




GPU PCIe Link Troubleshooting: Speed, Lanes and Errors

Plan and build private AI compute with GPU Smith. We connect workload requirements with GPU hardware, networking, deployment and operating decisions so your infrastructure fits the work it must perform.

Published by GPU Smith · 2026-10-15 · gpu pcie link troubleshooting · pcie link speed · pcie link width · pcie aer · nvidia-smi · dcgm · pcie replay counters · gpu bandwidth · server diagnostics

Original article: <https://gpusmith.com/articles/en/gpu-pcie-link-health-troubleshooting>



In this article

1. Executive Summary
 2. Introduction and Background
 3. Establish Expected Capability and Locate the Symptom
 4. Read Negotiated Speed and Width at Idle and Under Load
 5. Decode AER Severity and Preserve the Reporting Path
 6. Trend PCIe Replays Without Clearing the Evidence
 7. Use Controlled Transfers to Test the Hypothesis
 8. Isolate Safely and Prepare the OEM Evidence Bundle
 9. Data Analysis and Evidence
 10. Implications and Future Directions
 11. Frequently Asked Questions (FAQs)
 12. Conclusion
-

Executive Summary

GPU PCIe link troubleshooting, for a graphics processing unit (**GPU**), should establish what the installed platform is designed to provide, observe what the link actually negotiates, and then correlate errors with a controlled transfer. Peripheral Component Interconnect Express (**PCIe**) link capability and link status are different register views: **LnkCap** describes capability, while **LnkSta** describes the negotiated state. ([pciutils: capability](#); [status](#)) A low idle reading is insufficient evidence of degradation because NVIDIA documents that both current generation and width may fall while the GPU is unused. ([docs.nvidia.com](#)) The decisive comparison is the same endpoint and upstream path observed during a recorded workload.

For **PCIe Gen5 x16**, the calculated raw ceiling is **64 decimal gigabytes per second (GB/s) per direction**, before encoding or packet overhead. ([PCI-SIG: raw rate](#); [lane scaling](#)) Applying the documented encoding ratio produces an encoding-only ceiling of **63.015 GB/s**, which remains higher than useful application throughput. ([docs.amd.com](#)) These figures are explanatory limits, not acceptance criteria. Expected width also depends on the system configuration: Lenovo documents direct **x8** and switched **x16** GPU connections on its SR675 V3, with different subscription ratios. ([Lenovo: direct riser](#); [switched riser](#)) A full-width endpoint can therefore coexist with shared upstream capacity.

Advanced Error Reporting (AER) and replay telemetry answer different questions. Linux defines correctable conditions as hardware-corrected, while fatal errors make the link unreliable and require a different recovery path. ([Linux AER: correctable](#); [fatal](#)) A reporting bus/device/function identifier (**BDF**) is a location clue, not proof that the GPU is defective; the kernel also documents counters incrementing at a link partner. ([docs.kernel.org](#)) ([docs.kernel.org](#)) Preserve the kernel journal, endpoint and upstream status, and counter timestamps before resets or physical service. Volatile journals can disappear at reboot. ([documentation.suse.com](#))

The recommended decision ladder is **inventory, idle/active comparison, AER classification, replay deltas, a reserved transfer test, and isolation under original equipment manufacturer (OEM) guidance**. There is no universal replay cutoff in this report: Data Center GPU Manager (DCGM)'s documented replay limit belongs to its selected diagnostic run. ([docs.nvidia.com](#)) Compare matched observation intervals with a healthy

equivalent configuration, and treat a reset or counter decrease as a discontinuity. Escalate a persistent loaded-state mismatch or correlated error growth with the complete record. Replace hardware only after configuration expectations and the approved service procedure make the result interpretable.

Introduction and Background

A slow accelerator does not automatically have a slow PCIe link. The operational question is narrower: does the installed GPU path negotiate the capacity expected for this machine, and does it show transport errors during the workload that matters? PCIe link initialization negotiates lane counts and data rates; these are observable properties rather than conclusions about application performance. (www.keysight.com) This report develops an evidence ladder for server and GPU platform engineers who need to distinguish configuration, negotiation and error evidence before changing components.

The starting point is the installed machine, not a GPU datasheet considered alone. NVIDIA's reported maximum generation and width depend on the GPU **and system configuration**. (docs.nvidia.com) Linux's bandwidth helper similarly walks upstream to locate the minimum-bandwidth point. (docs.kernel.org) Consequently, a correct GPU endpoint reading can leave a narrower bridge or shared uplink unresolved. Inventory should identify that path without turning this procedure into a workload-affinity or purchasing guide.

The tools expose different layers. **PCI Utilities** supplies `lspci` for enumerating buses and connected devices. (mj.ucw.cz) Linux supplies AER evidence, while vendor interfaces provide accelerator telemetry; AMD SMI, for example, documents PCIe speed, width and replay count together. (rocm.docs.amd.com) These observations should be kept with their source, unit and timestamp. Combining them into an unexplained “healthy” score loses the distinction between unsupported telemetry, inactive links and observed recovery activity.

The engineering perspective is independent assessment against documented expectations. **GPU Smith** describes itself as an independent private AI compute engineering firm; its relevance here is that advisory perspective, rather than an alternative PCIe product. (gpusmith.com) The runbook below presents recommendations and calculations, not client results or a claimed benchmark dataset.

