



Independent Engineer Report for GPU Financing: Scope

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 · independent engineer report for gpu financing · gpu financing technical due diligence · lender technical adviser · gpu collateral assessment · data center due diligence · reliance letters · gpu acceptance evidence · project finance

Original article: <https://gpusmith.com/articles/en/independent-engineer-reports-gpu-financing>



In this article

1. Executive Summary
 2. Introduction and Background
 3. Roles and Opinion Boundaries
 4. Appointment, Independence and Reliance
 5. Proposed GPU Scope of Work
 6. Report Deliverables and Evidence Quality
 7. Acceptance Evidence and Post-Close Monitoring
 8. Data Analysis and Evidence
 9. Implications and Future Directions
 10. Conclusion
 11. References
-

Executive Summary

For **graphics processing unit (GPU) financing**, an **independent engineer (IE) report** is best commissioned as a defined technical opinion supporting a particular funding decision. The appointment should specify the assets and service obligations reviewed, the evidence cutoff, the work performed, the limitations and the named parties permitted to rely. The International Federation of Consulting Engineers (FIDIC) appointment guidance calls for tailored conditions; the Royal Institution of Chartered Surveyors (RICS) commercial-property guidance distinguishes report circulation through a data room from extending reliance to its readers. These are relevant analogies, not universal GPU-financing terms. ([fidic.org](https://www.fidic.org)) (www.rics.org)

The proposed pre-close scope should connect exact inventory and compatibility, power and cooling, compute and storage fabrics, delivery and commissioning, support and operations, and technical finance-model inputs. NVIDIA documents **eight GPUs** in a DGX B300 system, while its reference architecture separates networking functions. The engineering implication is that a server count is insufficient evidence of service capacity. ASHRAE also distinguishes nameplate ratings from actual operating draw, and Dominion describes staged data-center power delivery. (docs.nvidia.com) (docs.nvidia.com) (handbook.ashrae.org) (www.dominionenergy.com)

The required output should be a traceable report with a document register, methodology, findings, assumptions, unresolved-issue register and proposed conditions for the relevant funding stage. Evidence should distinguish observed results from sponsor forecasts and identify sampled, excluded or untested assets. Passing selected diagnostics cannot, on that evidence alone, establish every workload's performance: the documented Data Center GPU Manager (DCGM) plugin for NVIDIA Collective Communication Library (NCCL) tests is single-node, whereas the separate NCCL test project addresses performance and correctness across its configured participants. This report therefore treats acceptance as bounded evidence, not a guarantee of production throughput, uptime or repayment. (docs.nvidia.com) (github.com)

Professional eligibility, competence and specialist boundaries require attention before appointment. The National Council of Examiners for Engineering and Surveying (NCEES) notes that no single nationwide US engineering license exists; the National Society of Professional Engineers (NSPE) emphasizes competence in

the actual technical field. Counsel should settle reliance and contractual terms, with valuation and accounting work separately allocated. Finally, appoint post-close monitoring explicitly: the World Bank wind-project specimen separates pre-close evaluation from later construction, performance and operations work. This guide adapts that structure into a proposed GPU scope and blank technical-input worksheet, using public sources retrieved **October 15, 2026**. It offers neither universal certification nor transaction-specific legal forms. ([ncees.org](https://nces.ed.gov/ipeds/data/collegesanduniversities/2026/10152026/)) (www.nspe.org) (ppp.worldbank.org)

Introduction and Background

An **independent engineer report for GPU financing** should help a lender decide whether identified assets, infrastructure interfaces, delivery plans and operating assumptions support the technical case presented for funding. The appointment should define the questions the engineer will answer, the work undertaken to answer them, and the people entitled to rely on the resulting opinion. FIDIC's engineering-contract guidance emphasizes tailoring the agreement to the particular appointment. (fidic.org)

Conventional project-finance work supplies a useful starting point. The World Bank's publicly available **wind-project terms of reference** distinguish technical review before financial close from construction, performance-test and operations monitoring afterward. This article adapts that separation to graphics processing unit (GPU) infrastructure. Every GPU-specific scope, register and decision framework below is a **proposed scope**, not a World Bank GPU standard. (ppp.worldbank.org) (ppp.worldbank.org)

The adaptation matters because a cluster comprises more than accelerators in racks. NVIDIA's B300 reference architecture distinguishes compute, storage and management networking; its storage discussion explains that checkpoint activity can stop forward training progress. The engineering inference is that equipment identification, facility readiness and customer-service capability require different evidence. A report should connect them without treating one successful observation as proof of the others. (docs.nvidia.com) (docs.nvidia.com)

This guide addresses banks, private-credit and infrastructure funds, counsel and borrowers commissioning that work. It uses public primary sources retrieved **October 15, 2026**, rather than confidential engagement letters or unpublished transaction reports. GPU Smith's perspective here is that of an adjacent independent engineering advisor. The article concerns appointment and deliverables, with a separate [GPU financing due-diligence overview](#) available for commissioning context.

