

Common errors when running jobs on H100 and how to fix them

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: H100

Cat. No.: B15585327

[Get Quote](#)

H100 GPU Technical Support Center

This technical support center provides troubleshooting guides and frequently asked questions (FAQs) to help researchers, scientists, and drug development professionals resolve common errors when running jobs on NVIDIA **H100** GPUs.

Table of Contents

- [--INVALID-LINK--](#)
- [--INVALID-LINK--](#)
- [--INVALID-LINK--](#)
- [--INVALID-LINK--](#)
- [--INVALID-LINK--](#)

Performance Degradation

Question: My application is running slower than expected on the **H100 GPU**. How can I troubleshoot and improve its performance?

Answer:

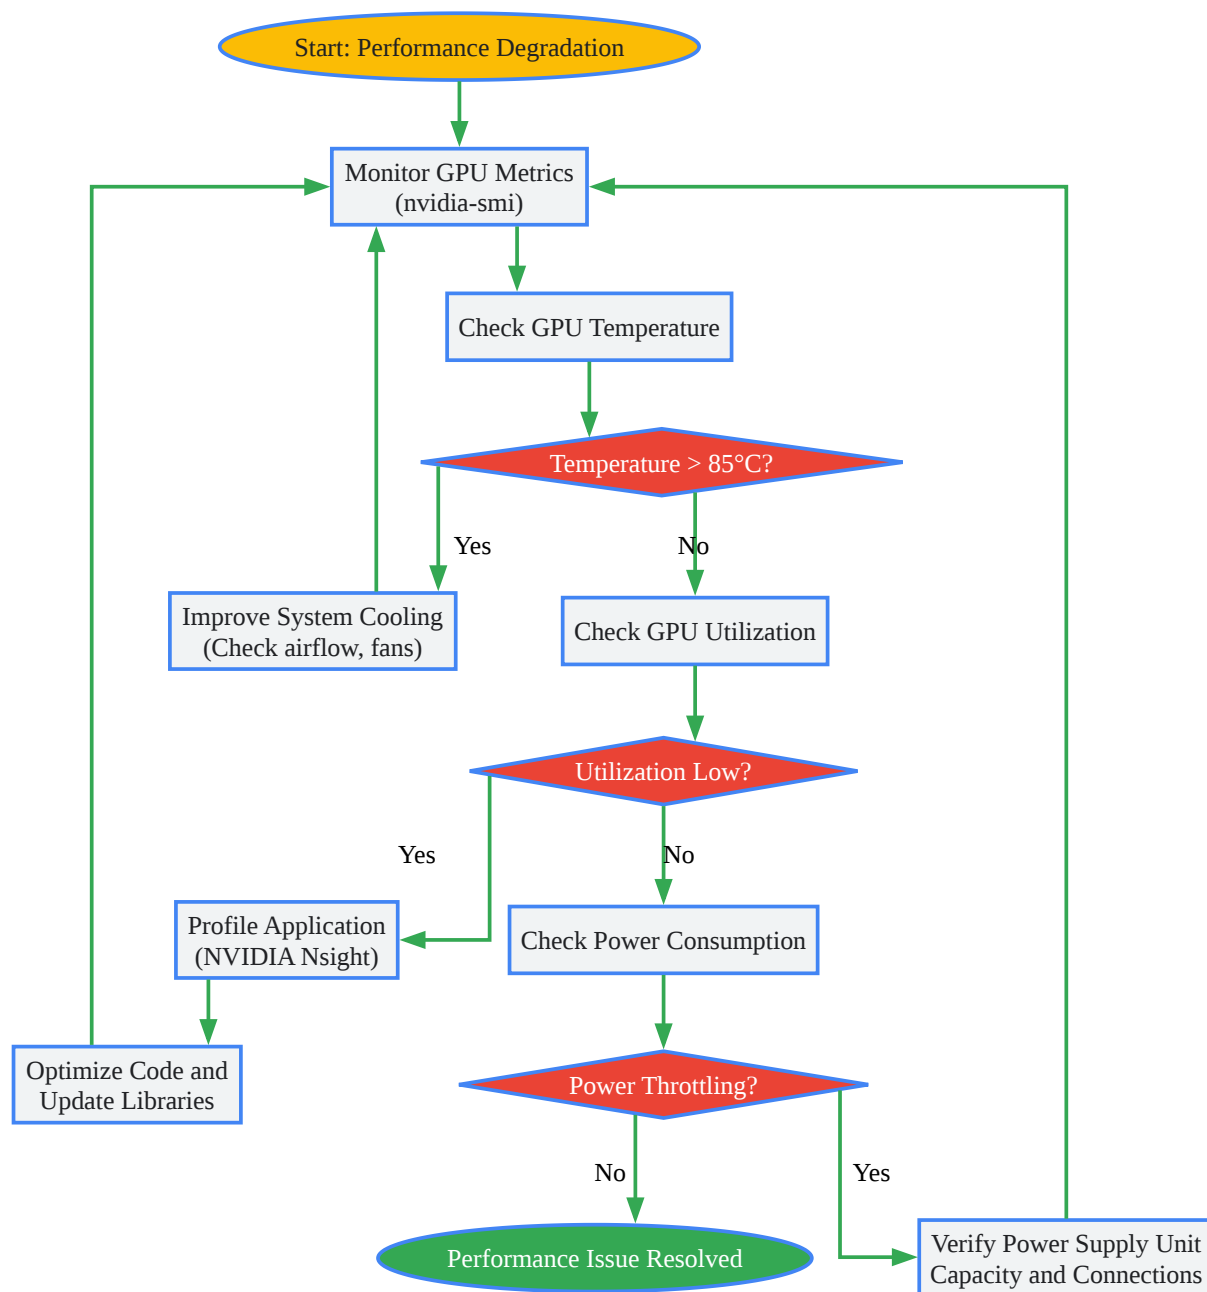
Performance degradation on **H100** GPUs can stem from various factors, including thermal throttling, inefficient workload distribution, and software bottlenecks.^{[1][2]} A systematic approach to identifying and resolving these issues is crucial for optimal performance.

