



GPUDirect Storage Requirements and Readiness Matrix

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 · gpudirect storage requirements · gpudirect storage compatibility · nvidia · cufire · gdscheck · gpu storage · rdma · nvme · linux kernel

Original article: <https://gpusmith.com/articles/en/gpudirect-storage-readiness-compatibility-matrix>



In this article

1. Executive Summary
 2. Introduction and Background
 3. Key Changes
 4. The End-to-End GPUDirect Storage Compatibility Ledger
 5. Implementation Considerations and Process Changes
 6. Qualification and Runtime Path Evidence
 7. Fallback and Failure Modes
 8. Data Analysis and Evidence
 9. Implications and Future Directions
 10. Conclusion
-

Executive Summary

GPUDirect Storage requirements form a support chain, not a single GPU or filesystem specification. **NVIDIA GPUDirect Storage (GDS)** provides a direct-memory-access (DMA) path from storage to graphics processing unit (GPU) memory that can avoid a bounce buffer in central processing unit (CPU) memory. Eligibility, documented support, and observed runtime behavior must be assessed separately. A suitable device does not establish the whole path: Linux peer-to-peer DMA also requires cooperating drivers. (docs.nvidia.com) (docs.kernel.org)

The purchasing rule proposed here is strict: call a deployment **supported** only when the exact component combination has applicable vendor documentation and an identified support owner. Generic integration statements belong in a preliminary shortlist. **NetApp ONTAP 9.10.1** introduced its documented GDS-enabling Network File System (NFS) over remote direct memory access (RDMA) capability, but protocol-specific release selection remains necessary. (docs.netapp.com) (docs.netapp.com) **BeeGFS** documents a CUDA compatibility boundary at **CUDA 12.2 release 2** and requires RDMA throughout its deployment. (doc.beegfs.io) (doc.beegfs.io) Those are different kinds of prerequisites, and neither certifies an arbitrary server build.

The ledger must identify the intended direct mechanism. NVIDIA documents Non-Volatile Memory Express (NVMe) Peripheral Component Interconnect Express (PCIe) peer-to-peer DMA, or **P2PDMA**, beginning with **CUDA 12.8**. Its overview specifies **OpenRM (open GPU kernel module) 570.x or later** for that eligible path without the traditional nvidia-fs dependency. (docs.nvidia.com) A distributed-filesystem deployment can have different driver and network requirements. **VAST** explicitly requires its client to match the selected OpenFabrics Enterprise Distribution (OFED) and exact operating system (OS)/kernel. (kb.vastdata.com) **Amazon Web Services (AWS) FSx for Lustre** adds service-specific requirements, including selecting Elastic Fabric Adapter (EFA) support when creating the filesystem. (docs.aws.amazon.com)

After installation, acceptance should establish input/output (I/O) correctness and path selection. A working application is insufficient: **WEKA** documents automatic fallback when acceleration requirements are unmet, and BeeGFS can transparently convert writes to Portable Operating System Interface (POSIX) I/O when GDS writes are disabled. (docs.weka.io) (doc.beegfs.io) Capture effective configuration, per-GPU activity, and mechanism-appropriate counters alongside the support records. IBM's Storage Scale FAQ Q2.39 states that

asynchronous CUDA I/O is not supported. ([IBM Storage Scale FAQ, Q2.39](#)) This report provides a dependency ledger, procurement worksheet, and change-impact matrix, not a benchmark or deployment certificate. Its publication anchor is **October 6, 2026**; live documentation was checked **October 15, 2026**. Undated live requirements are labeled accordingly, rather than backdated to publication.

Introduction and Background

The practical question is whether a particular application, on a particular mount and GPU, can use a documented direct storage path. A proposal that merely lists GDS-capable components leaves the integration decision unresolved. Buyers should require evidence connecting each component to the next, then preserve that evidence with the delivered configuration.

GDS concerns storage transfers; **GPUDirect RDMA** concerns the direct path between a network interface card (NIC) and GPU memory. The CPU still prepares and controls communication. Avoiding a CPU payload copy should therefore not be described as eliminating CPU involvement. ([docs.nvidia.com](#)) This distinction matters when an evaluation demonstrates network GPU communication but has not qualified the filesystem client.