Command guidance was checked against primary documentation and implementation sources for publication on **October 9, 2026**. This research environment had no GPU and lacked `lspci`, `nvidia-smi` and `dcmgi`, so no live server execution or bandwidth measurements are claimed. The installed kernel, driver, DCGM release and OEM maintenance documentation remain the execution authority on the target node.

Establish Expected Capability and Locate the Symptom

Begin with a configuration record that someone else can reproduce. A physical connector's size is insufficient: HPE's DL380 Gen11 QuickSpecs define bus width as the number of physical electrical lanes running to the connector. Check the installed riser variant and its slot-population table before setting the expected width. The cited revision is dated **May 4, 2026**. ([HPE QuickSpecs, V43](#)) Current model-specific population rules are the appropriate source for the node under examination.

For the **PowerEdge R760xa**, Dell distinguishes front configurations containing **four double-width Gen5 x16** GPU cards or **eight single-width Gen5 x8** cards. (www.dell.com) Lenovo's **SR675 V3** direct eight-GPU riser configuration provides **x8** connections with **1:1 subscription**, while its switched configuration provides **x16** connections with **1:2 subscription**. (Lenovo: [direct riser](#); [switched riser](#)) Therefore, “lower than the GPU maximum” and “lower than the installed platform's design” are different findings.

Capture the following before proposing changes:

- **Identity:** chassis model, serial, GPU model, GPU universally unique identifier (**UUID**) and full BDF. DCGM numeric IDs require correlation with a UUID or PCI bus ID. (docs.nvidia.com)
- **Population:** processor installation, slot assignment and riser variant. Dell's R760xa riser table assigns slots according to processor and riser configuration. (www.dell.com)
- **Path:** endpoint, intervening bridges or switches, and upstream port. The `lspci` tree view exposes these connections. (manpages.ubuntu.com)
- **Software:** operating-system build, kernel, driver and installed diagnostic releases; retain their reported version output.
- **Physical configuration:** approved riser and cable part numbers, recorded against the model-specific routing instructions. (pubs.lenovo.com)
- **Operating context:** idle or active workload, transfer direction, start/end times, and whether other jobs were running.

Use symptoms to choose the next observation, not the replacement part:

- **Low idle generation or width:** repeat during a recorded active interval because NVIDIA allows both readings to decrease when unused. (docs.nvidia.com)
- **Persistently low loaded width:** compare the endpoint and each upstream link with the populated-slot design, including intentional x8 configurations. (pubs.lenovo.com)
- **Expected endpoint state but slow transfers:** inspect upstream capacity and test conditions before concluding that negotiation failed. (docs.kernel.org)
- **Correctable messages without interruption:** retain the message stream and correlate it with counters; correction is distinct from a software recovery requirement. (docs.kernel.org)
- **Replay growth during transfers:** investigate the physical path as a candidate, without assigning blame from the counter alone. NVIDIA associates replays with PCIe transmission problems. (docs.nvidia.com)
- **Fatal error or active recovery:** preserve evidence and move to the platform's recovery process; Linux treats fatal errors as an unreliable link. (docs.kernel.org)

This symptom map is the report's diagnostic interpretation. It deliberately leaves open the possibility that the expected link state is correct and the performance limitation lies elsewhere.

Read Negotiated Speed and Width at Idle and Under Load

Collect a complete, unambiguous inventory

Use the full **domain:bus:device.function** address. `lspci` documents that address format and prints its numeric components in hexadecimal. ([lspci: address format](#); [hexadecimal notation](#)) The commands below are templates: replace `DOMAIN:BUS:DEVICE.FUNCTION` with the recorded GPU address. They inspect inventory and state; they do not request retraining.

`-D` retains the domain, `-nn` retains numeric identifiers with names, `-k` includes driver information, and `-vv` supplies detailed capability decoding. ([lspci: domains](#); [numeric identifiers](#); [drivers](#); [verbose output](#))

```
lspci -Dnnk
lspci -Dtv
sudo lspci -D -s DOMAIN:BUS:DEVICE.FUNCTION -vv
nvidia-smi -q -i DOMAIN:BUS:DEVICE.FUNCTION
```

The device-selection form follows `lspci -s`; the NVIDIA command uses the full query rather than an assumed PCI-only display filter. ([manpages.ubuntu.com](#)) ([docs.nvidia.com](#)) Some configuration-space details require elevated access. If they are inaccessible, record that limitation rather than treating the capability as absent. ([manpages.ubuntu.com](#))

Retain complete output before extracting **LnkCap**, **LnkSta** and AER blocks. The upstream decoder reads capability and status from separate registers. ([pciutils: capability register](#); [status register](#)) Compare GPU and upstream port output, preserving which device each block belongs to. A bridge's output is not interchangeable with the GPU endpoint's output.

Cross-check through sysfs

Linux's current link attributes read the Link Status register; maximum speed and width use capability information. ([Linux: current link state](#); [link capability](#)) Where present and readable, the following files give a compact second view:

```
cat /sys/bus/pci/devices/DOMAIN:BUS:DEVICE.FUNCTION/current_link_speed
cat /sys/bus/pci/devices/DOMAIN:BUS:DEVICE.FUNCTION/current_link_width
cat /sys/bus/pci/devices/DOMAIN:BUS:DEVICE.FUNCTION/max_link_speed
cat /sys/bus/pci/devices/DOMAIN:BUS:DEVICE.FUNCTION/max_link_width
```

Attribute availability and permissions must be checked on the node. Reading these files is different from writing PCI sysfs action interfaces, which can change device state. ([docs.kernel.org](#)) Likewise, `setpci` can configure devices, so it is outside this default capture procedure. ([man7.org](#))

Take an idle snapshot, an active snapshot while a known workload transfers data, and a final snapshot after it finishes. Record the observation times and workload state. An idle **Gen1** reading is not sufficient to call a GPU “stuck at Gen1”: NVIDIA explicitly allows reduced current generation and width when unused.

