



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

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

When training a neural network, what is the most common pattern of storage access?

- A. Random write
- B. Sequential read
- C. Sequential write

Answer: B

NEW QUESTION 3

Which of the following statements correctly differentiates between AI, Machine Learning, and Deep Learning?

- A. Machine Learning is a subset of AI, and AI is a subset of Deep Learning.
- B. AI and Deep Learning are the same, while Machine Learning is a separate concept.
- C. AI is a subset of Machine Learning, and Machine Learning is a subset of Deep Learning.
- D. Deep Learning is a subset of Machine Learning, and Machine Learning is a subset of AI.

Answer: D

NEW QUESTION 4

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 5

When monitoring a GPU-based workload, what is GPU utilization?

- A. The maximum amount of time a GPU will be used for a workload.
- B. The GPU memory in use compared to available GPU memory.
- C. The percentage of time the GPU is actively processing data.
- D. The number of GPU cores available to the workload.

Answer: C

NEW QUESTION 6

Which two components are included in GPU Operator? (Choose two.)

- A. Drivers
- B. PyTorch
- C. DCGM
- D. TensorFlow

Answer: AC

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 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 9

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 10

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 10

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 14

Which aspect of computing uses large amounts of data to train complex neural networks?

- A. Machine learning
- B. Deep learning
- C. Inferencing

Answer: B

NEW QUESTION 18

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 23

.....

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/>