An operator example makes the distinction concrete. **St. Jude Children's Research Hospital** describes separate compute and storage InfiniBand networks, with the storage network attached to a DDN AICD2 device and enabling GDS. Its resource page documents architecture, not a reproducible compatibility certificate or a GDS performance result. ([www.stjude.org](#)) ([www.stjude.org](#)) That is useful deployment context; the acceptance evidence for another cluster must come from that cluster.

The report adopts an engineering-advisor perspective. **GPU Smith** describes written acceptance criteria and as-built documentation as part of its infrastructure work; those concepts motivate the worksheet here, without making the consultancy a storage-product option or implying that it tested the combinations discussed. ([gpusmith.com](#)) No lab stack, customer measurements, or original throughput results were supplied.

The scope is filesystem-based cuFile readiness, including local NVMe and remote storage. Vendor integration announcements establish where further qualification is worthwhile. For example, **DDN EXAScaler 6** explicitly announced GDS support, ([www.ddn.com](#)) and **Pure Storage ATRI** documentation describes a GDS-enabled architecture using FlashBlade//S and NVIDIA DGX. ([www.purestorage.com](#)) Neither statement substitutes for the release-specific evidence requested below.

Key Changes

Replace a universal minimum with a path-specific requirement

Record the intended path before choosing package versions. Keep separate entries for upstream NVMe P2PDMA, nvidia-fs integration, and the filesystem client's own build requirements. The selected vendor's client qualification belongs alongside the GPU-driver qualification: VAST explicitly calls for matching client-build inputs. ([kb.vastdata.com](#)) This prevents a minimum copied from one mechanism becoming a blanket purchasing rule.

The distinction also changes installation checks. An absent `nvidia-fs` module cannot, by itself, disqualify a deployment intentionally using the eligible upstream NVMe path. Conversely, a filesystem integration that requires `nvidia-fs` callbacks cannot be qualified merely because an NVMe-only example works. Record the selected mechanism next to every mount in the ledger.

Treat release histories as scoped evidence

NVIDIA's release notes retain broad GPU capability language, but their **GDS 1.15** history explicitly removes **Pascal and Volta** support. They also do not guarantee functional or performant GDS with an input-output memory management unit (IOMMU) or Access Control Services (ACS) enabled on x86-64. (docs.nvidia.com) The purchasing question must name the GPU and selected software release, rather than rely on a compute-capability threshold copied from another section. When documentation differs in scope, the resolution should be an applicable release statement from the responsible vendor.

Packaging has independent version tracks. The live `nvidia-fs` guide identifies **2.30** as its latest published filesystem driver, specifies open kernel driver **535 or later**, lists **CUDA Toolkit 12.2 or later** as a packaging/build prerequisite, and describes separate `nvidia-fs` distribution beginning with **CUDA 13.4**. The same filesystem-driver guide lists XFS, plus EXT4 in ordered mode, on NVMe and NVMe-over-Fabrics (NVMe-oF). (docs.nvidia.com) These are live-document statements, not a reconstructed publication-day installer manifest. Keep library, CUDA, and filesystem-driver releases in separate fields.

Distinguish feature availability from complete qualification

Feature support evolves within an integration. IBM's historical note distinguishes accelerated writes supported in **Spectrum Scale 5.1.6** from their earlier technical-preview status. (www.ibm.com) That history demonstrates why an old claim of "GDS support" cannot answer every current operation-level question.

Likewise, **VDURA's September 30, 2026** announcement establishes V12 availability, while its separately fetched whitepaper explicitly names GDS. (www.vdura.com) (www.vdura.com) NVIDIA's newest visible release-notes section is **GDS 1.19.1**, with its PanFS support row dated **September 2026**. (docs.nvidia.com) These support product-level statements; an exact maintenance-release support chain still needs confirmation. This report does not infer a complete qualification from either document.

The End-to-End GPUDirect Storage Compatibility Ledger

Use three labels throughout procurement and acceptance:

- **Architecturally possible:** The selected design could provide the intended direct path, with no identified architectural exclusion.
- **Vendor supported:** Applicable documentation covers the proposed combination, operations, and deployment mode, with a support owner.
- **Locally verified:** The installed application exercised the intended path on the named mounts and GPUs, with preserved evidence.