The central decision is what evidence must exist at the relevant funding stage, which conclusions the IE can support, and which questions require other specialists. The European Investment Bank describes appraisal as spanning technical, financial, procurement and other disciplines. That breadth is a reason to allocate responsibilities explicitly rather than expect an engineering opinion to answer the entire credit decision. (www.eib.org)

Roles and Opinion Boundaries

An **independent engineer (IE)** or lender's technical adviser should be appointed to assess specified technical questions for the financing. That differs from responsibility for developing the asset on the sponsor's behalf. ORTECH's energy-oriented service description distinguishes independent financing review from an owner's

representative role concerned with development and construction. It illustrates different mandates, not a standardized GPU appointment. (www.ortechconsulting.com) (www.ortechconsulting.com)

An **owner's engineer** may develop or review design and support procurement and delivery. A **commissioning agent** organizes or witnesses defined verification work. The lender's IE can examine their outputs and, if appointed, witness selected tests. Under the proposed allocation here, those roles remain distinguishable even when one firm offers several services. ASHRAE's commissioning guidance describes integrated systems testing and operations-team involvement; neither activity automatically supplies an independent financing opinion. (www.ashrae.org) (www.ashrae.org)

Valuation requires its own mandate. The Appraisal Foundation describes personal-property appraisal as valuing movable tangible assets and certain intangible assets. A GPU collateral technical assessment should identify equipment, configuration, condition evidence, support dependencies and practical removal or redeployment constraints. A valuation specialist should determine the appropriate value premise and methodology. An engineering condition observation should not silently become an assumed resale price. (appraisalfoundation.org)

Likewise, an IE's technical-model review is distinct from a **financial-statement audit**. The Public Company Accounting Oversight Board (PCAOB) auditing standard describes the auditor's responsibility as expressing an opinion on financial statements. The proposed IE review checks the engineering basis of inputs such as service dates, capacity, maintenance and restoration assumptions. Accountants address accounting work within their own engagement; credit specialists assess borrower and customer credit; counsel addresses contractual and security questions. (pcaobus.org) (ppp.worldbank.org)

Match signatories to the actual work

Competence is specific to the discipline. The National Society of Professional Engineers (NSPE) calls for assignments to be undertaken by people qualified in the relevant technical fields and limits signing documents outside their competence or direction and control. A GPU financing may therefore need several specialists, with the report identifying who performed and took responsibility for each segment. (www.nspe.org) (www.nspe.org)

Professional requirements also depend on jurisdiction and the services actually supplied. The National Council of Examiners for Engineering and Surveying (NCEES) states that there is no single nationwide US engineering license. California's board, as a jurisdiction-specific example, identifies civil, electrical and mechanical engineering as practice-act branches. These sources do not establish that every GPU financing report needs the same license or seal. Confirm applicability with the relevant board and transaction counsel. (ncees.org) (www.bpelsg.ca.gov)

The proposed opinion boundary should consequently specify whether the work covers design compliance, installed configuration, witnessed tests, technical assumptions, or only selected documents. List excluded specialties and the reports relied upon. A statement that the technical case is supportable on defined evidence is different from guaranteeing repayment, production throughput or future availability. That distinction follows from the deliberately bounded appraisal, audit and engineering mandates discussed above.

Appointment, Independence and Reliance

Agree the **appointment before substantive diligence**. FIDIC's White Book guidance says particulars require specific drafting and that legal advice can help parties understand the agreement. For GPU financing, scope should match the lender's decision, the assets at that funding stage, and the evidence that can realistically be obtained. The document request list follows from that appointment; it does not replace it. ([fidic.org](https://www.fidic.org)) ([fidic.org](https://www.fidic.org))

The appointment should address **independence and conflicts** directly. NSPE calls for disclosure of known or potential conflicts affecting, or appearing to affect, judgment. Ask about design work on the same project, supplier relationships, referral arrangements and other roles that could affect the proposed review. Disclosure, agreed safeguards and acceptance are matters for the parties; the word independent should not stand alone as the entire explanation. (www.nspe.org)

