

ANNEXURE AB

TERMS OF REFERENCE

1 BACKGROUND & INTRODUCTION

The Government of Pakistan, under its broader Digital Pakistan Vision, aims to leverage digital and emerging technologies to drive inclusive economic growth, improve public service delivery, enhance innovation & governance, and position the country competitively in the global digital economy. Within this vision, Artificial Intelligence (AI) has been identified as a transformative general-purpose technology with the potential to significantly impact productivity, innovation, and national competitiveness across multiple sectors.

To operationalize this vision, the Government of Pakistan, through Ignite – National Technology Fund under the Ministry of Information Technology & Telecommunication (MoITT), is establishing the National AI Innovation Hub (NAIIH). Ignite’s mandate is to enable and strengthen Pakistan’s innovation and entrepreneurship ecosystem by supporting applied research, technology commercialization, entrepreneurship, and industry–academia collaboration. NAIIH will build directly on this mandate by serving as a national platform that translates AI research into real-world solutions, fosters AI-driven startups, and supports national priorities including healthcare, agriculture, governance, education, finance and climate resilience.

Despite growing interest and pockets of excellence, Pakistan’s current AI and innovation landscape faces significant structural challenges. The country lacks sovereign AI compute capacity and remains heavily dependent on foreign cloud infrastructure for advanced AI development and deployment. This dependence limits the ability to securely develop, scale, and govern AI systems, constrains long-term control over technology evolution, and weakens the development of local AI engineering and operations capabilities. At the same time, AI research and innovation efforts in Pakistan remain fragmented and insufficiently connected to commercialization pathways, resulting in a weak transition from research outputs to scalable products and public-sector solutions.

The absence of domestic high-performance AI compute capacity also exposes national innovation programs to external supply chain, pricing, regulatory, and geopolitical risks, reinforcing the strategic importance of sovereign AI infrastructure as a national capability.

These challenges are reflected in Pakistan’s performance in the Global Innovation Index (GII) 2025, published by the World Intellectual Property Organization (WIPO). Pakistan is ranked 99th out of 139 economies, with a particularly low ranking of 124th in the Innovation Input Sub-Index, which measures foundational enablers such as institutions, human capital, research, infrastructure, market sophistication, and business sophistication. By contrast, Pakistan’s Innovation Output Sub-Index ranking of 75th indicates that the country generates innovation outcomes disproportionate to its level of ecosystem support. This persistent input–output gap suggests substantial untapped potential that remains unrealized due to systemic constraints, especially in infrastructure, skills, and coordinated innovation platforms.

Global labor market and skills trends further reinforce the urgency of targeted intervention. The World Economic Forum’s Future of Jobs Report 2025 highlights AI and big data as among the fastest-growing job categories globally, with approximately 75% of companies worldwide expected to adopt AI technologies by 2028. Demand for roles such as AI and Machine Learning Specialists and Data Scientists is projected to grow by over 40%, yet only about 50% of the global workforce is currently considered ready for AI-driven job transformation, with readiness levels significantly lower in developing economies. Complementary global skills analyses, including large-scale online learning platforms, indicate rapid growth in demand for AI-related competencies, further underscoring the need for structured national capacity-building pipelines.

Within this context, the National AI Innovation Hub (NAIIH) is conceived as Pakistan's flagship, long-term, system-level intervention to address national AI readiness and innovation gaps through a unified and solution-driven framework. The NAIIH integrates skills development, applied research, commercialization, technology adoption, governance, and international collaboration within a single coordinated ecosystem, with a strategic emphasis on building sustainable pipelines for AI talent, research translation, and market-ready solutions rather than isolated or short-term initiatives. Central to this vision is the establishment of sovereign AI infrastructure as a foundational enabler, including government-owned AI compute capacity and an open, vendor-neutral AI software stack to support advanced research and experimentation, enable secure and scalable model training and inference, reduce long-term dependency on external platforms, and retain national control over the evolution of AI systems. This comprehensive approach is fully aligned with the National AI Policy 2025, particularly its priorities on AI innovation ecosystems, workforce readiness, secure and responsible AI, national infrastructure development, and global partnerships, collectively positioning NAIIH as a cornerstone for Pakistan's sustainable and sovereign AI advancement.