These labels are an analyst-proposed decision framework. They describe different evidence, not successive badges granted by one utility. Linux's documented restrictions on routing across PCIe hierarchy domains illustrate why compatible endpoint devices alone are insufficient. (docs.kernel.org)

Table 1 maps the support chain to the evidence a buyer should request. “Record” entries are worksheet requirements, not assertions that a listed component is universally supported.

Dependency	Record before purchase	Acceptance question and source boundary
GPU and GPU driver	Exact GPU, driver version, kernel-module flavor, server model	Does the chosen GDS release support that GPU, rather than an older family-level minimum? NVIDIA's release history changes architecture support. (docs.nvidia.com)
CUDA, cuFile, and GDS packages	Toolkit, library, tools, package origins, build identifiers	Can the installed manifest be reproduced? Keep AWS's named driver floor separate from a cuFile release number. (docs.aws.amazon.com)
nvidia-fs or NVMe P2PDMA	Selected mechanism per mount; module version if applicable	Does the vendor confirm the selected mechanism and matching client build? VAST's documented integration uses exact client-build inputs. (kb.vastdata.com)
Linux distribution and kernel	Full running kernel, package source, client-module build	Does the vendor qualify this exact kernel? VAST calls for exact OS/kernel matching; WEKA directs custom-kernel certification to its team. (kb.vastdata.com) (docs.weka.io)
Filesystem and client	Client/server versions, build options, mount configuration	Is GDS enabled in the client? BeeGFS documents client-module configuration and rebuilding, not just installing a server release. (doc.beegfs.io)
Storage software and protocol	Storage release, export, protocol version, enabled features	Does the protocol match the support statement? NetApp's NFS-over-RDMA support does not include Server Message Block (SMB) RDMA. (docs.netapp.com)
NIC, transport, and OFED	Adapter, firmware, network-driver distribution, RDMA path	Is remote RDMA enabled throughout? NetApp requires RDMA data interfaces; VAST requires RDMA end to end. (docs.netapp.com) (kb.vastdata.com)
NVMe and backend media	Local device versus remote storage, namespace and mount	Is NVMe actually the host transfer device? HPE's documented hybrid backend illustrates why backend media and host access path are separate questions. (www.hpe.com)
PCIe topology	GPU/NIC/NVMe placement, switches, root ports, NUMA affinity	Is the intended device pairing appropriate? Dell recommends grouping GPU and NIC placement by NUMA affinity. (infohub.delltechnologies.com)
IOMMU and ACS	Architecture-specific platform guidance and effective settings	Do not apply a universal platform-settings recipe. NVIDIA's Grace Hopper IOMMU guidance differs from bare-metal x86 assumptions. (docs.nvidia.com)
Container and runtime	Image digest, host modules, devices, mount visibility, privileges	Can the container access the qualified host path? NVIDIA says containers carry user-level components and rely on host-installed kernel components. (docs.nvidia.com)
cuFile configuration and application programming interface (API)	Effective configuration, environment overrides, operation types, alignment	Are writes and reads qualified separately? IBM's historical write explanation distinguishes compatibility copies from acceleration. (www.ibm.com)

The table makes missing evidence visible. A blank kernel-certification field is an unresolved support question, even when every component has a GDS logo. A positive read test should leave the write, asynchronous, and container fields unresolved until their own evidence exists. Procurement teams should preserve these unknowns rather than turn them into implied approval.

Supported filesystems are release-bound integrations

A filesystem name is too broad to serve as a compatibility matrix. NetApp's documentation requires selecting the ONTAP release appropriate to the NFS protocol version, and configuring RDMA data logical interfaces. Its network configuration includes restrictions that belong in the design review, not merely in post-installation troubleshooting. (docs.netapp.com) (docs.netapp.com) Both nodes in an NFS-over-RDMA high-availability pair must run the same version, and interface groups are unsupported. (docs.netapp.com)

WEKA's fetched networking documentation permits RDMA-capable InfiniBand or Ethernet interfaces and states that GDS is not supported for encrypted filesystems. That is a feature interaction worth resolving before an encryption requirement and a direct-I/O requirement enter the same purchase specification. The fetched pages are mutable and undated; their URL version label alone should not be treated as a dated qualification. (docs.weka.io) (docs.weka.io)