The following are **sample drafting questions for counsel and the advisor**, not legal forms or statements of universal market terms:

- **Instruction:** Who appoints the IE, who pays, and who can direct or change the scope?
- **Purpose:** Which financing decision and asset perimeter is the report intended to support?
- **Reliance:** Which legal entities may rely, for what purpose, and subject to which limitations?
- **Circulation:** Who may receive or read the report without becoming a reliance party?
- **Access:** Can the IE inspect the site, reach relevant staff and examine underlying records and raw logs?
- **Responsibility:** Which specialists sign or contribute, and whose separate opinions are incorporated?
- **Change:** What happens when equipment, designs, contracts or evidence change before close?

RICS's commercial-property technical-diligence guidance explicitly discusses avoiding unintended reliance by everyone reviewing a report in a data room. Applied here as an analogy, that means counsel should distinguish **permission to circulate** from **permission to rely**. A borrower uploading a report, or an agent distributing it to prospective lenders, should not be treated in this proposed process as settling the reliance question. (www.rics.org)

What a reliance letter needs to resolve

A proposed reliance-letter discussion should identify the precise report and version, permitted recipient entities, purpose, assumptions and exclusions, and any agreed conditions on onward transfer. Counsel should also resolve governing law, responsibility and liability terms, and whether the letter merely permits reliance on existing work or commissions additional procedures. Public sources reviewed for this guide do not provide a universal GPU-financing reliance form.

Date handling deserves separate attention. The **report issue date**, **information cutoff**, **site-visit date** and **test dates** may differ. NSPE calls for reports to bear a date indicating when their information was current. RICS notes that a data room can change after submission. The proposed appointment should require a controlled process for later information, with an addendum or further review where relevant. (www.nspe.org) (www.rics.org)

If closing is delayed or the financed inventory changes, the parties should ask what needs updating rather than assume a recently circulated report has become current. Record substitutions, design revisions and changes in support status against the original assumptions. Counsel should decide whether a reliance recipient needs the original report, an updated report, or an expressly defined supplemental opinion. This is a proposed control, grounded in the sources' emphasis on current information and bounded inspection scope.

(www.rics.org)

Proposed GPU Scope of Work

Start with the financed service and asset perimeter

Begin the proposed GPU scope with the **service obligation** and its measurable acceptance requirements. Ask the borrower to identify the workload class, customer acceptance requirements, operating environment and dependencies needed to make capacity usable. Training, inference and reserved bare-metal capacity should each have their own evidence questions. Kubernetes, for example, documents a scheduling mechanism requiring vendor drivers and device plugins; that mechanism is narrower than demonstrating a customer's complete service. (kubernetes.io)

The proposed service review should make these questions explicit:

- **Workload:** Which model, dataset, precision and software environment define the intended service?
- **Service unit:** Is the customer receiving devices, partitions, nodes or an end-to-end managed service?
- **Performance:** Which throughput or latency measure and observation conditions govern acceptance?
- **Availability:** What measurement period, service boundary and exclusions define the contractual target?
- **Data path:** Which storage, network and access dependencies must work for the customer?
- **Acceptance:** Who approves the result, and which evidence must be retained?

Next establish the **asset perimeter**: legal borrowing entity, equipment location, server and GPU identifiers, configuration, ownership documents and financed supporting assets. The IE should reconcile the physical and logical inventory within the agreed scope; counsel determines legal title and security treatment. Record compatibility evidence against the installed configuration. NVIDIA specifically states that ConnectX-8 firmware determines supported cables and adapters, illustrating why a generic bill of materials is insufficient technical identification. (docs.nvidia.com)

Table 1 maps conventional diligence questions to proposed GPU-specific additions. The evidence and deliverables are recommendations for the appointment, not requirements claimed to apply to every financing. The starting distinction between design review, technical-model review and later monitoring follows the wind-project specimen. (ppp.worldbank.org)

Conventional question	GPU-specific addition	Source evidence to request	Responsible specialist	Proposed deliverable and exclusion
What is being financed?	Which servers, GPUs and supporting assets form the financed cluster?	Serial register, invoices, delivery records, site reconciliation	IE for technical identity; counsel for title	Asset reconciliation; no legal-title opinion

