

Pass Guaranteed Quiz Marvelous NVIDIA NCP-AIN - NVIDIA-Certified Professional AI Networking Latest Exam Format



P.S. Free & New NCP-AIN dumps are available on Google Drive shared by ActualCollection: https://drive.google.com/open?id=1y_BbRK9LoPu2icyHd8yADkPDhe_TIoVP

At the ActualCollection, we guarantee that our customers will receive the best possible NCP-AIN study material to pass the NVIDIA-Certified Professional AI Networking (NCP-AIN) certification exam with confidence. Joining this site for the NCP-AIN exam preparation would be the greatest solution to the problem of outdated material. The NCP-AIN would assist applicants in preparing for the NVIDIA NCP-AIN Exam successfully in one go NCP-AIN would provide NCP-AIN candidates with accurate and real NVIDIA-Certified Professional AI Networking (NCP-AIN) Dumps which are necessary to clear the NCP-AIN test quickly. Students will feel at ease since the content they are provided with is organized rather than dispersed.

The second format ActualCollection also has a product support team available every time to help you out in any terms. And they will fix all of your problems on time. provides its users to study for Prepare for your NVIDIA-Certified Professional AI Networking (NCP-AIN) exam is web-based practice exam. This format has all the features of desktop practice exam software for NVIDIA NCP-AIN exam preparation.

>> NCP-AIN Latest Exam Format <<

Excellent NCP-AIN Preparation Materials: NVIDIA-Certified Professional AI Networking donate you the best Exam Simulation - ActualCollection

Worrying over the issue of passing exam has put many exam candidates under great stress. Many people feel on the rebound when they aimlessly try to find the perfect practice material. Our team will relieve you of tremendous pressure with passing rate of the NVIDIA-Certified Professional AI Networking prepare torrents up to 98 percent to 100 percent. Even we have engaged in this area over ten years, professional experts never blunder in their handling of the NCP-AIN Exam torrents. By compiling our NVIDIA-Certified Professional AI Networking prepare torrents with meticulous attitude, the accuracy and proficiency of them is nearly perfect. As the leading elites in this area, our NVIDIA-Certified Professional AI Networking prepare torrents are in concord with syllabus of the exam. They are professional backup to this fraught exam.

NVIDIA NCP-AIN Exam Syllabus Topics:

Topic	Details

Topic 1	<ul style="list-style-type: none"> AI Network Architecture: This section of the exam measures the skills of AI Infrastructure Architects and covers the ability to distinguish between AI factory and AI data center architectures. It includes understanding how Ethernet and InfiniBand differ in performance and application, and identifying the right storage options based on speed, scalability, and cost to fit AI networking needs.
Topic 2	<ul style="list-style-type: none"> Spectrum-X Configuration, Optimization, Security, and Troubleshooting: This section of the exam measures the skills of Network Performance Engineers and covers configuring, managing, and securing NVIDIA Spectrum-X switches. It includes setting performance baselines, resolving performance issues, and using diagnostic tools such as CloudAI benchmark, NCCL, and NetQ. It also emphasizes leveraging DPUs for network acceleration and using monitoring tools like Grafana and SNMP for telemetry analysis.
Topic 3	<ul style="list-style-type: none"> InfiniBand Configuration, Optimization, Security, and Troubleshooting: This section of the exam measures the skills of Data Center Network Administrators and covers the configuration and operational maintenance of NVIDIA InfiniBand switches. It includes setting up InfiniBand fabrics for multi-tenant environments, managing subnet configurations, testing connectivity, and using UFM to troubleshoot and analyze issues. It also focuses on validating rail-optimized topologies for optimal network performance.

NVIDIA-Certified Professional AI Networking Sample Questions (Q63-Q68):

NEW QUESTION # 63

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

Pick the 2 correct responses below

- A. NCCL_SHARP_AUTOINIT
- B. NCCL_COLLECTIVE_ENABLE=1
- C. NCCL_ALGO=CollNet
- D. NCCLSPECTRUM_ENABLE=1

Answer: A,B

Explanation:

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

* NCCL_COLLECTIVE_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_COLLECTIVE_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.