- **Monitor GPU Health and Performance Metrics:** The first step is to gather data on the GPU's operational state. Key metrics to monitor include GPU utilization, temperature, power consumption, and memory usage.^[1] You can use the NVIDIA System Management Interface (nvidia-smi) command-line utility for real-time monitoring.^{[1][2]}
 - **Experimental Protocol: Monitoring with nvidia-smi**
 1. Open a terminal or command prompt on the system with the **H100** GPU.
 2. To get a continuous, real-time update of GPU metrics, run the following command:
 3. Pay close attention to the following fields in the nvidia-smi output:
 - **Fan Speed:** Abnormal values may indicate cooling issues. ^[3] * **Temp:** Consistently high temperatures can indicate cooling problems. ^[3] * **Perf:** Performance state (P0 is the maximum performance state).
 - **Pwr:Usage/Cap:** Power usage versus capacity. Unusual power patterns can signal hardware issues. ^[3] * **Memory-Usage:** Unexpected memory consumption may indicate memory leaks. ^[3] * **GPU-Util:** GPU utilization. Mismatches between utilization and expected workload can signal issues. ^[3]
- **Check for Thermal Throttling:** High GPU temperatures can cause the GPU to automatically reduce its clock speeds to prevent damage, a phenomenon known as thermal throttling. ^[2]^[4] This can significantly degrade performance.
 - **Recommended Temperature Ranges:**

| Metric | Recommended Range |
|--------|-------------------|
|--------|-------------------|

| GPU Temperature | Keep below 85°C for optimal performance. Throttling may occur above this temperature. |

- Optimize Workload and Software: Inefficient code or outdated software libraries can lead to performance bottlenecks.
 - Profile your application: Use profiling tools like NVIDIA Nsight to identify bottlenecks in your code. [\[2\]](#)
 - * Update frameworks and libraries: Ensure you are using the latest versions of your deep learning frameworks (e.g., TensorFlow, PyTorch) and CUDA libraries, as they often include performance optimizations for the latest GPU architectures. [\[5\]](#)
 - * Optimize CUDA kernels: If you are developing custom CUDA kernels, ensure they are optimized for the **H100**'s Hopper architecture. [\[5\]](#)



[Click to download full resolution via product page](#)

Caption: Troubleshooting workflow for **H100** performance degradation.

Memory Errors

Question: My job is failing with an "out of memory" (OOM) error. How can I resolve this?