Conventional question	GPU-specific addition	Source evidence to request	Responsible specialist	Proposed deliverable and exclusion
Does the design support output?	Can compute, storage and management paths support the specified service?	Architecture, topology, workload requirements, vendor support matrix	Cluster, network and storage engineers	Interface assessment; no universal throughput guarantee
Is power available?	What staged IT load can be energized at the relevant racks?	Utility correspondence, connection agreement, electrical design, commissioning records	Electrical engineer; utility and counsel interfaces	Capacity/date reconciliation; no utility commitment invented
Is cooling suitable?	Are airflow, coolant, flow and pressure requirements configuration compatible?	Mechanical design, supplier limits, water-quality records, test logs	Mechanical/cooling engineer	Thermal-interface assessment; no substitute for design responsibility
Can delivery be completed?	Do server, fabric, storage and software milestones converge?	Procurement schedule, dependency plan, installation and acceptance procedures	IE with commissioning agent	Integrated readiness schedule; no contractor performance guarantee
Are technical costs plausible?	Are support, spares, restoration and replacement assumptions evidenced?	Support terms, operating plan, supplier lead-time evidence, budget inputs	IE; accountant and valuation specialist separately	Technical-input bridge; no accounting or residual-value opinion

The matrix assigns an **interface owner** as well as a reviewer. Dominion Energy describes data-center service as typically arriving in stages matching the load ramp, while National Grid distinguishes its works from the customer's works. These are utility-specific examples, not promises about another project. The proposed IE deliverable should connect the applicable utility evidence to the actual cluster stages.

(www.dominionenergy.com) (www.nationalgrid.com)

Review the interfaces and operating plan

Cooling review should address the actual configuration. ASHRAE identifies configuration-specific flow and pressure requirements and the importance of cooling-water quality. HPE also distinguishes cooling approaches by which components they cool. The proposed scope should therefore follow heat from the equipment through the installed facility interfaces, identifying any dependency whose evidence remains incomplete. (handbook.ashrae.org) (handbook.ashrae.org) (www.hpe.com)

The **support and operating plan** should identify who diagnoses faults, approves changes, restores failed nodes and obtains replacement equipment. Technical assumptions in the finance model should then trace to these plans and their evidence. This extends the World Bank specimen's review of technical inputs, maintenance, spares and operating assumptions into the GPU context, without assigning the IE responsibility for the model's commercial forecasts. (ppp.worldbank.org)

Report Deliverables and Evidence Quality

A useful IE report is a traceable set of opinions, not a folder-completeness score. Its document register should identify what was reviewed, which version was used, when it was received, and which conclusion depends on it. RICS's commercial-property diligence guidance discusses access limitations and documents added after a data-room review. The GPU application is to freeze the evidence set and assess later changes explicitly. (www.rics.org)

The proposed package should contain:

- **Appointment summary:** purpose, addressees, permitted reliance, reviewed assets and specialist boundaries.
- **Evidence register:** document versions, raw-log locations, site access, interviews and sponsor representations.
- **Methodology:** inspections, calculations, test witnessing and sampling, including work omitted.
- **Opinion and assumptions:** supported conclusions, conditions needed for them, and their information cutoff.
- **Issue register:** unresolved questions, technical consequences, owner, evidence required and closure status.
- **Decision schedule:** proposed pre-funding conditions, later monitoring items and changes requiring reconsideration.

These deliverables adapt the World Bank specimen's separate draft and final technical reports and its requirement to highlight outstanding items. They should not imply that every matter is resolved simply because a final report has been issued. A final document can contain a conditional opinion and clearly identified work that remains outside the completed scope. (ppp.worldbank.org)

Table 2 shows how proposed report sections support distinct credit decisions. Evidence cutoff is specific to the section where necessary; the report should identify the dates of the underlying observations alongside its issue date.

Report section	Credit decision supported	Evidence cutoff and boundary
Appointment and reliance	Whether the opinion is addressed to the intended financing parties	Executed instructions and named permitted users; counsel confirms reliance language
Asset and configuration reconciliation	Whether the technical perimeter matches the proposed funded assets	Inventory version and inspection date; unobserved assets listed
Facility and interface assessment	Whether the proposed deployment stages have evidenced dependencies	Design revisions, utility correspondence and commissioning state
Acceptance evidence assessment	Whether the agreed technical acceptance questions have been answered	Actual test timestamps, versions, population and exclusions
Technical-model reconciliation	Whether model inputs match the engineering evidence	Model version, source assumptions and unresolved sensitivities

Report section	Credit decision supported	Evidence cutoff and boundary
Findings and closeout schedule	Which technical matters require action before the relevant funding decision	Issue status at cutoff; subsequent closure evidence separately recorded

The table's purpose is to prevent a broad conclusion from outrunning a narrow observation. Utility preliminary planning, for example, is a different evidence state from completed energization requirements. Dominion offers preliminary infrastructure evaluation; National Grid describes technical, contractual and safety requirements for commissioning before energization. The report should name the state actually evidenced.