(docs.nvidia.com) A persistent loaded-state mismatch deserves further investigation only after the expected configuration has been established.

GT/s means gigatransfers per second, and **x8** or **x16** describes lane width, not generation. The underlying standard rates and lane arithmetic are separated in the data section. (pcisig.com) A capability maximum, a current negotiated value and a measured transfer rate should never occupy the same unlabeled worksheet cell.

Decode AER Severity and Preserve the Reporting Path

Capture before recovery changes the record

AER is a PCIe error-reporting mechanism, not a GPU-only log. Red Hat documents its kernel-log collection and optional rasdaemon database storage. (docs.redhat.com) For a Linux node, preserve kernel records before rebooting, resetting the device or taking physical service action. The journal may be volatile, and retained boot history varies with configuration. (SUSE: [volatile journal](#); [boot history](#))

The journal commands below select the current kernel log, list retained boots, and request the previous boot when available. ([journalctl: kernel messages](#); SUSE: [boot selection](#); [journalctl: previous boot](#))

```
sudo journalctl -k -b --no-pager > kernel-current-boot.txt
journalctl --list-boots
sudo journalctl -k -b -1 --no-pager > kernel-previous-boot.txt
sudo dmesg > dmesg-capture.txt
```

dmesg and journalctl -k are documented AER evidence locations. (access.redhat.com) Save full records, then create a filtered working copy. Keep surrounding messages, error status and mask values, the reporting BDF, severity, and workload timestamps. Filtering alone can remove the sequence needed to understand recovery.

Table 1 decodes the Linux severity classes. The action column is the report's operating recommendation, not a universal replacement rule.

Class	Documented meaning	Recommended evidence response
Correctable	Hardware corrects the condition without software recovery or data loss. (docs.kernel.org)	Preserve counts and timestamps; compare behavior during matched idle and transfer intervals.
Uncorrectable, non-fatal	The transaction is unreliable while the link remains functional; recovery need not always include a link reset. (docs.kernel.org)	Save the full event and recovery sequence; evaluate application impact with the platform owner.
Uncorrectable, fatal	The link is unreliable; Linux's fatal recovery path does not let drivers opt out of reset. (docs.kernel.org)	Stop adding diagnostic workload, preserve available evidence, and follow the approved recovery process.

Severity determines the recovery category. It does not establish which replaceable component is responsible, nor does a corrected event prove that measured performance meets requirements.

Interpret counters without equating different mechanisms

Requester ID in the kernel example identifies the device sending the error message to the root port. Treating it as a physical component diagnosis would go beyond that definition. (docs.kernel.org) Linux explicitly notes that an endpoint problem can increment its link partner's AER counters while the endpoint counters remain unchanged. (docs.kernel.org) Capture both sides and the upstream reporting path.

Read the AER files where available, retaining each device's address:

```
cat /sys/bus/pci/devices/DOMAIN:BUS:DEVICE.FUNCTION/aer_dev_correctable
cat /sys/bus/pci/devices/DOMAIN:BUS:DEVICE.FUNCTION/aer_dev_nonfatal
cat /sys/bus/pci/devices/DOMAIN:BUS:DEVICE.FUNCTION/aer_dev_fatal
```

The ABI documents those device files and separate root-port totals. (docs.kernel.org) `TOTAL_ERR_COR` need not equal the sum of category counts because an error message may report multiple conditions. (docs.kernel.org) A root-port correctable total counts reported error messages, not a GPU's PCIe replay metric. (docs.kernel.org) These are important distinctions when a log, AER counter and vendor counter appear to disagree.

BadTLP and **BadDLLP** refer to transaction-layer packets (**TLPs**) and data-link-layer packets (**DLLPs**), respectively. **ReplayTimerTimeout** and **ReplayNumRollover** are specific error indications, not synonyms for every successful replay. Microsoft's register definitions distinguish them. (Microsoft: [packet errors; replay timeout](#)) Windows Hardware Error Architecture (**WHEA**) also recognizes PCIe root-port error sources, but its event collection and recovery procedures should remain separate from this Linux runbook. (learn.microsoft.com)

Finally, an empty Linux AER log is not a comprehensive physical-health certificate. Kernel support and firmware ownership matter: Linux requires the relevant configuration and operating-system control through the Advanced Configuration and Power Interface (**ACPI**) `_OSC` operating-system capabilities method. (docs.kernel.org) Check coverage before interpreting absence. On systems already running `rasdaemon`, `ras-mc-ctl --summary` and `ras-mc-ctl --errors` can supplement the preserved record. (docs.redhat.com)

