Vendor: NVIDIA

Exam Code: NCP-AIN

Exam Name: AI Networking

Version: DEMO

QUESTION 1
You suspect there might be connectivity issues in your InfiniBand fabric and need to perform a comprehensive check. Which tool should you use to run a full fabric diagnostic and generate a report?

A. ibnetdiscover
B. perfquery
C. ibdiagnet
D. taping

Answer: C
Explanation:
The ibdiagnet utility is a fundamental tool for InfiniBand fabric discovery, error detection, and diagnostics. It provides comprehensive reports on the fabric's health, including error reporting, switch and Host Channel Adapter (HCA) configuration dumps, various counters reported by the switches and HCAs, and parameters of devices such as switch fans, power supply units, cables, and PCI lanes. Additionally, ibdiagnet performs validation for Unicast Routing, Adaptive Routing, and Multicast Routing to ensure correctness and a credit-loop-free routing environment.

QUESTION 2
When utilizing the ib_write_bw tool for performance testing, what does the -S flag define?

A. The burst size
B. The number of QP's
C. The maximum rate of sent packages
D. Which service level to use

Answer: D
Explanation:
From NVIDIA Performance Tuning Guide (ib_write_bw Tool Usage):
"-S <SL>: Specifies the Service Level (SL) to use for the InfiniBand traffic. SL is used for setting priority and mapping to virtual lanes (VLs) on the IB fabric."
This flag is useful when testing QoS-aware setups or validating SL/VL mappings.

QUESTION 3
When upgrading DOCA on a BlueField DPU, what command should first be run on the host?

A. sudo apt-get autoremove
B. /usr/sbin/ofed_uninstall.sh -force
C. sudo apt-get upgrade doca
D. sudo apt-get install doca

Answer: B
Explanation:
Before upgrading the DOCA SDK on a BlueField DPU, it is mandatory to uninstall the existing OFED drivers to prevent compatibility conflicts.

QUESTION 4
You are troubleshooting InfiniBand connectivity issues in a cluster managed by the NVIDIA Network Operator. You need to verify the status of the InfiniBand interfaces. Which command should you use to check the state and link layer of InfiniBand interfaces on a node?

A. rdma show devices
B. ibstat -d mlx5_X
C. ifconfig ib0
D. ip link show dev ib0

Answer: B
Explanation:
To check the status and link layer of InfiniBand interfaces, the ibstat command is used.
For example:
ibstat -d mlx5_0
This command provides detailed information about the InfiniBand device, including its state (e.g., Active), physical state (e.g., LinkUp), and link layer (e.g., InfiniBand).

QUESTION 5
You are concerned about potential security threats and unexpected downtime in your InfiniBand data center. Which UFM platform uses analytics to detect security threats, operational issues, and predict network failures in InfiniBand data centers?

A. Host Agent
B. Enterprise Platform
C. Cyber-AI Platform
D. Telemetry Platform

Answer: C
Explanation:
The NVIDIA UFM Cyber-AI Platform is specifically designed to enhance security and operational efficiency in InfiniBand data centers. It leverages AI-powered analytics to detect security threats, operational anomalies, and predict potential network failures. By analyzing real-time telemetry data, it identifies abnormal behaviors and performance degradation, enabling proactive maintenance and threat mitigation.
This platform integrates with existing UFM Enterprise and Telemetry services to provide a comprehensive view of the network's health and security posture. It utilizes machine learning algorithms to establish baselines for normal operations and detect deviations that may indicate security breaches or hardware issues.

QUESTION 6
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 full membership to both GPU servers and storage system.
B. Assign limited membership to both GPU servers and storage system.
C. Assign limited membership PKey to the shared storage system and full membership PKey to each tenant's GPU servers.
D. Assign full membership PKey to the shared storage system and limited membership PKey to each tenant's GPU servers.

Answer: D
Explanation:

To enforce strictmulti-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

QUESTION 7
Why is the InfiniBand LRH called a local header?

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

Answer: A
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.

QUESTION 8
In Cumulus Linux, which technology enables the ability to provide active-active redundancy to servers, without the need for direct inter-switch links?

A. MLAG
B. VSS
C. EVPN Multi-homing

Answer: C
Explanation:
EVPN Multi-homingenablesactive-active redundancy without inter-switch links by using overlay routing over VXLAN and distributed control plane using BGP EVPN.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.lead2pass.com/all-products.html

10% Discount Coupon Code: **ASTR14**