

# NVIDIA

## NCP-AII

### NVIDIA AI Infrastructure

For More Information – Visit link below:

<https://www.examsempire.com/>

**Product Version**

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/ncp-aii>

# Latest Version: 6.0

## Question: 1

What command is needed to measure BER (Bit Error Rate)?

- A. mlxconfig -d <device> q
- B. ethtool -S <device>
- C. mlxlink -d <device> -c -e
- D. mstflint -d <device> q full

**Answer: C**

Explanation:

In NVIDIA networking environments, specifically those utilizing InfiniBand or high-speed Ethernet via ConnectX adapters, monitoring the physical link quality is critical for preventing packet loss and RDMA retransmissions. The mlxlink tool is part of the NVIDIA Firmware Tools (MFT) package and is the primary utility for checking the status and health of the physical link. Using the -d flag specifies the device (e.g., /dev/mst/mt4123\_pciconf0), while the -c (counters) and -e (error counters/BER) flags provide a detailed readout of the link's performance. Bit Error Rate (BER) is a fundamental metric for signal integrity. NVIDIA systems typically distinguish between "Raw BER" (errors before Forward Error Correction) and "Effective BER" (errors remaining after FEC). A high BER often points to a failing transceiver, a dirty fiber connector, or a marginal DAC cable. While ethtool can show general statistics in Ethernet mode, mlxlink is the verified method for granular BER measurement across InfiniBand and high-speed fabrics, allowing engineers to determine if a link meets the "Error-Free" operation standards required for large-scale AI collective communications like NCCL.

## Question: 2

When updating the firmware on an NVLink switch transceiver, how can an engineer apply new firmware without interrupting the network?

- A. mlxfwreset -d -lid 27 reset --yes to reset the transceiver
- B. Physically disconnect and reconnect the transceiver.
- C. flint -d -lid 27 --linkx --linkx\_auto\_update --activate
- D. nv action reboot system to force immediate activation.

**Answer: C**

Explanation:

NVIDIA's LinkX optical transceivers and active copper cables often require firmware updates to ensure compatibility and performance optimizations. In a production DGX SuperPOD environment,

interrupting the NVLink fabric can cause GPU-to-GPU communication failures and crash training jobs. To mitigate this, NVIDIA utilizes the flint utility (part of MFT) with specific flags for "Live" or "Seamless" updates. The --linkx flag targets the transceiver or cable specifically, rather than the switch ASIC itself. The --linkx\_auto\_update flag automates the sequence, while the --activate flag ensures the new firmware is applied to the module's active memory without requiring a full system reboot or a manual flap of the network link. This "in-service" update capability is essential for largescale AI clusters where uptime is measured in weeks or months of continuous training. By using the -lid (Logical Identifier) target, an administrator can address specific modules across the fabric from a central management node, ensuring that the high-bandwidth NVLink mesh remains stable while maintaining the latest hardware optimizations.

### Question: 3

An infrastructure engineer in an AI factory has successfully replaced a power supply unit on an NVIDIA DGX H100. After installation, both the IN and OUT LEDs on the new power supply illuminate solid green. Which NVSM CLI command should the engineer use to quickly verify the overall system status and ensure it is operating as expected?

- A. nvsm show power
- B. nvsm show powermode
- C. nvsm show health
- D. nvsm show alerts

**Answer: C**

Explanation:

The NVIDIA System Management (NVSM) tool is the definitive CLI utility for monitoring the health of DGX platforms. While replacing a PSU (Power Supply Unit) is a common maintenance task, verifying that the new component is correctly integrated into the system's health model is mandatory. While nvsm show power would provide specific data regarding wattage and voltage for the PSU, the most comprehensive way to ensure the replacement hasn't caused secondary issues or that the system hasn't remained in a "Degraded" state is to run nvsm show health. This command performs a global check across all subsystems: GPUs, NVLink switches, storage, fans, and power. If the PSU replacement was successful and the system is back to full redundancy, nvsm show health will return a "Healthy" status. In an AI factory setting, where DGX H100 nodes pull significant power, ensuring that all 6 PSUs (in an N+N or N+1 configuration) are not only physically green but logically acknowledged by the Baseboard Management Controller (BMC) is critical for preventing unexpected shutdowns during high-load training iterations.

### Question: 4

A leaf switch shows "FW Version Mismatch" alerts for transceivers after cluster expansion. Which tool validates transceiver firmware against expected versions?

- A. flint
- B. iblinkinfo
- C. mlxconfig
- D. ethtool

**Answer: A**

Explanation:

Firmware consistency is a pillar of stable InfiniBand fabric performance. When a cluster is expanded, new transceivers or cables may arrive with newer or older firmware than the existing base, leading to "FW Version Mismatch" alerts in management consoles like UFM (Unified Fabric Manager). The flint tool (or mstflint) is the correct utility for querying the specific firmware levels embedded within the transceivers. While iblinkinfo provides data on link speeds and port states, it does not provide the deep hardware-level firmware telemetry required for version validation. flint allows the administrator to query the device, compare the current burn version against the target image, and perform the necessary updates to bring the cluster into a uniform state. In NVIDIA AI infrastructure, maintaining uniform firmware across the fabric ensures that features like Adaptive Routing and Congestion Control operate predictably. Without version parity, inconsistent behavior in Forward Error Correction (FEC) or link-up negotiation can lead to intermittent performance drops that are difficult to diagnose at the application (NCCL) level.

### Question: 5

A system administrator needs to install a GPU/DPU in a server. The server has a free PCI-e slot, there are enough free PCI-e lanes, and there is enough room for the card. Which procedure should be followed?

- A. Ensure the server has enough power. Verify compatibility of cables with server's platform. Make sure the server is down to remove cables safely. Do not wear an ESD bracelet.
- B. Ensure the server has enough power. Make sure the server is down to remove cables safely. Wear an ESD bracelet.
- C. Ensure the server has enough power. Make sure the server is up and running with attached cables. Wear an ESD bracelet.
- D. Ensure the server has enough power. Verify compatibility of cables with server's platform. Make sure the server is down to remove cables safely. Wear an ESD bracelet.

**Answer: D**

Explanation:

The physical installation of high-performance NVIDIA components, such as H100 PCIe GPUs or BlueField DPUs, requires strict adherence to data center safety and hardware preservation standards. Option D is the only "100% verified" procedure because it covers three critical pillars: Power, Compatibility, and Safety. First, high-end GPUs can draw up to 300W-450W individually; verifying the server's PDU and internal PSU capacity is essential to prevent over-current shutdowns. Second, verifying cable compatibility (such as 12VHPWR or specific PCIe power 8-pin layouts) is vital to avoid

electrical damage. Third, "Cold Service" (ensuring the server is powered down and cables are removed) is the standard for non-hot-plug PCIe components to prevent short circuits. Finally, wearing an ESD (Electrostatic Discharge) bracelet is non-negotiable when handling NVIDIA hardware, as static charges can destroy the sensitive HBM (High Bandwidth Memory) or the GPU die itself. Skipping ESD protection (as suggested in Option A) or performing the install while the system is "up and running" (as suggested in Option C) are leading causes of hardware infant mortality in AI infrastructure.

**Thank You for Trying Our Product**  
**Special 16 USD Discount Coupon: NSZUBG3X**

**Email:** [support@examsempire.com](mailto:support@examsempire.com)

**Check our Customer Testimonials and ratings  
available on every product page.**

**Visit our website.**

**<https://examsempire.com/>**