

# NCP-AII勉強の資料 & NCP-AII最新関連参考書

---

**NVIDIA NCP-AII Exam**

**NVIDIA Certified Professional AI Infrastructure**

<https://www.passquestion.com/ncp-aii.html>



Pass NVIDIA NCP-AII Exam with PassQuestion NCP-AII questions  
and answers in the first attempt.

<https://www.passquestion.com/>

---

P.S.JpshikenがGoogle Driveで共有している無料の2026 NVIDIA NCP-AIIダンプ: <https://drive.google.com/open?id=1R26711QjZPapEW66cDiPhJV6AdwpOPA5>

NCP-AIIガイドトレントの3つの異なるバージョンの中で最も普及しているのはPDFバージョンであり、NVIDIAは特に適切であり、若者に歓迎されていることは間違いありません。このバージョンにはいくつかの機能があります。まず、NCP-AII準備ガイドのPDFバージョンを紙に印刷できます。ただし、メモを作成して重要な試験ポイントを強調することができます。前述のように、NCP-AII試験トレントサポート無料のデモダウンロードに加えて、NCP-AII準備ガイドを十分に理解し、適切で満足できる場合は購入することが理想的です。

競争がますます激しいIT業種では、NVIDIAのNCP-AII試験の認定は欠くことができない認証です。Jpshikenを選んだら、君が一回でNVIDIAのNCP-AII認定試験に合格するのを保証します。もしJpshikenのNVIDIAのNCP-AII試験トレーニング資料を購入した後、学習教材は問題があれば、或いは試験に不合格になる場合は、私たちが全額返金することを保証いたします。

>> NCP-AII勉強の資料 <<

## NCP-AII最新関連参考書 & NCP-AII試験感想

NVIDIAのNCP-AII試験トレントを購入した後、10分以内にできるだけ早く製品をお届けすることを保証します。そのため、長時間待つ必要がなく、配達時間や遅延を心配する必要はありません。NCP-AII準備トレントをすぐにオンラインでお客様に転送します。このサービスは、JpshikenのNCP-AIIテストブレインダンプが人々

の心をつかむ理由でもあります。さらに、NCP-AIIトレーニングガイドで20~30時間だけ学習すれば、NVIDIA AI Infrastructure試験に自信を持って合格することができます。

## NVIDIA AI Infrastructure 認定 NCP-AII 試験問題 (Q109-Q114):

### 質問 # 109

You need to verify the NVLink connectivity between GPUs in a DGX server. Which command-line utility is the MOST reliable and provides detailed NVLink status?

- A. lspci
- B. gpustat
- C. nvlink\_info (Hypothetical command)
- D. **dcgmi diag -t 1004**
- E. nvidia-smi

正解: D

解説:

'dcgmi diag -t 1004' is the correct command. 'nvidia-smi' provides basic GPU information, but 'dcgmi diag -t 1004' (part of the Data Center GPU Manager) provides specific diagnostic tests for NVLink connectivity. 'lspci' lists PCIe devices, not specifically NVLink. 'gpustat' is a monitoring tool. 'nvlink\_info' is hypothetical.

### 質問 # 110

Which of the following techniques are effective for improving inter-GPU communication performance in a multi-GPU Intel Xeon server used for distributed deep learning training with NCCL?

- A. **Utilizing InfiniBand or RoCE interconnects if available.**
- B. **Configuring NCCL to use the correct network interface and transport protocol (e.g., IB, Socket).**
- C. Increasing the system RAM size to minimize data transfer to disk.
- D. **Enabling PCIe peer-to-peer transfers between GPUs.**
- E. Disabling CPU frequency scaling to maintain consistent performance.

正解: A、B、D

解説:

Improving inter-GPU communication involves optimizing the network used for transferring data between GPUs. PCIe peer-to-peer, InfiniBand/RoCE, and proper NCCL configuration all contribute to faster communication. Increasing RAM size helps with data caching but doesn't directly affect inter-GPU communication speed. Disabling CPU frequency scaling is about CPU performance stability, not inter-GPU communication directly.

### 質問 # 111

You are tasked with automating the BlueField OS deployment process across a large number of SmartNICs. Which of the following methods is MOST suitable for this task?