Reference: NCCL SHARP Plugin Guide & NCCL User Guide - Environment Variables Section

NEW QUESTION # 64

When creating a simulation 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?

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

Answer: B

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>":"<port>". 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.

NEW QUESTION # 65

You are optimizing an InfiniBand network for AI workloads that require low-latency and high-throughput data transfers. Which feature of InfiniBand networks minimizes CPU overhead during data transfers?

- A. SHARP
- B. PKey
- C. **Direct Memory Access (DMA)**
- D. TCP/IP Offloading

Answer: C

Explanation:

Direct Memory Access (DMA) in InfiniBand networks allows data to be transferred directly between the memory of two devices without involving the CPU. This capability significantly reduces CPU overhead, lowers latency, and increases throughput, making it ideal for AI workloads that demand efficient data transfers.

NEW QUESTION # 66

In Cumulus Linux, which technology enables the ability to provide active-active redundancy to servers, without the need for direct inter-switch links?

- A. MLAG
- **B. EVPN Multi-homing**
- C. VSS

Answer: B

Explanation:

EVPN Multi-homing enables active-active redundancy without inter-switch links by using overlay routing over VXLAN and distributed control plane using BGP EVPN.

From the official NVIDIA Cumulus Linux EVPN Multihoming Documentation:

"EVPN multihoming allows multiple Top-of-Rack (ToR) switches to connect to the same server while maintaining full layer-2 redundancy without the need for inter-switch links or traditional MLAG configuration." Key benefits:

- * Simplified topology (no ISL/peer-link needed)
- * BGP-based control plane
- * Fast convergence
- * Active-active links per host NIC

Incorrect Options:

- * MLAG requires ISL between switches and peer-link configuration.
- * VSS (Virtual Switching System) is a Cisco term, not supported in NVIDIA networking.

Reference: Cumulus Linux Docs - EVPN Multihoming

NEW QUESTION # 67

When utilizing the `ib_write_bw` tool for performance testing, what does the `-S` flag define?

- A. The burst size
- B. The maximum rate of sent packages
- C. The number of QP's
- **D. Which service level to use**

Answer: D

Explanation:

From NVIDIA Performance Tuning Guide (`ib_write_bw` Tool Usage):

"`-S <SL>`: Specifies the Service Level (SL) to use for the InfiniBand traffic. SL is used for setting priority and mapping to virtual lanes (VLs) on the IB fabric." This flag is useful when testing QoS-aware setups or validating SL/VL mappings.

Incorrect Options:

- * A- No such flag for burst size.
- * B- `-q` defines number of QPs.
- * C- `--rate` or `-R` is used for rate-limiting.

Reference: NVIDIA InfiniBand Performance Guide - `ib_write_bw` Options Section

NEW QUESTION # 68

.....

One NVIDIA certification will help you get highly favor of large enterprises, it will bring you better opportunities. NCP-AIN valid exam dumps PDF will be a stepping-stone for you to success. The most important method for passing exams is targeted learning and preparing. Programmatic learning may make you know professional knowledge better. But it will not only cost a lot of your time and energy but also can't guarantee you pass. Our NCP-AIN Valid Exam Dumps PDF can help you pass exam for sure.

NCP-AIN Questions: <https://www.actualcollection.com/NCP-AIN-exam-questions.html>

- 2026 NCP-AIN Latest Exam Format: NVIDIA-Certified Professional AI Networking - The Best NVIDIA NCP-AIN Questions Download NCP-AIN for free by simply searching on { www.vceengine.com } Valid NCP-AIN Test Notes

What's more, part of that ActualCollection NCP-AIN dumps now are free: https://drive.google.com/open?id=1y_BbRK9LoPu2icyHd8yAdkPDhe_TloVP