1 topology supports more compute node NICs per switch that the Class 3 topology.
B. The Class 1 topology scales to a larger number of compute node NICs than the Class 3 topology.
C. The Class 1 topology supports a higher ratio of L1 and L2 links versus L0 links than the Class 3 topology.
D. The Class 1 topology supports additional adaptive routing technologies beyond those supported with the Class 3 topology.

Answer: A
Explanation:
Class 1 is a smaller Slingshot dragonfly topology with a single switch per group, so more switch ports can be used as L0 host/NIC links. Class 3 uses larger switch groups, which consume more ports for L1/L2 switch-to-switch connectivity and scale the fabric larger, but with fewer compute NICs per switch.

QUESTION 4
A customer needs to deploy a NIM LLM with HPE Machine Learning Inference Software running

on HPE Private Cloud AI (AI Essentials 1.6.1). What is one option for the image for serving the model?

A. A custom image stored in the customer's choice of registry
B. An image available in the public NGC registry
C. An image stored in the HPE Private Cloud AI registry
D. A BentoML image provided by HPE Machine Learning Inference Software

Answer: C
Explanation:
For AI Essentials 1.6.1, NIM LLM serving images are handled as prepackaged/curated images in the HPE Private Cloud AI registry. Direct dynamic NGC delivery is associated with later AI Essentials releases, not 1.6.1.

QUESTION 5
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 the SUSE Rancher Operators catalog to find the HPE driver. Install the driver from there.
D. Install an HPE Data Ops Manager plug-in for SUSE Rancher. Use the plug-in to deploy the CSI driver.

Answer: C
Explanation:
For SUSE Rancher, the HPE CSI Driver is delivered as a Rancher partner app/chart in the official Partner repository and is installed through Rancher Apps/Charts rather than by manually downloading Helm charts or using Ansible.

QUESTION 6
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:
HPE Private Cloud AI can export observability data to an external OpenTelemetry endpoint, allowing organizations to retain Prometheus metrics long term outside the default in-cluster retention period.

QUESTION 7
HPE Private Cloud AI Users want to run interactive Spark applications from notebook servers running on HPE AI Essentials. What should you do?

A. Create a custom notebook image that supports magic spark commands.
B. Make sure that Livy is installed from the HPE AI Essentials Administration settings for

Tools & Frameworks.

C. Ask users to select the image with “spark” in the name when they run a notebook server.
D. Configure the Spark Operator framework from HPE AI Essentials Tools & Frameworks; enable interactive access.

Answer: B
Explanation:
Interactive Spark sessions from notebooks use Livy; HPE AI Essentials Spark magic commands connect to the Livy server to start and manage Spark sessions. Spark Operator is for Spark Applications, not interactive notebook sessions.

QUESTION 8
An HPE Private Cloud AI customer wants to eliminate the certificate errors that show when users access HPE AI Essentials. The customer has created a certificate signed by their CA to replace the current certificate. The customer has given you the certificate file, private key file, and CA certificate file. How should you install this certificate?

A. Use the HPE Private Cloud AI interface to access the VM that hosts the HPE AI Essentials UI. Install the certificate, private key, and CA certificate on that VM.
B. Create a CertificateIssuer that references the CA certificate in the HPE AI Essential cert-manager namespace. Then create an ingress-cert certificate based on the provided certificate and private key in the istio-system namespace.
C. Log into the Keycloak interface using the credentials stored in the HPE AI Essentials ezaddon-system namespace. Upload the certificate, private key, and CA certificate files in the “ua” realm.
D. Delete the current “ingress-cert” certificate and secret in the HPE AI Essentials istio-system namespace. Add a new “ingress-cert” secret in this namespace referencing the provided files.

Answer: D
Explanation:
For a customer-provided certificate, key, and CA certificate, you replace the existing Istio ingress gateway certificate material by recreating the ingress-cert secret in the istio-system namespace. Cert-manager/ClusterIssuer is used when you want cert-manager to issue or renew the certificate automatically.

QUESTION 9
You are demonstrating HPE AI Essentials to a customer. You want to show running a DAG (workflow). Which framework or page should you access?

A. NVIDIA AI Enterprise
B. MLflow
C. Ray
D. Airflow

Answer: D
Explanation:
In HPE AI Essentials, Airflow is the framework used to create, view, and run DAG-based workflows. MLflow tracks experiments, and Ray is for distributed compute/tuning.

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**