

## NCA-AIIO Dumps

# NVIDIA-Certified Associate AI Infrastructure and Operations

<https://www.certleader.com/NCA-AIIO-dumps.html>



**NEW QUESTION 1**

A company is implementing a new network architecture and needs to consider the requirements and considerations for training and inference. Which of the following statements is true about training and inference architecture?

- A. Training architecture and inference architecture have the same requirements and considerations.
- B. Training architecture is only concerned with hardware requirements, while inference architecture is only concerned with software requirements.
- C. Training architecture is focused on optimizing performance while inference architecture is focused on reducing latency.
- D. Training architecture and inference architecture cannot be the same.

**Answer: C**

**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**

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 5**

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 6**

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 7**

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 8**

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 9**

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 10**

In an AI cluster, what is the purpose of job scheduling?

- A. To gather and analyze cluster data on a regular schedule.
- B. To monitor and troubleshoot cluster performance.
- C. To assign workloads to available compute resources.
- D. To install, update, and configure cluster software.

**Answer:** C

**NEW QUESTION 10**

What is a significant benefit of using containers in an AI development environment?

- A. They increase the base accuracy of AI models by optimizing their algorithms.
- B. They ensure that AI applications run consistently across different computing environments.
- C. They can automatically generate AI datasets for machine learning model training.
- D. They directly increase the processing speed of GPUs used in AI computations.

**Answer:** B

**NEW QUESTION 14**

How many Mellanox ConnectX-6 Single Port VPI cards are in a DGX A100 system?

- A. 8
- B. 16
- C. 4

**Answer:** A

**NEW QUESTION 17**

In an AI cluster, what is the importance of using Slurm?

- A. Slurm is used for data storage and retrieval in an AI cluster.
- B. Slurm is responsible for AI model training and inference in an AI cluster.
- C. Slurm is used for interconnecting nodes in an AI cluster.
- D. Slurm helps with managing job scheduling and resource allocation in the cluster.

**Answer:** D

**NEW QUESTION 18**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your NCA-AIIO Exam with Our Prep Materials Via below:**

<https://www.certleader.com/NCA-AIIO-dumps.html>