



# **NVIDIA**

## **Exam Questions NCA-AIIO**

NVIDIA-Certified Associate AI Infrastructure and Operations

## About ExamBible

### *Your Partner of IT Exam*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

#### NEW QUESTION 1

Which of the following aspects have led to an increase in the adoption of AI? (Choose two.)

- A. Moore's Law
- B. Rule-based machine learning
- C. High-powered GPUs
- D. Large amounts of data

**Answer:** CD

#### NEW QUESTION 2

Which type of GPU core was specifically designed to realistically simulate the lighting of a scene?

- A. Tensor Cores
- B. CUDA Cores
- C. Ray Tracing Cores

**Answer:** C

#### NEW QUESTION 3

In training and inference architecture requirements, what is the main difference between training and inference?

- A. Training requires real-time processing, while inference requires large amounts of data.
- B. Training requires large amounts of data, while inference requires real-time processing.
- C. Training and inference both require large amounts of data.
- D. Training and inference both require real-time processing.

**Answer:** B

#### NEW QUESTION 4

Which of the following NVIDIA tools is primarily used for monitoring and managing AI infrastructure in the enterprise?

- A. NVIDIA NeMo System Manager
- B. NVIDIA Data Center GPU Manager
- C. NVIDIA DGX Manager
- D. NVIDIA Base Command Manager

**Answer:** D

#### NEW QUESTION 5

An IT professional is considering whether to implement an on-prem or cloud infrastructure. Which of the following is a key advantage of on-prem infrastructure?

- A. Lower upfront costs and capital expenditure.
- B. Scalability and flexibility.
- C. Ensure data security and sovereignty.
- D. Easy remote management.

**Answer:** C

#### NEW QUESTION 6

What NVIDIA tool should a data center administrator use to monitor NVIDIA GPUs?

- A. NVIDIA System Monitor
- B. NetQ
- C. DCGM

**Answer:** C

#### NEW QUESTION 7

Which solution should be recommended to support real-time collaboration and rendering among a team?

- A. A cluster of servers with NVIDIA T4 GPUs in each server.
- B. A DGX SuperPOD.
- C. An NVIDIA Certified Server with RTX-based GPUs.

**Answer:** C

#### NEW QUESTION 8

Which NVIDIA software provides the capability to virtualize a GPU?

- A. Horizon
- B. vGPU
- C. virtGPU

**Answer:** B

**NEW QUESTION 9**

Which phase of deep learning benefits the greatest from a multi-node architecture?

- A. Data Augmentation
- B. Training
- C. Inference

**Answer:** B

**NEW QUESTION 10**

When deploying high-density workloads in a data center, what are the three main resource constraints that need to be considered?

- A. Processing speed, storage capacity, and network connectivity.
- B. Power, cooling, and physical space.
- C. Bandwidth, security, and redundancy.

**Answer:** B

**NEW QUESTION 10**

Which NVIDIA parallel computing platform and programming model allows developers to program in popular languages and express parallelism through extensions?

- A. CUDA
- B. CUMML
- C. CUGRAPH

**Answer:** A

**NEW QUESTION 12**

Which feature of RDMA reduces CPU utilization and lowers latency?

- A. Increased memory buffer size.
- B. Network adapters that include hardware offloading.
- C. NVIDIA Magnum I/O software.

**Answer:** B

**NEW QUESTION 14**

Which is the best PUE value for a data center?

- A. PUE of 1.2
- B. PUE of 3.5
- C. PUE of 5.0
- D. PUE of 2.0

**Answer:** A

**NEW QUESTION 16**

What is a common tool for container orchestration in AI clusters?

- A. Kubernetes
- B. MLOps
- C. Slurm
- D. Apptainer

**Answer:** A

**NEW QUESTION 18**

Which of the following statements is true about GPUs and CPUs?

- A. GPUs are optimized for parallel tasks, while CPUs are optimized for serial tasks.
- B. GPUs have very low bandwidth main memory while CPUs have very high bandwidth main memory.
- C. GPUs and CPUs have the same number of cores, but GPUs have higher clock speeds.
- D. GPUs and CPUs have identical architectures and can be used interchangeably.

**Answer:** A

**NEW QUESTION 20**

What is one key advantage that Cloud GPU Infrastructure has over On-Prem GPU infrastructure?

- A. Lower cost barrier to entry.
- B. Reduced cost of I/O traffic.

C. Greater flexibility for hardware orchestration.

**Answer:** A

#### NEW QUESTION 21

How is out-of-band management utilized by network operators in an AI environment?

- A. It is used to remotely manage and troubleshoot network devices independently of the production network.
- B. It is used to directly manage the AI model's learning rate during training sessions.
- C. It is used to increase the computational power of AI models by adapting additional processing resources.
- D. It is used to manage the data throughput of AI applications by prioritizing network traffic.

**Answer:** A

#### NEW QUESTION 23

When using an InfiniBand network for an AI infrastructure, which software component is necessary for the fabric to function?

- A. Verbs
- B. MPI
- C. OpenSM

**Answer:** C

#### NEW QUESTION 25

In a data center, what is the purpose and benefit of a DPU?

- A. A DPU is responsible for providing backup and disaster recovery solutions.
- B. A DPU is used for managing physical infrastructure, such as power and cooling.
- C. A DPU is responsible for managing network connections and security.
- D. A DPU is designed to offload, accelerate, and isolate infrastructure workloads.

**Answer:** D

#### NEW QUESTION 29

A customer is evaluating an AI cluster for training and is questioning why they should use a large number of nodes. Why would multi-node training be advantageous?

- A. The model is too large to fit into GPU memory.
- B. The model is being used by a large number of users.
- C. The model is being used for large-scale inference workloads.

**Answer:** A

#### NEW QUESTION 30

Which GPUs should be used when training a neural network for self-driving cars?

- A. NVIDIA H100 GPUs
- B. NVIDIA L4 GPUs
- C. NVIDIA DRIVE Orin

**Answer:** A

#### NEW QUESTION 33

Which NVIDIA tool aids data center monitoring and management?

- A. NVIDIA Mellanox Insight
- B. NVIDIA Clara
- C. NVIDIA TensorRT
- D. NVIDIA DCGM

**Answer:** D

#### NEW QUESTION 34

What is the maximum number of MIG instances that an H100 GPU provides?

- A. 7
- B. 8
- C. 4

**Answer:** A

#### NEW QUESTION 39

What is a key value of using NVIDIA NIMs?

- A. They provide fast and simple deployment of AI models.
- B. They have community support.
- C. They allow the deployment of NVIDIA SDKs.

**Answer:** A

**NEW QUESTION 42**

What is an advantage of InfiniBand over Ethernet?

- A. InfiniBand always provides higher bandwidth than Ethernet.
- B. InfiniBand supports RDMA while Ethernet does not.
- C. InfiniBand offers lower latency than Ethernet.

**Answer:** C

**NEW QUESTION 44**

.....

## Relate Links

**100% Pass Your NCA-AIIO Exam with Exambible Prep Materials**

<https://www.exambible.com/NCA-AIIO-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>