


# Examinations NCP-AIO Actual Questions & Regualer NCP-AIO Update



What's more, part of that It-Tests NCP-AIO dumps now are free: <https://drive.google.com/open?id=1qZ76A13aShh3Q4fEpzvnFQHSiO8sf0VC>

Please believe that our It-Tests team have the same will that we are eager to help you pass NCP-AIO exam. Maybe you are still worrying about how to prepare for the exam, but now we will help you gain confidence. By constantly improving our dumps, our strong technical team can finally take proud to tell you that our NCP-AIO exam materials will give you unexpected surprises. You can download our free demo to try, and see which version of NCP-AIO Exam Materials are most suitable for you; then you can enjoy your improvement in IT skills that our products bring to you; and the sense of achievement from passing the NCP-AIO certification exam.

## NVIDIA NCP-AIO Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"><li>Workload Management: This section of the exam measures the skills of AI infrastructure engineers and focuses on managing workloads effectively in AI environments. It evaluates the ability to administer Kubernetes clusters, maintain workload efficiency, and apply system management tools to troubleshoot operational issues. Emphasis is placed on ensuring that workloads run smoothly across different environments in alignment with NVIDIA technologies.</li></ul>
Topic 2	<ul style="list-style-type: none"><li>Installation and Deployment: This section of the exam measures the skills of system administrators and addresses core practices for installing and deploying infrastructure. Candidates are tested on installing and configuring Base Command Manager, initializing Kubernetes on NVIDIA hosts, and deploying containers from NVIDIA NGC as well as cloud VMI containers. The section also covers understanding storage requirements in AI data centers and deploying DOCA services on DPU Arm processors, ensuring robust setup of AI-driven environments.</li></ul>
Topic 3	<ul style="list-style-type: none"><li>Troubleshooting and Optimization: This section of the exam measures the skills of AI infrastructure engineers and focuses on diagnosing and resolving technical issues that arise in advanced AI systems. Topics include troubleshooting Docker, the Fabric Manager service for NVIDIA NVlink and NVSwitch systems, Base Command Manager, and Magnum IO components. Candidates must also demonstrate the ability to identify and solve storage performance issues, ensuring optimized performance across AI workloads.</li></ul>

Topic 4	<ul style="list-style-type: none"> <li>Administration: This section of the exam measures the skills of system administrators and covers essential tasks in managing AI workloads within data centers. Candidates are expected to understand fleet command, Slurm cluster management, and overall data center architecture specific to AI environments. It also includes knowledge of Base Command Manager (BCM), cluster provisioning, Run.ai administration, and configuration of Multi-Instance GPU (MIG) for both AI and high-performance computing applications.</li> </ul>
---------	---

>> Examinations NCP-AIO Actual Questions <<

## Regualer NCP-AIO Update | NCP-AIO Exams Dumps

The product is made in three different formats to help customers with different preparation styles meet their needs. One of these formats is NVIDIA NCP-AIO Dumps PDF file which is printable and portable. Users can take NVIDIA NCP-AIO PDF Questions anywhere and use them anytime. They can print these real NCP-AIO questions to save them as paper notes.

### NVIDIA AI Operations Sample Questions (Q34-Q39):

#### NEW QUESTION # 34

You are managing an on-premises cluster using NVIDIA Base Command Manager (BCM) and need to extend your computational resources into AWS when your local infrastructure reaches peak capacity.

What is the most effective way to configure cloudbursting in this scenario?

- A. Use BCM's built-in load balancer to distribute workloads evenly between on-premises and cloud resources without any pre-configuration.
- B. Set up a standby deployment in AWS and manually switch workloads to the cloud during peak times.
- **C. Use BCM's Cluster Extension feature to automatically provision AWS resources when local resources are exhausted.**
- D. Manually provision additional cloud nodes in AWS when the on-premises cluster reaches its limit.

**Answer: C**

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

NVIDIA Base Command Manager (BCM) provides a Cluster Extension feature that enables automatic provisioning and scaling of cloud resources (e.g., AWS) when on-premises capacity is fully utilized. This cloudbursting capability allows seamless extension of computational resources without manual intervention, improving flexibility and reducing downtime during peak demand. Options A, B, and C involve manual or incomplete automation approaches that do not leverage BCM's integrated cluster extension functionality.