Answer:

"Out of memory" (OOM) errors are common when the GPU's memory is insufficient to hold the data and models required for your job. [6][7]The **H100** GPU has a large memory capacity, but large-scale models and high-resolution data can still exceed its limits.

- **Monitor Memory Usage:** Use `nvidia-smi` to monitor the GPU's memory consumption in real-time. [6]This will help you understand how much memory your application is using.
- **Reduce Batch Size:** A simple and often effective solution is to reduce the batch size in your training script. This reduces the amount of data loaded into the GPU memory at once.
- **Optimize Your Model:**
 - **Model Parallelism:** For very large models, consider using model parallelism techniques to distribute the model across multiple GPUs.
 - **Gradient Accumulation:** This technique allows you to simulate a larger batch size by accumulating gradients over several smaller batches before updating the model weights.
- **Use Mixed-Precision Training:** Training in mixed precision (using both 16-bit and 32-bit floating-point numbers) can significantly reduce memory usage and improve performance on **H100** GPUs.

If you suspect a memory leak (a gradual increase in memory usage over time without a corresponding increase in workload), you can use the following protocol:

- **Baseline Monitoring:** Run your application with a fixed workload and monitor the memory usage over a set period using `nvidia-smi`.
- **Code Profiling:** Use a memory profiler, such as the one included in NVIDIA Nsight, to trace memory allocations and deallocations in your code.

- Isolate the Leak: Systematically comment out or disable parts of your code to identify the section that is causing the memory leak.

Driver and Software Conflicts

Question: I am encountering driver installation failures or my H100 GPU is not being recognized. What should I do?

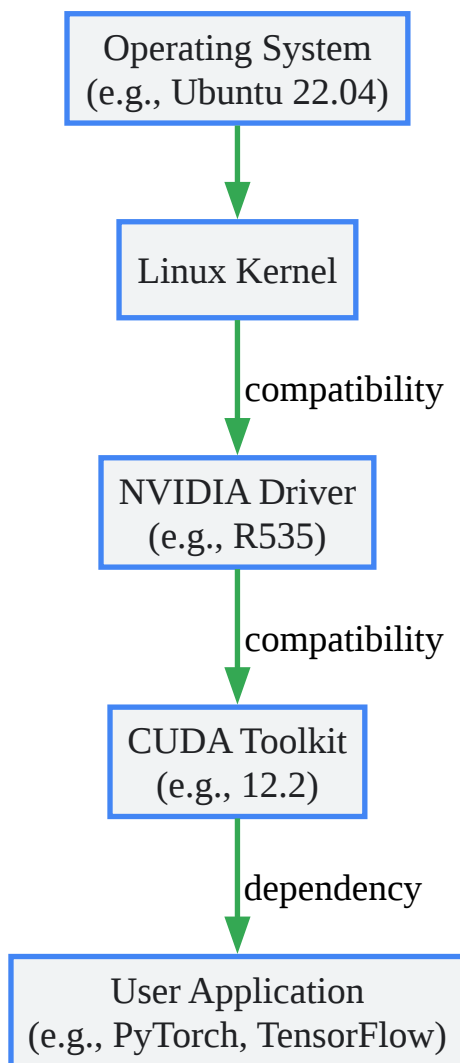
Answer:

Driver and software conflicts are a common source of issues with **H100** GPUs. [2][4] Ensuring compatibility between the operating system, NVIDIA driver, and CUDA toolkit is crucial for stable operation. [8]

| Component | Recommended Version for H100 |
|---------------|--|
| NVIDIA Driver | R535 or later (R535 provides production-stable performance) [9] |
| CUDA Toolkit | 11.8 or later (12.2 recommended for optimal performance) [9] |
| Linux Kernel | Check for compatibility with the specific NVIDIA driver version. Using unsupported kernel versions can lead to performance issues or the GPU not being recognized. [9] |

- Verify Driver and CUDA Compatibility: Ensure that the installed NVIDIA driver and CUDA toolkit versions are compatible with the **H100** GPU and your operating system. [8]2. Clean Installation of Drivers: If you are experiencing issues, it is recommended to perform a clean installation of the NVIDIA drivers. This involves completely removing any existing drivers before installing the new ones.
- Check for System Conflicts: Other hardware or software components may interfere with GPU detection. [10]Disable any onboard graphics in the system's BIOS to avoid conflicts. [10]4.