22.1 GROWING DOMESTIC DEMAND FOR SOVEREIGN AI COMPUTE AND SERVICES

Pakistan is already witnessing significant and demonstrable demand for sovereign AI infrastructure across academia, startups, public institutions, regulated industries, and compute-intensive national sectors. The National Centre of Artificial Intelligence (NCAI) consortium and leading universities, including NUST, COMSATS, LUMS, UET Lahore, and NED, increasingly require advanced accelerator infrastructure for sovereign Urdu and regional language large language models (LLMs), financial market prediction, climate intelligence, healthcare AI, judicial analytics, and advanced scientific research. However, current access to large-scale AI compute remains fragmented, expensive, and heavily dependent on overseas hyperscale cloud providers, resulting in high USD-denominated costs, limited accessibility, data sovereignty concerns, and latency challenges for Pakistani institutions and enterprises.

Pakistan's emerging AI startup ecosystem, including companies such as Metric, Metal, and Vyro Systems, increasingly relies on AI-native products, large-scale inference, multilingual LLMs, and accelerator-intensive generative AI systems. Yet, the absence of affordable domestic AI compute infrastructure constrains product scaling, weakens cost competitiveness, and limits the ability of startups to commercialise locally relevant AI solutions. Similarly, Pakistan's growing gaming, animation, and digital content industries are increasingly adopting AI for procedural content generation, Urdu-language NPCs, real-time inference, esports analytics, and advanced rendering workloads, all of which require scalable and affordable accelerator infrastructure.

Strategic national sectors including WAPDA, the Pakistan Meteorological Department (PMD), OGDCL, and SUPARCO increasingly require high-performance compute for flood forecasting, hydrological modelling, seismic inversion, reservoir simulation, satellite imagery analytics, crop monitoring, and disaster assessment while ensuring sensitive national datasets remain within Pakistan. The devastating impact of recurring climate events, energy security challenges, and the growing need for precision agriculture underscore the urgency of establishing sovereign compute capabilities for national resilience and evidence-based planning.

At the same time, regulated sectors including banking, financial services, telecommunications, healthcare, and critical infrastructure are increasingly seeking secure sovereign AI capabilities for credit scoring, fraud detection, anti-money laundering (AML), Urdu-language customer support, network optimisation, healthcare diagnostics, and enterprise automation. Organisations such as HBL, Telenor, and PTCL are already exploring AI-driven transformation while operating within increasingly stringent domestic data governance, confidentiality, and regulatory requirements. The inability to

securely process sensitive financial, citizen, telecom, healthcare, and enterprise data on foreign platforms further reinforces the need for sovereign AI infrastructure hosted on Pakistani soil.

Pakistan’s National AI Policy 2025 and broader digital transformation agenda recognise AI as a strategic national capability with the potential to transform agriculture, industry, services, governance, and productivity. However, achieving these ambitions requires foundational digital public infrastructure capable of providing trusted compute, sovereign data processing, secure inference, and scalable access to academia, startups, enterprises, and public institutions alike. These emerging use cases collectively demonstrate a strong, immediate, and cross-sectoral national demand for sovereign, Pakistan-hosted AI infrastructure capable of supporting secure, scalable, and locally governed AI innovation, commercialization, and public service transformation.

The growing national demand for sovereign AI infrastructure spans multiple sectors of Pakistan’s economy, including academia, startups, strategic public institutions, regulated industries, and emerging digital sectors. The table below provides an illustrative overview of key demand areas, representative institutions, and the corresponding AI and compute requirements that necessitate secure, sovereign, and scalable national AI infrastructure.

