

Desktop and Web-Based Practice Exams to Evaluate NCP-AIN Exam Preparation



What's more, part of that ExamCost NCP-AIN dumps now are free: https://drive.google.com/open?id=1R7yWfyn9TxYqE_M97Db1-bHB4GZw62IG

ExamCost have a professional IT team to do research for practice questions and answers of the NVIDIA NCP-AIN exam certification exam. They provide a very effective training tools and online services for your. If you want to buy ExamCost products, ExamCost will provide you with the latest, the best quality and very detailed training materials as well as a very accurate exam practice questions and answers to be fully prepared for you to participate in the NVIDIA Certification NCP-AIN Exam. Safely use the questions provided by ExamCost's products. Selecting the ExamCost is equal to be 100% passing the exam.

The users can instantly access the product after purchasing it from ExamCost NCP-AIN, so they don't have to wait to prepare for the NVIDIA NCP-AIN Exams. The 24/7 support system is available for the customers, so they can contact the support whenever they face any issue, and it will provide them with the solution. Furthermore, ExamCost offers up to 1 year of free updates and free demos of the product.

>> NCP-AIN Reliable Study Materials <<

Exact Inside NCP-AIN Reliable Study Materials Questions and Answers

In order to help these people who have bought the NCP-AIN study materials of our company, There is a team of expert in our company, which is responsible to renovate and update the NCP-AIN study materials provided by our company. We are going to promise that we will have a lasting and sustainable cooperation with customers who want to buy the NCP-AIN Study Materials from our company. If you decide to buy our NCP-AIN study materials, you will never miss any important information. In addition, we can promise the updating system is free for you.

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 (Q14-Q19):

NEW QUESTION # 14

You are using NVIDIA Air to simulate a Spectrum-X network for AI workloads. You want to ensure that your network configurations are optimal before deployment.

Which NVIDIA tool can be integrated with Air to validate network configurations in the digital twin environment?

- A. GPU Cloud
- **B. NetQ**
- C. Spectrum-X Manager
- D. DOCA

Answer: B

Explanation:

NVIDIA NetQ is a highly scalable network operations toolset that provides visibility, troubleshooting, and validation of networks in real-time. It delivers actionable insights and operational intelligence about the health of data center networks—from the container or host all the way to the switch and port—enabling a NetDevOps approach.

NetQ can be used as the functional test platform for the network CI/CD in conjunction with NVIDIA Air.

Customers benefit from testing the new configuration with NetQ in the NVIDIA Air environment ("digital twin") and fix errors before deploying to their production.

NEW QUESTION # 15

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

Answer: C

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 # 16

What is the purpose of WJH (What Just Happened)?

- A. Collate operating system logs and diagnose system crashes.
- B. Send notifications of failed login attempts to a pre-defined Slack channel.
- C. Identify potential cyberattacks or unusual traffic patterns across the cluster.
- **D. Provide contextual information regarding dropped packets in order to aid debugging.**

Answer: D

Explanation:

NVIDIA's What Just Happened (WJH) is a feature that provides real-time visibility into network problems by analyzing all packets passing through the switch and alerting on performance issues caused by packet drops, congestion, high latency, or misconfigurations.

WJH retains the last packets that were dropped from the switch with complete packet headers and the actual drop reason. This enhances the ability to debug network problems, identify affected flows, and decrease time- to-repair.

NEW QUESTION # 17

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?

- A. ibhosts
- B. ibswitches
- **C. ibnetdiscover**

- D. `ib_check_routes`

Answer: C

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.

Reference: NVIDIA InfiniBand Tools - `ibnetdiscover` Utility

NEW QUESTION # 18

You are tasked with troubleshooting a link flapping issue in an InfiniBand AI fabric. You would like to start troubleshooting from the physical layer.

What is the right NVIDIA tool to be used for this task?

- A. `tcpdump` tool
- **B. `mxlink` utility**
- C. `nvidia-smi` utility

Answer: B

Explanation:

The `mxlink` tool is used to check and debug link status and issues related to them. The tool can be used on different links and cables (passive, active, transceiver, and backplane). It is intended for advanced users with appropriate technical background.

Reference: `mxlink` Utility - NVIDIA Docs

NEW QUESTION # 19

.....

All these three NVIDIA NCP-AIN exam questions formats are easy to use and perfectly work with all devices, operating systems, and the latest web browsers. The NVIDIA-Certified Professional AI Networking (NCP-AIN) PDF dumps file is the collection of real and updated NVIDIA-Certified Professional AI Networking (NCP-AIN) exam questions that are being presented in PDF format. You can install NCP-AIN Pdf Dumps file on your desktop computer, laptop, tab, or even on your smartphone devices. Just install the NCP-AIN PDF dumps file and start NVIDIA-Certified Professional AI Networking (NCP-AIN) exam preparation anywhere and anytime.

NCP-AIN Test Questions Fee: <https://www.examcost.com/NCP-AIN-practice-exam.html>

- Marvelous NCP-AIN Reliable Study Materials, Ensure to pass the NCP-AIN Exam Search for 「 NCP-AIN 」 and obtain a free download on ⇒ www.prepawayete.com ⇐ NCP-AIN Dumps
- NCP-AIN Reliable Test Materials NCP-AIN Top Questions NCP-AIN Reliable Braindumps Sheet Open www.pdfvce.com enter ▷ NCP-AIN ◁ and obtain a free download NCP-AIN Study Materials Review
- 100% Pass Updated NCP-AIN - NVIDIA-Certified Professional AI Networking Reliable Study Materials Open 「 www.testkingpass.com 」 enter ➡ NCP-AIN and obtain a free download NCP-AIN Dumps
- 100% Pass Updated NCP-AIN - NVIDIA-Certified Professional AI Networking Reliable Study Materials Search for « NCP-AIN » and download it for free on (www.pdfvce.com) website NCP-AIN Test Testking
- Latest Real NCP-AIN Exam Customizable NCP-AIN Exam Mode NCP-AIN Dumps Copy URL 「 www.dumpsmaterials.com 」 open and search for ➡ NCP-AIN to download for free Reliable NCP-AIN Exam Online