Dell's published PowerScale report is more concrete about its demonstrated stack: it records storage, operating-system, kernel, driver, and GDS versions. It also enables NFS over RDMA at the network-pool and global levels. The useful lesson is the specificity of that evidence, not that its historical versions define today's recommended build. (infohub.delltechnologies.com) (infohub.delltechnologies.com)

Implementation Considerations and Process Changes

Obtain support evidence before issuing the purchase order

A request for proposal should ask for a completed support chain, including the boundaries of each vendor's responsibility. Product integration is a starting point: DDN's EXAScaler announcement and Pure's AIRI brief establish integrations, while Pure explicitly directs readers elsewhere for full technical details. (www.ddn.com) (www.purestorage.com)

Use the following questions as acceptance-contract inputs:

- **Release identity:** Which exact storage, client, kernel, GPU-driver, cuFile, and tools releases are covered?
- **Transport identity:** Which adapter, firmware, network-driver stack, and remote-storage transport are qualified?
- **Operation scope:** Are reads, writes, synchronous calls, asynchronous calls, and batch calls covered separately?
- **Feature interaction:** Which encryption, snapshot, caching, compression, or data-layout settings affect the qualification?
- **Deployment scope:** Is coverage for bare metal, a named container runtime, or a particular orchestration configuration?

- **Support owner:** Who receives the initial case and coordinates a problem spanning storage and server components?
- **Evidence delivery:** Which configuration snapshots, correctness checks, and path observations will accompany acceptance?
- **Change policy:** Which updates invalidate the qualification, and how will the replacement stack be approved?

The feature-interaction question has a documented practical basis: WEKA excludes encrypted filesystems from its GDS support, while Dell explicitly records configuration choices in its published test setup. These are vendor- and context-specific conditions, not settings to disable indiscriminately on another product.

(docs.weka.io) (infohub.delltechnologies.com)

Table 2 turns those questions into a bill-of-materials evidence worksheet. The blank identifiers should be populated for the proposed deployment; no row certifies a vendor combination.

Worksheet item	Required artifact	Proposed acceptance disposition
Named storage service and filesystem	Release document, export/mount identity, configuration snapshot	Hold if the requested service mode was not provisioned. AWS cannot add EFA to an existing non-EFA filesystem. (docs.aws.amazon.com)
Server and topology	As-built device map and GPU/NIC pairing	Hold if actual placement differs from the reviewed map. Dell's topology guidance supplies a relevant mapping approach. (infohub.delltechnologies.com)
Client software manifest	Installed packages, running kernel, module-build records	Hold if the client build differs from qualified OS/kernel/OFED inputs. (kb.vastdata.com)
Application I/O contract	Reads/writes, buffer alignment, offsets, synchronous/async usage	Qualify each behavior separately. BeeGFS's documented block-alignment requirement is an operation-level constraint. (doc.beegfs.io)
Container execution contract	Host-to-container device, mount, permission, and runtime map	Repeat qualification in the delivered runtime; account for documented non-root RDMA configuration. (developers.redhat.com)
Support and evidence packet	Vendor coverage records, named owner, test outputs, exceptions	Accept only the evidenced scope. HPE explicitly directs buyers to obtain network-compatibility confirmation. (www.hpe.com)

The worksheet separates missing support evidence from a failed runtime test. Either can prevent acceptance, but they require different remedies. A vendor clarification resolves documentation coverage; a corrected configuration and rerun resolve a demonstrated path problem. Neither should erase the other from the record.

Qualify the network-driver stack separately

For Mellanox OFED (**MOFED**), NVIDIA DOCA networking software, or an inbox distribution driver, request the storage vendor's answer for the selected path. An RDMA-capable adapter or a GPUDirect RDMA recipe does not establish the filesystem-client qualification. VAST requires matching OFED, operating-system (OS), and

kernel inputs; NetApp separately requires the appropriate NFS/RDMA configuration. ([kb.vastdata.com](https://kb.vastdata.com/docs.netapp.com)) (docs.netapp.com)

The live NVIDIA pages contain a specific unresolved difference: the dedicated `nvidia-fs` guide excludes **DOCA 3.3.x** for NVMe, while release notes require **DOCA 3.3.0 or newer** for that local-NVMe mechanism on **kernel 6.17 or newer**. Neither statement should become a universal minimum. Obtain a reconciled, exact-release answer before procurement. (docs.nvidia.com) (docs.nvidia.com)