Sector / Vertical	Illustrative Demand in Pakistan	Named Organisations / Beneficiaries	Key AI / Compute Requirement
Academia & Research	Sovereign Urdu/Pashto LLMs, climate modelling, healthcare AI, judicial analytics, financial market prediction	NCAI, NUST, COMSATS, LUMS, UET Lahore, NED, Punjab University	Multi-accelerator model training, sovereign research compute, inference APIs
AI Startups & Digital Economy	AI-native products, bilingual LLM inference, venture intelligence, creative AI	Metric, Metal, Vyro Systems, AI startups, SMEs	Affordable PKR-denominated inference, compute, low-latency APIs
Gaming & Animation	AI-powered NPCs, procedural content generation, esports analytics, rendering workloads	CEGA, Gamer Pakistan, gaming startups	Accelerator burst compute, real-time inference, rendering clusters
Climate, Water & Disaster Management	Flood prediction, monsoon forecasting, hydrological and climate simulation	WAPDA, PMD, NDMA	HPC-scale accelerator simulation, air-gapped sovereign processing
Agriculture & Space	Satellite imagery analysis, crop disease detection, yield forecasting, disaster assessment	SUPARCO, NARC, Provincial Agriculture Departments	Large-scale imagery analytics and inference
Oil & Gas / Energy	Seismic inversion, reservoir simulation, geological modelling	OGDCL, PPL, MPCL	High-performance confidential compute and air-gapped storage
Banking & Finance	SME credit scoring, fraud detection, AML/KYC, regulatory reporting	HLB, UBL, Meezan Bank, SBP, PSX	Confidential AI, explainable models, secure inference

Telecom & Digital Services	Urdu conversational AI, churn prediction, network optimisation, fraud detection	PTCL, Telenor, Jazz, Zong	Low-latency inference and sovereign language models
Healthcare & Pharma	Medical imaging, patient outcome prediction, Urdu clinical NLP	Aga Khan, Shaukat Khanum, PIMS, Ferozsons	Confidential AI and secure healthcare inference
Manufacturing, Logistics & Retail	Predictive maintenance, route optimisation, demand forecasting, customer intelligence	Lucky Cement, Packages Ltd, TCS, Daraz, Leopards	AI inference, analytics, computer vision, enterprise AI

2 OVERVIEW OF THE PROPOSAL

Establishing sovereign Artificial Intelligence (AI) capabilities is critical for Pakistan’s long-term economic growth, public-sector transformation, and global competitiveness. AI-driven innovation has the potential to significantly enhance productivity, enable high-value entrepreneurship, and support evidence-based decision-making across key national sectors. However, the absence of sovereign AI infrastructure and reliance on foreign cloud platforms currently limits Pakistan’s ability to securely develop, scale, and govern advanced AI solutions.

The NAIH initiative is aligned with the United Nations Sustainable Development Goals (SDGs), particularly Goal 4 (Quality Education), Goal 8 (Decent Work and Economic Growth), Goal 9 (Industry, Innovation and Infrastructure), Goal 16 (Strong Institutions), and Goal 17 (Partnerships for the Goals). By enabling access to advanced AI compute, strengthening local capabilities, and reducing dependency on external platforms, NAIH aims to build a resilient and inclusive national AI ecosystem while contributing to sustainable economic development.

The salient features of the NAIH AI Infrastructure are given below:

23.1 Provision of Sovereign AI Compute Infrastructure

There shall be provision of AI compute infrastructure, including associated AI accelerators, CPUs, memory, storage, and high-speed networking, capable of supporting advanced AI training, fine-tuning, and inference workloads. All AI compute infrastructure procured under this project shall be owned by Ignite under the Ministry of IT and Telecom.

