

# New Trustworthy NCP-AII Practice 100% Pass | Reliable NCP-AII: NVIDIA AI Infrastructure 100% Pass



BONUS!!! Download part of PracticeTorrent NCP-AII dumps for free: <https://drive.google.com/open?id=1K6KJjqyz6I7r4sxepScrAQX0hgXK0USy>

On PracticeTorrent website, you can easily prepare NCP-AII exam, also can avoid some common mistakes. Our IT elite team take advantage of their professional knowledge and experience, and probe into the IT industry development status by trial and error, finally summarizes PracticeTorrent's NVIDIA NCP-AII Exam Training materials. It is very accurate, authoritative. PracticeTorrent's NVIDIA NCP-AII exam dumps will be your best choice.

If you are preparing for an exam, it may spend lots of time, but don't worry, if you are preparing for the NCP-AII exam, the product of our company will help you save your time. The product of our company will list the major key points of the NCP-AII exam, and you can grasp the knowledge points as quickly as possible, therefore the time is saving. Besides, the product for NCP-AII Exam also provide specific training materials for the exam. And the PDF version is convenient to read, and support printing, while the software version stimulate the real environment of the NCP-AII exam. The APP online version is slao available of the product, you can learn at any time and at any place. Choosing our product, it wil help you.

>> Trustworthy NCP-AII Practice <<

## Exam NCP-AII Papers, NCP-AII Reliable Test Voucher

Most of our clients found our NCP-AII exam questions and answers amazing. All they learned from PracticeTorrent is that the NVIDIA NCP-AII practice test questions were accurately similar to the actual questions they faced on their NVIDIA AI Infrastructure exam. It made them utterly confident to go through the whole process of the NVIDIA AI Infrastructure. Feel free to compare our quality of NVIDIA NCP-AII Exam Questions dumps with other courses. Nothing can help people pass their NVIDIA NCP-AII certification exam more than we do. Even people who were on their first time taking NVIDIA Target NCP-AII certification can pass their NVIDIA AI Infrastructure exam with PracticeTorrent's help.

## NVIDIA NCP-AII Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> <li>Physical Layer Management: Covers configuring BlueField network platform devices and setting up Multi-Instance GPU (MIG) partitioning for AI and HPC workloads.</li> </ul>

Topic 2	<ul style="list-style-type: none"> <li>• Control Plane Installation and Configuration: Covers deploying the software stack including Base Command Manager, OS, Slurm</li> <li>• Enroot</li> <li>• Pyxis, NVIDIA GPU and DOCA drivers, container toolkit, and NGC CLI.</li> </ul>
Topic 3	<ul style="list-style-type: none"> <li>• System and Server Bring-up: Covers end-to-end physical setup of GPU-based AI infrastructure, including BMC</li> <li>• OOB</li> <li>• TPM configuration, firmware upgrades, hardware installation, and power and cooling validation to ensure servers are workload-ready.</li> </ul>
Topic 4	<ul style="list-style-type: none"> <li>• Troubleshoot and Optimize: Covers identifying and replacing faulty hardware components such as GPUs, network cards, and power supplies, along with performance optimization for AMD</li> <li>• Intel servers and storage.</li> </ul>
Topic 5	<ul style="list-style-type: none"> <li>• 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.</li> </ul>

## NVIDIA AI Infrastructure Sample Questions (Q67-Q72):

### NEW QUESTION # 67

You are attempting to install NGC CLI on a CentOS 7 system, but the 'pip install nvidia-cli' command fails with a 'Could not find a version that satisfies the requirement nvidia-cli' error. You have confirmed that 'pip' is installed and working. What could be the cause of this issue?

- A. The system's package manager (YUM) is not configured correctly, preventing 'pip' from finding the NGC CLI package.
- B. The CentOS 7 system does not have the required Python version installed. NGC CLI requires Python 3.6 or later.
- C. The 'pip' version is outdated and incompatible with the NGC CLI package. Upgrade 'pip' using 'pip install -upgrade pip'.
- D. The system's firewall is blocking access to the Python Package Index (PyPI). CentOS 7 is not supported by NGC CLI.

**Answer: B,C**

Explanation:

A likely reason is an outdated Python version (A), as NGC CLI requires Python 3.6 or later. Another potential issue is an outdated 'pip' version (C) which could be incompatible with the NGC CLI package. Confirming the correct python version and up to date pip usually resolves this issue. Option E is incorrect, CentOS 7 is supported with correct configuration.

