

NCP-AIIトレーニングサンプル、NCP-AII基礎訓練



BONUS!!! JPTestKing NCP-AIIダンプの一部を無料でダウンロード: https://drive.google.com/open?id=1kPFcLYadZs2CmxE_m1ZhyAUQaWfmbcjS

天帝様は公平ですから、人間としての一人一人は完璧ではないです。私のように、以前が努力しなかったの
で、今は無駄に悩んでいます。現在のIT領域で競争が激しくなっていることは皆は良く知っていますから、み
んなはIT認証を通じて自分の価値を高めたいです。私もそう思いますが、IT認証は私にとって大変難しいで
す。でも、幸い私はインターネットでJPTestKingのNVIDIAのNCP-AII試験トレーニング資料を見つけました。そ
れを手に入れてから私は試験に合格する自信を持つようになります。JPTestKingのNVIDIAのNCP-AII試験トレ
ーニング資料のカバー率がとても高いですから、自分で勉強するよりずっと効率が高いです。あなたもIT業種の
一人としたら、ためらわずにJPTestKingのNVIDIAのNCP-AII試験トレーニング資料をショッピングカートに入れ
ましょう。JPTestKingはきっとあなたが成功への良いアシスタントになります。

短い時間に最も小さな努力で一番効果的にNVIDIAのNCP-AII試験の準備をしたいのなら、JPTestKingのNVIDIA
のNCP-AII試験トレーニング資料を利用することができます。JPTestKingのトレーニング資料は実践の検証に合
格したもので、多くの受験生に証明された100パーセントの成功率を持っている資料です。JPTestKingを利用
したら、あなたは自分の目標を達成することができ、最良の結果を得ます。

>> NCP-AIIトレーニングサンプル <<

NVIDIA NCP-AII Exam | NCP-AIIトレーニングサンプル - 返金保証 NCP-AII基礎訓練

NVIDIA NCP-AII認証試験に合格することが簡単ではなくて、NVIDIA NCP-AII証明書は君にとってはIT業界に入
るの一つの手づるになるかもしれません。しかし必ずしも大量の時間とエネルギーで復習しなくて、弊社が丹
精にできあがった問題集を使って、試験なんて問題ではありません。

NVIDIA AI Infrastructure 認定 NCP-AII 試験問題 (Q126-Q131):

質問 # 126

A BlueField-2 DPU is being used to accelerate storage access via NVMe-oF. You notice high CPU utilization on the host server, despite the DPU. What is the MOST likely cause and the MOST effective initial troubleshooting step?

- A. Firmware mismatch between the host NIC driver and the BlueField DPU firmware; Upgrade both to the latest compatible versions.
- **B. The host is still handling a significant portion of the I/O processing due to incorrect offload settings; Check the NVMe-oF initiator settings and verify that offload is enabled.**
- C. Insufficient memory allocated to the BlueField DPIJ; Increase memory allocation via the BMC.
- D. Incorrect RoCE configuration; Verify RoCE version and MTU settings using 'roce_util'.
- E. Incorrect NVMe-oF target configuration on the storage array; Check target NQN and discovery settings.

正解: B

解説:

The MOST likely cause of high host CPU utilization with a DPIJ accelerating NVMe-oF is that the host is still handling significant I/O processing. The initial step is to verify NVMe-oF initiator offload settings to ensure the DPIJ is actually handling the workload. Checking target configuration or RoCE, while important, wouldn't directly address high host CPU if the offload isn't working in the first place.

質問 # 127

You are deploying a multi-node NVIDIA GPU cluster for distributed deep learning. Each node has a different ambient operating temperature due to varying airflow patterns within the data center. To ensure optimal performance and longevity of the GPUs across all nodes, which approach is MOST effective for managing GPU power limits?

- A. Disable power capping altogether to allow GPUs to operate at their maximum potential performance.
- B. Manually adjust the fan speeds of each GPU to ensure they are all running at maximum RPM.
- **C. Implement dynamic power management using NVIDIA's Data Center GPU Manager (DCGM) to adjust power limits on a per-GPU basis, taking into account real-time temperature readings and workload characteristics.**
- D. Rely on the default power management settings provided by the GPU driver.
- E. Set a uniform power limit for all GPUs across the entire cluster based on the GPU's Thermal Design Power (TDP) specification.

正解: C

解説:

Option C, using DCGM for dynamic power management, is the most effective approach. It allows for per-GPU power limit adjustments based on real-time conditions, optimizing performance while ensuring thermal safety and longevity across nodes with different operating temperatures. A uniform power limit (A) might be too restrictive for some nodes or insufficient for others. Disabling power capping (B) risks overheating and damage. Default settings (D) may not be optimal. Manually adjusting fan speeds (E) can help, but doesn't address power limits directly.

