NCP-AIN Preparation, Reliable NCP-AIN Dumps Book

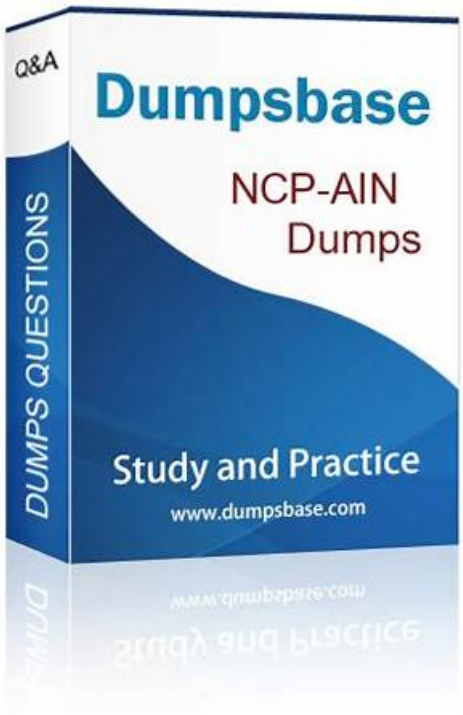

2026 Latest PracticeVCE NCP-AIN PDF Dumps and NCP-AIN Exam Engine Free Share: https://drive.google.com/open?id=1uWIp_d0yLKUN7cPowblOaNJmj_Y9s-kj

NVIDIA certification is one of the best golden-content certifications in IT expert field all over the world, and it is also the necessary condition of choosing talents standard in large enterprises. NCP-AIN exam questions answers is useful for candidates who are eager to go through the examination. There are thousands of companies recognized and valued the certification in the world. NCP-AIN Exam Questions Answers will make you pass exam easily.

There are many ways to help you pass NVIDIA certification NCP-AIN exam and selecting a good pathway is a good protection. PracticeVCE can provide you a good training tool and high-quality reference information for you to participate in the NVIDIA certification NCP-AIN exam. PracticeVCE's practice questions and answers are based on the research of NVIDIA certification NCP-AIN examination Outline. Therefore, the high quality and high authoritative information provided by PracticeVCE can definitely do our best to help you pass NVIDIA certification NCP-AIN exam. PracticeVCE will continue to update the information about NVIDIA certification NCP-AIN exam to meet your need.

>> NCP-AIN Preparation <<

2026 NCP-AIN Preparation : NVIDIA-Certified Professional AI Networking Realistic NCP-AIN 100% Pass

We understand our candidates have no time to waste, everyone wants an efficient learning. So we take this factor into consideration, develop the most efficient way for you to prepare for the NCP-AIN exam, that is the real questions and answers practice mode, firstly, it simulates the real NVIDIA-Certified Professional AI Networking test environment perfectly, which offers greatly help to our customers. Secondly, it includes printable PDF Format, also the instant access to download make sure you can study anywhere and anytime. All in all, high efficiency of NCP-AIN Exam Material is the reason for your selection.

NVIDIA NCP-AIN Exam Syllabus Topics:

Topic	Details
Topic 1	• 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 2	• 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.
Topic 3	• 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.

NVIDIA-Certified Professional AI Networking Sample Questions (Q22-Q27):

NEW QUESTION # 22
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?

- A. "spine-01":*swp01" - *gpu-leaf-01":"swp41"
- B. "spine-01":"swp1" to "gpu-leaf-01":"swp41"
- C. "spine-01":"eth1" - "gpu-leaf-01":"eth41"
- D. "spine-01 'eth1" to "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>":"<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 # 23
You are troubleshooting an InfiniBand network issue and need to check the status of the InfiniBand interfaces. Which command should you use to display the state, physical state, and link layer of InfiniBand interfaces?

