

#	Technical Clarification Question	Brief Reason for Asking	Why We Need Clarification
1	DR / Colocation Requirements; Draft SLA §11–12, pp. 234 onward	The RFP requires Primary and DR facilities with defined recovery targets. Kindly clarify the intended overall resilience architecture: should both sites provide equivalent AI compute, storage, network and platform capacity, or may bidders propose an appropriately sized DR environment capable of restoring the required critical services within the prescribed RPO/RTO?	Defines the complete Primary/DR architecture, capacity, replication model and major cost baseline. The SLA defines recovery targets but not equivalent infrastructure capacity.
2	DR / Colocation Requirements; Technical Evaluation – DR, p. 81	As part of the end-to-end DR solution, should bidders determine whether the two sites operate Active/Active or Active/Passive, together with the required inter-site bandwidth, provided the committed availability, RPO and RTO are achieved? If not, Kindly provide IGNITE's preferred operating model and minimum inter-site connectivity requirements.	Architecture, WAN capacity, replication technology and DR investment all depend on this decision.
3	Layer 1 – Accelerator / AI Fabric Requirements	The solution is expected to support future AI capacity growth in addition to the initial accelerator deployment. Kindly clarify whether the required 25% growth headroom means only that the architecture must be expandable, or whether Day-1 facility capacity—including rack space, power, cooling and network ports—must also be reserved for this expansion.	Ensures compute, network and colocation facilities are engineered as one scalable solution rather than independently sized components.
4	Layer 1 – Accelerator Requirements	Given the rapid evolution of AI accelerator technologies, Kindly clarify at what project milestone the requirement for “latest-generation, in-production” accelerators will be assessed—bid submission, contract award, procurement or delivery. May the bidder upgrade to a newer equivalent/superior generation during execution, subject to compatibility and approval?	Prevents the overall solution becoming technologically outdated during a potentially long procurement and deployment cycle.
5	Data-Pipeline / General-Purpose Compute Tier, p. 92	The RFP requires a dedicated data-pipeline compute tier with ≥1,500 physical CPU cores, ≥1 TB memory per node and ≥500 TB flash-class working storage. Kindly clarify whether IGNITE's objective is physical separation of compute and storage resources, or guaranteed dedicated performance and isolation. Accordingly, may bidders integrate the 500 TB scratch requirement with the high-performance parallel-storage architecture if the required ≥100 GB/s read/write performance and workload isolation are demonstrably achieved?	Allows the complete data path—ingestion → processing → storage → GPU training—to be optimized as one architecture rather than creating potentially duplicated storage islands.
6	Parallel Storage System, §26.1 / Technical Evaluation Section D, pp. 93 onward	For the parallel-storage solution, Kindly confirm whether the ≥1 PB hot-tier expansion requirement represents usable capacity after data protection. Further, will IGNITE prescribe a standard benchmark methodology for validating ≥380 GB/s read, ≥190 GB/s write and ≥6 GB/s concurrent delivery per accelerator node, or should bidders propose an OEM-validated test methodology for acceptance?	Capacity and performance must be evaluated together to correctly engineer the storage nodes, media, protection scheme and storage fabric.
7	Parallel Storage §26.1.4 – Direct Storage-to-Accelerator DMA	The RFP requires a direct data path from parallel storage to accelerator memory. Is the functional objective to achieve CPU-bypass/high-performance accelerator data movement, allowing any OEM-supported standards-based or equivalent implementation, rather than requiring a specific proprietary technology?	Maintains interoperability and allows the storage, accelerator and network ecosystem to be selected as an integrated solution rather than locking the architecture to one implementation.
8	AI Fabric / Networking Requirements; Scope Overview	For clarity of the complete network architecture, Kindly confirm whether bidder responsibility extends across all required connectivity domains—AI accelerator fabric, parallel-storage fabric, service/data network, management/OOB network, security perimeter, Internet/external connectivity and Primary-to-DR connectivity—or identify any connectivity that will be provided by IGNITE or another nominated party.	Establishes the end-to-end network boundary and prevents interface gaps between different partners/OEMs. The Prime Bidder otherwise carries end-to-end integration accountability.
9	§26.1.6 – Container Orchestration, pp. 151–152	The RFP intentionally permits a CNCF-conformant Kubernetes platform without mandating a vendor. To support long-term technology sovereignty and portability, Kindly confirm whether a production-grade open-source Kubernetes architecture with enterprise-grade 24×7 support is acceptable, or whether IGNITE expects commercial OEM subscriptions for the complete Managed Operations Term.	Determines the overall platform strategy, long-term licensing model, support responsibility and degree of vendor independence.
10	§26.1.8 – DR & Backup; Draft SLA §11	The RFP defines different RPO/RTO targets for critical, standard and bulk/archival services. Should bidders therefore design a policy-based data-protection architecture in which datasets, model artefacts, platform state and other data are protected according to their service criticality, or is the complete parallel-storage dataset expected to be replicated at the DR facility?	This is fundamental to DR storage capacity, backup sizing, replication bandwidth and overall solution economics.
11	§26.1.8 – Backup & DR; Draft SLA §11	To engineer the complete data-protection solution, Kindly confirm whether IGNITE will prescribe backup retention, backup frequency, immutability and archival requirements, or whether bidders should propose an appropriate policy aligned with the RFP-defined service tiers and RPO/RTO objectives.	Without the protection policy, backup capacity, technology and long-term storage requirements cannot be accurately sized.