(www.dominionenergy.com) (www.nationalgrid.com)

Express confidence without an invented standardized score

Use descriptive **assertion status**: observed by the IE, independently corroborated, supported by supplied documents, sponsor forecast, or not assessed. These are proposed reporting labels, not industry certification grades. Attach a rationale to each. A forecast may be reasonable for planning while remaining an assumption rather than an observed operating result. FIDIC's White Book guidance says the Particular Conditions must be drafted to suit each individual agreement and type of service. (fidic.org)

An unresolved-issue entry should state the affected asset or service, technical consequence, missing evidence, responsible party and decision deadline. Severity should reflect that consequence and the relevant funding stage, with an explanation. Avoid converting unlike issues into a synthetic universal score. The wind-project specimen's express treatment of outstanding closing matters provides the underlying project-finance analogy. (ppp.worldbank.org)

Proposed **pre-funding conditions** are recommendations to the lender and counsel, not binding loan terms issued by the IE. An item may need closure before an equipment draw, before customer service begins, or during an agreed monitoring period. The report should explain that timing judgment and identify the professional or contractual decision that remains with another party. (www.eib.org)

Acceptance Evidence and Post-Close Monitoring

Witness tests with a declared perimeter

Acceptance evidence should answer agreed questions about identified assets under recorded conditions. NVIDIA's Data Center GPU Manager (DCGM) documentation describes readiness diagnostics with stated exclusions. Its documented NCCL plugin runs only single-node tests; if its required test-binary path is unset, the plugin is disabled and absent from diagnostic output. A general diagnostic result therefore needs inspection at the test and entity level. (docs.nvidia.com) (docs.nvidia.com) (docs.nvidia.com)

The official **NVIDIA Collective Communication Library (NCCL)** test project separately assesses collective-operation performance and correctness across configured processes, threads and devices. That provides bounded communication evidence. The proposed IE assessment should state the topology and participating assets actually exercised, and whether those observations address the financed service. It should not extrapolate a selected collective test into a customer-workload or availability guarantee. (github.com) (github.com)

For witnessed acceptance, request:

- **Population:** A reconciliation of financed, installed, selected and actually tested assets.
- **Environment:** Hardware configuration, firmware, drivers, libraries, containers and scheduler settings.
- **Procedure:** Agreed test purpose, topology, parameters, duration and acceptance criteria.
- **Raw evidence:** Timestamped logs, output files, measurement boundaries and configuration records.
- **Exceptions:** Failed, skipped, unavailable, interrupted and untested items, each stated explicitly.
- **Disposition:** Retest evidence, remaining limitations and the conclusion supported at the cutoff.

These are proposed evidence requirements, not a repetition of tool execution protocols. Consult the linked official diagnostics and test documentation for those details. ASHRAE's commissioning guidance reinforces a complementary principle: receipt and review of test reports should remain a hold point. A notification that testing is complete is a weaker evidence state than reviewed results. (www.ashrae.org)

Partial scope needs explicit exclusions

A site visit to selected racks should identify the selection and unvisited population. A document-only appointment should say that installation was not observed. Missing records need an assumption, exclusion or unresolved issue, with the effect on the opinion explained. RICS's property-diligence guidance provides a parallel approach to stating limitations and treating unavailable material; this is an adaptation to GPU evidence. (www.rics.org) (www.rics.org)

Software and monitoring limitations also matter. Slurm documents gaps in utilization accounting for Multi-Instance GPU (MIG) devices and differences in what autodetection modes detect. Thus, **missing telemetry is not automatically zero utilization**, and an inventory export should be interpreted against the configured collection method. The proposed report should identify these measurement limitations before using operational data in a technical conclusion. (slurm.schedmd.com) (slurm.schedmd.com)

Appoint post-close work separately

The World Bank wind-project specimen separates pre-close review from later construction, performance-test and operations monitoring. Adapt that distinction into a separate GPU mandate with its own access, cadence, deliverables and permitted reliance. Do not imply that commissioning a pre-close report also purchases continuing oversight. (ppp.worldbank.org)

A proposed monitoring schedule should define:

- **Cadence:** Reporting dates tied to construction, deployment or operating stages.
- **Baseline:** The approved asset, configuration, evidence and assumption versions.
- **Triggers:** Substitutions, relocations, material software changes, facility changes or relevant outages.
- **Retesting:** Which changes require additional technical evidence and who accepts it.
- **Reporting:** Exception logs, remediation status and the technical scope of any draw review.

