# HP.HPE7-S02.v2026-11-14.q20

| | |
|---|---|
| **Exam Code:** | HPE7-S02 |
| **Exam Name:** | Advanced HPE Compute Integrator Solutions Written Exam |
| **Certification Provider:** | HP |
| **Free Question Number:** | 20 |
| **Version:** | v2026-11-14 |
| **# of views:** | 109 |
| **# of Questions views:** | 200 |
| https://www.freepdfdumps.com/HP.HPE7-S02.v2026-11-14.q20.html | |

**NEW QUESTION: 1**
A customer has fine-tuned a model and wants to find out whether the model is producing any toxic content.
Which NVIDIA microservice best meets this need?
**A.** NeMo Customizer
**B.** NeMo Curator
**C.** NeMo Guardrails
**D.** NeMo Evaluator
**Answer: (SHOW ANSWER)**
D is correct. NVIDIA NeMo Evaluator is designed to measure model and application quality against defined evaluation criteria, including safety-related dimensions. Current NeMo evaluation tooling includes safety evaluations for harmful-content generation, toxicity, bias, hate speech, and related robustness concerns. A customer who has fine-tuned a model and wants to determine whether it produces toxic output needs an evaluation service rather than a data-preparation or runtime-policy component.
NeMo Customizer is used to adapt or fine-tune models. NeMo Curator prepares training data through curation, filtering, and deduplication. NeMo Guardrails is used to apply runtime controls to conversational AI applications, such as restricting unsafe topics or validating inputs and outputs. Guardrails can help prevent unacceptable responses in production, but the question asks which microservice is used to evaluate whether the fine-tuned model is producing toxic content.
A sound AI lifecycle therefore separates customization, evaluation, and enforcement: customize the model, evaluate its behavior, and then apply appropriate guardrails and deployment controls.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 5 "NVIDIA NeMo deep dive"; NVIDIA NeMo Evaluator documentation, safety and security evaluation capabilities.

**NEW QUESTION: 2**
You are deploying several HPE ProLiant Gen12 servers to act as VMware vSAN nodes.
How should you configure the vSAN drives on the servers' storage controller?
**A.** In passthrough mode
**B.** In Erasure Coding (ECC) mode
**C.** With RAID 5
**D.** With RAID 1
**Answer: (SHOW ANSWER)**
A is correct. VMware vSAN is designed to manage its storage devices directly so that it can create disk groups or storage pools, apply its own data-placement policies, and monitor device health and performance.
On HPE ProLiant servers using a supported storage controller, vSAN data drives should therefore be presented in passthrough/HBA mode where the controller and vSAN design support it. This exposes the individual drives rather than hiding them behind a hardware RAID virtual disk.
RAID 1 and RAID 5 in options D and C would move redundancy and data placement into the controller, preventing vSAN from managing the devices as intended. Option B incorrectly treats erasure coding as a storage-controller mode. vSAN can apply software-defined RAID/erasure-coding policies at the datastore layer, but that is separate from configuring the physical controller.
The exact controller, firmware, driver, and device configuration must match the VMware Compatibility Guide and HPE vSAN design guidance. Passthrough is the architectural principle; supportability still depends on the validated hardware stack.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 8 "VMware vSAN on HPE solutions"; HPE and VMware vSAN configuration guidance.

**NEW QUESTION: 3**
A customer requires accessing the HPE AI Essentials KubeRay cluster interactively (with ray.init()) from an external client.
What should you explain?
**A.** External access is enabled by default at the Ray dashboard UI. The external client should direct the ray.
init() requests to TCP port 10001.
**B.** External access is not supported with the built-in KubeRay cluster. The customer should deploy a new cluster to achieve this access.
**C.** External access is not available by default, but can be achieved by, for example, changing the kuberay- head-svc to a NodePort service.
**D.** External access is enabled by default at the Ray dashboard UI. The external client should direct the ray.init() requests to TCP port 8265.
**Answer: (SHOW ANSWER)**

