

NVIDIA Xid triage reference card

Preserve the first event, classify the impact, and change one variable before recovery.

First response

1	Drain only the affected workload or node when the event can affect peers. Do not erase evidence with an immediate reboot.
2	Record the first Xid line, PCI address, GPU UUID, affected process or rank, driver version, GPU model, and the preceding 100 log lines.
3	Capture <code>nvidia-smi -q</code> , <code>journalctl -k</code> , DCGM diagnostics when safe, and <code>nvidia-bug-report.sh</code> before recovery.
4	Use the Xid as a starting point. A repeated controlled test decides whether the fault follows the workload, node, device, or configuration.

Common codes and safe first actions

Xid	Meaning	Safe first action
13	Graphics engine exception	Preserve the channel and application context. Run DCGM or field diagnostics, then debug the application if hardware checks pass.
31	GPU memory page fault	Capture the faulting process and address. Reproduce with Compute Sanitizer before treating the board as faulty.
45	Preemptive channel removal	Treat as informational in most cases. Find the earlier Xid or shutdown event that caused channel removal.
48	Double-bit ECC error	Drain the workload. If followed by 63 or 64, reset the reporting GPU after work completes. Otherwise collect field diagnostics.
63/64	Page retirement or row-remap event	Interpret with the GPU generation and adjacent 48 or 94 event. Schedule the reset or reboot required to activate repair.
74	NVLink error	Drain affected work, preserve the full Xid payload, inspect link counters and topology, then check physical links if supported.
79	GPU fell off the bus	Drain the node and collect the support packet. Investigate PCIe, power, thermals, seating, and the platform before replacement.
94	Contained memory error	Restart the affected application. Other applications may continue. Reset the GPU at a convenient maintenance point.
95	Uncontained memory error	Drain affected work. Reset the GPU before applications restart, or reboot the node when the platform cannot isolate the reset.
119/120	GSP RPC timeout or GSP error	Collect the bug report and driver context. Follow the reported recovery action, then escalate if it recurs after recovery.

Evidence and escalation checklist

Collect before reset, reboot, driver reload, or device replacement

1. System and GPU model, serial or asset ID, OS, kernel, driver, CUDA, firmware, and Fabric Manager versions.
2. Complete first-failure log context with timestamp, PCI address, GPU UUID, rank or process, and workload identifier.
3. `nvidia-smi -q`, `nvidia-bug-report.log.gz`, DCGM diagnostic output, and relevant Fabric Manager logs.
4. Every controlled test performed, what changed, what stayed fixed, and whether the Xid followed the workload or hardware.
5. Recovery action, observation window, recurrence count, and the vendor's requested RMA or support case fields.

Decision rule

Do not turn one Xid into a root-cause or RMA claim. Repeat a production-shaped control after the targeted change. Record whether the same first Xid recurs, whether it follows the workload or hardware, and how long the corrected system remains under observation.

Official sources

[NVIDIA Xid Errors](https://docs.nvidia.com/deploy/xid-errors/latest/index.html)

<https://docs.nvidia.com/deploy/xid-errors/latest/index.html>

[NVIDIA GPU Debug Guidelines](https://docs.nvidia.com/deploy/gpu-debug-guidelines/)

<https://docs.nvidia.com/deploy/gpu-debug-guidelines/>

[NVIDIA Xid Catalog analysis workflow](https://docs.nvidia.com/deploy/xid-errors/analyzing-xid-catalog.html)

<https://docs.nvidia.com/deploy/xid-errors/analyzing-xid-catalog.html>

This card summarizes public NVIDIA guidance for incident triage. It is not a substitute for model-specific vendor instructions, safety procedures, or an RMA decision.