

Updated and Error-free BraindumpStudy NCP-AII Exam Practice Test Questions

 **Get the Premier NVIDIA NCP-AII Exam PDF Questions**
— Succeed on Your First Shot!

Are you aiming to clear the NVIDIA NCP-AII exam questions on your very first try? Certification exams in 2025 are harder than ever, demanding extensive knowledge and practical experience. To enhance your chances of accomplishment, you need reliable and latest learning tools. HelloDumps provides [NVIDIA NCP-AII dumps](#) and expert-reviewed study materials to help you get ready effectively. The premier part? You can now enjoy an exclusive 35% promotion on all study materials!

Why Choose HelloDumps for Your NCP-AII Exam Preparation?

 **Legitimate NCP-AII Exam Questions – Obtain an Superiority Over Others**

HelloDumps offers actual exam-like questions, closely matched with the latest NCP-AII composition, difficulty level, and syllabus of the NVIDIA Certified Professional AI Infrastructure certification. Practicing these questions helps you understand the test structure and enhances your output.

 **Retrieve NVIDIA Certified Professional AI Infrastructure Exam Questions 2025 Here: <https://hellodumps.com/ncp-aii-dumps.html>**

 **Three Practical Learning Formats**

We provide study materials in various formats to suit different learning styles:

-  **PDF Layout** – Study on any device, anytime, anywhere.
-  **Desktop Training Test** – Simulates real exam conditions with customizable mock tests.
-  **Web-Based Exercise Test** – No installation required—practice directly from your browser.

BONUS!!! Download part of BraindumpStudy NCP-AII dumps for free: https://drive.google.com/open?id=1PFmVLTJm2BxFzJgQdUCFPni7tVK57_rv

There are three versions of our NCP-AII exam questions: the PDF, Software and APP online. Now I want to introduce the online version of our NCP-AII learning guide to you. The most advantage of the online version is that this version can support all electronic equipment. If you choose the online version of our NCP-AII Study Materials, you can use our products by your any electronic equipment. We believe it will be very convenient for you, such as IPAD, phone and laptop.

The NVIDIA NCP-AII mock tests are specially built for you to evaluate what you have studied. These NVIDIA AI Infrastructure (NCP-AII) practice exams (desktop and web-based) are customizable, which means that you can change the time and questions according to your needs. Our NVIDIA AI Infrastructure (NCP-AII) practice tests teach you time management so you can pass the NVIDIA AI Infrastructure (NCP-AII) certification exam.

>> **New NCP-AII Exam Labs** <<

Free PDF Quiz 2026 NCP-AII: NVIDIA AI Infrastructure – High-quality New Exam Labs

The client only needs 20-30 hours to learn our NCP-AII learning questions and then they can attend the test. Most people may devote their main energy and time to their jobs, learning or other important things and can't spare much time to prepare for the NCP-AII test. But if clients buy our NCP-AII Training Materials they can not only do their jobs or learning well but also pass the NCP-AII test smoothly and easily because they only need to spare little time to learn and prepare for the NCP-AII test.

NVIDIA AI Infrastructure Sample Questions (Q33-Q38):

NEW QUESTION # 33

A system administrator needs to install a container toolkit and successfully run the following commands:

```
sudo apt-get update
sudo apt-get install -y nvidia-container-toolkit
sudo nvidia-ctk runtime configure --runtime docker
What step should be taken next to finish the installation?
```

- A. apt-get remove nvidia-container-toolkit
- B. apt-get install cuda-drivers
- C. systemctl restart docker
- D. dpkg -i doca-host-repo-ubuntu<version>_amd64.deb

Answer: C

Explanation:

The nvidia-ctk runtime configure command is a crucial step that modifies the Docker daemon configuration file (/etc/docker/daemon.json) to register the nvidia runtime. However, the Docker daemon only reads this configuration file during its initialization phase. Even though the toolkit is installed and the configuration file is updated, Docker will not be able to spawn GPU-accelerated containers until the service is refreshed.

Executing sudo systemctl restart docker (or the equivalent for your container engine) is the mandatory final step. This forces Docker to reload its settings and recognize the NVIDIA Container Runtime as a valid option.

Without this restart, attempting to run a container with the --gpus all flag will result in an error stating that the "nvidia" runtime is not found or is unconfigured. This is a common point of failure in automated AI infrastructure deployments where the configuration script finishes, but the service state remains stale.

NEW QUESTION # 34

During BCM cluster setup, an engineer must configure bonded network interfaces on DGX nodes for high availability. Which cmsh command sequence properly configures a bond0 interface with two physical NICs?

