



Vendor: NVIDIA

Exam Code: NCA-AIIO

Exam Name: AI Infrastructure and Operations

Version: DEMO

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

Explanation:

Training architectures are designed to maximize computational throughput and accelerate model convergence, often by leveraging distributed systems with multiple GPUs or specialized accelerators to process large datasets efficiently. This focus on performance ensures that models can be trained quickly and effectively. In contrast, inference architectures prioritize minimizing response latency to deliver real-time or near-real-time predictions, frequently employing techniques such as model optimization (e.g., pruning, quantization), batching strategies, and deployment on edge devices or optimized servers. These differing priorities mean that while there may be some overlap, the architectures are tailored to their specific goals--performance for training and low latency for inference.

QUESTION 2

For which workloads is NVIDIA Merlin typically used?

- A. Recommender systems
- B. Natural language processing
- C. Data analytics

Answer: A

Explanation:

NVIDIA Merlin is a specialized, end-to-end framework engineered for building and deploying large-scale recommender systems. It streamlines the entire pipeline, including data preprocessing (e.g., feature engineering, data transformation), model training (using GPU-accelerated frameworks), and inference optimizations tailored for recommendation tasks. Unlike general-purpose tools for natural language processing or data analytics, Merlin is optimized to handle the unique challenges of recommendation workloads, such as processing massive user-item interaction datasets and delivering personalized results efficiently.

QUESTION 3

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

- A. CUDA
- B. CUMM
- C. CUGRAPH

Answer: A

Explanation:

CUDA (Compute Unified Device Architecture) is NVIDIA's foundational parallel computing

platform and programming model. It enables developers to harness GPU parallelism by extending popular languages such as C, C++, and Fortran with parallelism-specific constructs (e.g., kernel launches, thread management). CUDA also provides bindings for languages like Python (via libraries like PyCUDA), making it versatile for a wide range of developers. In contrast, CUMML and CUGRAPH are higher-level libraries built on CUDA for specific machine learning and graph analytics tasks, not general-purpose programming models.

QUESTION 4

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

Explanation:

The surge in AI adoption is driven by two key enablers: high-powered GPUs and large amounts of data. High-powered GPUs provide the massive parallel compute capabilities necessary to train complex AI models, particularly deep neural networks, by processing numerous operations simultaneously, significantly reducing training times. Simultaneously, the availability of large datasets-spanning text, images, and other modalities-provides the raw material that modern AI algorithms, especially data-hungry deep learning models, require to learn patterns and make accurate predictions. While Moore's Law (the doubling of transistor counts) has historically aided computing, its impact has slowed, and rule-based machine learning has largely been supplanted by data-driven approaches.

QUESTION 5

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

Explanation:

The primary distinction between training and inference lies in their operational demands. Training necessitates large amounts of data to iteratively optimize model parameters, often involving extensive datasets processed in batches across multiple GPUs to achieve convergence. Inference, however, is designed for real-time or low-latency processing, where trained models are deployed to make predictions on new inputs with minimal delay, typically requiring less data volume but high responsiveness. This fundamental difference shapes their respective architectural designs and resource allocations.

QUESTION 6

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

Explanation:

GPUs and CPUs are architecturally distinct due to their optimization goals. GPUs feature thousands of simpler cores designed for massive parallelism, excelling at executing many lightweight threads concurrently--ideal for tasks like matrix operations in AI. CPUs, conversely, have fewer, more complex cores optimized for sequential processing and handling intricate control flows, making them suited for serial tasks. This divergence in design means GPUs outperform CPUs in parallel workloads, while CPUs excel in single-threaded performance, contradicting claims of identical architectures or interchangeable use.

QUESTION 7

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

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

Answer: AC

Explanation:

The NVIDIA GPU Operator is a tool for automating GPU resource management in Kubernetes environments. It includes two key components: GPU drivers, which provide the necessary software to interface with NVIDIA GPUs, and the NVIDIA Data Center GPU Manager (DCGM), which offers health monitoring, telemetry, and diagnostics for GPU clusters. Frameworks like PyTorch and TensorFlow are separate AI development tools, not part of the GPU Operator, which focuses on infrastructure rather than application layers.

QUESTION 8

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

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

Answer: B

Explanation:

Training is the deep learning phase that benefits most from a multi-node architecture. It involves compute-intensive operations-forward and backward passes, gradient computation, and synchronization-across large datasets and complex models. Distributing these tasks across multiple nodes with GPUs accelerates processing, reduces time to convergence, and enables handling models too large for a single node. While data augmentation and inference can leverage multiple nodes, their gains are less pronounced, as they typically involve lighter or more localized computation.

QUESTION 9

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

Explanation:

The Transformer model is the foundational architecture for modern large language models (LLMs). Introduced in the paper "Attention is All You Need," it uses stacked layers of self-attention mechanisms and feed-forward networks, often in encoder-decoder or decoder-only configurations, to efficiently capture long-range dependencies in text. While BERT (a specific Transformer-based model) and attention mechanisms (a component of Transformers) are related, the Transformer itself is the core concept. State space models are an alternative approach, not the primary basis for LLMs.

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

Explanation:

NVIDIA NIMs (NVIDIA Inference Microservices) are pre-built, GPU-accelerated microservices with standardized APIs, designed to simplify and accelerate AI model deployment across diverse environments--clouds, data centers, and edge devices. Their key value lies in enabling fast, turnkey inference without requiring custom deployment pipelines, reducing setup time and complexity. While community support and SDK deployment may be tangential benefits, they are not the primary focus of NIMs.

QUESTION 11

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

Explanation:

DGX OS, the operating system powering NVIDIA DGX systems, is built on Ubuntu Linux, specifically the Long-Term Support (LTS) version. It integrates Ubuntu's robust base with NVIDIA-specific enhancements, including GPU drivers, tools, and optimizations tailored for AI and high-performance computing workloads. Neither Red Hat nor CentOS serves as the foundation for DGX OS, making Ubuntu the correct choice.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14