Trend PCIe Replays Without Clearing the Evidence

Correlate GPU identity and field identity

Start with the full `nvidia-smi -q` output captured earlier. Preserve its PCIe replay reading, if supported, with the GPU BDF and timestamp; **N/A** is an unsupported or unavailable observation, not a zero. (docs.nvidia.com) For longer observations, **Data Center GPU Manager (DCGM)** provides finite monitoring through `dcmgi dmon`, including a locally discoverable field catalogue. (DCGM `dmon`: [field catalogue](#); [sample interval](#))

dcgmi discovery -l lists active GPUs and supplies identity information; use it to map a DCGM GPU identifier to the recorded UUID and BDF. (docs.nvidia.com) The field example below must be checked against dcgmi dmon -l on the installed release.

```
dcgmi discovery -l
dcgmi dmon -l
dcgmi dmon -i gpu:GPU_ID -e 202,235,236,237,238 -d 1000 -c 60
```

In the verified catalogue, **field 202** is the PCIe replay total; **235/236** describe maximum generation/width and **237/238** current generation/width. (DCGM: [field catalogue](https://docs.nvidia.com); [numeric identifiers](https://docs.nvidia.com)) The example's sample interval and count are observation settings, not acceptance thresholds. -d specifies milliseconds and -c bounds the output count. (docs.nvidia.com) Numeric entity identifiers belong to the running host engine, so retain the UUID/BDF mapping with the output. (docs.nvidia.com)

Compare defined intervals, not unexplained totals

Use **replay delta = end count minus start count** only when both samples refer to the same GPU and an uninterrupted counter history. Retain collection timestamps, because DCGM can present retained samples rather than new hardware queries. (docs.nvidia.com) If a count decreases, identity changes, a restart occurs, or data becomes unavailable, mark that interval discontinuous rather than computing a reassuring negative rate.

The proposed observation record should include:

- **Baseline:** the initial total, its collection time and the GPU identity.
- **Final sample:** the final total and collection time for the same observation interval.
- **Workload:** idle, host-to-device, device-to-host, or an identified peer transfer.
- **Link state:** current generation and width during the interval, not only before it.
- **Error context:** contemporaneous endpoint and upstream AER evidence.
- **Continuity:** driver reloads, resets, reboots, host-engine changes or unsupported samples.

These are proposed recordkeeping fields. They do not define a universal healthy replay rate. A counter delta observed during a specific transfer is more interpretable than a lifetime total with unknown history.

For **GPU PCIe replay counter increasing**, the next question is reproducibility: does growth recur on the same path under matched conditions, and does it coincide with other evidence? NVIDIA describes replays as transmission-problem indicators. (docs.nvidia.com) Its PCIe diagnostic replay parameter is scoped to the plugin run and GPU, so importing it as a production cutoff would change its meaning. (docs.nvidia.com)

Cross-vendor observations also need explicit definitions. AMD separates total replays from replay rollover counts, and labels instantaneous bandwidth differently from link speed. (AMD SMI: [replay count](https://docs.amd.com); [bandwidth units](https://docs.amd.com)) Preserve those units and names rather than combining unlike metrics. The decision should rest on a documented local baseline and OEM guidance, with unresolved differences left visible.

Use Controlled Transfers to Test the Hypothesis

A transfer test should answer a specific question: does the observed path deliver repeatable transfers under known conditions, and does the error evidence change while it runs? It should follow evidence capture and use scheduler-reserved resources. This procedure does not prescribe a stress duration or replace the platform's approved acceptance suite.

Select a current, bounded host-device test

NVIDIA's official CUDA Samples changelog removed **bandwidthTest in CUDA 12.9** and directs platform bandwidth testing to **NVBandwidth**. (github.com) Some CUDA guidance still discusses the older sample, so an operational recipe should follow the installed tool and current repository instructions rather than assume that historical sample exists. CUDA Samples also warns that its examples are not comprehensive validation or performance benchmarks. (github.com)

For an approved NVBandwidth installation on a reserved node, list available cases and select host-device transfer directions:

```
nvbandwidth -l
nvbandwidth -t host_to_device_memcpy_ce device_to_host_memcpy_ce
```

The repository documents case selection and configurable buffer size and sample count. (github.com) Record the exact executable version or commit, visible GPUs, complete command, options and output. The example selects test cases, not a promise of single-GPU scope: current source enumerates all visible GPUs in single-node mode, and its pair filter does not restrict host-only cases. (**NVBandwidth: visible GPUs; pair-filter scope**) Its default measurement process uses **three repetitions** and reports the median, which should be retained as a tool-specific method rather than mistaken for a platform qualification rule. (github.com)

Host allocation matters. NVIDIA identifies pinned, or page-locked, host memory as the transfer mode that attains the highest bandwidth. (docs.nvidia.com) Pageable memory introduces staging through pinned memory, so those modes are not interchangeable observations. (developer.nvidia.com) Use the following proposed checklist for every approved transfer:

- **Scope:** record participating GPUs and the application's visible-device set.
- **Direction:** distinguish host-to-device, device-to-host and peer results.
- **Allocation:** retain the test's pinned or pageable memory mode.
- **Transfer settings:** retain buffer sizes and configured repetition counts.
- **Host context:** retain placement, competing work and relevant memory configuration.
- **Result:** save the complete output and its reported units and aggregation method.

Non-uniform memory access (**NUMA**) information belongs in the evidence record as test context; this article does not prescribe an [affinity tuning procedure](#).

Interpret a slow result cautiously. Link capacity is one constraint, while packet size, buffering and latency also affect achieved throughput. (www.synopsys.com) Compare only like-for-like tool configurations on matched systems. Capture active link state and replay/AER changes around the same run, rather than using an

unrelated utilization snapshot.