Review System Logs: Check system logs for any error messages related to the NVIDIA driver. On Linux, you can use `dmesg | grep -i nvidia` to look for relevant messages.



[Click to download full resolution via product page](#)

Caption: Dependency hierarchy of software components for **H100** GPU operation.

Hardware and Power Issues

Question: My system is not detecting the H100 GPU, or it is crashing under load. What could be the cause?

Answer:

Hardware and power-related issues can manifest as the GPU not being detected by the system or unexpected crashes during operation. [2][4] These problems often stem from improper installation, insufficient power, or faulty hardware.

- Verify Physical Installation:
 - Ensure the **H100** GPU is securely seated in the PCIe slot. [2][11] * Check that all necessary power cables are firmly connected to the GPU. [2][11]
- Check Power Supply:
 - The **H100** GPU has high power requirements (up to 700W). [12] Ensure that your power supply unit (PSU) has sufficient capacity to power the entire system, including the GPU under full load. [4][13] * Use high-quality power cables and avoid using adapters or splitters that may not provide stable power.
- Inspect for Physical Damage: Carefully inspect the GPU for any signs of physical damage to the card or the PCIe connector.
- Test in Another System: If possible, test the **H100** GPU in another compatible system to rule out issues with the motherboard or other system components. [10]

Overheating Issues

Question: My H100 GPU is running at very high temperatures. How can I prevent it from overheating?

Answer:

Overheating is a primary cause of performance degradation and can lead to hardware failure if not addressed. [2][4][12] The **H100** GPU generates a significant amount of heat, and proper cooling is essential for its stable operation.

- Monitor Temperatures: Use nvidia-smi to continuously monitor the GPU's temperature.
- Ensure Proper Airflow:

- Make sure the server or workstation has adequate airflow. [1] * Ensure that no cables or other components are obstructing the fans on the GPU or the case fans.
- Regularly clean any dust buildup from the GPU's heatsink and fans, as this can significantly impede cooling efficiency. [13]
- Check Ambient Temperature: The ambient temperature of the room or data center can affect GPU temperatures. Ensure the environment is within the recommended operating temperature range for your system.
- Consider Advanced Cooling Solutions: For high-density deployments or systems running under continuous heavy load, consider using more advanced cooling solutions such as liquid cooling.

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- 1. How do I monitor and troubleshoot NVIDIA H100 GPU performance and reliability issues in a data center? - Massed Compute [massedcompute.com]
- 2. How do I troubleshoot common issues with NVIDIA H100 GPUs in a data center? - Massed Compute [massedcompute.com]
- 3. support.exxactcorp.com [support.exxactcorp.com]
- 4. What are the common causes of GPU-related errors with NVIDIA H100 GPUs? - Massed Compute [massedcompute.com]
- 5. How do I troubleshoot common issues with H100 GPU performance in a large-scale HPC cluster? - Massed Compute [massedcompute.com]
- 6. What are some common memory-related issues that can occur on the H100 and how can I troubleshoot them? - Massed Compute [massedcompute.com]
- 7. youtube.com [youtube.com]
- 8. How do I troubleshoot compatibility issues with NVIDIA GPU drivers on the A100 and H100 GPUs? - Massed Compute [massedcompute.com]

- 9. docs.hyperstack.cloud [docs.hyperstack.cloud]
- 10. How do I resolve the 'NVIDIA H100 NVL GPU not detected' issue during driver installation? - Massed Compute [massedcompute.com]
- 11. How do I troubleshoot HBM2 memory issues on the NVIDIA H100 PCIe GPU? - Massed Compute [massedcompute.com]
- 12. cyfuture.cloud [cyfuture.cloud]
- 13. What are the common causes of NVIDIA H100 GPU errors in a data center environment? - Massed Compute [massedcompute.com]
- To cite this document: BenchChem. [Common errors when running jobs on H100 and how to fix them]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b15585327#common-errors-when-running-jobs-on-h100-and-how-to-fix-them]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [[Contact our Ph.D. Support Team for a compatibility check](#)]

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd
Ontario, CA 91761, United States
Phone: (601) 213-4426
Email: info@benchchem.com