

NCP-AIN Download - Vce NCP-AIN Test Simulator



BONUS!!! Download part of Prep4SureReview NCP-AIN dumps for free: <https://drive.google.com/open?id=1AN5laTnhbMLc4BdOus8B6QzgDiY4gsv3>

If you are still struggling to prepare for passing NVIDIA real exam at this moment, our Prep4SureReview NCP-AIN vce dumps can help you preparation easier and faster. Our website can provide you Valid NCP-AIN Exam Cram with high pass rate to help you get certification, and then you will become a good master of certification exam.

We are aimed to improve customer satisfaction and always put customers first. Our experts check daily whether there is an update to the NVIDIA-Certified Professional AI Networking torrent prep, and if there is an update system, we will automatically send it to you. So it can guarantee latest knowledge and keep up with the pace of change. Many people are worried that online shopping electronics have viruses. But you don't have to worry about our products. Our NCP-AIN Exam Questions are absolutely safe and virus-free. If you have any questions during the installation process, we will arrange professional staff on guidance of your installation and use. We always put your needs first.

[**>> NCP-AIN Download <<**](#)

Pass Guaranteed Quiz 2026 Useful NCP-AIN: NVIDIA-Certified Professional AI Networking Download

The NVIDIA NCP-AIN desktop practice exam software simulates a real test environment and familiarizes you with the actual test format. This NVIDIA NCP-AIN practice exam software tracks your progress and performance, allowing you to see how much you've improved over time. We frequently update the NVIDIA NCP-AIN Practice Exam software with the latest NVIDIA NCP-AIN DUMPS PDF.

NVIDIA NCP-AIN Exam Syllabus Topics:

Topic	Details
Topic 1	<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 2	<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.
Topic 3	<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.

NVIDIA-Certified Professional AI Networking Sample Questions (Q16-Q21):

NEW QUESTION # 16

What command sequence is used to identify the exact name of the server that runs as the master SM in a multi-node fabric?

- A. `sminfo`
`smpquery ND <lid>`
- B. `lstat`
`sminfo <lid>`
- C. `ibis`
`ibsim <lid>`
- D. `sminfo`
`smpquery NI <lid>`

Answer: A

Explanation:

To identify the active Subnet Manager (SM) node in an InfiniBand fabric, the correct command sequence is:

* `sminfo`
* Displays general information about the active SM in the fabric, including its LID.
* `smpquery ND <lid>`

* Resolves the Node Description (ND) at the given LID, revealing the exact hostname or label of the SM server.

From the InfiniBand Tools Guide:

"The `sminfo` utility provides the LID of the master SM. Use `smpquery ND <lid>` to resolve the node name hosting the SM." This two-step approach is standard for locating and validating the SM identity in fabric diagnostics.

Incorrect Options:

* B (NI) is an invalid query type.
* C and D do not identify SMs.

Reference: InfiniBand SM Tools - `sminfo` & `smpquery` Usage

NEW QUESTION # 17

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?

- A. `ping`
- B. `ib_write_bw`
- C. `ibdiagnet`
- 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.

Reference: `ib_write_bw` - NVIDIA Enterprise Support Portal

NEW QUESTION # 18

You are tasked with configuring multi-tenancy using partition key (PKey) for a high-performance storage fabric running on InfiniBand. Each tenant's GPU server is allowed to access the shared storage system but cannot communicate with another tenant's GPU server.

Which of the following partition key membership configurations would you implement to set up multi-tenancy in this environment?

- A. Assign limited membership to both GPU servers and storage system.
- B. Assign full membership to both GPU servers and storage system.
- **C. Assign full membership PKey to the shared storage system and limited membership PKey to each tenant's GPU servers.**
- D. Assign limited membership PKey to the shared storage system and full membership PKey to each tenant's GPU servers.

Answer: C

Explanation:

To enforce strict multi-tenancy, where:

- * Tenant A's GPU cannot talk to Tenant B's GPU
- * But both can access shared storage

The correct solution is:

- * Storage system # Full PKey membership
- * Each tenant's GPU # Limited PKey membership

From the NVIDIA InfiniBand P_Key Partitioning Guide:

"A port with limited membership can only communicate with full members of the same PKey. It cannot communicate with other limited members, even within the same partition." This isolates tenants from each other, while allowing shared access to storage.

Incorrect Options:

- * A permits tenant-to-tenant communication.
- * B isolates everything, including access to storage.
- * C prevents GPU access to storage.

Reference: NVIDIA InfiniBand - Multi-Tenant PKey Partitioning Design

NEW QUESTION # 19

When designing a multi-tenancy East/West (E/W) fabric using Unified Fabric Manager (UFM), which method should be used?

- **A. Partition / PKey**
- B. VLAN
- C. ROMA
- D. VXLAN

Answer: A

Explanation:

In InfiniBand networks, Partitioning using Partition Keys (PKeys) is the standard method for implementing multi-tenancy and traffic isolation. PKeys allow administrators to define logical partitions within the fabric, ensuring that traffic is confined to designated groups of nodes. This mechanism is essential for creating secure and isolated environments in multi-tenant architectures.

The Unified Fabric Manager (UFM) leverages PKeys to manage these partitions effectively, enabling administrators to assign and control access rights across different tenants. This approach ensures that each tenant's traffic remains isolated, maintaining both security and performance integrity within the shared fabric.

Reference: NVIDIA UFM Enterprise User Manual v6.15.6-4

NEW QUESTION # 20

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

Answer: D

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 # 21

.....

In this career advancement NVIDIA-Certified Professional AI Networking (NCP-AIN) certification journey you can get help from valid, updated, and real NCP-AIN Dumps questions which you can instantly download from Prep4SureReview. At this platform, you will get the top-rated and Real NCP-AIN Exam Questions that are ideal study material for quick NVIDIA NCP-AIN exam preparation.

Vce NCP-AIN Test Simulator: <https://www.prep4surereview.com/NCP-AIN-latest-braindumps.html>

- Fast-Download NCP-AIN Download - Pass NCP-AIN Once - First-Grade Vce NCP-AIN Test Simulator Open

DOWNLOAD the newest Prep4SureReview NCP-AIN PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1AN5laTnhbMLc4BdOus8B6QzgDiY4gsv3>