C is correct. The built-in KubeRay head service is normally exposed only inside the Kubernetes cluster, so an external client cannot simply call ray.init() against it by default. One supported way to provide interactive external access is to change kuberay-head-svc from ClusterIP to NodePort. After the service is exposed, the administrator identifies the assigned node port for the Ray client endpoint and connects to the cluster using a ray:// address.
Option A incorrectly states that external access is enabled by default. Port 10001 is the Ray client service port inside the service definition, but an external client still requires a reachable service exposure mechanism.
Option D confuses the Ray dashboard port, commonly 8265, with the Ray client endpoint.
Option B is too restrictive; a separate Ray cluster is not required merely to permit external interactive access.
Because NodePort expands network exposure, the change should be combined with appropriate firewalling, authentication, and access controls rather than treated as a convenience-only setting.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 4 "Using Ray"; HPE AI Essentials Software documentation, "Connecting to Ray Cluster."

**NEW QUESTION: 4**
What is one way that NVIDIA Spectrum-X meets the needs for AI networks?
**A.** It replaces IP ECN with PFC to provide more sophisticated congestion control for large AI flows.
**B.** It uses ECMP to provide fine-grained load balancing and to saturate redundant links.
**C.** It uses dedicated per-port buffers to ensure losslessness for RDMA traffic.
**D.** It uses targeted Flow Metering to manage congestion for bursty AI traffic.
**Answer: (SHOW ANSWER)**
D is correct. NVIDIA Spectrum-X is engineered for Ethernet-based AI fabrics where synchronized GPU traffic can create short, extremely large bursts. Its congestion-management architecture uses high-frequency telemetry and flow metering to identify the sources contributing to congestion and react with much finer control than conventional static Ethernet behavior. This is especially valuable for RDMA/RoCE AI traffic, where packet loss, tail latency, and uneven use of parallel paths can materially reduce GPU utilization.
Option A reverses the relationship between ECN and PFC. Spectrum-X does not simply replace ECN with Priority Flow Control; modern AI Ethernet designs combine congestion signaling and control mechanisms to avoid relying on broad pause behavior. Option B is also inaccurate because static ECMP hashing is not the defining fine-grained balancing mechanism; Spectrum-X uses adaptive routing and telemetry-aware decisions. Option C describes dedicated per-port buffering as the basis for losslessness, which is not the principal Spectrum-X innovation.
The exam objective is to understand why an AI network differs from ordinary data-center Ethernet: the network must react rapidly to elephant flows and collective-communication bursts while keeping accelerated nodes productive.

References/topics: Advanced HPE Compute Solutions, Rev. 26.21, AI/HPC networking objectives; NVIDIA Spectrum-X Ethernet networking documentation, congestion control and telemetry.

**NEW QUESTION: 5**
How do workloads running within HPE AI Essentials authenticate each other?
**A.** Using Kubernetes secrets auto-generated by Kyverno mutating admission controllers when workloads spin up
**B.** Using SVIDs that an Istio proxy helps them obtain from a built-in SPIRE server
**C.** Using certificates issued by the Istio ingress gateway built into HPE AI Essentials
**D.** Using self-signed certificates generated when HPE installs the solution
**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 6**
A hospital wants to fine-tune an LLM on their own dataset. However, the data needs to have sensitive information redacted.
Duplicate records also need to be cleaned up.
What can the customer use to achieve this goal?
**A.** Katib
**B.** NeMo Customizer
**C.** NeMo Curator
**D.** MLflow
**Answer: C (LEAVE A REPLY)**

**NEW QUESTION: 7**
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 (LEAVE A REPLY)**
C is correct. The Alerts view in HPE AI Essentials is designed to move an administrator from a detected condition to the evidence needed for diagnosis. From an alert, the administrator can open the logs for the pod associated with that alert. That is materially different from taking a remediation action such as restarting a pod: the Alerts interface provides observability and troubleshooting context, while restart or recovery actions are performed through the appropriate Kubernetes or platform-management workflow. Alert forwarding is separately configured in the platform's observability settings rather than as a per-alert action, and HPE GreenLake does not provide an alert-correlation function from this local Alerts view.

