

# 使用高質量的考試新版NCP-AII題庫準備您的NVIDIA NCP-AII考試，當然通過



BONUS!!! 免費下載NewDumps NCP-AII考試題庫的完整版: [https://drive.google.com/open?id=1B1VklD-5QrD71EtWKbBJ-4DZpa0\\_yNdz](https://drive.google.com/open?id=1B1VklD-5QrD71EtWKbBJ-4DZpa0_yNdz)

在這個人才濟濟的社會，人們不斷提高自己的知識想達到更高的水準，但是國家對尖端的IT人員需求量還在不斷擴大，國際上更是如此。所以很多人想通過NVIDIA的NCP-AII考試認證，但想通過並非易事。其實只要你們選擇一個好的培訓資料完全通過也不是不可能，我們NewDumps NVIDIA的NCP-AII考試認證培訓資料完全擁有這個能力幫助你們通過認證，NewDumps網站的培訓資料是通過許多使用過的考生實踐證明了的，而且在國際上一直遙遙領先，如果你要通過NVIDIA的NCP-AII考試認證，就將NewDumps NVIDIA的NCP-AII考試認證培訓資料加入購物車吧！

NewDumps擁有NVIDIA NCP-AII 認證考試的特殊培訓工具，能使你不用花費大量的時間和金錢就可以短時間獲得很多IT技術知識來提升你的技術，很快就能在IT行業中證明你的專業知識和技術。NewDumps的培訓課程是NewDumps的專家團隊利用自己的知識和經驗為NVIDIA NCP-AII 認證考試而研究出來的。

>> 新版NCP-AII題庫 <<

## 最新NVIDIA NCP-AII考證 - NCP-AII考試指南

當你在為準備NCP-AII考試而努力學習並且感到很累的時候，你知道別人都在幹什麼嗎？看一下你周圍跟你一樣要參加IT認證考試的人。為什麼當你因為考試惴惴不安的時候，他們卻都一副自信滿滿、悠然自得的樣子呢？是你的能力不如他們高嗎？當然不是。那麼想知道為什麼別人很輕鬆就可以通過NCP-AII考試嗎？那就是使用NewDumps的NCP-AII考古題。只用學習這個考古題就可以輕鬆通過考試。不相信嗎？覺得不可思議嗎？那就快點來試一下吧。你可以先體驗一下考古題的demo,這樣你就可以確認這個資料的品質了。快点击NewDumps的網站吧。

### NVIDIA NCP-AII 考試大綱：

主題	簡介
主題 1	<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>

主題 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>
主題 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>
主題 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>
主題 5	<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>

## 最新的 NVIDIA-Certified Professional NCP-AII 免費考試真題 (Q115-Q120):

### 問題 #115

Which protocol is commonly used in Spine-Leaf architectures for dynamic routing and load balancing across multiple paths?

- A. BGP (Border Gateway Protocol)
- B. STP (Spanning Tree Protocol)
- C. VRRP (Virtual Router Redundancy Protocol)
- **D. ECMP (Equal-Cost Multi-Path)**
- E. OSPF (Open Shortest Path First)

答案: D

解題說明:

ECMP (Equal-Cost Multi-Path) is crucial for efficiently utilizing the multiple paths available in a Spine-Leaf architecture. It allows traffic to be distributed across these paths, improving throughput and reducing congestion. OSPF and BGP can be used for routing but do not inherently provide per-packet load balancing. STP is used to prevent loops, and VRRP provides router redundancy, neither of which directly address load balancing across multiple equal-cost paths.

### 問題 #116

You are installing four NVIDIA A100 GPUs into a server designed for AI training. The server motherboard has multiple PCIe Gen4 x16 slots. However, the server's power supply unit (PSU) only has three 8-pin PCIe power connectors available. What is the BEST course of action to ensure all GPUs receive adequate power?

- A. Underclock the GPUs significantly to reduce their power consumption below the available PSU capacity.
- B. Connect the GPUs using the motherboard's internal SATA power connectors.
- C. Use a PCIe power splitter cable on one of the 8-pin connectors to power two GPUs.
- **D. Replace the existing PSU with a higher wattage PSU that has at least four 8-pin PCIe power connectors.**
- E. Install only three GPUs and leave the fourth unpowered.

答案: D

解題說明:

Using a splitter can overload the original connector and cause instability or damage. Leaving a GPU unpowered defeats the purpose. SATA connectors are not designed for the high power requirements of GPUs. Underclocking might work, but is not a reliable long-term solution. Replacing the PSU is the safest and most reliable solution.

### 問題 #117

You want to automate the NGC CLI installation process across multiple hosts in your infrastructure. What are the best practices to achieve this?

- A. Distribute the './.ngc/config.json' file to all hosts.
- B. Create a custom script that downloads the NGC CLI package, installs it using 'pip', and configures the API key.
- C. Use a configuration management tool like Ansible or Chef to automate the installation and configuration of the NGC CLI on all hosts.
- D. Manually install the NGC CLI on each host, as automation is not recommended for security reasons.
- E. Use a Dockerfile to create a container image with the NGC CLI pre-installed and configured.

答案: B,C,E

解題說明:

Automation is highly recommended. Configuration management tools (A), custom scripts (B), and containerization (D) are all viable options for automating the NGC CLI installation process. Manually installing on each host is inefficient and error-prone. Distributing the config.json (E) could be a security risk.

### 問題 #118

After a firmware upgrade on a DGX H100, the administrator notices that one GPU is not detected by the system. Which troubleshooting step should be performed first to identify the root cause?

- A. Immediately re-run the firmware upgrade on all system components.
- B. Ignore the issue and proceed with production workloads if the other GPUs are operational.
- C. Review firmware update logs and run `nvsm show health` to check for hardware or firmware errors on the affected GPU.
- D. Remove the GPU from the system and replace it with a new one before any diagnostics.

答案: C

解題說明:

The first step is to review the firmware update logs and run `nvsm show health`. After a DGX H100 firmware upgrade, a missing GPU can result from incomplete firmware activation, failed component update, PCIe enumeration failure, GPU tray communication issues, BMC inventory mismatch, or an actual hardware fault.

NVSM is the correct DGX platform-level health tool because it checks hardware state across GPUs, NVSwitch components, PCIe devices, storage, power, cooling, and system sensors. Firmware logs are equally important because they show whether each update completed successfully and whether a reboot, cold power cycle, BMC reset, or AC power cycle is still required. Replacing the GPU immediately is premature and may cause unnecessary downtime. Ignoring the issue is unsafe because production AI workloads expect all GPUs to be visible and healthy. Re-running firmware across all components without diagnosis can hide the original failure or introduce more risk. Proper bring-up practice is to collect evidence, verify hardware health, confirm firmware activation state, and then decide whether reseating, power cycling, reapplying firmware, or service escalation is required.

### 問題 #119

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

答案: C

解題說明:

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.

