**Vendor:** NVIDIA

**Exam Code:** NCP-AI

**Exam Name:** Nutanix Certified Professional - Artificial Intelligence v6.10

**Version:** DEMO

**QUESTION 1**
The administrator observes a "Pod Network Unreachable" health check failure in the NAI, specifically affecting several application pods on a newly added worker node. Existing pods on other nodes are functioning normally.

Upon initial investigation, kubectl get pods -o wide shows the affected pods stuck in ContainerCreating or Pending status, and kubectl describe pod <pod-name> events include messages, such as:

```
"FailedCreatePodSandBox: failed to create sandbox: rpc error: code =
Unknown desc = failed to set up network for sandbox..." or "network is
not ready"
```

What is the reason for this error and the appropriate next step that the administrator should take?

A. The Docker/containerd runtime on the worker node is corrupted; reinstall the container runtime.
B. The kube-proxy service is misconfigured; restart all kube-proxy pods on the cluster.
C. The underlying Nutanix network segmentation is incorrect; verify VLAN or IP pool configurations are correct in Prism Central.
D. The NKP CNI plugin on the affected worker node is not functioning correctly; inspect the logs of the CNI pod.

**Answer:** D

**QUESTION 2**
An AI/ML admin is testing access to an endpoint using Open AI compatible clients, but is unable to successfully access the endpoint. What could be the issue?

A. API key has been created by another AI/ML admin.
B. Endpoint has not been configured with the right LLM.
C. API key has been removed from the endpoint.
D. Endpoint as been created by another AI/ML admin.

**Answer:** C

**QUESTION 3**
What security requirement must a private container registry meet before pushing Nutanix AI mages?

A. Serve images over HTTPS with Basic-authentication only
B. Serve images over SSH with key-pair authentication required
C. Serve images over HTTP without trusted TLS/SSL certificates
D. Serve images over HTTPS with trusted TLS/SSL certificates

**Answer:** D

**QUESTION 4**
Which task is an AI/ML User unable to perform in Nutanix Enterprise AI?

A. Create an API key to access an AI inference endpoint.
B. Delete an endpoint that they previously created.

C. Test an AI inference endpoint that they previously created.
D. View a LLM that was imported by a different user.

**Answer:** D

**QUESTION 5**
What prerequisites must be skipped when adding GPU nodes to a managed Kubernetes service on an Azure AKS cluster using Azure CLI?

A. Install the GPU driver
B. Install the Worker driver
C. Install the Worker Connector
D. Install the GPU Connector

**Answer:** A

**QUESTION 6**
Nutanix Enterprise AI supports the NVIDIA GPU Operator when it is configured to use which GPU-virtualization mode?

A. GPU paravirtualization
B. SR-IOV
C. vGPU MIG
D. GPU Passthrough only

**Answer:** D

**QUESTION 7**
An administrator is monitoring the performance of a deployed Large Language Model within the Nutanix Enterprise AI platform. After initial deployment, users report slow inference response times and occasional timeouts when accessing the model through its API endpoint.

The administrator reviews the performance metrics available in the NAI Dashboard and notes the following:

CPU usage is consistently high across all inference-serving containers.

Memory utilization is nearing the allocated limits for the model service.

The request latency graph shows increasing average inference times during peak usage.

Which action should the administrator take to improve performance and reduce latency?

A. Restart the model container to clear memory cache and allow the system to rebalance performance.
B. Scale out the number of instances and allocate additional CPU and memory resources.
C. Disable logging temporarily to reduce resource consumption during peak load periods.
D. Increase the number of API keys assigned to the endpoint to allow more concurrent access.

**Answer:** B

**QUESTION 8**
An AI/ML administrator has been contacted by an endpoint owner who advises that one of their endpoints is not working properly anymore suggesting the API key has been removed. The AI/ML administrator can see that the API key is no longer listed under the endpoint. What filter under the Audit Events page can the AI/ML administrator use to help understand what has happened?

A. Action = Log Out
Entity Type = API Key
B. Action = Delete
Entity Type = API Key
C. Action = Log Out
Entity Type = Endpoint
D. Action = Delete
Entity Type = Endpoint

**Answer:** B

**QUESTION 9**
An administrator is configuring a new endpoint in Nutanix Enterprise AI (NAI) using an NVIDIA NIM model. During the configuration, the administrator notices that the Use GPU checkbox is selected and cannot be modified. What is the reason the checkbox is locked in this state?

A. Setting is locked due to a cluster-level policy restriction.
B. All endpoints in NAI require using GPUs.
C. System has already validated the endpoint configuration.
D. These models require GPU resources to run.

**Answer:** D

**QUESTION 10**
An administrator has been ask to troubleshoot the import of a new NIM LLM in NAI. The import Status is Pending. What could be the most common issue?

A. LLM repository read access restriction.
B. Unauthorized Hugging Face token.
C. Required storage space is not available.
D. System is downloading the LLM.

**Answer:** D

**QUESTION 11**
An administrator notices increased model inference latency and frequent timeout errors in a Nutanix AI deployment during peak usage. What is the most effective action to troubleshoot and resolve the performance issue?

A. Reduce the number of API keys to limit external access to the model.
B. Restart the Nutanix Kubernetes cluster to clear cached memory and reset service states.
C. Disable logging and monitoring tools to free up system resources.
D. Analyze resource metrics and scale out the model service to handle increased load.

**Answer:** D

**QUESTION 12**
How should a non-text-generation LLM endpoint be tested?

A. Copy the Sample API Code and URL and run it in the application.
B. Select Custom Request and upload the provided test file, and click Test.
C. Create a Sample Request, choose a default prompt, and click Test.
D. Open the URL on a Browser on the HTTPS web page.

**Answer:** A

**QUESTION 13**
Where can an API key be deactivated?

A. Settings
B. Endpoints
C. Dashboard
D. Admin Configuration

**Answer:** B

**QUESTION 14**
A developer with experience using Open AI compatible clients is exploring using OpenAI Compatible endpoints within Nutanix Enterprise AI.

An administrator has provided the developer with a list of OpenAI compatible endpoints that Nutanix Enterprise AI supports but one of them is unsupported.

Which option shows the unsupported endpoint type?

A. /v1/systems
B. /v1/embeddings
C. /v1/chat/completions
D. /v1/models

**Answer:** A

**QUESTION 15**
An administrator needs to download the Access token for installing Nutanix Enterprise AI. Which method should the administrator take to satisfy this task?

A. Download from Docker hub account.
B. Download the token from Prism Central marketplace.
C. Download from Nutanix portal.
D. Download from the NAI helm chart bundle.

**Answer:** C

# 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**