This is an important operational distinction for the exam. HPE Private Cloud AI troubleshooting follows an evidence-first workflow: identify the alert, inspect the affected workload and its logs, and then determine the corrective action. Pod logs can expose image-pull failures, scheduling problems, resource exhaustion, application exceptions, or service communication errors without immediately changing system state.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 7 "Troubleshooting HPE Private Cloud AI," especially alerts and solution-provided tools; HPE AI Essentials Software documentation,
"Viewing Alerts and Notifications."

**NEW QUESTION: 8**
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 (LEAVE A REPLY)**
B is correct. HPE Slingshot is Ethernet-based and can interoperate with selected standard Ethernet adapters, so existing 200 GbE NICs do not automatically have to be replaced. The key requirement is that the adapters be on the HPE-qualified support list for the Slingshot switches and intended configuration. Qualification confirms interoperability for link behavior, firmware, drivers, performance, and supported operating modes.
Option A is incorrect because HPE Slingshot is not an InfiniBand fabric and therefore does not require NDR InfiniBand adapters. Option C is too absolute; HPE Slingshot NICs are not the only possible edge adapters where supported Ethernet NICs are qualified. Option D confuses storage networking with host attachment.
LNet routers are used to bridge Lustre networking domains or provide routed access in certain HPC storage designs; they are not a general requirement for connecting 200 GbE server NICs to Slingshot top-of-rack switches.
The integrator's task is therefore to validate the exact adapter, firmware, cable/transceiver, and switch-port combination against HPE support documentation before deployment.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 10 "HPE HPC networking solutions"; HPE Slingshot QuickSpecs and qualified Ethernet-adapter guidance.

**NEW QUESTION: 9**
An HPE Private Cloud AI customer wants to import a custom framework. The customer wants to ensure that only authorized users access that framework's endpoint URL. You have created an authorization policy for this purpose.

