

Downloaded from <http://ajphaphapublications.org/> on November 10, 2015

--	--

Thema 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.
Thema 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.

>> NCP-AIN Prüfungsfragen <<

NCP-AIN Ausbildungsressourcen - NCP-AIN Fragen Antworten

Die NVIDIA NCP-AIN Zertifizierungsprüfung ist heutzutage sehr beliebt. PrüfungFrage wird Ihnen helfen, die NCP-AIN Prüfung zu bestehen, und bietet Ihnen einen einjährigen kostenlosen Update-Service. Dann wählen Sie doch PrüfungFrage, um Ihren Traum zu verwirklichen. Um Erfolg zu erringen, ist Ihnen weise, PrüfungFrage zu wählen. Wählen Sie PrüfungFrage, Sie werden der nächste IT-Elite sein.

NVIDIA-Certified Professional AI Networking NCP-AIN Prüfungsfragen mit Lösungen (Q13-Q18):

13. Frage

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. nvidia-smi utility
- **B. mlxlink utility**
- C. tcpdump tool

Antwort: B

Begründung:

The mlxlink 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: mlxlink Utility - NVIDIA Docs

14. Frage

In order to configure RoCE on a Cumulus switch, which command should be used?

- A. nv roce qos enable on
- B. nv set roce qos enable on
- C. nv qos roce enable on
- **D. nv set qos roce enable on**

Antwort: D

Begründung:

To enable RDMA over Converged Ethernet (RoCE) on a Cumulus switch, the correct command is:

nv set qos roce enable on

This command configures the Quality of Service (QoS) settings to support RoCE, ensuring that the necessary parameters for lossless Ethernet are applied.

Reference: NVIDIA Cumulus Linux Documentation - RDMA over Converged Ethernet (RoCE)

15. Frage

A financial services company is planning to implement an AI infrastructure to support real-time fraud detection and risk assessment. They need a solution that can handle both training and inference workloads while maintaining data privacy and security. Which NVIDIA reference architecture component would be most appropriate to address the data privacy and security concerns in this AI networking setup?

- A. NVIDIA Spectrum switches
- B. NVIDIA Magnum IO
- C. NVIDIA CUDA-X AI libraries
- **D. NVIDIA BlueField DPUs**

Antwort: D

Begründung:

NVIDIA BlueField Data Processing Units (DPUs) are integral to securing AI infrastructures, especially in environments requiring stringent data privacy and security measures. BlueField DPUs offload and accelerate critical infrastructure tasks such as encryption, firewall enforcement, and intrusion detection, thereby isolating sensitive data paths from potential threats.

In the context of AI workloads, BlueField DPUs enable secure and efficient data movement between GPUs and storage systems, ensuring that sensitive information, like financial data, is protected during both training and inference processes. Their integration into NVIDIA's reference architectures provides a hardware root of trust, essential for maintaining data integrity and compliance with security standards.

Reference: NVIDIA BlueField Networking Platform

16. Frage

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

Antwort: A

Begründung:

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.

17. Frage

Which service on Cumulus switches can monitor layer 1, layer 2, layer 3, tunnel, buffer, and ACL related issues?

- A. BGP
- B. ONIE
- C. WJH
- D. NCLU

Antwort: C

Begründung:

The "What Just Happened" (WJH) service on Cumulus switches provides real-time visibility into network problems by monitoring various layers and components, including layer 1, layer 2, layer 3, tunnel, buffer, and Access Control List (ACL) related issues. WJH streams detailed and contextual telemetry data, enabling administrators to diagnose and troubleshoot network problems effectively.

Reference Extracts from NVIDIA Documentation:

* "WJH can monitor layer 1, layer 2, layer 3, tunnel, buffer and ACL related issues."

* "The WJH service enables you to diagnose network problems by looking at dropped packets."

18. Frage

.....

Heute steigt sich alles außer dem Gehalt sehr schnell. Wollen Sie nicht einen Durchbruch machen? Sie können Ihr Gehalt verdoppeln. Das ist sehr wahrscheinlich. Wenn Sie nur die NVIDIA NCP-AIN Zertifizierungsprüfung bestehen können, können Sie bekommen, wie Sie wollen. Die Dumps von PrüfungFrage wird Ihnen helfen, die NVIDIA NCP-AIN Prüfung 100% zu bestehen, was uns sehr wundert. Das ist echt, Sie sollen keine Zweifel haben.

NCP-AIN Ausbildungsressourcen: <https://www.pruefungfrage.de/NCP-AIN-dumps-deutsch.html>

- NCP-AIN Simulationsfragen ☐ NCP-AIN Exam Fragen ☐ NCP-AIN Online Prüfung ☐ Öffnen Sie die Webseite ☀ de.fast2test.com ☀☀ und suchen Sie nach kostenloser Download von (NCP-AIN) ☐ NCP-AIN Online Prüfung
- NCP-AIN Schulungsangebot, NCP-AIN Testing Engine, NVIDIA-Certified Professional AI Networking Trainingsunterlagen ☐ Öffnen Sie die Website [www.itzert.com] Suchen Sie 《 NCP-AIN 》 Kostenloser Download ☐ NCP-AIN Demotesten
- NCP-AIN Praxisprüfung ☐ NCP-AIN Vorbereitung ☐ NCP-AIN Zertifikatsfragen ☐ Suchen Sie auf [www.pruefungfrage.de] nach kostenlosem Download von ➡ NCP-AIN ☐ ☘ NCP-AIN Ausbildungsressourcen
- NCP-AIN Examengine ☐ NCP-AIN Unterlage ☐ NCP-AIN Übungsmaterialien ☐ Suchen Sie auf { www.itzert.com

- [illegible]