



NVIDIA NCP-AIO Practice Questions

Shared by Franks on 21-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

An administrator is troubleshooting issues with an NVIDIA Unified Fabric Manager Enterprise (UFM) installation and notices that the UFM server is unable to communicate with InfiniBand switches.

What step should be taken to address the issue?

Options:

- A- Reboot the UFM server to refresh network connections.
- B- Install additional GPUs in the UFM server to boost connectivity.
- C- Disable the firewall on the UFM server to allow communication.
- D- Verify the subnet manager configuration on the InfiniBand switches.

Answer:

D

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

Communication issues between UFM server and InfiniBand switches often result from misconfigured or missing subnet manager configuration on the switches. The subnet manager controls fabric membership and routing, so verifying and correcting its setup is essential for proper UFM operation. Rebooting, adding GPUs, or disabling firewalls are less likely to resolve fabric-level communication problems.

Question 2

Question Type: MultipleChoice

What must be done before installing new versions of DOCA drivers on a BlueField DPU?

Options:

- A- Uninstall any previous versions of DOCA drivers.

- B- Re-flash the firmware every time.
- C- Disable network interfaces during installation.
- D- Reboot the host system.

Answer:

A

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

Before installing new versions of DOCA drivers on NVIDIA BlueField DPUs, it is required to uninstall any previous versions of DOCA drivers to prevent conflicts and ensure a clean upgrade. This ensures that the new installation is not affected by leftover files or configurations from earlier versions. Re-flashing firmware or disabling network interfaces is not always required before every driver installation. Rebooting the host system might be recommended after installation but is not a prerequisite before installing drivers.

Question 3

Question Type: MultipleChoice

You are configuring cloudbursting for your on-premises cluster using BCM, and you plan to extend the cluster into both AWS and Azure.

What is a key requirement for enabling cloudbursting across multiple cloud providers?

Options:

- A- You only need to configure credentials for one cloud provider, as BCM will automatically replicate them across other providers.
- B- You need to set up a single set of credentials that works across both AWS and Azure for seamless integration.
- C- You must configure separate credentials for each cloud provider in BCM to enable their use in the cluster extension process.
- D- BCM automatically detects and configures credentials for all supported cloud providers without requiring admin input.

Answer:

C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

When configuring BCM for cloudbursting across multiple cloud providers such as AWS and Azure, it is necessary to configure separate credentials for each cloud provider within BCM. This allows BCM to authenticate and manage resources appropriately in each distinct cloud environment. BCM does not automatically replicate or detect credentials, nor can a single credential set typically work across providers.

Question 4

Question Type: MultipleChoice

You are managing a Kubernetes cluster running AI training jobs using TensorFlow. The jobs require access to multiple GPUs across different nodes, but inter-node communication seems slow, impacting performance.

What is a potential networking configuration you would implement to optimize inter-node communication for distributed training?

Options:

- A- Increase the number of replicas for each job to reduce the load on individual nodes.
- B- Use standard Ethernet networking with jumbo frames enabled to reduce packet overhead during communication.
- C- Configure a dedicated storage network to handle data transfer between nodes during training.
- D- Use InfiniBand networking between nodes to reduce latency and increase throughput for distributed training jobs.

Answer:

D

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

For distributed AI training jobs that require fast inter-node communication, such as those using TensorFlow across multiple GPUs and nodes, InfiniBand networking is the preferred solution. InfiniBand provides ultra-low latency and high bandwidth, reducing communication delays

significantly and increasing overall training throughput. While jumbo frames on Ethernet can help, they do not match the performance of InfiniBand. Dedicated storage networks or increasing replicas do not directly address inter-node communication latency.

Question 5

Question Type: MultipleChoice

A system administrator needs to scale a Kubernetes Job to 4 replicas.

What command should be used?

Options:

- A- `kubecthl stretch job --replicas=4`
- B- `kubecthl autoscale deployment job --min=1 --max=10`
- C- `kubecthl scale job --replicas=4`
- D- `kubecthl scale job -r 4`

Answer:

C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

The correct command to scale a Kubernetes Job to a specific number of replicas is `kubecthl scale job --replicas=4`. This explicitly sets the number of desired pod instances for the Job resource. The other commands are either invalid (`stretch`), apply to Deployments rather than Jobs (`autoscale deployment`), or use incorrect syntax (`-r`).

Question 6

Question Type: MultipleChoice

An administrator needs to submit a script named "my_script.sh" to Slurm and specify a custom output file named "output.txt" for storing the job's standard output and error.

Which 'sbatch' option should be used?

Options:

- A- --o output.txt
- B- --e output.txt
- C- --output-output output.txt

Answer:

A

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

The correct sbatch option to specify a custom output file for both standard output and error is -o output.txt (or --output=output.txt). This option directs Slurm to write the job's standard output and error streams to the specified file. The -e option is for standard error only, and -output-output is not a valid option.

Question 7

Question Type: MultipleChoice

What two (2) platforms should be used with Fabric Manager? (Choose two.)

Options:

- A- HGX
- B- L40S Certified
- C- GeForce Series
- D- DGX

Answer:

A, D

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

NVIDIA Fabric Manager is designed to manage and optimize fabric resources like NVLink and NVSwitch in enterprise-class platforms such as HGX and DGX systems. These platforms have the necessary hardware fabric components. The L40S Certified and GeForce series are either not compatible or do not require Fabric Manager.



To Get Premium Files for NCP-AIO Visit

<https://www.p2pexams.com/products/ncp-aio>

For More Free Questions Visit

<https://www.p2pexams.com/nvidia/pdf/ncp-aio>

20%
DISCOUNT

P2P
exams