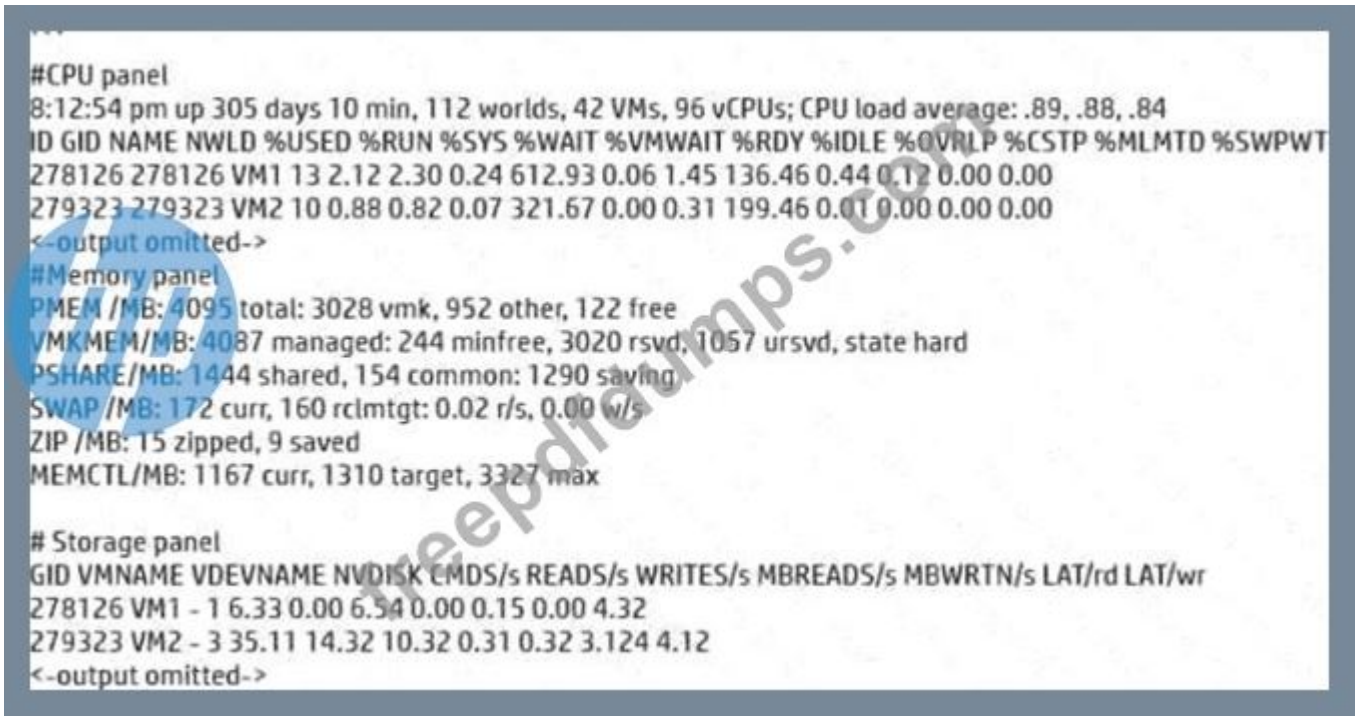
In which namespace should you add the policy?
**A.** istio-system
**B.** the namespace associated with your user account
**C.** ezapp-system
**D.** the namespace in which you will import the framework
**Answer: A (LEAVE A REPLY)**
A is correct. For an imported framework exposed through HPE AI Essentials, the authorization policy that controls access to the framework endpoint is placed in the istio-system namespace. The reason is architectural: access to the application is enforced at the Istio service-mesh and ingress layer, so the AuthorizationPolicy must be associated with the namespace where the relevant Istio gateway and policy enforcement context operate. The policy then limits which authenticated identities can reach the imported framework.
A user namespace is where many user-scoped workloads and resources run, but it is not the correct location for the ingress authorization policy in this workflow. The ezapp-system namespace contains platform application components and is not the documented target for this policy. Likewise, placing the policy only in the framework's application namespace would not implement the HPE AI Essentials ingress authorization pattern described for imported frameworks.
Current Rev. 26.21 training also introduces AutoHelm and updated framework-import behavior, but it continues to treat access control as an Istio authorization-policy function rather than a GreenLake user-setting shortcut.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 3 "Understanding and managing communications" and governance policies; HPE AI Essentials Software documentation, SSO/authorization for imported frameworks.

**NEW QUESTION: 10**
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 (LEAVE A REPLY)**
B is correct. HPE Cray Supercomputing Storage Systems C500 exposes management through the Redfish
/Swordfish REST API model. Redfish provides a standardized, HTTP-based management interface, and Swordfish extends the model for storage resources. This makes it suitable for scripted inventory, configuration, health, and lifecycle operations without requiring an administrator to automate an interactive management console.

HPE GreenLake is not the device-level scripting API for the C500. HPE OneView is a composable- infrastructure and server-management platform and is not the native C500 storage-management API. Slurm is a workload manager and scheduler for HPC jobs; it allocates compute resources and schedules work but does not replace the storage system's management API.
The spelling is also significant: the standard is "Redfish," not "RedFish." In an HPC environment, administrators commonly combine specialized tools rather than expecting one product to perform every management function. HPE Performance Cluster Manager or Slurm can participate in cluster operations, while the storage platform remains manageable through its supported Redfish/Swordfish interface.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Modules 9 and 10, HPE Cray management and HPC storage; HPE Cray Supercomputing Storage Systems C500 Redfish/Swordfish REST API documentation.