Draw the topology that the workload will actually use

The topology diagram should connect application-visible paths to physical devices. A generic server block diagram is insufficient for reviewing which GPU can reach which storage-facing adapter. Dell's guidance explicitly groups GPUs and NICs by non-uniform memory access (**NUMA**) affinity and uses `ibdev2netdev` to relate RDMA identifiers to interface names. (infohub.delltechnologies.com) (infohub.delltechnologies.com)

Include these fields in the diagram checklist:

- **GPU identity:** Application index, physical device identity, and NUMA association.
- **Storage identity:** Filesystem mount, export or namespace, and storage-facing interface.
- **PCIe route:** Endpoint, upstream switch, root-port relationship, and applicable platform restrictions.
- **RDMA identity:** Linux interface, RDMA device, physical port, and intended storage network.
- **Platform settings:** Recorded input-output memory management unit (IOMMU) and Access Control Services (ACS) policy.
- **Container mapping:** The mount and device identities seen inside the application runtime.

Use platform-specific guidance to evaluate that map. Linux documents how ACS affects routing inside a PCIe hierarchy; it also normally blocks unsupported routing between hierarchy domains. A topology review should therefore evaluate the installed platform rather than assume every [GPU-to-NIC pairing](#) is equivalent. (docs.kernel.org) (docs.kernel.org)

Qualification and Runtime Path Evidence

Separate installation, correctness, and path selection

The proposed acceptance procedure has separate outputs for exposed capability, correct I/O, and the observed mechanism. A platform report is necessary evidence about configuration; it is not an application trace. BeeGFS's own procedure starts with **gdscheck -p**, then tests a file using **gdscheck -f** and looks for successful read verification on each GPU. (doc.beegfs.io)

A [reproducible qualification sequence](#) is:

1. **Freeze the build:** Export the software manifest and capture the running configuration before testing.
2. **Inspect platform capability:** Run the installed `gdscheck` utility, preserving the complete output and timestamp.
3. **Exercise a test file:** Use the vendor-documented file check on the intended mount and named GPUs.
4. **Check correctness:** Run documented read/write verification with representative operation shapes.

5. **Observe the path:** Capture workload-specific cuFile and vendor counters during an active application run.
6. **Repeat in production context:** Use the intended identity, container, configuration, and mount visibility.

NVIDIA documents **gdscheck.py -p** under the installed toolkit's GDS tools directory and reports filesystem/device support and ACS state. Use the actual installation path rather than copying a toolkit version from an example. Its correctness tooling also has mode-dependent behavior: **gdsio -V** verification is ignored in timed **-T** mode. A throughput run should not be relabeled as a correctness check. (docs.nvidia.com) (docs.nvidia.com)

For an nvidia-fs integration, VAST recommends observing changing, nonzero counters in **/proc/driver/nvidia-fs/stats**. That observation must be scoped to that mechanism. NVIDIA explicitly excludes compatibility-mode and P2PDMA activity from those kernel statistics, so zero nvidia-fs activity alone cannot establish a CPU bounce path on an intended P2PDMA deployment. (kb.vastdata.com) (docs.nvidia.com)

Capture effective cuFile behavior

Archive the configuration the application actually used. A systemwide JavaScript Object Notation (JSON) file is not sufficient evidence when **CUFILE_ENV_PATH_JSON** selects application-specific parameters. The cuFile reference also documents compatibility fallback controls and operation-dependent behavior. (docs.nvidia.com) (docs.nvidia.com) The acceptance record should connect configuration to the running process, not simply to a package installation.

Useful distinctions include RDMA availability versus GDS use, and successful completion versus direct transfer. WEKA documents **weka cluster processes** for confirming RDMA usage; that establishes a network condition. Its documented fallback behavior means the same deployment still needs application-specific evidence for the storage path. (docs.weka.io) (docs.weka.io)

Retain these evidence-packet elements:

- **Configuration provenance:** Effective JSON, relevant environment settings, application build, and selected runtime.
- **Operation coverage:** Separate read/write and synchronous/asynchronous results with sizes and alignment recorded.
- **Path observations:** Per-process or per-GPU counters, vendor observations, and the observation interval.
- **Correctness outputs:** Test inputs, successful verification results, exit status, and any exclusions.
- **Support records:** Source URLs, release applicability, vendor confirmations, and unresolved exceptions.