12	Cross-Cutting Security / IAM / Security Operations	Should NAIH be designed with its own complete cybersecurity operations environment—including IAM, MFA/PAM, SIEM/SOC and associated security controls—or are any existing IGNITE/Government identity or SOC platforms expected to be integrated with the proposed ecosystem?	Determines whether we build a self-contained security ecosystem or an integrated security architecture and avoids duplicate technology investment.
13	§26.1.7 – Infrastructure Monitoring & Telemetry, p. 151; Cross-Cutting Monitoring/Logging	Does IGNITE expect a unified observability and operations platform covering compute, accelerators, network, storage, facilities, Kubernetes and platform services, or should the solution integrate with any existing NOC/SOC/ITSM platforms?	The ecosystem requires end-to-end operational visibility; this clarifies whether monitoring is a complete new solution or an integration requirement.
14	Colocation Requirements; Draft SLA §12, pp. 234 onward	Considering the high-density accelerator infrastructure, Kindly confirm whether bidders may propose liquid-cooled/high-density AI systems together with the necessary CDU, manifolds and facility integration. Additionally, must both proposed facilities already hold operational Tier III/TIA-942 Rated-3 certification at bid submission, or may certification be completed before deployment/Operational Acceptance?	AI system selection, colocation selection, rack density, power and cooling must be engineered together. The RFP requires operationally certified, concurrently maintainable facilities.
15	Technical Evaluation B9-2 – Core Component Implementation & Testing	The RFP requires OEM involvement for implementation and testing of “Core Components.” Kindly clarify which technology domains are considered Core Components and whether implementation by OEM-authorized/certified partners under OEM supervision/support is acceptable, while the Prime Bidder retains end-to-end integration accountability.	Critical for defining the delivery model across GPU, compute, storage, networking and platform OEMs and avoiding overlapping implementation responsibilities.
16	§31.1 Acceptance; Draft SLA §5, pp. 230 onward	Since Operational Acceptance depends on the integrated operation of compute, storage, network, orchestration, software and management layers, will IGNITE and the Successful Bidder jointly finalize an end-to-end Acceptance Test Plan during detailed design, defining workload profiles, benchmarks, test conditions and measurable pass/fail criteria before deployment?	Ensures the entire ecosystem is accepted based on agreed end-to-end outcomes rather than isolated component specifications.
17	End-to-End Capacity / Sizing	Are the capacities stated in the RFP intended as fixed minimum Day-1 capacities, or should bidders perform end-to-end sizing and may propose higher capacities where necessary to ensure that compute, GPU, storage, network and platform layers operate as a balanced solution?	Prevents us from treating individual minimum specifications as independently sized components.
18	Technology / OEM Interoperability	Does IGNITE expect bidders to propose a best-of-breed multi-vendor architecture where appropriate, provided end-to-end interoperability and single-point accountability are demonstrated, or is preference given to a more tightly integrated OEM ecosystem?	Very important for deciding NVIDIA/compute/storage/network/platform vendor strategy and consortium structure. The RFP already emphasizes interoperability and support accountability in areas such as storage.
19	Management & Control Infrastructure	Should management, control-plane and platform-support services be deployed on dedicated infrastructure independent of the AI/GPU workload clusters, and should equivalent management capability also be maintained at DR?	Defines the “brains” of the ecosystem and avoids using expensive GPU resources for management/platform services.
20	Day-1 vs. Future Expansion	For requirements specifying future scalability or expansion, Kindly clarify which capacities must be physically delivered and licensed at Day-1 versus capabilities that only need to be architecturally supported for future expansion.	This has major implications for BOQ, licensing, facility reservation and commercial pricing.
21	Prime vs Consortium Liability	Where the solution is delivered through a consortium involving a Prime Bidder, global AI/technology partner, local SI and multiple OEMs, please clarify whether contractual responsibility remains entirely with the Prime Bidder or whether specific obligations, liabilities and acceptance responsibilities may be contractually allocated to respective consortium members.	A prospective parent/prime needs to know whether it carries 100% liability for work actually delivered by other parties.
22	Responsibility Boundaries	Will IGNITE allow the proposal and subsequent contract to establish clearly defined responsibility boundaries for infrastructure, AI platform, colocation, managed operations, OEM support and other major workstreams while maintaining overall programme governance through the Prime Bidder?	Allows a workable consortium structure instead of placing every technical risk on one entity.
23	Phased Acceptance & Payment	Given the scale and multi-layer nature of the programme, can infrastructure, platform, DR and operational capabilities be commissioned through defined stage-gates with independent acceptance and corresponding commercial/payment milestones rather than making overall acceptance dependent upon completion of the entire ecosystem?	Probably one of the most important questions for partner risk. It limits accumulated delivery and cash-flow exposure.

24	Dependency / Third-Party Risk	Where achievement of a project milestone depends upon facilities, connectivity, approvals, data, access, integrations or services to be provided by IGNITE or another nominated stakeholder, how will delays outside the Successful Bidder's control be treated for schedule, SLA and acceptance purposes?	Prevents the prime from absorbing delays it cannot control.
25	Technology Evolution / Obsolescence	Given the project duration and rapid evolution of AI technologies, will a formal technology-refresh/change-control mechanism allow equivalent or superior technologies to be introduced during Build and Operate phases without being treated as deviation from the originally accepted solution?	Protects the consortium against accelerator/platform obsolescence during a multi-year programme.
26	Performance Responsibility	Where end-to-end performance depends on components from multiple OEMs and service providers, will compliance be assessed against the integrated solution's committed performance, with component-level responsibilities governed through the approved responsibility matrix and OEM support arrangements?	Important for determining who owns performance failures across GPU-network-storage-platform boundaries.
27	BOT Risk Boundary	Please clarify the point at which operational responsibility and associated service risk transfer from Build/Implementation into the Managed Operations phase, and subsequently from the Successful Bidder to IGNITE at completion of BOT.	Establishes when each major category of delivery and operational liability begins and ends.