

Free PDF Quiz 2026 NCP-AIN: NVIDIA-Certified Professional AI Networking–The Best Exams Collection



P.S. Free & New NCP-AIN dumps are available on Google Drive shared by BraindumpStudy: https://drive.google.com/open?id=1HKOSz1k5dGv_f0OYHDsp46SI3DB9dmpr

Several advantages we now offer for your reference. On the one hand, our NCP-AIN learning questions engage our working staff in understanding customers' diverse and evolving expectations and incorporate that understanding into our strategies, thus you can 100% trust our NCP-AIN Exam Engine. On the other hand, the professional NCP-AIN study materials determine the high pass rate. According to the research statistics, we can confidently tell that 99% candidates have passed the NCP-AIN exam.

NVIDIA NCP-AIN Exam Syllabus Topics:

Section	Weight	Objectives
Topic 1: Troubleshooting Tools and Techniques	10-15%	- Diagnostic and monitoring tools • 1. NetQ, CloudAI, NCCL analysis • 2. Performance benchmarking and validation - Common issues and resolution • 1. Error detection and correction • 2. Connectivity and performance problems
Topic 2: AI Data Center Design and Optimization	20-25%	- AI data center architecture principles • 1. Network topology design for GPU clusters • 2. Distributed training and communication requirements - Performance optimization strategies • 1. Congestion control and flow management • 2. Latency and bandwidth optimization
Topic 3: NVIDIA InfiniBand Networking	20-25%	- UFM (Unified Fabric Manager) operations • 1. Fabric monitoring and diagnostics • 2. Performance tuning and optimization - InfiniBand fabric deployment • 1. Multi-tenant and partition configuration • 2. Subnet management and routing

Topic 4: Kubernetes Integration	10-15%	- Networking for AI workloads in Kubernetes • 1. RDMA and SR-IOV integration • 2. Network policy and service management
Topic 5: NVIDIA Spectrum Networking	20-25%	- Spectrum switch configuration and management • 1. RoCEv2 and RDMA configuration • 2. Quality of Service (QoS) settings - Spectrum-X solutions and features • 1. Security and access control • 2. Rail-optimized topologies
Topic 6: Automation and Configuration Management	10-15%	- Infrastructure as Code and automation • 1. API usage and integration • 2. Configuration templates and scripts

>> NCP-AIN Exams Collection <<

NVIDIA NCP-AIN Examcollection Vce | Examinations NCP-AIN Actual Questions

The learning material is open in three excellent formats; NVIDIA NCP-AIN dumps PDF, a desktop NVIDIA NCP-AIN dumps practice test, and a web-based NVIDIA NCP-AIN dumps practice test. NVIDIA NCP-AIN dumps is organized by experts while saving the furthest down-the-line plan to them for the NVIDIA NCP-AIN Exam. The sans bug plans have been given to you all to drift through the NVIDIA-Certified Professional AI Networking certificate exam.

NVIDIA-Certified Professional AI Networking Sample Questions (Q25-Q30):

NEW QUESTION # 25

You're troubleshooting a Spectrum-X network and notice that the System Status LED on a switch is blinking for more than 5 minutes. What is the most likely cause of this issue?

- A. The Onyx software did not boot properly
- B. The switch is overheating
- C. The power supply unit is failing

Answer: A

Explanation:

According to the NVIDIA Spectrum-X Switch Operating System (SX_OS) Troubleshooting Guide, the System Status LED behavior is a critical indicator of the switch's internal operational state.

From the document:

"The System Status LED will blink green during system initialization. If the LED continues blinking for more than 5 minutes, it indicates that the Onyx OS has failed to load properly. The system may be stuck in the boot process, or the file system may be corrupted." This blinking LED beyond normal initialization time indicates that the system has either encountered a failure during software boot or is unable to transition from bootloader to the OS runtime environment (i.e., Onyx).

Key causes include:

- * Corrupted or missing system files.
- * Failed firmware or OS upgrade attempts.
- * Boot device (e.g., eMMC or SSD) issues or corrupted partitions.