23.2 Co-Location of AI Infrastructure in an Existing Data Center

The AI compute infrastructure, through partnership, shall be co-located in an existing Tier-III or above data center facility in Pakistan, with adequate power, cooling, physical security, and network connectivity. The co-location partner must have initiated the Cloud office accreditations process under Pakistan’s Cloud First Policy. The co-location arrangement shall strictly be limited to hosting services, while ownership and control of AI systems shall remain with the Company under the ambit of NAIH.

23.3 Deployment of Open and Vendor-Neutral AI Software Stack

The AI infrastructure shall be deployed with an open, modular, and vendor-neutral AI software stack to support orchestration, training, inference, monitoring, and governance of AI workloads. The

software stack shall ensure portability, replaceability, and long-term flexibility, avoiding proprietary lock-in.

23.4 Funding and Monitoring by the Company

Funding and monitoring of the NAIH AI Infrastructure shall be conducted by the Company (Ignite – National Technology Fund) or its authorized company for monitoring & evaluation of the project in accordance with its mandate, approved budgets, and applicable public procurement rules.

23.5 Procurement, Deployment, and Initial Operations by the Successful Bidder

Procurement, deployment, commissioning, and initial operations support of the AI infrastructure shall be carried out by the Successful Bidder. The Successful Bidder shall also be responsible for providing structured capacity building and knowledge transfer to NAIH technical teams during the operations support period.

23.6 Nature of the Setup as National AI Infrastructure

The established setup shall function as national AI infrastructure, not as a commercial cloud or managed service. The infrastructure shall support government-led AI initiatives, startups, applied research, innovation programs, and sectoral AI use cases under NAIH.

23.7 Policy-Governed and Non-Discriminatory Access

The NAIH AI Infrastructure shall be an open platform for eligible users, including public-sector entities, research institutions, startups, and ecosystem partners, subject to defined access policies, security requirements, and capacity constraints approved by the Company. Access to the AI infrastructure may be governed through quota-based allocation, priority policies, and usage-based cost allocation models as defined by the Company.

23.8 Support for AI Workloads and Innovation Programs

The AI infrastructure shall support onboarding and execution of multiple concurrent AI workloads, including experimentation, training, inference, pilots, and proof-of-concept deployments, aligned with NAIH objectives and national priorities.

23.9 User-Facing Access Portal

A user-facing access platform (NAIH Access Portal) shall be developed to enable controlled access to AI infrastructure, services, datasets, and models for approved users. The platform shall operate entirely within Pakistan, integrate with Government identity systems, and be delivered with full source code ownership, without dependency on external or vendor-controlled services.

23.10 AI Innovation Marketplace

A sovereign NAIH Innovation Marketplace shall be developed to support innovation, ecosystem engagement, and commercialization of AI solutions built on NAIH infrastructure. All platform data, and operations shall remain under full control of the Company while the solutions build will belong to the innovator.

23.11 Data Sovereignty and Air-Gapped Operation

All AI workloads and platform operations shall be executed within sovereign infrastructure in Pakistan. The AI stack shall support fully air-gapped deployment with no external connectivity, telemetry, or remote license enforcement, and no Government data shall be transmitted offshore.

23.12 Billing, Metering, and Cost Allocation

The AI platform shall support resource accounting, usage metering, and reporting to enable quota-based allocation and cost management. All such mechanisms shall operate within sovereign infrastructure under full Company control.

23.13 Capacity Building and Knowledge Transfer

A core objective of the project is to build local AI engineering and operations capacity. The Successful Bidder shall ensure structured training, hands-on operational exposure, and progressive transfer of responsibilities to NAIH teams.

23.14 Progress Reporting Requirements

The Successful Bidder shall be required to submit regular progress and performance reports covering infrastructure deployment, operations, utilization, incidents, and capacity-building activities to the Company on a regular basis.

Within this context, the present RFP represents a foundational step in operationalizing NAIH by establishing sovereign AI compute infrastructure coupled with an open, vendor-neutral AI software stack, supported by initial operations and capacity building. This infrastructure will serve as the backbone for NAIH's research, innovation, commercialization, and adoption activities, while remaining fully aligned with Pakistan's long-term national AI and digital sovereignty objectives.