#### NEW QUESTION # 35

When using GPUDirect RDMA for inter-GPU communication, what component MUST be supported by the network interface card (NIC) to ensure optimal performance?

- A. Jumbo Frames
- **B. Remote Direct Memory Access (RDMA)**
- C. TCP Offload Engine (TOE)
- D. Ethernet Flow Control
- E. Quality of Service (QOS)

**Answer: B**

Explanation:

GPUDirect RDMA requires RDMA support on the NIC. RDMA enables direct memory access between GPUs without CPU intervention, significantly reducing latency and improving bandwidth. While other features like TOE, QOS, flow control, and Jumbo Frames can contribute to overall network performance, they are not fundamental requirements for GPUDirect RDMA to function.

#### NEW QUESTION # 36

A system administrator needs to optimize the delivery of their AI applications to the edge.

What NVIDIA platform should be used?

- A. NetQ
- B. Base Command Platform
- C. Base Command Manager
- D. **Fleet Command**

**Answer: D**

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

NVIDIAFleet Command is the platform designed specifically to optimize and manage the deployment and delivery of AI applications at the edge. It enables secure and scalable orchestration of AI workloads across distributed edge devices, providing lifecycle management, remote monitoring, and updates. Fleet Command facilitates running AI applications closer to where data is generated (edge), improving latency and operational efficiency.

\* Base Command Platform and Base Command Manager primarily target data center and AI cluster management for configuration, monitoring, and troubleshooting.

\* NetQ is focused on network telemetry and network state monitoring rather than application delivery.

Therefore, for AI application delivery and optimization at the edge, Fleet Command is the recommended NVIDIA platform.

### NEW QUESTION # 37

You are tasked with configuring MIG on an NVIDIA A100 GPU for a mixed AI/HPC workload. You need to create two instances: one for a deep learning training job (requiring high memory bandwidth) and another for a molecular dynamics simulation (requiring high compute throughput). Which is the MOST optimal MIG configuration to create based on these workload requirements?

- A. One instance of 4g.20gb for deep learning and one instance of 3g.20gb for molecular dynamics.
- B. Two instances of 7g.40gb. This provides maximum performance for both workloads.
- C. One instance of 2g.10gb for deep learning and one instance of 1g.5gb for molecular dynamics.
- D. Two instances of 1g.5gb. This ensures balanced resource allocation.
- E. **One instance of 3g.20gb for deep learning and one instance of 4g.20gb for molecular dynamics simulation. This configuration dedicates larger memory and compute resources to each task based on the workload.**

**Answer: E**

Explanation:

Deep learning training typically benefits from larger memory capacities and bandwidth. While molecular dynamics often leverages compute throughput. Therefore, allocating 3g.20gb for deep learning, with focus on memory, and 4g.20gb for molecular dynamics will better utilize computational resources based on the workload characteristics. The 1g.2g options are too small, and 7g option might overcommit resources that other processes or users could need on the same node.

### NEW QUESTION # 38

Which of the following Magnum IO components would be MOST beneficial for accelerating data loading in a deep learning training pipeline that reads data directly from NVMe drives?

- A. GPUDirect RDMA
- B. NVSHMEM
- C. InfiniBand
- D. **GPUDirect Storage**
- E. CUDA-Aware MPI

**Answer: D**

Explanation:

GPUDirect Storage is specifically designed to allow direct memory access between NVMe drives and GPU memory, bypassing the CPU. This dramatically accelerates data loading and reduces CPU utilization. NVSHMEM is for inter-GPU shared memory.

GPUDirect RDMA is for network communication. CUDA-Aware MPI is for distributed processing. InfiniBand is a network technology but GPUDirect Storage utilizes it most efficiently in this data loading scenario.

## NEW QUESTION # 39

Everyone wishes to spend their career at one level. Obtaining a NVIDIA AI Operations NCP-AIO certificate is the reason that many people join the NVIDIA NCP-AIO exam. They can be sure of earning promotions and higher pay at their current job with this credential. While attempting career growth is crucial, you can only do so after clearing the NVIDIA AI Operations NCP-AIO Exam.

**Regualer NCP-AIO Update:** <https://www.it-tests.com/NCP-AIO.html>

DOWNLOAD the newest It-Tests NCP-AIO PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1qZ76A13aShh3Q4fEpzvnFQHSiO8sf0VC>