Technically, during power-on:

- * The switch performs POST (Power-On Self Test).
- * Then the Onyx OS attempts to load from the boot partition.
- * If the Onyx OS kernel or root filesystem is invalid, the system halts boot, and the LED remains in a blinking state, as no successful OS load confirmation is triggered.

Remediation Steps (as per NVIDIA guide):

- * Access the switch through console and monitor boot logs.
- * Use ONIE recovery or re-flash a stable Onyx OS version.
- * Check system storage integrity using built-in diagnostics.

Exact Extract Reference:

Source: NVIDIA SX_OS 3.9.3000 Documentation

Topic: Troubleshooting System Status LED

Extract: "If the LED blinks for more than 5 minutes and the switch is not accessible via CLI, the Onyx software failed to load properly and recovery procedures must be initiated."

NEW QUESTION # 26

Why is the InfiniBand LRH called a local header?

- A. It allows traffic on a local link only.
- B. It provides the parameters for each local HCA.
- C. It provides the LIDs from the local subnet manager.
- D. It is used for routing traffic between nodes in the local subnet.

Answer: D

Explanation:

The Local Route Header (LRH) in InfiniBand is termed "local" because it is used exclusively for routing packets within a single subnet. The LRH contains the destination and source Local Identifiers (LIDs), which are unique within a subnet, facilitating efficient routing without the need for global addressing. This design optimizes performance and simplifies routing within localized network segments. InfiniBand is a high-performance, low-latency interconnect technology widely used in AI and HPC data centers, supported by NVIDIA's Quantum InfiniBand switches and adapters. The Local Routing Header (LRH) is a critical component of the InfiniBand packet structure, used to facilitate routing within an InfiniBand fabric. The question asks why the LRH is called a "local header," which relates to its role in the InfiniBand network architecture.

According to NVIDIA's official InfiniBand documentation, the LRH is termed "local" because it contains the addressing information necessary for routing packets between nodes within the same InfiniBand subnet. The LRH includes fields such as the Source Local Identifier (SLID) and Destination Local Identifier (DLID), which are assigned by the subnet manager to identify the source and destination endpoints within the local subnet. These identifiers enable switches to forward packets efficiently within the subnet without requiring global routing information, distinguishing the LRH from the Global Routing Header (GRH), which is used for inter-subnet routing.

Exact Extract from NVIDIA Documentation:

"The Local Routing Header (LRH) is used for routing InfiniBand packets within a single subnet. It contains the Source LID (SLID) and Destination LID (DLID), which are assigned by the subnet manager to identify the source and destination nodes in the local subnet. The LRH is called a 'local header' because it facilitates intra-subnet routing, enabling switches to forward packets based on LID-based forwarding tables."

-NVIDIA InfiniBand Architecture Guide

This extract confirms that option A is the correct answer, as the LRH's primary function is to route traffic between nodes within the local subnet, leveraging LID-based addressing. The term "local" reflects its scope, which is limited to a single InfiniBand subnet managed by a subnet manager.

Reference: LRH and GRH InfiniBand Headers - NVIDIA Enterprise Support Portal

NEW QUESTION # 27

When upgrading Cumulus Linux to a new version, which configuration files should be migrated from the old installation?

Pick the 2 correct responses below.

- A. All files in /etc/network
- B. All files in /etc/mix
- C. All files in /etc/cumulus/ac1
- D. All files in /etc

Answer: A,C

Explanation:

Before upgrading Cumulus Linux, it's essential to back up configuration files to a different server. The /etc directory is the primary location for all configuration data in Cumulus Linux. Specifically, the following files and directories should be backed up:

- * /etc/frr/ - Routing application (responsible for BGP and OSPF)
- * /etc/hostname - Configuration file for the hostname of the switch
- * /etc/network/ - Network configuration files, most notably /etc/network/interfaces and /etc/network/interfaces.d/
- * /etc/cumulus/acl - Access control list configurations

Cumulus Linux is a network operating system used on NVIDIA Spectrum switches, including those in the Spectrum-X platform, to provide a Linux-based environment for Ethernet networking in AI and HPC data centers. When upgrading Cumulus Linux to a new version, it's critical to migrate specific configuration files to preserve network settings and ensure continuity. The question asks for the two configuration file locations that should be migrated from the old installation during an upgrade.

