NVIDIA
NCP-AIOL
NVIDlA-Certified Professional: Al Operations

For More Information – Visit link below:
https://www.examsempire.com/

Product Version

1. **Up to Date products, reliable and verified.**
2. **Questions and Answers in PDF Format.**



Latest Version: 6.0

Question: 1

An operator needs a single pane of glass to watch node health, GPU utilization, and job status across a large NVIDIA cluster in real time.
Which tool provides this cluster-monitoring interface?

A. A spreadsheet updated by hand
B. Base Command Manager's Base View interface
C. A single node's local system log
D. The building's badge-access system

Answer: B

Question: 2

A team wants to stand up Kubernetes across the NVIDIA cluster's nodes in a consistent, managed way rather than installing it by hand on each host.
Which approach fits the blueprint?

A. Install and initialize Kubernetes on the NVIDIA hosts using BCM
B. Email each node's owner to install it themselves
C. Skip orchestration entirely
D. Run all workloads on a single laptop

Answer: A

Question: 3

A single powerful GPU must be partitioned into several smaller, isolated instances so multiple jobs can share it without interfering with each other.
Which NVIDIA feature enables this?

A. Overclocking the GPU
B. Disabling ECC memory
C. Multi-Instance GPU (MIG)
D. Uninstalling the driver

Answer: C

Question: 4

An HPC-style batch of training jobs must be queued and scheduled across many nodes with fair-share priorities on the cluster.
Which workload manager is designed for this?

A. A word processor
B. A web browser
C. A spreadsheet macro
D. Slurm

Answer: D

Question: 5

A team needs to fractionalize and schedule GPU resources across many data-science users, giving each a share while maximizing overall utilization.
Which NVIDIA tool is built for GPU orchestration and allocation?

A. NVIDIA Run:ai
B. A screen-recording app
C. A PDF reader
D. A calculator

Answer: A

Question: 6

An operator needs GPU-optimized, prebuilt container images for common AI frameworks rather than building every image from scratch.
Where should they pull these containers from?

A. A random public forum
B. An unsigned email attachment
C. The NGC catalog
D. A personal USB drive

Answer: C

Question: 7

A low-latency inference service must scale across pods and be managed declaratively on the cluster.
Which platform is the right fit for deploying the inference workload?

A. A single cron job on one node
B. Kubernetes
C. A static HTML page
D. A screensaver

Answer: B

Question: 8

On a multi-GPU system, peer-to-peer GPU communication over the high-speed interconnect has failed, and the operator suspects the service that manages NVLink and NVSwitch.
Which service should they investigate?

A. The email server
B. The office printer queue
C. The building HVAC controller
D. The fabric manager service for NVLink and NVSwitch

Answer: D

Question: 9

A containerized GPU job fails to start, and the operator needs to inspect the container runtime and image to find the cause.
Which component are they troubleshooting?

A. The parking-lot gate
B. Docker
C. The coffee machine
D. The desk lamp

Answer: B

Question: 10

A cluster must run the latest firmware and a consistent software image across all nodes after a maintenance window.
Which BCM capability handles this?

A. Manually reflashing each node from a USB stick
B. Ignoring updates until something breaks
C. Applying patches, updating firmware, and synchronizing software images across nodes using BCM
D. Rebuilding every node from scratch by hand

Answer: C

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

Check our Customer Testimonials and ratings available on every product page.

Visit our website.

https://examsempire.com/