Treat peer tests as path-specific evidence

Peer-to-peer (P2P) results require a recorded GPU pair and transport. NVIDIA's PCIe diagnostic can use **NVLink** between GPUs when available, so a successful peer transfer is not automatically proof of host PCIe performance. (docs.nvidia.com) NVBandwidth supports restricting applicable peer tests to a specified UUID pair; check the installed list of cases before selecting one. (github.com)

Reserve participating resources and avoid unrelated transfers. NVIDIA's plugin documentation specifically calls for idle GPUs so competing work does not distort measurements. (docs.nvidia.com) This means idle before the test reserves the devices, not that the link should stay in an idle power state while the transfer runs.

Do not add a long diagnostic suite merely because the node is slow. DCGM warns that longer suites can reset state or conflict with applications, and a named diagnostic can include software checks beyond the named test. (DCGM: [longer suites](#); [named-test software checks](#)) Such procedures need their own approved scope. A host bandwidth measurement also is not a PCIe compliance certification: formal compliance covers physical-layer, data-link and interoperability behavior. (www.keysight.com)

Isolate Safely and Prepare the OEM Evidence Bundle

The proposed isolation order preserves a causal record: establish expectations, reproduce observations, change a documented variable, and repeat the same measurement. It is an analytical method subject to the original equipment manufacturer's (**OEM's**) approved configuration, not a universal service instruction.

- **Confirm workload conditions:** repeat the observation with known transfer direction and controlled competing work before changing hardware. (docs.nvidia.com)
- **Confirm population and configuration:** verify processor, riser, distribution board and slot rules against the exact model. (www.dell.com) (pubs.lenovo.com)
- **Review software and firmware:** compare installed versions with the approved platform stack; Lenovo notes that component replacement can require firmware or driver updates. (pubs.lenovo.com)
- **Preserve the initial record:** export logs and measurements before a reboot, reset or service step; volatile journals may be lost. (documentation.suse.com)
- **Plan physical isolation:** use the model's service procedure and approved components, recording the initial slot and cable arrangement. (pubs.lenovo.com) (www.supermicro.com)
- **Repeat a matched observation:** after an approved change, collect the same link, error and transfer evidence and describe exactly what changed.

Physical service is a separate maintenance operation. Dell states that R760xa expansion slots are **not hot-swappable**. (www.dell.com) Supermicro's SYS-212GT-DNAF GPU procedure similarly starts with an operating-system shutdown and power-cord disconnection. (www.supermicro.com) Neither document authorizes live reseating merely because the diagnostic commands were read-only.

Riser and cable work must follow routing details. Dell's R760xa GPU riser procedure disconnects specified power, signal and other cables, illustrating why “reseat the riser” is incomplete guidance. (www.dell.com) Supermicro names **CBL-PWEX-1316-30** for the H100 power connection in the cited SYS-212GT-DNAF configuration. (www.supermicro.com) That is a model-specific rule, not a generally interchangeable GPU cable recommendation.

Table 2 is an escalation worksheet. Fill one record per GPU and observation interval, attaching raw files rather than replacing them with a verdict.

Evidence field	What to retain	Why it matters
Identity and location	GPU model, UUID, full BDF, chassis, slot, riser and installed software.	Device selection must remain unambiguous. (manpages.ubuntu.com)
NUMA and root path	Reported NUMA node or unknown value, root complex and intervening bridges/switches.	Upstream capacity can constrain the endpoint. (docs.kernel.org)
Expected capability	OEM population rule and approved configuration, with document revision.	A documented x8 configuration may be intentional. (pubs.lenovo.com)
Negotiated state	Idle and active generation/width for endpoint and relevant upstream links.	Current and maximum readings answer different questions. (raw.githubusercontent.com)
AER status	Severity, reporting IDs, masks, device counters and surrounding kernel messages.	Reporting-path interpretation needs both link partners. (docs.kernel.org)
Replay delta	Start/end values, collection times, interval and continuity notes.	Sample time is part of the telemetry meaning. (docs.nvidia.com)
Workload state	Transfer direction, concurrency, allocation mode and reserved-resource scope.	Pinned and pageable observations differ. (developer.nvidia.com)
Measured test	Exact command, tool version/commit, buffers, repetitions, units and full output.	Tool settings belong to the measurement method. (github.com)
OEM action	Case reference, approved change, pre/post record and support bundle.	Dell documents locally exportable SupportAssist support collection. (www.dell.com)

The worksheet links a symptom to an observable path and a repeatable intervention. It should make “expected,” “observed,” and “proposed” distinguishable even when the final cause remains unknown.

For Dell systems, the Integrated Dell Remote Access Controller (**iDRAC**) can save a **Technical Support Report (TSR)** locally; operating-system and application data require the running iDRAC Service Module. ([Dell: support report; OS and application data](#)) Check bundle completeness rather than assuming the management controller included the same logs preserved from Linux.

Stop adding test workload after a fatal event, during active recovery, or when the platform owner cannot provide an approved test/service scope. The fatal-link interpretation is documented; the conservative stop recommendation is this report's procedure. (docs.kernel.org) [DCGM diagnostics](#) do not replace NVIDIA's offline field diagnostics or determine return eligibility. (docs.nvidia.com) The evidence bundle should support the OEM decision rather than pre-empt it.