**NEW QUESTION: 11**
Which correctly characterizes HPE ProLiant Compute XD685 servers?
**A.** Support for 4 processors and 4 Lovelace GPUs per 5U chassis, combined with a choice of iLO or BMC management
**B.** Support for up to 8 NVIDIA Blackwell GPUs per 5U chassis, combined with iLO management
**C.** Scale-out design with up to 8 single-processor nodes per 5U chassis
**D.** Scale-up design with 8 processors and 8TB of memory per chassis
**Answer: (SHOW ANSWER)**
B is correct. HPE ProLiant Compute XD685 is a purpose-built accelerated system for large AI training and tuning workloads. Current HPE configurations support eight NVIDIA Blackwell-family GPUs, including B200 and Blackwell Ultra options, in a dense direct-liquid-cooled 5U chassis. The platform also uses HPE Integrated Lights-Out (iLO) for secure server-level management, giving it the familiar HPE management and security foundation expected in enterprise deployments.
Option A is incorrect in processor count, GPU generation, and management characterization. The XD685 is not a four-processor/four-GPU Lovelace design with a choice between unrelated management controllers.
Option C describes a multi-node scale-out chassis architecture that does not match the XD685; the 5U system is a single accelerated server platform with eight GPUs. Option D similarly invents an eight-processor, 8 TB scale-up design.
The exam objective is not merely to memorize a model name. The XD685 is positioned where workload requirements demand very high GPU density, large-model training capability, fast accelerator interconnects, and enterprise serviceability.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 1 "HPE solutions for AI"; HPE ProLiant Compute XD685 product specifications and current HPE Blackwell platform materials.

**NEW QUESTION: 12**
You are helping a customer troubleshoot performance issues on an HPE ProLiant DL server acting as a VMware ESXi host.
You have used the esxtop command, and the output is shown below. (In this output, the details for many VMs have been omitted; however, other VMs have similar stats to the ones that are shown.)




Based on this output, what should you explain?
**A.** The host needs more memory, or some VMs need to be moved to other hosts.
**B.** The performance issues derive from the storage backend, which might require a controller upgrade.
**C.** The host needs processors with more cores, or some VMs need to be moved to other hosts.
**D.** The performance issues derive from the storage network, which needs higher bandwidth.
**Answer: A (LEAVE A REPLY)**
A is correct. The esxtop output indicates memory pressure rather than CPU or storage saturation. The host has only about 122 MB of free physical memory, the VMkernel memory state is shown as hard, ballooning is active, and swap is in use. Those are strong indicators that ESXi is reclaiming memory because demand is approaching or exceeding comfortable physical capacity. Reclamation can increase latency sharply, especially when guest memory is swapped to storage.
The CPU figures do not show the classic signature of CPU contention. The displayed VMs have low %RDY values, so they are not spending excessive time waiting for physical CPU scheduling. The storage panel also shows relatively low read/write latency for the displayed workloads, which does not support a storage- controller or storage-network bottleneck as the primary cause.

The appropriate remediation is therefore to add physical memory, reduce VM memory commitments, or evacuate some VMs to other hosts. The screenshot is a good example of why VMware troubleshooting should correlate multiple esxtop panels rather than infer the bottleneck from utilization alone.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 8 "Troubleshooting HPE running VMware"; VMware ESXi esxtop memory-reclamation concepts.

