



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

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 3

Which of the following statements is true about Kubernetes orchestration?

- A. It is bare-metal based but it supports containers.
- B. It has advanced scheduling capabilities to assign jobs to available resources.
- C. It has no inferencing capabilities.
- D. It does load balancing to distribute traffic across containers.

Answer: BD

NEW QUESTION 4

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 5

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 6

The foundation of the NVIDIA software stack is the DGX OS. Which of the following Linux distributions is DGX OS built upon?

- A. Ubuntu
- B. Red Hat
- C. CentOS

Answer: A

NEW QUESTION 7

Which architecture is the core concept behind large language models?

- A. BERT Large model
- B. State space model
- C. Transformer model
- D. Attention model

Answer: C

NEW QUESTION 8

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 9

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 10

What is a key benefit of using NVIDIA GPUDirect RDMA in an AI environment?

- A. It increases the power efficiency and thermal management of GPUs.
- B. It reduces the latency and bandwidth overhead of remote memory access between GPUs.
- C. It enables faster data transfers between GPUs and CPUs without involving the operating system.
- D. It allows multiple GPUs to share the same memory space without any synchronization.

Answer: C

NEW QUESTION 10

Which NVIDIA software provides the capability to virtualize a GPU?

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

Answer: B

NEW QUESTION 15

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 19

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 23

What is the name of NVIDIA's SDK that accelerates machine learning?

- A. Clara
- B. RAPIDS
- C. cuDNN

Answer: C

NEW QUESTION 24

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

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

Answer: A

NEW QUESTION 27

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 30

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 34

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 37

Which are three key features of InfiniBand networking technology?

- A. High reliability, high latency, and CPU offloads.
- B. High latency, high reliability, and high bandwidth.
- C. GPU offloads, low latency, high reliability.
- D. Low latency, high bandwidth, and CPU offloads.

Answer: D

NEW QUESTION 40

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 43

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 48

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

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

Answer: A

NEW QUESTION 52

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 55

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 58

.....

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/>