質問 # 128

Consider a scenario where you are setting up a high-performance computing cluster with several GPU-accelerated nodes using Slurm as the resource manager. You want to ensure that jobs requesting GPUs are only scheduled on nodes with the appropriate NVIDIA drivers and CUDA toolkit installed. How can you achieve this within Slurm?

- **A. Use Slurm's node features to tag nodes with the 'Feature=' keyword in 'slurm.conf'. For example, tag nodes with GPUs as 'Feature=gpu'. Jobs can then request nodes with the 'gpu' feature using the option.**
- B. Use Slurm's 'GresTypeS' configuration option in 'slurm.conf' to define a generic resource type called 'gpu' and then configure each node to advertise the available GPUs. Slurm will automatically ensure that jobs requesting GPUs are only scheduled on nodes with the 'gpu' resource.
- C. Utilize Slurm's Prolog and Epilog scripts to dynamically install the necessary NVIDIA drivers and CUDA toolkit on each node before and after a job runs. This ensures that the required software is always available.
- D. Install the NVIDIA Data Center GPU Manager (DCGM) on each node and configure Slurm to query DCGM for GPU availability and health. Slurm will then only schedule jobs on healthy and available GPUs.
- E. Create a custom Slurm script that checks for the presence of the NVIDIA driver and CUDA toolkit before submitting a

job to a node. If the requirements are not met, the job is rejected.

正解: A

解説:

Using Slurm's node features is the most straightforward and recommended approach for tagging nodes with specific capabilities. The '-constraint' option allows jobs to request nodes with particular features. GresTypeS can be used, but node features provide more flexibility and control. Installing drivers dynamically is impractical and inefficient. DCGM is primarily for monitoring, not core scheduling requirements.

質問 # 129

Which configuration file should be modified to blacklist the 'nouveau' driver on a system running systemd to prevent conflicts with the NVIDIA driver?

- A. /etc/modules
- B. /etc/modprobe.d/blacklist-nouveau.conf
- C. /etc/systemd/system.conf
- D. /boot/grub/grub.cfg
- E. /etc/default/grub

正解: B

解説:

The file '/etc/modprobe.d/blacklist-nouveau.conf' is the standard location to blacklist the 'nouveau' driver. This prevents it from loading and conflicting with the NVIDIA driver. '/etc/modules' is a list of modules to load at boot, but not for blacklisting. '/boot/grub/grub.cfg' is the GRUB configuration file, not for module blacklisting. '/etc/systemd/system.conf' is the systemd configuration file, not module related. '/etc/default/grub' is where you configure Grub settings and update GRUB, but it is not the right location for blacklisting modules.

質問 # 130

You are designing a network for a distributed training job utilizing multiple GPUs across multiple nodes. Which network characteristic is MOST critical for minimizing training time?

- A. Large MTU
- B. High packet loss rate
- C. Low latency
- D. Low cost
- E. High bandwidth

正解: A、C、E

解説:

Distributed training involves frequent data exchange between GPUs. High bandwidth is necessary to move large datasets quickly. Low latency reduces the round-trip time for communication. Large MTU (Maximum Transmission Unit) can reduce overhead by transmitting larger packets. High packet loss is detrimental to performance. Low cost should not be prioritized over performance in a distributed training environment.

質問 # 131

.....

あなたは君の初めてのNVIDIAのNCP-AII認定試験を受ける時に認定試験に合格したいか。JPTestKingでは、私たちは君のすべての夢を叶えさせて、君の最も早い時間でNVIDIAのNCP-AII認定試験に合格するということを保証します。JPTestKingのNVIDIAのNCP-AII試験トレーニング資料は豊富な経験を持っているIT専門家が研究したもので、問題と解答が緊密に結んでいるものです。JPTestKingを選ぶなら、絶対に後悔させません。

NCP-AII基礎訓練: <https://www.jpctestking.com/NCP-AII-exam.html>

お客様に最高のサービスをお楽しみいただくために、当社のNCP-AII試験準備はすべて、何百人もの経験豊富な

蓮が、いつか普通の恋愛をして、俺から離れて行くんじゃないかって、明らかに上等な部類の男、お客様に最高のサービスをお楽しみいただくために、当社のNCP-AII試験準備はすべて、何百人もの経験豊富な専門家によって設計されました。

購入した後、我々はあなたがNCP-AII試験にうまく合格するまで細心のヘルプをずっと与えます、したがって、レビュープランを調整するために、NCP-AIIの各練習問題を要約することが不可欠です、すべてのコンテンツは、NCP-AII試験の規制に準拠しています。

[illegible]

BONUS !!! JPTesKing NCP-AIIダンプの一部を無料でダウンロード: <https://drive.google.com/open?id=1kPFcLYadZs2CmxEm1ZhyAUQaWfnbcjS>