For active-process inspection, NVIDIA documents **gds_stats -p PID -l 3** and per-GPU POSIX-operation counters. Preserve the associated configuration enabling the requested statistics. The exact installed tool's help output should accompany a reusable test recipe. (docs.nvidia.com) (docs.nvidia.com)

Fallback and Failure Modes

The most useful troubleshooting question is “which layer changed the intended path?” A functional workload can remain functional after that change. IBM's historical explanation of compatibility writes explicitly describes copying the GPU buffer into client system memory before the filesystem write. This explains the distinction without assuming that a particular current IBM deployment uses that route. (www.ibm.com)

Track these conditions in the exception register:

- **Compatibility policy:** A permissive policy may allow an operation to complete on another path; record whether fallback is acceptable.
- **Write asymmetry:** Read-path success does not establish accelerated writes. BeeGFS explicitly documents transparent POSIX writes under its write-support setting. (doc.beegfs.io)
- **Network prerequisites:** A mounted filesystem alone does not show the required remote transport. NetApp requires RDMA-enabled data interfaces. (docs.netapp.com)
- **Kernel drift:** A newly booted kernel needs renewed client-build and support checks. WEKA recommends preventing automatic upgrades to unsupported kernels. (docs.weka.io)
- **Runtime permissions:** A container's RDMA setup needs its own review, including non-root configuration and memory-lock limits described by Red Hat. (developers.redhat.com)
- **Service provisioning:** Some design choices precede installation. FSx's EFA creation requirement cannot be repaired by changing cuFile configuration. (docs.aws.amazon.com)

Async CUDA I/O needs its own support answer

An application programming interface (**API**) can expose an asynchronous operation without establishing an accelerated implementation on every filesystem. NVIDIA's live cuFile reference says its IBM Spectrum Scale and WekaFS async and batch operations use compatibility-mode I/O. Treat that as NVIDIA's documented API behavior, not as a current IBM release certification. (docs.nvidia.com)

IBM's Storage Scale FAQ Q2.39 states that **asynchronous CUDA I/O is not supported**. ([IBM Storage Scale FAQ, Q2.39](#)) Record this as IBM's documented support restriction in the ledger. Keep NVIDIA's compatibility-mode API behavior described above separate from IBM's support policy, and require IBM's exact-release confirmation before accepting a deployment that needs asynchronous CUDA I/O.

The same discipline applies to filesystem flags. NVIDIA describes **O_DIRECT** as a hint and documents eligible direct paths for some non-O_DIRECT cases. Opening a file with a particular flag is therefore insufficient proof of the runtime transfer mechanism. (docs.nvidia.com) Qualification should use the application's actual call pattern and preserve the resulting observations.

Data Analysis and Evidence

The quantitative evidence here consists of documented release boundaries and configuration constraints. These are not benchmark measurements and should not be combined into a synthetic readiness score. A higher version in one component does not establish that the combination is covered by another component's qualification.

Several numerical anchors illustrate the different scopes:

- **ONTAP 9.10.1:** NetApp identifies this as the starting release for its documented NFS-over-RDMA GDS capability; protocol-specific selection remains necessary. (docs.netapp.com)
- **CUDA 12.2 release 2:** BeeGFS explicitly names this boundary for CUDA 12 support, a finer distinction than a toolkit major version. (doc.beegfs.io)
- **4 KB alignment:** BeeGFS documents block-aligned request offsets and lengths, making workload shape part of readiness. (doc.beegfs.io)
- **Driver 2.24.2 or higher:** AWS's live FSx setup names this NVIDIA GDS driver floor; it should not be entered as a cuFile library version. (docs.aws.amazon.com)
- **Ubuntu 22.04 or newer, kernel 6.8 or newer:** AWS's live client documentation gives this Ubuntu qualification, not a universal Linux kernel minimum. (docs.aws.amazon.com)
- **V12 availability on September 30, 2026:** VDURA's dated announcement establishes timing; its undated GDS whitepaper supplies a separate integration claim. (www.vdura.com) (www.vdura.com)

Dell's published test bill of materials offers a useful granularity example: **OneFS 9.5.0.0**, **CUDA 12.1**, GPU driver **530.30.02**, and **GDS 1.6.1.9** appear together with the operating-system and kernel details. (infohub.delltechnologies.com) Those values describe a historical test combination, not authorization to replace individual components with arbitrary newer releases.