According to NVIDIA's official Cumulus Linux documentation, the key directories containing configuration files that should be migrated during an upgrade are /etc/cumulus/acl (for access control list configurations) and /etc/network (for network interface configurations). These directories store critical network settings that define the switch's behavior, such as ACL rules and interface settings, which must be preserved to maintain network functionality after the upgrade.

Exact Extract from NVIDIA Documentation:

"When upgrading Cumulus Linux, you must back up and migrate specific configuration files to ensure continuity of network settings. The following directories should be included in the backup:

- * /etc/cumulus/acl: Contains access control list (ACL) configuration files that define packet filtering and security policies.
 - * /etc/network: Contains network interface configuration files, such as interfaces and ifupdown2 settings, which define the network interfaces and their properties. Back up these directories before upgrading and restore them after the new version is installed to maintain consistent network behavior."
- NVIDIA Cumulus Linux Upgrade Guide This extract confirms that options A and B are the correct answers, as /etc/cumulus/acl and /etc/network contain essential configuration files that must be migrated during a Cumulus Linux upgrade. These files ensure that ACL policies and network interface settings are preserved, which are critical for Spectrum-X configurations in AI networking environments.

Reference:Upgrading Cumulus Linux - NVIDIA Docs

NEW QUESTION # 28

A cloud service provider is deploying the NVIDIA Spectrum-X Ethernet platform in a multi-tenant environment. To ensure the security and isolation of each tenant's AI workload, the provider wants to implement a feature that prevents unauthorized access to the network. Which of the following features of the Spectrum-X platform should the provider implement?

- A. Congestion Control
- B. Streaming Telemetry
- C. Adaptive Routing
- **D. Traffic Isolation**

Answer: D

Explanation:

In multi-tenant AI cloud environments, ensuring that each tenant's workloads are isolated and secure is paramount. The NVIDIA Spectrum-X platform addresses this need through its Traffic Isolation capabilities. This feature ensures that network resources are partitioned effectively, preventing unauthorized access and interference between tenants. By implementing Traffic Isolation, the provider can maintain strict boundaries between different tenant environments, ensuring both security and performance consistency.

NEW QUESTION # 29

In a multi-tenant InfiniBand environment managed by UFM, you need to configure access controls to prevent unauthorized users from altering the fabric configuration. Which method is used within UFM to manage user access and ensure authorized modifications only?

- A. Virtual Network Segmentation (VNS)
- B. Network Access Control (NAC)
- **C. Role-Based Access Control (RBAC)**
- D. Digital Certification Management (DCM)

Answer: C

Explanation:

Role-Based Access Control (RBAC) is implemented within NVIDIA's Unified Fabric Manager (UFM) to manage user permissions effectively. RBAC allows administrators to assign roles to users, each with specific permissions, ensuring that only authorized individuals can make changes to the fabric configuration. This structured approach to access control enhances security by limiting the

potential for unauthorized modifications and streamlines the management of user privileges across the network.
Reference: Role-Based Access Control (RBAC) - One Identity

NEW QUESTION # 30

• • • • •

BraindumpStudy's NVIDIA NCP-AIN exam training materials are the best training materials of all the Internet training resources. Our visibility is very high, which are results that obtained through many candidates who have used the BraindumpStudy's NVIDIA NCP-AIN exam training materials. If you also use BraindumpStudy's NVIDIA NCP-AIN Exam Training materials, we can give you 100% guarantee of success. If you do not pass the exam, we will refund the full purchase cost to you. For the vital interests of the majority of candidates, BraindumpStudy is absolutely trustworthy.

NCP-AIN Examcollection Vce: https://www.braindumpstudy.com/NCP-AIN_braindumps.html

- [illegible]

P.S. Free & New NCP-AIN dumps are available on Google Drive shared by BraindumpStudy: https://drive.google.com/open?id=1HKOSzlK5dGv_f0OYHDsp46SI3DB9dmpr