### NEW QUESTION # 68

You are building a cloud-native application that uses both CPU and GPU resources. You want to optimize resource utilization and cost by scheduling CPU-intensive tasks on nodes without GPUs and GPU-intensive tasks on nodes with GPUs. How would you achieve this node selection and workload placement in Kubernetes?

- A. Use taints and tolerations to dedicate nodes without GPUs to CPU-intensive tasks and nodes with GPUs to GPU-intensive tasks.
- B. Use labels to identify the CPU and GPU-intensive nodes.
- C. Use node affinity rules to schedule CPU-intensive tasks on nodes with GPUs and GPU-intensive tasks on nodes without GPUs.
- D. Use node affinity rules to schedule CPU-intensive tasks on nodes without GPUs and GPU-intensive tasks on nodes with GPUs.
- E. Use resource quotas to limit the CPU resources available on nodes with GPUs and the GPU resources available on nodes without GPUs.

**Answer: D**

Explanation:

Node affinity rules allow you to specify constraints on which nodes a pod can be scheduled on. By using node affinity rules, you can

ensure that CPU-intensive tasks are scheduled on nodes without GPUs and GPU-intensive tasks are scheduled on nodes with GPUs. This optimizes resource utilization and cost. Taints and tolerations can be used, but affinity is more flexible. Resource quotas limit resource usage but do not control placement.

#### NEW QUESTION # 69

You are planning the network infrastructure for a DGX SuperPOD. You need to ensure that the network fabric can handle the high bandwidth and low latency requirements of AI training workloads. Which network technology is the RECOMMENDED choice for interconnecting the DGX nodes within the SuperPOD, and why?

- A. Gigabit Ethernet, because it's widely available and inexpensive.
- B. 10 Gigabit Ethernet, for a balance between cost and performance.
- C. Token Ring, because it's a reliable and deterministic networking protocol.
- D. Wi-Fi 6, for wireless connectivity and flexibility.
- E. InfiniBand, due to its high bandwidth, low latency, and RDMA support.

**Answer: E**

Explanation:

InfiniBand is the recommended network technology for DGX SuperPODs due to its high bandwidth, low latency, and support for RDMA (Remote Direct Memory Access). RDMA allows GPUs to directly access each other's memory without involving the CPU, significantly reducing latency and improving performance for distributed AI training workloads. Ethernet, even at higher speeds, generally doesn't offer the same level of performance and RDMA capabilities as InfiniBand.

#### NEW QUESTION # 70

Which of the following statements regarding the benefits of using a BlueField DPU for network offload are TRUE? (Select TWO)

- A. Automatic compatibility with all existing network protocols without requiring software updates.
- B. Increased network throughput due to hardware acceleration.
- C. Reduced CPU utilization on the host server for network-related tasks.
- D. Simplified network configuration compared to traditional NICs.
- E. Elimination of the need for a dedicated network interface card (NIC).

**Answer: B,C**

Explanation:

BlueField DPUs offload network tasks, freeing up CPU resources and improving network performance through hardware acceleration. Simplified configuration is not always guaranteed, as DPUs often require more configuration due to their advanced features. DPUs do not eliminate the need for a NIC, they are a NIC with advanced processing capabilities.

#### NEW QUESTION # 71

You are designing a storage solution for a new AI inference cluster that requires extremely low latency for model serving. Which storage technology and configuration would be MOST suitable to meet this stringent latency requirement?

- A. A distributed file system deployed on spinning HDDs with a large read-ahead cache.
- B. NVMe-oF (NVMe over Fabrics) using RDMA over Converged Ethernet (RoCE) connected to a cluster of NVMe drives.
- C. A traditional Fibre Channel SAN with a dedicated storage array.
- D. Amazon S3 object storage accessed over a high-bandwidth internet connection.
- E. A software-defined storage (SDS) solution running on commodity hardware with SATA SSDs.

**Answer: B**

Explanation:

NVMe-oF with RoCE (option B) provides the lowest latency due to the combination of NVMe's low-latency access to flash storage and RDMA's ability to bypass the CPU for data transfer. HDDs (A) are too slow. SATA SSDs (C) are slower than NVMe. Object storage (D) introduces network latency. Fibre Channel (E) is a viable option, but NVMe-oF typically offers lower latency and greater flexibility.

