NVIDIA NCP-AIN Practice Questions

Shared by Lopez on 21-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

[InfiniBand Troubleshooting]

A user has requested confirmation that the InfiniBand network is performing optimally and is not limiting the speed of a training run. To verify this, you would like to measure the RDMA throughput rate between two endpoints.

Which tool should be used?

Options:

A- ibdiagnet
B- ib_write_bw
C- ping
D- iperf

Answer:

B

Explanation:

The ib_write_bw tool is part of the Perftest package and is specifically designed to measure the bandwidth of RDMA write operations between two InfiniBand endpoints. It provides accurate assessments of RDMA throughput, which is crucial for verifying the performance of InfiniBand networks in high-performance computing and AI training environments.

Question 2

Question Type: MultipleChoice

[InfiniBand Troubleshooting]

You are troubleshooting connectivity issues in your InfiniBand network and need to test basic connectivity between nodes. Which command should you use to test basic connectivity between InfiniBand nodes?

Options:

A- ibping
B- traceroute
C- ibnetdiscover
D- ping

Answer:

A

Explanation:

The tool specifically designed for testing InfiniBand connectivity is **ibping**. It functions similarly to the traditional ping utility but is optimized for InfiniBand fabrics.

From the NVIDIA InfiniBand Diagnostic Utilities Documentation:

'ibping tests the connectivity of InfiniBand nodes by sending management datagrams (MADs) and verifying the response from the destination LID or GUID.'

Tests basic node-to-node reachability

Supports testing via LID, GUID, or port number

Helps verify subnet manager routing and fabric health

Incorrect Options:

ping and traceroute are IP-based, not fabric-aware.

ibnetdiscover maps topology but doesn't test live connectivity.

Question 3

Question Type: MultipleChoice

[Spectrum-X Optimization]

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 best for creating performance baseline results in this Spectrum-X environment?

Options:

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.

Question 4

Question Type: MultipleChoice

[InfiniBand Optimization]

Which of the following NCCL environment variables enable SHARP aggregation with NCCL when using the NCCL-SHARP plugin?

Pick the 2 correct responses below

Options:

A- NCCL_COLLNET_ENABLE=1
B- NCCL_ALGO=CollNet
C- NCCLSPECTRUM_ENABLE=1
D- NCCL_SHARP_AUTOINIT

Answer:

A, D

Explanation:

To enable SHARP (Scalable Hierarchical Aggregation and Reduction Protocol) aggregation using the NCCL-SHARP plugin, the following two environment variables are required:

NCCL_COLLNET_ENABLE=1

Enables NCCL's support for CollNet (Collective Network) operations, including SHARP.

NCCL_SHARP_AUTOINIT=1

Automatically initializes the SHARP plugin when available, activating SHARP-based collectives.

From the NVIDIA NCCL User Guide -- SHARP Plugin Section:

'NCCL_COLLNET_ENABLE must be set to enable collective network acceleration features.'

'NCCL_SHARP_AUTOINIT enables automatic SHARP plugin integration at NCCL runtime.'

Incorrect Options:

B . NCCL_ALGO=CollNet -- This variable controls the algorithm used for collectives but does not enable SHARP.

C . NCCLSPECTRUM_ENABLE -- This is not a documented NCCL variable.

Question 5

Question Type: MultipleChoice

[Spectrum-X Optimization]

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

Options:

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.

Question 6

Question Type: MultipleChoice

[Spectrum-X Configuration]

When upgrading Cumulus Linux to a new version, which configuration files should be migrated from the old installation?

Pick the 2 correct responses below.

