



We offer free update service for one year!
<https://www.certqueen.com>

Exam : NCP-AII

**Title : NVIDIA Certified
Professional AI
Infrastructure**

Version : DEMO

1.A system administrator noticed a failure on a DGX H100 server. After a reboot, only the BMC is available.

What could be the reason for this behavior?

- A. A boot disk has failed.
- B. The network card has no link / connection.
- C. There are more than two failed power supplies.
- D. Multiple GPUs have failed.

Answer: A

2.When updating the firmware on an NVLink switch transceiver, how can an engineer apply new firmware without interrupting the network?

- A. nv action reboot system to force immediate activation.
- B. mlxfwreset -d -lid 27 reset --yes to reset the transceiver
- C. Physically disconnect and reconnect the transceiver.
- D. flint -d -lid 27 --linkx --linkx_auto_update --activate

Answer: D

3.You are evaluating the integration of NVIDIA BlueField DPUs into your data center's storage architecture to optimize AI workloads. The storage solution chosen has incorporated BlueField DPUs to enhance performance and efficiency.

Which of the following benefits directly results from this integration?

- A. Reduced CPU load by offloading data processing tasks to DPUs.
- B. Enhanced I/O performance with NVMe storage access speeds.
- C. Elimination of latency issues in data processing tasks.
- D. Unlimited scalability by adding more DPUs without architectural changes.

Answer: A

4.During a 48-hour NeMo question-answering model burn-in test, GPU memory errors occur when processing large datasets.

Which configuration strategy prevents Out-of-Memory (OOM) errors while maintaining processing efficiency?

- A. Set blocksize="1GB" for data loading and enable RMM asynchronous allocation.
- B. Disable add_filename for Parquet files to reduce metadata.
- C. Increase files_per_partition to 1000 for larger batch processing.
- D. Switch from FP16 to FP32 precision for numerical stability.

Answer: A

5.A 24-hour HPL burn-in fails with "illegal value" errors during the first iteration.

Which **initial troubleshooting step** resolves this without compromising burn-in validity?

- A. Reduce test duration to 12 hours.
- B. Disable GPU affinity.
- C. Switch from FP64 to FP32 precision.
- D. Verify the matrix size is divisible by block size.

Answer: D