- A. Manually flashing each SmartNIC using the 'bfboot' utility on a workstation.
- B. Utilizing a custom-built python script to flash each individual card, controlled from a central server. This method supports parallel flashing.
- C. **Using a network boot (PXE) server to deploy the BlueField OS image over the network. This allows centralized management and scalability.**
- D. Utilizing the 'dd' command to directly copy the image to each SmartNIC's flash memory.
- E. Creating a custom ISO image with the BlueField OS and booting each SmartNIC from a USB drive.

正解: C

解説:

PXE boot allows for automated and scalable OS deployment over the network, making it the most suitable option for managing a large number of SmartNICs. Manually flashing or using USB drives is not practical at scale, and using 'dd' directly can be risky and error-prone without proper checks.

### 質問 # 112

You are using the NVIDIA Container Toolkit in a Kubernetes environment with multiple GPUs per node. You want to ensure that pods can request specific GPUs on a node, rather than simply requesting 'any' GPU. Which Kubernetes feature, in conjunction with the NVIDIA Device Plugin, allows you to achieve this fine-grained GPU resource allocation?

- A. Device Plugins API
- B. Taints and Tolerations
- C. Resource Quotas
- D. Node Affinity
- E. Topology Manager

正解: E

解説:

The Kubernetes Topology Manager (C) allows you to align resource allocations (including GPUs) with specific NUMA nodes. This is critical for performance when dealing with multiple GPUs per node. Resource Quotas (A) limit resource usage but don't control specific GPU selection. Node Affinity (B) selects nodes based on labels, not specific GPUs. The Device Plugins API (D) enables GPU discovery, but the Topology Manager is needed for fine-grained allocation within a node. Taints and Tolerations (E) are used to prevent pods from being scheduled on certain nodes unless they have the corresponding toleration, and does not directly allow for the selection of a particular GPU.

### 質問 # 113

You are configuring a BlueField DPU to run a custom packet processing application. You want to ensure that the application has exclusive access to certain CPU cores on the DPU. Which mechanism is best suited for isolating CPU cores for your application on the Bluefield DPU?

- A. Modifying the DPU's bootloader configuration to disable the cores you want to reserve.
- B. Using CPU affinity settings within the application code itself.
- C. Utilizing cgroups (control groups) to create a dedicated cgroup for the application and limit its CPU usage to specific cores.
- D. Using 'taskset' command to pin the application's processes to specific cores.
- E. Adjusting the kernel's scheduler parameters to prioritize the application's threads on the desired cores.

正解: C

解説:

Cgroups provide a robust and flexible way to isolate and manage resources, including CPU cores, for applications. They allow you to create a dedicated cgroup for your application and limit its CPU usage to specific cores. 'taskset' is a viable option, but cgroups offer more comprehensive resource management capabilities. Modifying the bootloader is not a practical or recommended approach. CPU affinity settings in the application code depend on the application's design and may not be as reliable. Adjusting kernel scheduler parameters can be complex and affect other processes.

### 質問 # 114

.....

NCP-AII準備試験では、国内および海外の専門家と学者を取り入れた専門家のチームを集めて、関連する試験銀行の調査と設計を行い、受験者がNCP-AII試験に合格するのを支援します。ほとんどの専門家は長年プロの分野で勉強しており、NCP-AII練習問題で多くの経験を蓄積しています。当社は才能の選択にかなり慎重であり、夢のNCP-AII認定の取得を支援するために、専門知識とスキルを備えた従業員を常に雇用しています。

NCP-AII最新関連参考書: [https://www.jpshiken.com/NCP-AII\\_shiken.html](https://www.jpshiken.com/NCP-AII_shiken.html)

NVIDIA NCP-AII勉強の資料 これは本当のことです、試験に合格したお客様は「NCP-AIIオンラインテストエンジンを利用して、模擬試験を繰り返して受けました、だから、弊社のNCP-AII試験参考書はいろいろな資料の中で目立っています、NCP-AIIスタディガイドから多くのメリットを得ることができます、もちろん、NCP-AII問題集が一番いいです、弊社には、NCP-AII試験資料の合格率について、記載があります、NVIDIA NCP-AII勉強の資料 その後、試験に合格するのに十分な自信があります、NCP-AII試験の準備をするとき、がむしゃらにITに関連する知識を学ぶのは望ましくない勉強法です、NVIDIA NCP-AII 勉強の資料 弊社が提供した製品は一部