Similarly, HPE's **September 2, 2025** ClusterStor E1000 QuickSpecs describe a Lustre implementation based on **2.15** and explicitly request network-compatibility confirmation. (www.hpe.com) (www.hpe.com) A filesystem lineage and a dated specification are useful evidence, but neither is a complete CUDA/client-driver qualification. For the backend-media question, the same QuickSpecs describe combining flash and disk in one filesystem; that is distinct from qualifying a local NVMe host path. (www.hpe.com)

The proposed support-chain rule is conjunctive: every required relationship must have applicable support evidence before the deployment receives the supported label. Local verification adds evidence about installed behavior; it cannot supply missing vendor coverage. Unknown fields remain unknown. This avoids giving a mostly documented stack a misleading numerical pass mark.

No performance-gain figure is reported. The evidence does not include a controlled workload comparison, repeated runs, or a supplied laboratory system. An integration announcement, a published version floor, and an operator's architecture description answer different questions and should remain separate in technical diligence.

Implications and Future Directions

GDS readiness should become a maintained configuration record. Treat the procurement ledger as the beginning of that record, then attach the installed manifest, runtime observations, and approved exceptions. GPU Smith's first-party description of architecture records, port maps, operating procedures, and acceptance reports illustrates the documentation perspective used here; it is not evidence that the consultancy has qualified the third-party combinations in this report. (gpusmith.com)

Table 3 proposes refresh triggers. These are engineering acceptance recommendations derived from the dependency chain, rather than a vendor-mandated universal procedure.

Change	Evidence to refresh	Reason and disposition
GPU driver or kernel-module flavor	Support coverage, module build, capability report, operation tests	NVIDIA documents rebuilding nvidia-fs when a driver update changes exported P2P symbols. (docs.nvidia.com)
Linux kernel or filesystem client	Vendor qualification, build inputs, mount-level tests	VAST's exact OS/kernel/OFED client-build requirement makes these inputs material. (kb.vastdata.com)
CUDA, cuFile, or application API usage	Package manifest, effective configuration, operation-level path evidence	BeeGFS's CUDA release boundary illustrates why a toolkit label alone is insufficient. (doc.beegfs.io)
Storage release, protocol, or enabled feature	Release-specific coverage, export configuration, read/write acceptance	NetApp requires protocol-appropriate release selection; WEKA documents its encryption restriction. (docs.netapp.com) (docs.weka.io)
NIC firmware, network drivers, or topology	Device map, RDMA verification, per-GPU path tests	BeeGFS requires RDMA throughout; Linux routing depends on the PCIe hierarchy. (doc.beegfs.io) (docs.kernel.org)
Container image or runtime	Mount/device/permission mapping and application-context tests	Red Hat's RDMA example separately configures the container engine for non-root use. (developers.redhat.com)

The matrix defines which evidence an update puts back into question. It does not imply every change produces fallback. An update can be acceptable after its affected support links and runtime results have been re-established.

Future buying decisions should ask vendors to make these links easier to audit: immutable release-qualified matrices, explicit operation coverage, and clear ownership for cross-vendor cases. Product briefs remain useful introductions. Pure's own brief directs readers to reference architectures for full details; the procurement process should carry that distinction through to acceptance. (www.purestorage.com)

Conclusion

A defensible answer to “what are the GPUDirect Storage requirements?” begins with the named workload and intended storage path. The answer must connect the GPU, driver, CUDA/cuFile packages, optional path-specific kernel module, Linux kernel, filesystem client, storage release, network, topology, runtime, and effective configuration.

The three evidence labels should remain separate. Architectural possibility establishes a plausible design. Vendor support establishes documented coverage of the actual combination. Local verification establishes what the installed workload did under the recorded conditions. No one label should silently stand in for the others.

The practical deliverable is a completed ledger and acceptance packet. Before purchase, request exact release identities, covered operations, applicable feature restrictions, and a support owner. After installation, record the build, check exposed capability, verify correctness, and observe the intended path while the real execution context is active. Repeat affected checks when a material dependency changes.