**NEW QUESTION: 13**
A hospital wants to fine-tune an LLM on their own dataset. However, the data needs to have sensitive information redacted.
Duplicate records also need to be cleaned up.
What can the customer use to achieve this goal?
**A.** MLflow
**B.** NeMo Customizer
**C.** Katib
**D.** NeMo Curator
**Answer: D (LEAVE A REPLY)**
D is correct. NVIDIA NeMo Curator is the data-curation component intended to prepare large datasets for generative-AI training and fine-tuning. Its pipeline capabilities include filtering, quality processing, deduplication, and privacy-oriented data preparation. That directly matches both requirements in the scenario:
duplicate records must be removed and sensitive information must be redacted before the hospital's dataset is used to adapt the LLM. Performing these controls before fine-tuning reduces data leakage risk and prevents repeated records from biasing training.
The distractors address different stages of the AI lifecycle. MLflow is principally used for experiment tracking, model lifecycle management, and related MLOps functions; it is not the primary tool for cleansing the raw training corpus. NeMo Customizer performs model customization or fine-tuning after an appropriate dataset has been prepared. Katib is a Kubernetes-native hyperparameter-optimization component and does not provide the required privacy redaction and deduplication workflow.
The sequencing matters: curate the data first, then customize and evaluate the model. That separation is part of the software-stack knowledge expected for HPE Private Cloud AI with NVIDIA.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 5 "Using NVIDIA AI Enterprise and Serving Models," NeMo deep dive; NVIDIA NeMo Curator documentation.

**NEW QUESTION: 14**
An AI team is using Spark Applications to analyze and format data for AI fine-tuning. The team would like to automate running the applications in a series.
Which component of HPE AI Essentials can help?
**A.** Katib

**B.** Airflow
**C.** Ray Tune
**D.** MLflow
**Answer: B (LEAVE A REPLY)**
B is correct. Airflow is the workflow-orchestration component in HPE AI Essentials. It is designed to define directed acyclic graphs (DAGs) whose tasks execute in a controlled sequence based on dependencies. If a team has several Spark applications that prepare, transform, and validate data before model fine-tuning, Airflow can schedule those applications as stages in a repeatable pipeline and coordinate retries, dependencies, and execution order. Katib and Ray Tune are focused on hyperparameter optimization rather than general-purpose workflow orchestration. MLflow tracks experiments, runs, artifacts, and model lifecycle information; it is not the scheduler that executes an ordered series of Spark jobs. Airflow complements these components rather than replacing them.
The key exam distinction is between orchestration and execution engines. Spark performs distributed data processing. Airflow controls when the Spark applications run and how one stage depends on another. That separation permits the data-preparation pipeline to be automated without embedding orchestration logic inside each Spark application.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 4 "Using Spark" and "Airflow"; HPE AI Essentials Software Airflow documentation.

**NEW QUESTION: 15**
An HPE Private Cloud AI User wants to use KServe to run a NIM that is available in NGC. What is one prerequisite?
**A.** Creating an NGC registry in HPE Machine Learning Inference Software that includes credentials for a NVIDIA AI Enterprise account
**B.** Importing NVIDIA AI Enterprise into HPE AI Essentials using a Helm chart
**C.** Adding KServe decorators to the code that launches the NIM
**D.** Creating an NGC API Key and specifying it in a secret in the user's namespace
**Answer: D (LEAVE A REPLY)**
D is correct. A NIM image hosted in NVIDIA NGC is protected by NVIDIA registry authentication. To allow KServe in a user namespace to pull that image, the user needs a valid NGC API key and a Kubernetes secret containing the corresponding registry credentials in the namespace where the InferenceService will run. The workload can then reference that secret for the image pull or other required NGC access.
Option A describes registering NGC with HPE Machine Learning Inference Software. That is relevant when MLIS is the serving mechanism, but the scenario explicitly uses KServe. Option B is not how NVIDIA AI Enterprise is made available to HPE AI Essentials; it is not imported as a user Helm chart. Option C is also invalid because KServe does not require special "decorators" in application code to launch a NIM.
Namespace placement matters because Kubernetes secrets are namespace-scoped. A secret created elsewhere is not automatically available to the user's KServe deployment.

References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 5 "Other options for serving models on HPE Private Cloud AI" and NVIDIA AI Enterprise; HPE AI Essentials KServe guidance; NVIDIA NGC authentication documentation.