Data Analysis and Evidence

Calculate capacity with explicit units and assumptions

Table 3 presents **per-direction** theoretical ceilings, calculated for this report as of **October 9, 2026**. The raw signaling inputs come from PCI-SIG; AMD's documented encoding taxonomy supplies the encoding-only conversion. (pcisig.com) (docs.amd.com) (docs.amd.com) **GB/s** uses decimal gigabytes. Values are calculations, not measured accelerator performance.

Generation	Rate per lane, GT/s	Encoding	Raw x1, GB/s	Raw x8, GB/s	Raw x16, GB/s	x16 after encoding only, GB/s
Gen1	2.5	8b/10b	0.3125	2.5	5	4.000 (pcisig.com) (docs.amd.com)
Gen2	5	8b/10b	0.625	5	10	8.000 (pcisig.com) (docs.amd.com)
Gen3	8	128b/130b	1	8	16	15.754 (pcisig.com) (docs.amd.com)
Gen4	16	128b/130b	2	16	32	31.508 (pcisig.com) (docs.amd.com)
Gen5	32	128b/130b	4	32	64	63.015 (pcisig.com) (docs.amd.com)

The table explains what a generation or width change does to a capacity ceiling. It does not supply a minimum acceptable benchmark rate: protocol overhead and the measurement conditions remain outside the encoding-only column.

The calculation is **raw GB/s = rate per lane × lane count ÷ 8**. Apply **8/10** for Gen1/Gen2 or **128/130** for Gen3/Gen4/Gen5 to obtain the last column. (pcisig.com) (docs.amd.com) (docs.amd.com) PCIe offers bandwidth simultaneously in both directions; a one-way host-to-device result belongs against the one-way ceiling, not a combined bidirectional advertisement. (pcisig.com)

For example, calculated **Gen4 x8** and **Gen3 x16** both yield **16 GB/s raw** and approximately **15.754 GB/s after encoding** in one direction. (pcisig.com) (docs.amd.com) This arithmetic explains why a bandwidth result alone cannot identify the negotiated generation and width. Conversely, an expected Gen5 x16 capability does not guarantee an application achieves the encoding-only maximum.

Encoding factors are mapped by generation: AMD documents 8b/10b for Gen1/Gen2 and 128B/130B for Gen3/Gen4/Gen5. (docs.amd.com) (docs.amd.com) PCI-SIG separately explains the byte-to-10-bit mapping and the more efficient encoding introduced with PCIe 3.0. ([PCI-SIG: 8b/10b](https://pcisig.com); [128b/130b](https://pcisig.com)) These references establish encoding behavior; they do not establish the capabilities of an installed GPU or processor.

Keep newer-generation calculations separate

PCIe 6.0 changes the model: PCI-SIG specifies **64 GT/s** and **256 GB/s maximum bidirectional x16 bandwidth**, together with fixed-size flow control units (**FLITs**), pulse-amplitude modulation with four levels (**PAM4**) and forward error correction (**FEC**). (PCI-SIG: [PCIe 6.0 bandwidth](#); [FLIT encoding](#)) Synopsys explains the fixed-packet requirement for FEC and the changed FLIT encoding. ([www.synopsys.com](#)) ([www.synopsys.com](#))

PCI-SIG describes **236 TLP bytes within a 256-byte FLIT** and notes that an enabled FLIT link retains that mode at lower speeds. ([pcisig.com](#)) ([pcisig.com](#)) Transaction-layer packets (**TLPs**) and error-control overhead therefore need their own accounting. Applying the Gen3-to-Gen5 multiplier indiscriminately would produce a misleading payload estimate.

Implications and Future Directions

The operational benefit of this evidence ladder is a narrower, reproducible escalation. A capability mismatch, a negotiated-state mismatch, an error trend and a transfer limitation are different findings. Their combination may justify an approved isolation step; their separation may show that the physical link has not yet been implicated.

Make the evidence record durable enough to compare before and after maintenance. The proposed process should retain raw outputs, configuration revisions and the assumptions behind each calculation. GPU Smith's stated documentation scope includes architecture records, port maps, operating procedures and acceptance reports; those records provide useful context for this kind of independent assessment. ([gpusmith.com](#)) No consulting relationship or delivered outcome is assumed.

Future platforms will require generation-aware interpretation. PCI-SIG's **PCIe 7.0** specification identifies **128 GT/s** and an official member release on **June 11, 2025**; a specification's publication does not establish that a particular installed GPU path supports it. ([pcisig.com](#)) Device, server and firmware capabilities still need direct verification.

Likewise, replay observations should be interpreted with the appropriate link mode. Synopsys describes FLIT replay following a failed cyclic redundancy check (**CRC**). ([www.synopsys.com](#)) That mechanism helps explain the role of recovery telemetry, but does not turn a count into an application-impact measurement or a universal service threshold.

When basic observations cannot resolve the path, specialist protocol capture may be appropriate under an OEM-approved plan. Keysight distinguishes an analyzer that captures exchanges from an exerciser that generates protocol activity. ([www.keysight.com](#)) These are different interventions with different scopes, beyond the default read-only ladder. The escalation record should explain what remains uncertain so that deeper testing answers a specific question.

Frequently Asked Questions (FAQs)

How to check GPU PCIe link speed and width?