The specimen includes review of progress reports and disbursement support. For GPU assets, the proposed extension is to reconcile each relevant technical milestone with the assets and dependencies actually delivered. Counsel and the lender retain responsibility for the funding decision and contractual conditions. The

monitoring report should identify what changed since the baseline, what remains unresolved and what its inspection did not establish. (ppp.worldbank.org)

Data Analysis and Evidence

Published quantities are reference inputs, not project observations

The quantitative task is to reconcile the borrower's technical model, not import a generic financing ratio. The wind-project specimen explicitly includes technical data inputs and calculations in financial-model review. For GPUs, the proposed adaptation asks where capacity, service dates, power consumption, restoration time and replacement assumptions originate, and which remain forecasts. No universal advance rate, residual value or engineering fee is proposed here. (ppp.worldbank.org)

Published specifications illustrate why units and configuration matter. NVIDIA documents **eight B300 GPUs per DGX B300 system**. (docs.nvidia.com) AMD's MI355X page lists **288 gigabytes (GB)** of third-generation enhanced High Bandwidth Memory (HBM3E) and **8 terabytes per second (TB/s)** memory bandwidth. (www.amd.com) Its cluster-support documentation states a **ROCm 7.0.1 or higher** requirement for GPU applications. (instinct.docs.amd.com) These are vendor specifications and a support requirement, not measurements of a financed cluster or comparable customer-workload performance.

Measurement boundaries require equal care. Lenovo's enclosure-management documentation distinguishes enclosure, node and power-supply consumption. ASHRAE cautions that a nameplate rating does not represent actual operating draw. The US Department of Energy (DOE) defines **power usage effectiveness (PUE)** using total annual facility energy divided by annual IT-equipment energy. An annual energy ratio should therefore not silently become a measured rack-demand or peak-capacity figure. (pubs.lenovo.com) (handbook.ashrae.org) (www.energy.gov)

Blank project reconciliation worksheet

Table 3 is a proposed calculation appendix. Fields are deliberately blank. Populate them with actual project records, units, versions and observation dates; a published reference specification should remain labeled as such.

Model input	Project value	Proposed reconciliation	Evidence and limitation
Available service capacity	[enter value and unit]	Reconcile installed, accepted and allocatable assets; apply actual workload constraints	Serial inventory, acceptance population, scheduler configuration; schedulable resources alone do not establish service output
Staged service dates	[enter dates by stage]	Reconcile delivery, installation, utility, cooling, fabric, storage and customer acceptance dependencies	Versioned schedule and supporting records; distinguish commitments from forecasts
IT power and facility energy	[enter units and boundary]	Sum project loads on a consistent boundary; reconcile energy separately from peak demand	Meter records and configured load assumptions; no generic PUE multiplier for all purposes

Model input	Project value	Proposed reconciliation	Evidence and limitation
Restoration time	[enter project assumptions]	Reconcile diagnosis, parts access, repair, reinstallation and validation; explain overlaps	Support arrangements and incident/restoration evidence; prospective values remain assumptions
Replacement timing and cost inputs	[enter project evidence]	Trace technical substitution, compatibility, removal, installation and retesting assumptions	Supplier evidence and operating plan; separate appraisal and accounting responsibilities
Sensitivity inputs	[enter identified cases]	Vary unresolved technical dependencies and show which model input changes	Explicit engineering assumptions; lender/model owner assesses resulting credit effects

The worksheet supports **input reconciliation**, not a repayment opinion. Kubernetes exposes schedulable GPU resources after plugin installation; that is a useful inventory and allocation signal, not the same quantity as contract-ready service capacity. Dominion's staged-service description similarly supports checking each deployment stage rather than applying one site-wide readiness date. (kubernetes.io) (www.dominionenergy.com)

For a separate planning aid, see the [GPU power calculator](#). The financing report should retain its project-specific units and measurement boundary. Where data are unavailable, show the blank field or assumption and its consequence. Arithmetic becomes auditable when another reviewer can reproduce the reconciliation from the cited project records, even when the underlying values remain confidential.