**NEW QUESTION: 16**
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: (SHOW ANSWER)**
B is correct. NVIDIA AI Blueprints are packaged reference workflows that combine models, NIM microservices, retrieval components, guardrails, and application logic for specific generative-AI use cases.
Selecting a suitable blueprint from the NVIDIA catalog can shorten a proof-of-concept cycle because the team starts from a validated architecture and deployable workflow instead of assembling every component from scratch. That is particularly useful when demonstrating a complex security-oriented AI application to a regulated customer.
Sample Kubeflow pipelines can accelerate individual machine-learning workflow steps, but they do not provide the same end-to-end application blueprint. NVIDIA Base Command Manager is an infrastructure and cluster-management product; it is not the development accelerator being asked for. HPE AI Essentials tutorials are useful instructional resources, but a blueprint is explicitly intended to speed implementation of a working use case.
Current Rev. 26.21 training has a dedicated module on using NVIDIA Blueprints and updates the recommended installation process for validated blueprints, reinforcing their role in demonstrations and PoCs.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 6 "Using an NVIDIA Blueprint to demo an AI Chatbot"; NVIDIA AI Blueprints catalog and deployment documentation.

**Valid HPE7-S02 Dumps** shared by Actual4test.com for Helping Passing HPE7-S02 Exam! Actual4test.com now offer the **newest HPE7-S02 exam dumps**, the Actual4test.com HPE7-S02 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com HPE7-S02 dumps with Test Engine here: https://www.actual4test.com/HPE7-S02_examcollection.html (**72** Q&As Dumps, **30%OFF** **Special Discount: Freepdfdumps**)

**NEW QUESTION: 17**
How do workloads running within HPE AI Essentials authenticate each other?
**A.** Using certificates issued by the Istio ingress gateway built into HPE AI Essentials
**B.** Using SVIDs that an Istio proxy helps them obtain from a built-in SPIRE server
**C.** Using Kubernetes secrets auto-generated by Kyverno mutating admission controllers when workloads spin up
**D.** Using self-signed certificates generated when HPE installs the solution
**Answer: B (LEAVE A REPLY)**
B is correct. HPE AI Essentials uses service identities based on SPIFFE concepts. Workloads receive SPIFFE Verifiable Identity Documents (SVIDs) from the platform's SPIRE infrastructure, and the Istio service mesh uses those identities to establish authenticated service-to-service communication. This gives workloads cryptographically verifiable identities that are independent of network location and supports mutual authentication between services inside the platform.
Option A incorrectly attributes workload identity to the Istio ingress gateway. The ingress gateway handles north-south access, but it is not the certificate authority that independently issues identities to every internal workload. Option C describes Kubernetes secrets generated by Kyverno, which is not the HPE AI Essentials workload-authentication architecture. Option D is too generic; static self-signed certificates created at installation would not provide the dynamic, workload-specific identity and rotation model that SPIRE supplies.
For integrators, the important design point is that the platform treats workload identity as a first-class security control. Istio provides the service-mesh enforcement and mTLS capabilities, while SPIRE supplies trusted identities used by workloads and proxies.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 3 "Understanding and managing communications"; HPE Private Cloud AI/HPE AI Essentials core-concepts documentation for Istio, SPIRE, and service identity.

**NEW QUESTION: 18**
An HPE Private Cloud AI Administrator has imported a custom framework. Users can click the framework's endpoint link and access the framework successfully. However, sometimes users will try to go to the framework's page and they get an error.
What is one thing to investigate?
**A.** Whether the framework's virtual service uses the correct service port number
**B.** Whether the framework's ezua.virtualService.endpoint value uses the correct domain name
**C.** Whether the users have been authorized to access the framework in the HPE GreenLake cloud User Management settings
**D.** Whether the correct redirect URL has been added for the framework to the Keycloak UA client
**Answer: (SHOW ANSWER)**

References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 3 imported-framework communications and troubleshooting; HPE7-S02 Release Notes, June 15, 2026, framework-import updates and automatic Keycloak redirect URL handling.

