

100% Pass NCP-AII - NVIDIA AI Infrastructure—Reliable Exam Dumps



P.S. Free 2026 NVIDIA NCP-AII dumps are available on Google Drive shared by Exam4PDF: <https://drive.google.com/open?id=1tTR11dNxAGKxAPbtSGFUOVd39EDb3vcA>

As is known to all, NCP-AII practice test simulation plays an important part in the success of exams. By simulation, you can get the hang of the situation of the real exam with the help of our free demo of NCP-AII exam questions. Just as an old saying goes, knowing the enemy and yourself, you can fight a hundred battles with no danger of defeat. Simulation of our NCP-AII Training Materials make it possible to have a clear understanding of what your strong points and weak points are and at the same time, you can learn comprehensively about the NCP-AII exam and pass it easily.

NVIDIA NCP-AII Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> • Troubleshoot and Optimize: Covers identifying and replacing faulty hardware components such as GPUs, network cards, and power supplies, along with performance optimization for AMD • Intel servers and storage.
Topic 2	<ul style="list-style-type: none"> • Cluster Test and Verification: Covers full cluster validation through HPL and NCCL benchmarks, NVLink and fabric bandwidth tests, cable and firmware checks, and burn-in testing using HPL, NCCL, and NeMo.
Topic 3	<ul style="list-style-type: none"> • Physical Layer Management: Covers configuring BlueField network platform devices and setting up Multi-Instance GPU (MIG) partitioning for AI and HPC workloads.
Topic 4	<ul style="list-style-type: none"> • System and Server Bring-up: Covers end-to-end physical setup of GPU-based AI infrastructure, including BMC • OOB • TPM configuration, firmware upgrades, hardware installation, and power and cooling validation to ensure servers are workload-ready.
Topic 5	<ul style="list-style-type: none"> • Control Plane Installation and Configuration: Covers deploying the software stack including Base Command Manager, OS, Slurm • Enroot • Pyxis, NVIDIA GPU and DOCA drivers, container toolkit, and NGC CLI.

Up to one year of Free NVIDIA NCP-AII Exam Questions Updates

Though the content of our NCP-AII practice guide is the same, the varied formats indeed bring lots of conveniences to our customers. The PDF version of NCP-AII exam materials can be printed so that you can take it wherever you go. And the Software version can simulate the real exam environment and support offline practice. Besides, the APP online can be applied to all kind of electronic devices. No matter who you are, I believe you can do your best to achieve your goals through our NCP-AII Preparation questions!

NVIDIA AI Infrastructure Sample Questions (Q35-Q40):

NEW QUESTION # 35

You have installed the NVIDIA Container Toolkit and are attempting to run a container with GPU support. However, the 'docker run' command fails with an error indicating that the NVIDIA runtime is not found. You have already verified that the NVIDIA Container Toolkit is installed, and the Docker daemon has been restarted. What is the most likely cause of this error?

- A. The 'nvidia-container-runtime' package is not installed.
- B. The '/etc/docker/daemon.json' file is missing or has incorrect configuration settings related to the NVIDIA runtime.
- C. The system doesn't have a GPU.
- D. The NVIDIA driver version is incompatible with the CUDA version specified in the container image.
- E. The container image is corrupted and needs to be rebuilt.

Answer: B

Explanation:

The most likely cause is an issue with the 'S/etc/docker/daemon.json' file (A). This file configures Docker's runtime settings, including specifying the NVIDIA runtime. If the file is missing or has incorrect entries, Docker will not be able to find the NVIDIA runtime. While driver incompatibility (B) can cause issues, it typically manifests as runtime errors within the container, not a failure to find the runtime itself. 'nvidia- container-runtime' might be a required package depending on the installation method. A missing GPU is unlikely since the Toolkit would likely fail to install, although this is also an error that can prevent the NVIDIA runtime from being started.

NEW QUESTION # 36

After installing the NGC CLI, you attempt to run 'ngc config set' and encounter the following error: 'Error: API key is invalid or missing'.

What are the most likely causes of this issue and how can you resolve them?