Implications and Future Directions

The practical implication is a **versioned technical baseline** that can survive changes between diligence, close and operation. Public documentation itself illustrates why version control matters: NVIDIA describes CUDA minor-version compatibility within a major-release family with feature limitations, while AMD specifies a minimum software version for MI355X applications. The proposed appointment should require the project's exact support evidence rather than a general statement that the software is compatible. (docs.nvidia.com) (instinct.docs.amd.com)

Facility and service evidence should also remain distinct as they are brought into one report. Uptime Institute's topology standard concerns site-infrastructure functionality, capacity and expected availability. ASHRAE emphasizes configuration-specific cooling interfaces and dynamic loads. The engineering implication is to ask which facility observations support which cluster-service assumptions, then identify the additional tests or operating evidence needed for the latter. (uptimeinstitute.com) (handbook.ashrae.org)

Operations introduce further interfaces. The National Institute of Standards and Technology (NIST) Cybersecurity Framework organizes cybersecurity risk management around functions including governance, protection, detection, response and recovery. A proposed IE operating-plan review can identify dependencies on those capabilities while leaving a detailed cybersecurity assessment to an explicitly appointed specialist. The financing package should state whether that specialist's work was reviewed, incorporated or excluded. (www.nist.gov)

For procurement of the advisor, compare the proposed work and deliverables directly:

- **Coverage:** Does the team address facility, compute, fabric, storage and software interfaces?
- **Access:** Can it examine raw evidence and relevant sites within the available timetable?
- **Accountability:** Are specialist responsibilities, conflicts and signatories clearly disclosed?
- **Usability:** Will the lender receive supported conclusions, unresolved issues and an update process?

These are service-evaluation questions, not claims that one credential covers every layer. NCEES directs attention to jurisdictional licensing, and NSPE emphasizes competence in the actual technical fields. Lenders and borrowers should resolve both professional eligibility and GPU-system expertise when appointing the team. ([ncees.org](https://www.ncees.org)) (www.nspe.org)

Future monitoring can build on the same evidence baseline without implying that automated telemetry replaces judgment. Slurm's documented collection limitations show why a reviewer must understand what a metric measures and what is absent. The durable contribution of the IE is the connection between an assertion, its evidence, its limitations and the financing decision it is intended to inform. (slurm.schedmd.com)

Conclusion

An independent engineer report for GPU financing should leave the lender with a precise account of what was reviewed, which technical conclusions the evidence supports, and what remains conditional or unassessed. Appointment, reliance and deliverables deserve the same attention as the borrower's equipment list. Defined responsibilities keep the engineering opinion connected to the broader legal, valuation, accounting and credit work.

The proposed scope in this guide bridges conventional project-finance diligence and the additional layers of a GPU service: asset configuration, compute and storage fabrics, management software, facility interfaces, acceptance evidence and operating dependencies. Its central discipline is traceability. Each material assertion should identify the underlying evidence, observation date, limitations and responsible specialist. A forecast should remain distinguishable from an observation.

Before the relevant funding decision, the parties should agree the report version, information cutoff, named reliance recipients, unresolved issues and any proposed conditions requiring further evidence. Partial inspection, sampled tests and missing logs should stay visible. A final report should communicate its actual opinion boundary rather than imply universal certification of throughput, availability or repayment.

Post-close monitoring should have a separate mandate that explains what is monitored, when reports are delivered, which changes trigger further review and who may rely on the results. That makes the original diligence useful as a controlled baseline throughout the financing process.

For commissioning context, see [independent pre-close GPU technical due diligence](#).

References

Sources below were retrieved on October 15, 2026. Vendor documents establish their own documented specifications and tool scope. Wind-project and commercial-property guidance provide explicitly labeled analogies; the GPU scope and worksheets are proposed adaptations.