**NEW QUESTION: 19**
An HPE Private Cloud AI customer wants to replace the self-signed certificate that users see when they access HPE AI Essentials.
The customer wants to reduce maintenance efforts by having the solution auto-renew the certificate.
What is part of the implementation process?
**A.** Creating a ClusterIssuer object in HPE AI Essentials cert-manager namespace
**B.** Creating an authorization policy in HPE AI Essentials cert-manager namespace
**C.** Creating a secret in HPE AI Essentials istio-system namespace that includes a certificate, private key, and CA certificate
**D.** Creating an ingress gateway in HPE AI Essentials istio-system namespace
**Answer: A (LEAVE A REPLY)**
A is correct. When the objective is automated renewal of the HPE AI Essentials ingress certificate, cert- manager should control the certificate lifecycle. A ClusterIssuer is part of that implementation because it defines the certificate-authority or ACME issuing service that cert-manager can use cluster-wide. Once the issuer and the corresponding Certificate resource are configured correctly, cert-manager can request, replace, and renew the certificate rather than requiring repeated manual secret replacement.
Option C describes a manually supplied secret containing certificate material. That can be appropriate when an organization already has a certificate and wants to install it manually, but by itself it does not provide automated renewal. An AuthorizationPolicy controls traffic authorization, not certificate issuance, so B is unrelated. The Istio ingress gateway already exists as part of HPE AI Essentials; creating another gateway is not the required certificate-renewal step.
Rev. 26.21 specifically updates certificate material for HPE AI Essentials 1.9 and later, making it important to distinguish automated cert-manager workflows from manual ingress-cert replacement.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Module 3 "Replacing the Istio ingress gateway certificate"; HPE AI Essentials documentation, "Using cert-manager to Generate and Replace SSL Certificate."

**NEW QUESTION: 20**
A customer has HPE Cray XD systems, an HPE Slingshot network that includes HPE Slingshot Fabric Manager, HPE Performance Cluster Manager (HPCM), HPE Slingshot AIOps Software, and SchedMD Slurm.
Which tool provides heatmaps and interactive topologies that help you find trouble spots in the network performance?

**A.** HPE Slingshot AIOps software
**B.** HPE Slingshot Fabric Manager
**C.** SchedMD Slurm
**D.** HPCM
**Answer: A (LEAVE A REPLY)**
A is correct. HPE Slingshot AIOps Software is the analytics and visualization layer designed to expose network performance behavior across a Slingshot fabric. It provides topology-oriented views and telemetry heatmaps that allow an administrator to identify regions with abnormal bandwidth use, congestion, or other performance hot spots. Those visualizations are particularly useful when correlating network behavior with application or job activity.
HPE Slingshot Fabric Manager has a different primary role: fabric configuration, management, discovery, and operational control. SchedMD Slurm schedules jobs and allocates compute resources; although job information can be correlated with infrastructure telemetry, Slurm does not provide the Slingshot network heatmap interface described. HPE Performance Cluster Manager manages cluster provisioning and lifecycle operations rather than serving as the specialized Slingshot performance-analytics application.
The tools are complementary. Fabric Manager controls the network, HPCM manages the cluster, Slurm manages workload scheduling, and Slingshot AIOps helps diagnose how traffic is behaving on the fabric.
References/topics: Advanced HPE Compute Solutions, Rev. 26.21, Modules 9 and 10, HPE Cray management and HPC networking; HPE Slingshot AIOps Software User Guide, system-wide and telemetry heatmap views.

**Valid HPE7-S02 Dumps** shared by Actual4test.com for Helping Passing HPE7-S02 Exam! Actual4test.com now offer the **newest HPE7-S02 exam dumps**, the Actual4test.com HPE7-S02 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com HPE7-S02 dumps with Test Engine here: https://www.actual4test.com/HPE7-S02_examcollection.html (**72** Q&As Dumps, **30%OFF** **Special Discount: Freepdfdumps**)