- A. The NGC CLI configuration file is corrupted. Delete the file (A.ngc/config.json) and reconfigure the CLI.
- B. The NGC service is down. Check the NVIDIA NGC status page for any known outages.
- C. The host does not have network access to NGC.
- D. The NGC API key is incorrect or has expired. Verify the API key in your NVIDIA account and update the configuration using 'ngc config set'.
- E. The NGC CLI is not properly installed. Reinstall the package using 'pip install -upgrade nvidia-cli'

Answer: A,C,D

Explanation:

The most likely cause is an invalid API key (B) or a corrupted configuration file (C), or the host lacks network access (E). Reinstalling the package (A) might not resolve the issue if the problem lies with the API key or config file. While NGC service outages (D) are possible, they are less common.

NEW QUESTION # 37

What is the role of GPUDirect RDMA in an NVLink Switch-based system, and how does it improve performance?

- A. It encrypts data transmitted between GPUs, enhancing security.
- B. It facilitates the virtualization of GPUs, allowing multiple virtual machines to share a single physical GPU.

- C. It provides a mechanism for GPUs to offload compute-intensive tasks to the CPU, improving overall system throughput.
- **D. It allows GPUs to directly access each other's memory without involving the CPU, reducing latency and CPU overhead.**
- E. It enables direct communication between GPUs and storage devices, bypassing the network interface.

Answer: D

Explanation:

GPUDirect RDMA enables direct memory access between GPUs, bypassing the CPU and reducing latency. This significantly improves performance for applications that require frequent data transfers between GPUs. Other options describe functionalities that are not associated with RDMA in this context.

NEW QUESTION # 38

You're setting up a cluster with 8 NVIDIA A100 GPUs. Each GPU needs to read 4GB/s from storage to keep it fully utilized. The network connecting the storage and compute nodes has a bandwidth of 25GB/s. What is the maximum number of GPUs that can be simultaneously saturated with data without exceeding the network bandwidth?

- A. 0
- **B. 1**
- C. 2
- D. 3
- E. 4

Answer: B

Explanation:

Each GPU requires 4GB/s of bandwidth. The network has a bandwidth of 25GB/s. Therefore, the maximum number of GPUs that can be saturated is $25\text{GB/s} / 4\text{GB/s per GPU} = 6.25$ GPUs. Since you can't have a fraction of a GPU, the answer is 6 GPUs.

NEW QUESTION # 39

After replacing a faulty NVIDIA GPU, the system boots, and 'nvidia-smi' detects the new card. However, when you run a CUDA program, it fails with the error "no CUDA-capable device is detected". You've confirmed the correct drivers are installed and the GPU is properly seated. What's the most probable cause of this issue?

- A. The 'LD LIBRARY PATH*' environment variable is not set correctly.
- B. The new GPU is incompatible with the existing system BIOS.
- C. The user running the CUDA program does not have the necessary permissions to access the GPU.
- **D. The GPII is not properly initialized by the system due to a missing or incorrect ACPI configuration.**
- E. The CUDA toolkit is not properly configured to use the new GPU.

Answer: D

Explanation:

The error "no CUDA-capable device is detected", even when 'nvidia-smi' sees the GPII, points to a lower-level system issue that prevents CUDA from accessing the card. In such scenarios, ACPI (Advanced Configuration and Power Interface) misconfiguration is frequently the culprit. ACPI handles device initialization and power management. If ACPI doesn't properly configure the new GPU, CUDA programs won't be able to access it. Checking and correcting ACPI configuration would be the first line of action, which includes ensuring proper settings in the system BIOS/IJEFI related to PCI devices, especially those related to GPU/accelerators. LD LIBRARY PATH would affect runtime linking of CUDA libraries, but not the base device detection. User permissions are less likely to be the cause since 'nvidia-smr works.

NEW QUESTION # 40

.....

The language of our NCP-AII study materials is simple. The learners may come from many social positions and their abilities to master our NCP-AII study materials are varied. Based on this consideration we apply the most simple and easy-to-be-understood language to help the learners no matter he or she is the students or the in-service staff, the novice or the experienced employee which have worked for many years. NCP-AII Study Material use the simple language to explain the answers and detailed knowledge points and the concise words to show the complicated information about the NCP-AII study material.

