

NVIDIA NCP-AIN Reliable Cram Materials - Practice NCP-AIN Tests

KILLTEST
20% de descuento en todos los productos

Preguntas del examen
NCP-AIN
NVIDIA-Certified
Professional AI
Networking

The diagram shows a whiteboard with the text above. Below the whiteboard is a stylized house with a blue roof and a wooden frame, built into a shelf. The shelf is filled with books and a small potted plant.

BONUS!!! Download part of Prep4away NCP-AIN dumps for free: <https://drive.google.com/open?id=1HkarvavSbRcqYzKx8FxygiWOfn8fNYc>

We assure that you can not only purchase high-quality NCP-AIN prep guide but also gain great courage & trust from us. A lot of online education platform resources need to be provided by the user registration to use after purchase, but it is simple on our website. We provide free demo of NCP-AIN Guide Torrent, you can download any time without registering. Fast delivery—after payment you can receive our NCP-AIN exam torrent no more than 10 minutes, so that you can learn fast and efficiently. What are you waiting for? Just come and buy our NCP-AIN exam questions!

For the peace of your mind, you can also try a free demo of NVIDIA NCP-AIN Dumps practice material. You will not find such affordable and latest material for NVIDIA certification exam anywhere else. Don't miss these incredible offers. Order real NVIDIA NCP-AIN Exam Questions today and start preparation for the certification exam.

>> NVIDIA NCP-AIN Reliable Cram Materials <<

NVIDIA NCP-AIN PDF Questions – Ideal Material for Quick Preparation

First and foremost, you can get the latest version of our NCP-AIN study materials for free during the whole year. Second, our responsible after sale service staffs are available in twenty four hours a day, seven days a week, so if you have any problem after

purchasing NCP-AIN study materials, you can contact our after sale service staffs anywhere at any time. Finally, we have installed the most advanced operation machines in our website, so you can use credit for payment in the process of trading and register your personal information under a safe payment environment. Do not waver any more, the most effective and the Latest NCP-AIN Study Materials is right here waiting for you.

NVIDIA-Certified Professional AI Networking Sample Questions (Q66-Q71):

NEW QUESTION # 66

You are optimizing an AI workload that involves multiple GPUs across different nodes in a data center. The application requires both high-bandwidth GPU-to-GPU communication within nodes and efficient communication between nodes. Which combination of NVIDIA technologies would best support this multi-node, multi-GPU AI workload?

- A. PCIe for intra-node GPU communication and RoCE for inter-node communication.
- B. NVLink for both intra-node and inter-node GPU communication.
- C. InfiniBand for both intra-node and inter-node GPU communication.
- D. NVLink for intra-node GPU communication and InfiniBand for inter-node communication.

Answer: D

Explanation:

For optimal performance in multi-node, multi-GPU AI workloads:

- * NVLink provides high-speed, low-latency communication between GPUs within the same node.
- * InfiniBand offers efficient, scalable communication between nodes in a data center. Combining these technologies ensures both intra-node and inter-node communication needs are effectively met.

Reference:NVIDIA NVLink & NVSwitch: Fastest HPC Data Center Platform

NEW QUESTION # 67

You are deploying a Kubernetes cluster for AI workloads using NVIDIA Spectrum-X switches. You need to automate the deployment and management of networking components in this environment.

Which NVIDIA tool is specifically designed to automate the deployment and management of networking components in a Kubernetes cluster with Spectrum-X switches?

- A. Container Runtime
- B. GPU Operator
- C. Network Operator
- D. Mellanox OFED

Answer: C

Explanation:

The NVIDIA Network Operator is designed to simplify and automate the deployment and management of networking components in Kubernetes clusters, particularly those utilizing NVIDIA Spectrum-X switches. It manages the installation and configuration of necessary drivers, plugins, and other networking resources to enable features like RDMA and GPUDirect RDMA, which are essential for high-performance AI workloads.

By leveraging Kubernetes Custom Resource Definitions (CRDs) and the Operator Framework, the Network Operator ensures that networking components are consistently and correctly configured across the cluster, reducing manual intervention and potential configuration errors.

Reference:NVIDIA Network Operator Documentation

NEW QUESTION # 68

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

Answer: D

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 # 69

How does Spectrum-X achieve network isolation for multiple tenants?

- A. By assigning unique IP address ranges to each tenant.
- B. Using manual configuration of access control lists (ACLs).
- C. By implementing physical network segmentation.
- D. By implementing a Layer 3 Virtual Network Identifier (L3VNI) per VRR

Answer: D

Explanation:

Spectrum-X achieves network isolation in multi-tenant environments by implementing Layer 3 Virtual Network Identifiers (L3VNIs) per Virtual Routing and Forwarding (VRF) instance. This approach allows each tenant to have a separate routing table and network segment, ensuring that traffic is isolated and secure between tenants.

Reference Extracts from NVIDIA Documentation:

* "Spectrum-X enhances multi-tenancy with performance isolation to ensure tenants' AI workloads perform optimally and consistently."

NEW QUESTION # 70

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":*swp01" - *gpu-leaf-01":"swp41"
- C. "spine-01":"eth1" - "gpu-leaf-01":"eth41"
- D. "spine-01 'eth1" to "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 # 71

.....

Our NCP-AIN certification material is closely linked with the test and the popular trend among the industries and provides all the information about the NCP-AIN test. The answers and questions seize the vital points and are verified by the industry experts.

Diversified functions can help you get an all-around preparation for the test. Our online customer service replies the clients' questions about our NCP-AIN Certification material at any time. So our NCP-AIN learning file can be called perfect in all aspects.

Practice NCP-AIN Tests: <https://www.prep4away.com/NVIDIA-certification/braindumps.NCP-AIN.ete.file.html>

We provide 24/7 customer service for all of you, please feel free to send us any questions about NVIDIA NCP-AIN test pdf through email or online chat, and we will always try our best to keeping our customer satisfied, Our NCP-AIN learning guide is for you to improve your efficiency and complete the tasks with a higher quality, We do not hope that you spend all your time on learning the NCP-AIN certification materials.

Managing the life cycle of all Incidents, where an Incident is an unplanned interruption NCP-AIN to an IT service or a reduction in the quality of an IT service, Do not wait longer than one month from finishing the class until taking the exam.

NVIDIA NCP-AIN Exam Questions in Convenient PDF Format

We provide 24/7 customer service for all of you, please feel free to send us any questions about NVIDIA NCP-AIN Test Pdf through email or online chat, and we will always try our best to keeping our customer satisfied.

Our NCP-AIN learning guide is for you to improve your efficiency and complete the tasks with a higher quality, We do not hope that you spend all your time on learning the NCP-AIN certification materials.

If you are looking for a good learning site that can help you to pass the NVIDIA NCP-AIN exam, Prep4away is the best choice,

With exemplary quality, we boast that our NCP-AIN exam products are the best of all.

DOWNLOAD the newest Prep4away NCP-AIN PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1HkarvavSbRcqYzKx8lFxygjWOfn8fNYc>