Options:

A- All files in /etc/cumulus/acl
B- All files in /etc/network
C- All files in /etc
D- All files in /etc/mix

Answer:

A, B

Explanation:

Before upgrading Cumulus Linux, it's essential to back up configuration files to a different server. The /etc directory is the primary location for all configuration data in Cumulus Linux. Specifically, the following files and directories should be backed up:

/etc/frr/ - Routing application (responsible for BGP and OSPF)

/etc/hostname - Configuration file for the hostname of the switch

/etc/network/ - Network configuration files, most notably /etc/network/interfaces and /etc/network/interfaces.d/

/etc/cumulus/acl - Access control list configurations

Cumulus Linux is a network operating system used on NVIDIA Spectrum switches, including those in the Spectrum-X platform, to provide a Linux-based environment for Ethernet networking in AI and HPC data centers. When upgrading Cumulus Linux to a new version, it's critical to migrate specific configuration files to preserve network settings and ensure continuity. The question asks for the two configuration file locations that should be migrated from the old installation during an upgrade.

According to NVIDIA's official Cumulus Linux documentation, the key directories containing configuration files that should be migrated during an upgrade are /etc/cumulus/acl (for access control list configurations) and /etc/network (for network interface configurations). These directories store critical network settings that define the switch's behavior, such as ACL rules and interface settings, which must be preserved to maintain network functionality after the upgrade.

Exact Extract from NVIDIA Documentation:

''When upgrading Cumulus Linux, you must back up and migrate specific configuration files to ensure continuity of network settings. The following directories should be included in the backup:

/etc/cumulus/acl: Contains access control list (ACL) configuration files that define packet filtering and security policies.

/etc/network: Contains network interface configuration files, such as interfaces and ifupdown2 settings, which define the network interfaces and their properties.

Back up these directories before upgrading and restore them after the new version is installed

to maintain consistent network behavior.''

--- NVIDIA Cumulus Linux Upgrade Guide

This extract confirms that options A and B are the correct answers, as /etc/cumulus/acl and /etc/network contain essential configuration files that must be migrated during a Cumulus Linux upgrade. These files ensure that ACL policies and network interface settings are preserved, which are critical for Spectrum-X configurations in AI networking environments.

Question 7

Question Type: MultipleChoice

[AI Network Architecture]

Which of the following statements are true about AI workloads and adaptive routing?

Pick the 2 correct responses below.

Options:

A- AI workloads are made of a small number of volumetric flows called elephant flows.
B- AI workloads have very high entropy that helps spread traffic evenly without congestion.
C- Flow-based load balancing mechanisms increase congestion risk.
D- ECMP-based load balancing works best for AI workloads.

Answer:

A, C

Explanation:

AI workloads, particularly in large-scale training scenarios, are characterized by a small number of high-bandwidth, long-lived flows known as 'elephant flows.' These flows can dominate network traffic and are prone to causing congestion if not managed effectively.

Traditional flow-based load balancing mechanisms, such as Equal-Cost Multipath (ECMP), distribute traffic based on flow hashes. However, in AI workloads with low entropy (i.e., limited variability in flow characteristics), ECMP can lead to uneven traffic distribution and congestion on certain paths.

Adaptive routing techniques, which dynamically adjust paths based on real-time network conditions, are more effective in managing AI traffic patterns and mitigating congestion risks.

Question 8

Question Type: MultipleChoice

[Spectrum-X Optimization]

You have implemented adaptive routing in your Spectrum-X network to optimize AI workload performance. You need to verify the effectiveness of this configuration and monitor its impact on network congestion. Which tool would be best for monitoring and analyzing the adaptive routing performance in your Spectrum-X environment?

Options:

A- NetQ
B- Ansible
C- CloudAI Benchmark
D- MLNXOS

Answer:

A

Explanation:

NVIDIA NetQ is a comprehensive network operations tool designed to provide real-time visibility into the health and performance of NVIDIA networking environments, including Spectrum-X. It offers detailed telemetry and analytics, allowing administrators to monitor adaptive routing behaviors, detect congestion, and analyze traffic patterns. By leveraging NetQ, you can ensure that adaptive routing is functioning as intended and that the network is optimized for AI workloads.

Reference Extracts from NVIDIA Documentation:

'The NVIDIA NetQ network validation and ASIC monitoring tool set provide visibility into the network health and behavior. The NetQ flow telemetry analysis shows the paths that data flows take as they traverse the network, providing network latency and performance insights.'

