

Exam : NCP-AIN

**Title : NVIDIA-Certified
Professional AI Networking**

Version : DEMO

1.You are designing a new AI data center for a research institution that requires high-performance computing for large-scale deep learning models. The institution wants to leverage NVIDIA's reference architectures for optimal performance.

Which NVIDIA reference architecture would be most suitable for this high-performance AI research environment?

- A. NVIDIA Base Command Platform
- B. NVIDIA DGX Cloud
- C. NVIDIA LaunchPad
- D. NVIDIA DGX SuperPOD

Answer: D

Explanation:

The NVIDIA DGX Super POD is a turnkey AI supercomputing infrastructure designed for large-scale deep learning and high-performance computing workloads. It integrates multiple DGX systems with high-speed networking and storage solutions, providing a scalable and efficient platform for AI research institutions. The architecture supports rapid deployment and is optimized for training complex models, making it the ideal choice for environments demanding top-tier AI performance.

Reference: DGX SuperPOD Architecture VIDIA Docs

2.What is the basic functionality of an IB Router?

- A. Connecting Ethernet switches to an InfiniBand fabric
- B. Connecting SAN to an InfiniBand fabric
- C. Connecting NVLink domains to an InfiniBand fabric
- D. Connecting two (or more) InfiniBand fabrics

Answer: D

Explanation:

AnInfiniBand (IB) Router connects two or more InfiniBand subnets, making it possible for nodes in different subnets to communicate through route-managed communication.

From the official NVIDIA InfiniBand Routers Documentation:

"An InfiniBand router provides connectivity between two or more InfiniBand subnets, enabling communication between hosts that are not on the same subnet while preserving isolation and scalability."

Ensures fabric scalability by allowing subnet segmentation.

Uses LID routing across subnet managers (SMs).

Essential in large clusters with thousands of nodes.

Incorrect Options:

AandBare incorrect: InfiniBand does not connect directly to Ethernet or SANs without a gateway.

Cis unrelated: NVLink is a GPU interconnect, not tied to InfiniBand routers.

Reference: NVIDIA InfiniBand Routers Guide

3.Which of the following scenarios would the Network Traffic Map in UFM be least useful for troubleshooting?

- A. When investigating reports of network congestion or latency problems.
- B. After making changes to network configuration.
- C. When troubleshooting a single node's hardware failure.

D. When optimizing job placement and workload distribution across the cluster.

Answer: C

Explanation:

The Network Traffic Map in NVIDIA's Unified Fabric Manager (UFM) provides a visual representation of the network topology and traffic flows, which is particularly useful for identifying congestion points, verifying network configurations, and optimizing workload distribution.

However, when troubleshooting a single node's hardware failure, the Network Traffic Map is less effective, as it focuses on network-level issues rather than individual hardware components.

4.You have recently implemented NVIDIA Spectrum-X in your data center to optimize AI workloads. You need to verify the performance improvements and create a baseline for future comparisons.

Which tool would be most appropriate for creating performance baseline results in this Spectrum-X environment?

- A. NetQ
- B. CloudAI Benchmark
- C. MLNX-OS
- D. Ansible

Answer: B

Explanation:

The CloudAI Benchmark is designed to evaluate and establish performance baselines in AI-optimized networking environments like NVIDIA Spectrum-X. It assesses various performance metrics, including throughput and latency, ensuring that the network meets the demands of AI workloads. This benchmarking is essential for validating the benefits of Spectrum-X and for ongoing performance monitoring.

Reference: NVIDIA Spectrum-X Validated Solution Stack

5.Which tool would you use to gather telemetry data in a SpectrumX network?

- A. NVIEW
- B. UFM
- C. NetQ
- D. BCM

Answer: C

Explanation:

The NVIDIA Spectrum-X networking platform is an Ethernet-based solution optimized for AI workloads, combining Spectrum-4 switches, BlueField-3 SuperNICs, and advanced software to deliver high performance and low latency. Gathering telemetry data is critical for optimizing Spectrum-X networks, as it provides visibility into network performance, congestion, and potential issues. The question asks for the tool used to collect telemetry data in a Spectrum-X network.

According to NVIDIA's official documentation, NVIDIA NetQ is the primary tool for gathering telemetry data in Ethernet-based networks, including those running on Spectrum-X platforms with Cumulus Linux or SONiC. NetQ is a network operations toolset that provides real-time monitoring, telemetry collection, and analytics for network health, enabling administrators to optimize performance, troubleshoot issues, and validate configurations. It collects detailed telemetry data such as link status, packet drops, latency, and congestion metrics, which are essential for Spectrum-X optimization.

Exact Extract from NVIDIA Documentation:

“NVIDIA NetQ is a highly scalable network operations tool that provides telemetry-based monitoring and analytics for Ethernet networks, including NVIDIA Spectrum-X platforms. NetQ collects real-time telemetry data from switches and hosts, offering insights into network performance, congestion, and connectivity. It supports Cumulus Linux and SONiC environments, making it ideal for optimizing Spectrum-X networks by providing visibility into key metrics like latency, throughput, and packet loss.”

—NVIDIA NetQ User Guide

This extract confirms that option C, NetQ, is the correct tool for gathering telemetry data in a Spectrum-X network. NetQ’s integration with Spectrum-X switches and its ability to collect and analyze telemetry data make it the go-to solution for network optimization tasks.