

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

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 4

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 5

How is the architecture different in a GPU versus a CPU?

- A. A GPU acts as a PCIe controller to maximize bandwidth.
- B. A GPU is architected to support massively parallel execution of simple instructions.
- C. A GPU is a single large and complex core to support massive compute operations.

Answer: B

NEW QUESTION 6

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

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

Answer: B

NEW QUESTION 7

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 8

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

- A. CUDA
- B. CUML

C. CUGRAPH

Answer: A

NEW QUESTION 9

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 10

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

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

Answer: C

NEW QUESTION 10

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

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

Answer: A

NEW QUESTION 12

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 13

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 16

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 20

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 23

How many 1 Gb Ethernet in-band network connections are in a DGX H100 system?

- A. 1

B. 2

Answer: C

NEW QUESTION 26

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 28

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 31

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

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

Answer: A

NEW QUESTION 33

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 34

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 35

.....

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>