

TESTPASSPORT

PRACTICE QUESTIONS



H i g h Q u a l i t y

B e t t e r S e r v i c e

<https://www.testpassport.com>

Free updates for one year

Exam : NCP-AII

**Title : NVIDIA Certified
Professional AI
Infrastructure**

Version : DEMO

1. A GPU in your AI server consistently overheats during inference workloads. You've ruled out inadequate cooling and software bugs.

Running 'nvidia-smi' shows high power draw even when idle.

Which of the following hardware issues are the most likely causes?

- A. Degraded thermal paste between the GPU die and the heatsink.
- B. A failing voltage regulator module (VRM) on the GPU board, causing excessive power leakage.
- C. Incorrectly seated GPU in the PCIe slot, leading to poor power delivery.
- D. A BIOS setting that is overvolting the GPU.
- E. Insufficient system RAM.

Answer: B

Explanation:

Degraded thermal paste loses its ability to conduct heat effectively. A failing VRM can cause excessive power draw and heat generation. An incorrectly seated GPU can cause instability and poor power delivery, leading to overheating. Overvolting in BIOS will definitely cause overheating. While insufficient RAM can cause performance issues, it is less likely to lead to overheating.

2. You are monitoring a server with 8 GPUs used for deep learning training. You observe that one of the GPUs reports a significantly lower utilization rate compared to the others, even though the workload is designed to distribute evenly. 'nvidia-smi' reports a persistent "XID 13" error for that GPU.

What is the most likely cause?

- A. A driver bug causing incorrect workload distribution.
- B. Insufficient system memory preventing data transfer to that GPU.
- C. A hardware fault within the GPU, such as a memory error or core failure.
- D. An incorrect CUDA version installed.
- E. The GPU's compute mode is set to 'Exclusive Process'.

Answer: C

Explanation:

XID 13 errors in 'nvidia-smi' typically indicate a hardware fault within the GPU. Driver bugs or memory issues would likely cause different error codes or system instability across multiple GPUs. CUDA version mismatch might prevent the application from running altogether, but is less likely to lead to a specific XID error on a single GPU. Exclusive Process mode will lead to it being used by a different process but not necessarily cause that XID error.

3. You notice that one of the fans in your GPU server is running at a significantly higher RPM than the others, even under minimal load. 'ipmitool sensor' output shows a normal temperature for that GPU.

What could be the potential causes?

- A. The fan's PWM control signal is malfunctioning, causing it to run at full speed.
- B. The fan bearing is wearing out, causing increased friction and requiring higher RPM to maintain airflow.
- C. The fan is attempting to compensate for restricted airflow due to dust buildup.
- D. The server's BMC (Baseboard Management Controller) has a faulty temperature sensor reading, causing it to overcompensate.
- E. A network connectivity issue is causing higher CPU utilization, leading to increased system-wide heat.

Answer: A, C

Explanation:

A malfunctioning PWM control signal, worn fan bearings, or restricted airflow can all cause a fan to run at higher RPMs. While a faulty BMC sensor could be a cause, the question states that 'ipmitool sensor' shows a normal temperature. Network connectivity issues are less likely to cause an isolated fan to run high, if the GPU temperature is normal.

4. After upgrading the network card drivers on your A1 inference server, you experience intermittent network connectivity issues, including packet loss and high latency. You've verified that the physical connections are secure.

Which of the following steps would be most effective in troubleshooting this issue?

- A. Roll back the network card drivers to the previous version.
- B. Check the system logs for error messages related to the network card or driver.
- C. Run network diagnostic tools like 'ping', 'tracert', and 'iperf3' to assess the network performance.
- D. Reinstall the operating system.
- E. Update the server's BIOS.

Answer: A

Explanation:

Rolling back drivers is a quick way to revert to a known working state. Checking system logs will provide valuable information about driver errors or network issues. Network diagnostic tools will quantify the network performance and help isolate the problem. Reinstalling the OS is drastic and should be a last resort. Updating the BIOS is unlikely to resolve driver-related network issues unless specifically recommended for the network card.

5. Your deep learning training job that utilizes NCCL (NVIDIA Collective Communications Library) for multi-GPU communication is failing with "NCCL internal error, unhandled system error" after a recent CUDA update. The error occurs during the 'all reduce' operation.

What is the most likely root cause and how would you address it?

- A. Incompatible NCCL version with the new CUDA version. Update NCCL to a version compatible with the installed CUDA version.
- B. Insufficient shared memory allocated to the CUDA context. Increase the shared memory limit using 'cudaDeviceSetLimit(cudaLimitSharedMemory, new_limit)'.
`cudaDeviceSetLimit(cudaLimitSharedMemory, new_limit);`
- C. Firewall rules blocking inter-GPU communication. Configure the firewall to allow communication on the NCCL-defined ports (typically 8000-8010).
- D. Faulty network cables used for inter-node communication (if the training job spans multiple servers). Replace the network cables with certified high-speed cables.
- E. GPU Direct RDMA is not properly configured. Check 'dmesg' for errors and ensure RDMA is enabled.

Answer: A

Explanation:

NCCL relies on specific CUDA versions. An incompatibility after a CUDA update is the most probable cause. Insufficient shared memory is less likely to cause a system error within NCCL. Firewall rules usually manifest as connection refused errors. Faulty network cables affect inter-node communication, not intra-node. While RDMA issues can cause problems, they typically don't present as 'unhandled system error' immediately after a CUDA update, and are more likely if RDMA was working previously.