- A. sudo ibnodes -C mlx5_0
- B. ibv_devices -c mlx5_0
- C. ibstat -d mlx5_X
- D. cat /proc/net/ib/device

Answer: C

Explanation:
The ibstat command is utilized to display the operational status of InfiniBand Host Channel Adapters (HCAs).
It provides detailed information, including the state (e.g., Active, Down), physical state (e.g., LinkUp, Polling), and link layer (e.g., InfiniBand, Ethernet) of each port on the HCA. This information is crucial for diagnosing connectivity issues and ensuring that the InfiniBand interfaces are functioning correctly.
Reference Extracts from NVIDIA Documentation:
* "The ibstat command displays the status of the host channel adapters (HCAs) in your InfiniBand fabric.
The status includes the HCAs' state, physical state, and link layer."
* "For proper operation, you are looking for 'State: Active' and 'Physical State: LinkUp'."

NEW QUESTION # 24
You are automating the deployment of a Spectrum-X network using Ansible. You need to ensure that the playbooks can handle different switch models and configurations efficiently.
Which feature of the NVIDIA NVUE Collection helps simplify the automation by providing pre-built roles for common network configurations?

- A. Collection plugins
- B. Collection libraries
- C. Collection modules
- D. Collection roles

Answer: D

Explanation:
The NVIDIA NVUE Collection for Ansible includes pre-built roles designed to streamline automation tasks across various switch models and configurations. These roles encapsulate common network configurations, allowing for efficient and consistent deployment.
By utilizing these roles, network administrators can:
* Apply standardized configurations across different devices.

* Reduce the complexity of playbooks by reusing modular components.
* Ensure consistency and compliance with organizational policies.
This approach aligns with Ansible best practices, promoting maintainability and scalability in network automation.
Reference: NVIDIA NVUE Collection Documentation - Ansible Roles

NEW QUESTION # 25
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. Spectrum-X Manager
- B. NetQ
- C. DOCA
- D. GPU Cloud

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 # 26
What are the two general user account types in MLNX-OS?
Pick the 2 correct responses below:

- A. monitor
- B. enable
- C. viewer
- D. admin

Answer: A,D

Explanation:
MLNX-OS, the operating system for NVIDIA's networking devices, defines two primary user account types: adminandmonitor. Theadminaccount has full administrative privileges, allowing for complete configuration and management of the system. Themonitoraccount, on the other hand, is designed for users who need to view system configurations and statuses without making any changes. This separation ensures a clear distinction between users who manage the system and those who monitor its operations.
Reference Extracts from NVIDIA Documentation:
* "There are two user roles or account types: admin and monitor. As 'admin', the user is privileged to run all the available commands. As 'monitor', the user can run commands that show system configuration and status, or set terminal settings." MLNX-OS is the network operating system used on NVIDIA's Mellanox Ethernet switches, including the Spectrum family (e.g., Spectrum-4 switches in the Spectrum-X platform), designed for high-performance Ethernet networking in AI and HPC data centers. MLNX-OS provides a command-line interface (CLI) for configuring and managing switch operations, with user accounts controlling access to various commands and functions. The question asks for the two general user account types in MLNX-OS, which define the primary privilege levels for user access.
According to NVIDIA's official MLNX-OS documentation, the two general user account types in MLNX-OS are:
* monitor: This account type has read-only access, allowing users to view configurations, status, and logs but not modify settings. It is used for monitoring and troubleshooting without risking unintended changes.
* admin: This account type has full read-write access, enabling users to view and modify all configurations, execute commands, and manage the switch's operations. It is intended for administrators with complete control over the system.
These two account types represent the primary privilege levels in MLNX-OS, providing a clear distinction between read-only monitoring and full administrative access.
Exact Extract from NVIDIA Documentation:
"MLNX-OS supports two primary user account types for managing switch operations:

* monitor: Users with monitor privileges have read-only access to the system. They can view configuration details, system status, and logs but cannot make changes to the configuration.
* admin: Users with admin privileges have full read-write access, allowing them to configure, manage, and troubleshoot all aspects of the switch, including executing privileged commands.These account types ensure secure and controlled access to the switch's management functions."-NVIDIA MLNX-OS User Manual This extract confirms that options B (monitor) and C (admin) are the correct answers. These account types are the standard privilege levels in MLNX-OS, used to manage access for monitoring and administrative tasks on Spectrum switches, including those in Spectrum-X deployments.

NEW QUESTION # 27
.....