'By leveraging telemetry from Spectrum Ethernet switches and BlueField-3 SuperNICs, NVIDIA NetQ can detect network issues proactively and troubleshoot network issues faster for optimal use of network capacity.'

Question 9

Question Type: MultipleChoice

[InfiniBand Security]

You are concerned about potential security threats and unexpected downtime in your InfiniBand data center.

Which UFM platform uses analytics to detect security threats, operational issues, and predict network failures in InfiniBand data centers?

Options:

A- Host Agent
B- Enterprise Platform
C- Cyber-AI Platform
D- Telemetry Platform

Answer:

C

Explanation:

The NVIDIA UFM Cyber-AI Platform is specifically designed to enhance security and operational efficiency in InfiniBand data centers. It leverages AI-powered analytics to detect security threats, operational anomalies, and predict potential network failures. By analyzing real-time telemetry data, it identifies abnormal behaviors and performance degradation, enabling proactive maintenance and threat mitigation.

This platform integrates with existing UFM Enterprise and Telemetry services to provide a comprehensive view of the network's health and security posture. It utilizes machine learning algorithms to establish baselines for normal operations and detect deviations that may indicate security breaches or hardware issues.

Question 10

Question Type: MultipleChoice

[InfiniBand Troubleshooting]

A fabric administrator added new servers to a 40-port edge switch. The administrator now needs to gather and map the newly added ports' LIDs and LINK SPEED. Which of the following commands can be used for that purpose?

Options:

A- ib_check_routes
B- ibnetdiscover
C- ibhosts
D- ibswitches

Answer:

B

Explanation:

The correct utility is ibnetdiscover.

From the official NVIDIA InfiniBand Utilities Guide:

'ibnetdiscover scans the fabric and returns a topology of all switches and end nodes, including their GUIDs, LIDs, port numbers, and link speeds.'

It generates a fabric map with node-to-port relationships and shows:

GUIDs

LIDs (Local IDs)

Link speeds and widths

Switch-to-host connections

This is essential for network topology validation and mapping physical port additions.

Incorrect Options:

ib_check_routes -- for routing table diagnostics.

ibhosts -- shows host information but not switch-level port mapping.

ibswitches -- shows switch info, but lacks port-level LID/link speed mapping.

Question 11

Question Type: MultipleChoice

[Spectrum-X Configuration]

When creating a simu-lation in NVIDIA AIR, what syntax would you use to define a link between port 1 on spine-01 and port 41 on gpu-leaf-01?

Options:

A- 'spine-01':*swp01' - *gpu-leaf-01':'swp41'
B- 'spine-01':'swp1' to 'gpu-leaf-01':'swp41'
C- 'spine-01 'eth1' to 'gpu-leaf-01':'eth41'
D- 'spine-01':'eth1' - 'gpu-leaf-01':'eth41'

Answer:

A

Explanation:

NVIDIA AIR (AI-Ready Infrastructure) is a cloud-based simulation platform designed to model and validate data center network deployments, including Spectrum-X Ethernet networks, using realistic topologies and configurations. When creating a custom topology in NVIDIA AIR, users can define network links between devices (e.g., spine and leaf switches) using a DOT file format, which is based on the Graphviz graph visualization software. The question asks for the correct syntax to define a link between port 1 on a spine switch (spine-01) and port 41 on a leaf switch (gpu-leaf-01) in a NVIDIA AIR simulation.

According to NVIDIA's official NVIDIA AIR documentation, the DOT file format is used to specify network topologies, including nodes (devices) and links (connections between ports). The syntax for defining a link in a DOT file uses a double dash (--) to indicate a connection between two ports, with each port specified in the format '<node>':''. For Spectrum-X networks, which typically use Cumulus Linux or SONiC on NVIDIA Spectrum switches, ports are commonly labeled as swpX (switch port X) rather than ethX (Ethernet interface), especially for switch-to-switch connections in a leaf-spine topology. The correct syntax for the link between port 1 on spine-01 and port 41 on gpu-leaf-01 is:

'spine-01':'swp01' -- 'gpu-leaf-01':'swp41'

This syntax uses swp01 and swp41 to denote switch ports, consistent with Cumulus Linux conventions, and the double dash (--) to indicate the link, as required by the DOT file format.