This report supplies that process without certifying an unprovided stack or inventing performance results. Historical test configurations are examples of evidence granularity, not current shopping lists. Undated live pages require release confirmation. IBM's Storage Scale FAQ Q2.39 documents that asynchronous CUDA I/O is not supported; retain exact-release confirmation as a deployment acceptance requirement. ([IBM Storage Scale FAQ, Q2.39](#)) A purchase justified by GDS should proceed on a documented combination and a reviewable acceptance method, with unresolved fields carried forward visibly.

Source links

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

1. <https://docs.nvidia.com/gpudirect-storage/overview-guide/index.html>
2. <https://docs.kernel.org/driver-api/pci/p2pdma.html>
3. <https://docs.netapp.com/us-en/ontap/concepts/rdma-concept.html>
4. <https://docs.netapp.com/us-en/ontap/nfs-rdma/index.html>
5. https://doc.beegfs.io/latest/advanced_topics/gds_support.html
6. <https://kb.vastdata.com/docs/gpu-direct-storage-installation-and-operations-guide-2025>
7. <https://docs.aws.amazon.com/fsx/latest/LustreGuide/efa-file-systems.html>
8. <https://docs.weka.io/4.4/weka-system-overview/networking-in-wekaio>
9. <https://www.ibm.com/docs/en/STXKQY/pdf/gpfsclustersfaq.pdf>
10. <https://docs.nvidia.com/gpudirect-storage/design-guide/index.html>
11. <https://www.stjude.org/research/why-st-jude/data-tools/research-computing/research-computing-resources.html>
12. <https://gpusmith.com/>
13. <https://www.ddn.com/press-releases/ddn-boosts-ai-leadership-with-exascaler-6/>
14. <https://www.purestorage.com/content/dam/pdf/en/solution-briefs/sb-high-performance-ai-federal-agencies.pdf>
15. <https://docs.nvidia.com/gpudirect-storage/release-notes/index.html>
16. https://docs.nvidia.com/gpudirect-storage/nvidia_fs/index.html
17. <https://www.ibm.com/support/pages/ibm-spectrum-scale-gpudirect-storage-gds-accelerated-writes-available-technical-preview-feature-spectrum-scale-515>
18. <https://www.vdura.com/2026/09/30/vdura-data-platform-v12-now-generally-available-the-hyperscaler-storage-playbook-to-ai-factories-and-neoclouds-on-supermicro/>
19. <https://www.vdura.com/whitepaper/vdura-v12-ai-factory-storage-whitepaper/>
20. <https://docs.aws.amazon.com/fsx/latest/LustreGuide/configure-efa-clients.html>
21. <https://docs.weka.io/4.4/planning-and-installation/prerequisites-and-compatibility>
22. <https://www.hpe.com/psnow/downloadDoc/Cray%20ClusterStor%20E1000%20Storage%20Systems%20QuickSpecs-a00062172enw.pdf?deepLink=&form=false&hf=regular&id=a00062172enw.pdf&isFutureVersion=true&originalObjectName=&prelaunchSection=&preview=false&print=&r=§ion=&softrollSection=&ver=20>
23. <https://infohub.delltechnologies.com/en-us//dell-powerscale-and-nvidia-gpudirect-performance-report/pcie-topology/>
24. <https://docs.nvidia.com/gpudirect-storage/best-practices-guide/index.html>

25. <https://infohub.delltechnologies.com/en-us//dell-powerscale-and-nvidia-gpudirect-performance-report/bill-of-materials-116/>
26. <https://infohub.delltechnologies.com/en-us//dell-powerscale-and-nvidia-gpudirect-performance-report/powerscale-configuration-16/>
27. <https://developers.redhat.com/articles/2025/04/29/accelerate-model-training-openshift-ai-nvidia-gpudirect-rdma>
28. <https://gpusmith.com/articles/en/gpu-cpu-nic-affinity-rank-binding>
29. <https://gpusmith.com/articles/en/gpudirect-storage-gdsio-qualification-runbook>
30. <https://docs.nvidia.com/gpudirect-storage/troubleshooting-guide/index.html>
31. <https://docs.nvidia.com/gpudirect-storage/configuration-guide/index.html>
32. <https://docs.nvidia.com/gpudirect-storage/api-reference-guide/>
33. <https://docs.nvidia.com/gpudirect-storage/o-direct-guide/>

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.