Record the full BDF, inspect `lspci -D -s DOMAIN:BUS:DEVICE.FUNCTION -vv`, and compare `LnkCap` with `LnkSta`. Those fields decode capability and negotiated status separately. (raw.githubusercontent.com) Repeat during a recorded active workload before deciding that an idle reading is degraded.

Why is GPU PCIe link width lower than expected?

First verify the installed slot/riser configuration. Lenovo documents intentional x8 operation with direct GPU distribution boards on the SR675 V3. (pubs.lenovo.com) A GPU's advertised maximum is not a substitute for that configuration record.

What if GPU PCIe link speed is stuck at Gen1?

Capture it under an active transfer as well as idle, then inspect the upstream path. NVIDIA permits reduced generation and width while unused. (docs.nvidia.com) "Stuck" should describe a repeated active-state mismatch against the installed platform expectation.

Are PCIe correctable errors on NVIDIA GPUs a failure verdict?

No universal replacement verdict follows from the class. Linux defines correctable conditions as hardware-corrected without software recovery. (docs.kernel.org) Preserve their trend and reporting path, then compare with workload and transfer evidence.

What does an increasing nvidia-smi PCIe replay counter mean?

Keep a timestamped delta for a continuous observation interval, correlate with AER, and check an equivalent baseline. AER replay-timeout and rollover indications have specific meanings and should not be equated with every replay total. (learn.microsoft.com) Use the installed tool's documented field definitions rather than an invented cutoff.

Conclusion

GPU PCIe link troubleshooting is most useful when it produces a verifiable sequence of observations. Establish the installed configuration's expected capability, preserve device identity, inspect the endpoint and upstream negotiated states during idle and active conditions, and keep AER severity separate from vendor replay totals.

The next step should follow the evidence. A documented narrower configuration calls for correcting the expectation. A repeated active-state mismatch calls for configuration and path checks. Correlated error growth calls for preserving the reporting sequence and planning controlled isolation. An expected link with low measured throughput calls for examining test conditions and upstream sharing before selecting a replacement part.

Raw and encoding-only bandwidth calculations explain capacity changes; they do not certify useful transfer performance. Controlled host-device and peer tests must retain direction, allocation mode, participating GPUs, transport, tool settings and error deltas. Physical service and longer diagnostic suites belong to their own approved maintenance scope.

The final deliverable is an evidence bundle that another engineer can review: what was expected, what was observed, under which workload, what changed, and what remains uncertain. That record supports a more precise OEM escalation and a defensible decision about the next intervention, without inventing replay thresholds or presenting an unmeasured hardware diagnosis as a conclusion.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://raw.githubusercontent.com/pciutils/pciutils/master/lscaps.c>
2. <https://docs.nvidia.com/deploy/nvidia-smi/index.html>
3. <https://pcisig.com/what-bit-rates-does-pcie-50-specification-support-and-how-does-it-compare-prior-pcie-generations>
4. <https://docs.amd.com/r/en-US/am002-versal-gty-transceivers/Features>
5. <https://lenovopress.lenovo.com/lp1611-thinksystem-sr675-v3-server>
6. <https://docs.kernel.org/PCI/pcieaer-howto.html>
7. <https://docs.kernel.org/admin-guide/abi-testing-files.html>
8. <https://documentation.suse.com/sles/15-SP6/html/SLES-all/cha-journalctl.html>
9. <https://docs.nvidia.com/datacenter/dcgm/latest/reference/diagnostics/plugins/pcie.html>
10. <https://www.keysight.com/blogs/en/tech/educ/protocol-analyzer>
11. <https://docs.kernel.org/driver-api/pci/pci.html>
12. <https://mj.ucw.cz/sw/pciutils/>
13. <https://rocmdocs.amd.com/projects/amdsmi/en/latest/how-to/amdsmi-cli-tool.html>
14. <https://gpusmith.com/>
15. <https://www.hpe.com/psnow/downloadDoc/HPE%20ProLiant%20DL380%20Gen11%20QuickSpecs-a50004307enw.pdf?deepLink=&form=false&hf=regular&id=a50004307enw.pdf&isFutureVersion=true&originalObjectName=&prelaunchSection=&preview=false&print=&r=§ion=Product+Documentation&softrollSection=&ver=43>
16. https://www.dell.com/support/manuals/en-us/poweredge-r760xa/r760xa_ism_pub/expansion-card-riser-specifications?guid=guid-e4f2c1d5-1a15-4a17-ac9b-814c3a8bd6cc&lang=en-us
17. <https://docs.nvidia.com/datacenter/dcgm/latest/learn/getting-started-for-system-administrators/index.html>
18. https://www.dell.com/support/manuals/en-us/poweredge-r760xa/r760xa_ism_pub/expansion-card-installation-guidelines?guid=guid-ac875fab-425e-40d6-9f84-e1ad4c8784b1&lang=en-us
19. <https://manpages.ubuntu.com/manpages/noble/man8/lspci.8.html>
20. https://pubs.lenovo.com/sr675-v3/install_a_pcie_riser
21. https://pubs.lenovo.com/sr675-v3/server_specifications_technical
22. <https://docs.nvidia.com/fleet-intel/data-collection-rationale/>
23. <https://manpages.ubuntu.com/manpages/stonking/man8/lspci.8.html>