- A. device use dgx001 ; interfaces add bond bond0 ; append interfaces enp225s0flnp1 enp97s0flnp1 ; set mode 1 ; set network internalnet
- B. device use dgx001 ; interfaces add vlan vlan100 ; set parent bond0 ; set mode 1 ; set network internalnet
- C. device use dgx001 ; interfaces set enp225s0flnp1 network internalnet ; interfaces set enp97s0flnp1 network internalnet
- D. device use dgx001 ; interfaces delete enp225s0flnp1 ; interfaces delete enp97s0flnp1

Answer: A

Explanation:

In NVIDIA Base Command Manager (BCM), the management and storage traffic often requires redundancy via NIC bonding (Link Aggregation). The cmsh utility uses a hierarchical command structure to modify node configurations. To create a bond, the administrator must first navigate to the specific node's interface configuration (device use <node>). The correct sequence involves adding a new interface of type "bond" (interfaces add bond bond0). After the bond object is created, the physical slave interfaces must be associated with it using the append interfaces command. Finally, the bonding mode (e.g., Mode 1 for Active-Backup or Mode 4 for LACP) and the logical network assignment must be defined. Option B correctly follows this logic.

Assigning networks directly to physical ports (Option C) would prevent the use of a unified bond IP, and Option A incorrectly attempts to add a VLAN before the underlying bond is established. This configuration is essential for ensuring the control plane remains reachable even if a single management cable or switch port fails.

NEW QUESTION # 35

You are tasked with validating a newly installed NVIDIA A100 Tensor Core GPU within a server. You need to confirm the GPU is correctly recognized and functioning at its expected performance level. Describe the process, including commands and tools, to verify the following aspects: 1) GPU presence and basic information, 2) PCIe bandwidth and link speed, and 3) Sustained computational performance under load.

- A. 1) Check BIOS settings for GPU detection. 2) Use 'lspci -vv' to check PCIe speed. 3) Run a PyTorch ImageNet training script.
- B. 1) Use 'nvidia-smi' for presence and basic info. 2) Use 'nvidia-smi -q -d pcie' for bandwidth/speed. 3) Run a CUDA-

based matrix multiplication benchmark (e.g., using cuBLAS) with increasing matrix sizes and monitor performance.

- C. 1) Use 'lspci | grep NVIDIA' for presence, 'nvidia-smi' for basic info. 2) Use 'nvidia-smi -q -d pcie' for bandwidth/speed. 3) Run a TensorFlow ResNet50 benchmark.
- D. 1) Use 'nvidia-smi' for presence and basic info. 2) PCIe speed is irrelevant. 3) Run the 'nvprof' profiler during a CUDA application.
- E. 1) Use 'nvidia-smi' for presence and basic info. 2) Use 'nvlk-monitor' for bandwidth/speed. 3) Run a CPU-bound benchmark to avoid GPU bottlenecks.

Answer: B

Explanation:

'nvidia-smi' is the primary tool for NVIDIA GPU information. 'nvidia-smi -q -d pcie' provides PCIe details. A CUDA-based benchmark isolates GPU performance. Other options have elements of truth but aren't complete or optimally targeted (e.g., ResNet50 relies on other frameworks). Using a CPU-bound benchmark wouldn't test the GPU's capabilities.

NEW QUESTION # 36

You are tasked with setting up a secure environment for running GPU-accelerated machine learning workloads in Docker containers.

The security requirements dictate that containers should have minimal privileges and access only the necessary resources. Which of the following security measures are most relevant when using NVIDIA GPUs with Docker?

- A. Regularly scan Docker images for vulnerabilities using tools like Clair or Trivy and rebuild images with patched dependencies.
- B. Run the Docker daemon in rootless mode to reduce the risk of privilege escalation.
- C. Use AppArmor or SELinux profiles to restrict the capabilities of the Docker containers, limiting their access to system resources.
- D. Implement network segmentation and firewalls to isolate the Docker containers from other services and the internet.
- E. Grant the Docker containers direct access to the host's hardware devices, including the GPU, to maximize performance.

Answer: A,B,C,D

Explanation:

Security is paramount, and minimizing privileges is key. Running Docker in rootless mode (A) reduces the attack surface. AppArmor/SELinux (B) confines container capabilities. Regular vulnerability scanning (C) helps prevent attacks based on known weaknesses. Network segmentation (E) limits the impact of a compromised container. Granting direct hardware access (D) increases the risk of privilege escalation and should be avoided in a secure environment. The NVIDIA Container Toolkit facilitates GPU access without requiring direct device passthrough, adhering to principle of least privilege.

NEW QUESTION # 37

You're deploying BlueField OS to multiple SmartNICs with varying hardware revisions. How can you ensure that the correct device tree is loaded for each specific SmartNIC during the boot process?

- A. Manually specify the device tree file in the bootloader configuration for each SmartNIC
- B. Use a bootloader (e.g., U-Boot) that can detect the hardware revision and load the appropriate device tree based on a predefined mapping.
- C. Embed the device tree within the kernel image. The kernel will automatically select the correct one during boot.
- D. Create a single device tree that is compatible with all SmartNIC revisions. The kernel will automatically handle compatibility.
- E. Relocate the device tree after the OS is running, but before services are started using a custom script.

Answer: A

Explanation:

A bootloader like U-Boot is designed to handle hardware detection and conditional loading of resources like device trees. It can identify the SmartNIC revision and load the corresponding DTB file. Creating a single compatible DTB is difficult and may not fully utilize hardware capabilities. Manually specifying the DTB for each NIC is not scalable. Embedding the DTB in the kernel is uncommon. Attempting to modify the device tree at runtime could lead to instability.