1. World Bank, wind-project technical-diligence specimen. (ppp.worldbank.org)
2. FIDIC, client-consultant appointment guidance. (fidic.org)
3. ORTECH, financing review and owner's representative services. (www.ortechconsulting.com)
4. European Investment Bank, project appraisal cycle. (www.eib.org)
5. NSPE, engineering ethics, competence and current reports. (www.nspe.org)
6. NCEES, jurisdiction-specific engineering licensure. (ncees.org)
7. California engineering board, licensure categories. (www.bpelsg.ca.gov)
8. RICS, commercial-property technical due diligence. (www.rics.org)
9. Appraisal Foundation, personal-property valuation. (appraisalfoundation.org)
10. PCAOB, financial-statement auditor responsibilities. (pcaobus.org)
11. NVIDIA, B300 reference network fabrics. (docs.nvidia.com)
12. NVIDIA, B300 reference storage architecture. (docs.nvidia.com)
13. NVIDIA, DGX B300 system user guide. (docs.nvidia.com)
14. NVIDIA, DCGM diagnostic scope and exclusions. (docs.nvidia.com)
15. NVIDIA, official NCCL test project. (github.com)
16. NVIDIA, CUDA minor-version compatibility. (docs.nvidia.com)
17. AMD, MI355X product specifications. (www.amd.com)
18. AMD, cluster hardware and software support. (instinct.docs.amd.com)
19. Lenovo, enclosure and node power measurements. (pubs.lenovo.com)
20. HPE, liquid-cooling configuration descriptions. (www.hpe.com)
21. ASHRAE, data-center thermal and power considerations. (handbook.ashrae.org)
22. ASHRAE, commissioning and performance validation. (www.ashrae.org)
23. US Department of Energy, data-center energy and water efficiency. (www.energy.gov)
24. Uptime Institute, site-infrastructure topology standard. (uptimeinstitute.com)
25. Dominion Energy, data-center power requests and staged service. (www.dominionenergy.com)
26. National Grid, connection and commissioning process. (www.nationalgrid.com)
27. NIST, Cybersecurity Framework overview. (www.nist.gov)

Source links

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

1. <https://fidic.org/books/clientconsultant-model-services-agreement-5th-ed-2017-white-book>
2. https://www.rics.org/content/dam/ricsglobal/documents/standards/Technical_Due_Diligence_Of_Commercial_Property_1st_Edition.pdf
3. <https://docs.nvidia.com/dgx/dgxb300-user-guide/introduction-to-dgxb300.html>
4. <https://docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-b300/latest/network-fabrics.html>

5. https://handbook.ashrae.org/Handbooks/A23/SI/A23_Ch20/a23_ch20_si.aspx
6. <https://www.dominionenergy.com/virginia/large-business-services/data-center-requests>
7. <https://docs.nvidia.com/datacenter/dcgm/latest/user-guide/dcgm-diagnostics.html>
8. <https://github.com/NVIDIA/nccl-tests>
9. <https://ncees.org/licensure/international-professionals/>
10. <https://www.nspe.org/career-growth/nspe-code-ethics-engineers>
11. <https://ppp.worldbank.org/ppp-overview/practical-tools/sample-ppp-terms-reference/tor-technical-due-diligence-wind-power>
12. <https://docs.nvidia.com/dgx-superpod/reference-architecture/scalable-infrastructure-b300/latest/storage-architecture.html>
13. <https://gpusmith.com/gpu-financing-due-diligence>
14. <https://www.eib.org/en/projects/cycle/index.htm>
15. <https://www.ortechconsulting.com/due-diligence-and-peer-review-services>
16. <https://www.ashrae.org/technical-resources/topics-and-initiatives/ai-data-center/commissioning-performance-validation>
17. <https://appraisalfoundation.org/pages/personal-property-appraisal>
18. <https://pcaobus.org/oversight/standards/auditing-standards/details/as-1000--general-responsibilities-of-the-auditor-in-conducting-an-audit>
19. <https://www.bpelsg.ca.gov/applicants/appintrope.shtml>
20. <https://kubernetes.io/docs/tasks/manage-gpus/scheduling-gpus/>
21. <https://www.nationalgrid.com/electricity-transmission/connections/detailed-connections-process>
22. <https://www.hpe.com/us/en/solutions/direct-liquid-cooling.html>
23. <https://slurm.schedmd.com/gres.html>
24. <https://www.amd.com/en/products/accelerators/instinct/mi350/mi355x.html>
25. <https://instinct.docs.amd.com/projects/cluster-documentation/latest/reference/hardware-support.html>
26. https://pubs.lenovo.com/mgt_tools_smm2/c_power_overview
27. <https://www.energy.gov/cmei/femp/cooling-water-efficiency-opportunities-federal-data-centers>
28. <https://gpusmith.com/tools/gpu-power-calculator>
29. <https://docs.nvidia.com/deploy/cuda-compatibility/minor-version-compatibility.html>
30. <https://uptimeinstitute.com/resources/asset/tier-standard-topology>
31. <https://www.nist.gov/news-events/news/2024/02/nist-releases-version-20-landmark-cybersecurity-framework>

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.