Exact Extract from NVIDIA Documentation:

''You can create custom topologies in Air using a DOT file, which is the file type used with the

open-source graph visualization software, Graphviz. DOT files define nodes, attributes, and connections for generating a topology for a network. The following is an example of a link definition in a DOT file:

'leaf01':'swp31' -- 'spine01':'swp1'

This specifies a connection between port swp31 on leaf01 and port swp1 on spine01. Port names typically follow the switch port naming convention (e.g., swpX) for Cumulus Linux-based switches.''

--- NVIDIA Air Custom Topology Guide

This extract confirms that option A is the correct answer, as it uses the proper DOT file syntax with swp01 and swp41 for port names and the double dash (--) for the link, aligning with NVIDIA AIR's topology definition process for Spectrum-X simulations.

Analysis of Other Options:

B . 'spine-01':'swp1' to 'gpu-leaf-01':'swp41': This option uses the correct port naming convention (swp1 and swp41) but incorrectly uses the word to as the connector instead of the double dash (--). The DOT file format requires -- to define links, making this syntax invalid for NVIDIA AIR.

C . 'spine-01':'eth1' to 'gpu-leaf-01':'eth41': This option uses ethX port names, which are typically used for host interfaces (e.g., servers) rather than switch ports in Cumulus Linux or SONiC environments. Switch ports in Spectrum-X topologies are labeled swpX. Additionally, the use of to instead of -- is incorrect for DOT file syntax, making this option invalid.

D . 'spine-01':'eth1' - 'gpu-leaf-01':'eth41': This option uses a single dash (-) instead of the required double dash (--) and incorrectly uses ethX port names instead of swpX. The ethX naming is not standard for switch ports in Spectrum-X, and the single dash is not valid DOT file syntax, making this option incorrect.

Why 'spine-01':'swp01' -- 'gpu-leaf-01':'swp41' is the Correct Answer:

Option A correctly adheres to the DOT file syntax used in NVIDIA AIR for defining network links:

Node and Port Naming: The nodes spine-01 and gpu-leaf-01 are specified with their respective ports swp01 and swp41, following the swpX convention for switch ports in Cumulus Linux-based Spectrum-X switches.

Link Syntax: The double dash (--) is the standard connector in DOT files to indicate a link between two ports, as required by Graphviz and NVIDIA AIR.

Spectrum-X Context: In a Spectrum-X leaf-spine topology, connections between spine and leaf switches (e.g., Spectrum-4 switches) use switch ports labeled swpX, making swp01 and swp41 appropriate for this simulation.

This syntax ensures that the NVIDIA AIR simulation accurately models the physical connection between spine-01 port 1 and gpu-leaf-01 port 41, enabling validation of the Spectrum-X network topology. The DOT file can be uploaded to NVIDIA AIR to generate the topology, as described in the documentation.

Question 12

Question Type: MultipleChoice

[Spectrum-X Optimization]

Which service on Cumulus switches can monitor layer 1, layer 2, layer 3, tunnel, buffer, and ACL related issues?

Options:

A- WJH
B- ONIE
C- NCLU
D- BGP

Answer:

A

Explanation:

The 'What Just Happened' (WJH) service on Cumulus switches provides real-time visibility into network problems by monitoring various layers and components, including layer 1, layer 2, layer 3, tunnel, buffer, and Access Control List (ACL) related issues. WJH streams detailed and contextual telemetry data, enabling administrators to diagnose and troubleshoot network problems effectively.

Reference Extracts from NVIDIA Documentation:

'WJH can monitor layer 1, layer 2, layer 3, tunnel, buffer and ACL related issues.'

'The WJH service enables you to diagnose network problems by looking at dropped packets.'