24. <https://man7.org/linux/man-pages/man8/lspci.8.html>
25. <https://raw.githubusercontent.com/torvalds/linux/master/drivers/pci/pci-sysfs.c>
26. <https://docs.kernel.org/PCI/sysfs-pci.html>
27. <https://man7.org/linux/man-pages/man8/setpci.8.html>
28. https://docs.redhat.com/en/documentation/red_hat_enterprise_linux/8/html/managing_monitoring_and_updating_the_kernel/assembly_advanced-error-reporting_managing-monitoring-and-updating-the-kernel
29. <https://man7.org/linux/man-pages/man1/journalctl.1.html>
30. <https://access.redhat.com/solutions/7128107>
31. <https://docs.kernel.org/admin-guide/abi-testing.html>
32. https://learn.microsoft.com/en-us/windows-hardware/drivers/ddi/miniport/ns-miniport-_pci_express_correctable_error_status
33. <https://learn.microsoft.com/en-us/windows-hardware/drivers/whea/hardware-errors-and-error-sources>
34. <https://docs.nvidia.com/datacenter/dcgm/latest/reference/command-line-reference/dcgmi/dcgmi-dmon.html>
35. <https://docs.nvidia.com/datacenter/dcgm/latest/user-guide/feature-overview.html>
36. <https://docs.nvidia.com/datacenter/dcgm/latest/reference/field-identifiers.html>
37. <https://rocmdocs.amd.com/projects/amdsmi/en/latest/reference/amdsmi-cpp-api/types.html>
38. <https://github.com/NVIDIA/cuda-samples/blob/master/CHANGELOG.md>
39. <https://github.com/NVIDIA/cuda-samples>
40. <https://github.com/NVIDIA/nvbandwidth>
41. <https://github.com/NVIDIA/nvbandwidth/blob/main/nvbandwidth.cpp>
42. <https://docs.nvidia.com/cuda/cuda-c-best-practices-guide/index.html>
43. <https://developer.nvidia.com/blog/how-optimize-data-transfers-cuda-cc/>
44. <https://gpusmith.com/articles/en/gpu-cpu-nic-affinity-rank-binding>
45. <https://www.synopsys.com/articles/designing-effective-use-pcie6-bandwidth.html>
46. <https://docs.nvidia.com/datacenter/dcgm/latest/learn/modules/dcgm-diagnostics.html>
47. <https://www.keysight.com/blogs/en/tech/bench/2022/04/21/pcie-60-the-next-evolution-of-pcie>
48. <https://www.supermicro.com/en/support/manuals/product/system/sys-212gt-dnaf/Content/component-installation/expansion-cards/gpu-installation-sys-212gt-dnaf.htm>
49. https://www.dell.com/support/manuals/en-us/poweredge-r760xa/r760xa_ism_pub/removing-the-gpu-riser-module?guid=guid-03d32433-8317-4eb2-a81d-bab2eb7317c9&lang=en-us
50. <https://www.supermicro.com/en/support/manuals/product/system/sys-212gt-dnaf/Content/component-installation/cable-routing/system-cable-routing-h100-gpu-12vhpwr-sys-212gt-dnaf.htm>
51. <https://www.dell.com/support/kbdoc/en-us/000126308/export-a-supportassist-collection-via-idrac9>
52. <https://gpusmith.com/articles/en/gpu-cluster-acceptance-testing-runbook>
53. <https://docs.amd.com/r/en-US/pg343-pcie-versal/Status-Signals>
54. https://pcisig.com/faq?field_category_value%5B%5D=pci_express_3.0
55. <https://pcisig.com/blog/pcie%C2%AE-60-specification-released-members-double-bandwidth-next-generation-applications>
56. <https://www.synopsys.com/articles/pcie-6-designs.html>

- 57. <https://pcisig.com/blog/pcie%20AE-60-specification-webinar-qa-deeper-dive-flit-mode-pam4-and-forward-error-correction-fec>
- 58. <https://pcisig.com/specifications/pcie-70-specification-version-03-now-available-members>
- 59. <https://www.synopsys.com/blogs/chip-design/pcie-6-verification-fec-crc.html>

THE PUBLISHER

About GPU Smith

GPU Smith provides independent engineering and research for private AI infrastructure. We help technical buyers and operators reason about GPU workload sizing, hardware selection, procurement, deployment and inference operations, from the compute system to the networking, power and cooling requirements around it.

Start with the workload

Useful infrastructure decisions begin with the models, data, throughput, latency and operating constraints a team actually has. GPU Smith helps connect those requirements with system design and deployment choices. The goal is a privately controlled AI environment whose capacity and operational demands are understood before a purchasing decision.

Hardware and supplier research

Our public reference library covers GPUs, complete systems and networking components. The server-vendor directory and manufacturing research help teams investigate suppliers, compare options and follow sources behind technical claims. We distinguish manufacturer specifications from measured benchmarks and advertised capabilities from tested configurations. Reference pages are research resources, not a live inventory listing or a binding equipment quote.

Deployment and operations

GPU Smith's engineering scope includes infrastructure integration, networking, power, cooling and the practical operation of inference workloads. Our published articles explain the tradeoffs behind deployment, procurement and ongoing operations so teams can ask better questions and document decisions.

Work with GPU Smith

Explore the [hardware library](#), [GPU references](#), [networking references](#) and [vendor research](#). [Contact GPU Smith](#) with your workloads and deployment constraints to discuss a project and confirm current scope, pricing and availability.

Inclusion of a manufacturer or product in our research does not imply a partnership, certification or endorsement.

Website: <https://gpusmith.com>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.