3 PROGRAM OBJECTIVES & GUIDING PRINCIPLES

24.1 PROGRAM OBJECTIVES

The Ecosystem is established to deliver the following seven programme objectives. These objectives are binding statements of intent and shall govern the design, configuration, pricing model, and operation of every layer of the solution. Each requirement elsewhere in these TORs shall be read as being in service of one or more of these objectives.

Objective 1 — Affordable, Subsidised Inference for the Nation

Once the required infrastructure and compute are deployed, the Ecosystem shall provide AI inference services at rates cheaper than prevailing international market rates. The solution shall additionally provide subsidised (concessional) access to the open-weight model catalogue for Academia & Researchers and for Startups, available immediately upon operational readiness of the relevant infrastructure. The solution shall include the metering, quota, and billing capabilities necessary to operate affordable and subsidised access across the beneficiary tiers.

Objective 2 — Enterprise Hosting of Locally-Developed AI Models

The Ecosystem shall provide enterprise-grade hosting capability for AI models developed by Startups, Researchers, and Public- and Private-Sector organisations, enabling such third-party models to be served through the national inference layer. This capability shall extend and enrich the inference layer

beyond the open-weight catalogue, providing a managed, secure, multi-tenant home for home-grown models with appropriate isolation, service levels, and intellectual-property protection for model owners.

Objective 3 — Compute-as-a-Service Below International Rates

The Ecosystem shall offer Compute-as-a-Service (access to accelerated and general-purpose compute) to eligible beneficiaries at rates cheaper than prevailing international market rates. The Successful Bidder shall provide the metering, quota, allocation, and billing capabilities necessary to operate this service across the four beneficiary tiers.

Objective 4 — Sovereign LLM Fine-Tuned for Urdu and Local Context

The Ecosystem shall develop, fine-tune, and deliver a fully sovereign large language model adapted to the Urdu language, other local and regional languages, and the cultural context of Pakistan. The sovereign model and all associated artefacts shall be owned by Ignite under an unbroken permissive license chain, free of any foreign veto, usage cap, geographic, field-of-use, or output restriction.

Objective 5 — A National Developer & Innovation Ecosystem

The Ecosystem shall enable a national developer and innovation community by providing developer environments, tooling, self-service platforms, documentation, training, and the means for developers across the country to build, test, deploy, and scale AI solutions on national infrastructure.

Objective 6 — An AI Marketplace Connecting Solutions, Adopters & Investors

The Ecosystem shall operate a national AI marketplace in which AI solutions are listed with their complete technical detail, including standards-based interface documentation (e.g. OpenAPI-style specifications or equivalent) and working demonstrations, enabling discovery, evaluation, and adoption by public- and private-sector buyers. The marketplace shall additionally provide a mechanism to connect solution developers with investors, supporting commercialisation, funding, and growth of the national AI sector.

Objective 7 — Impact, Performance & Economic Measurement

The Ecosystem shall provide a comprehensive dashboard and feedback mechanism to measure adoption, performance, and the economic and sectoral impact of the national investment. This capability shall enable Ignite to monitor service performance and to analyse and report the return and impact of the programme across the economy and across sectors.

24.2 GUIDING PRINCIPLES

Sovereignty & data-residency principles

All data, models designated sovereign, and primary and disaster-recovery hosting shall reside within Pakistani territory and under Pakistani jurisdiction. Ownership of the sovereign model and all associated artefacts shall vest in Ignite.

Vendor- and technology-neutrality principles

The solution shall be procured and evaluated on fair, vendor-neutral terms, free from requirements that mandate or favour any particular manufacturer, product, or technology. The solution shall remain under the purchaser's control and free from dependence on any foreign-controlled service, licence,

or remote facility for its continued operation. Open, published standards shall be preferred at integration boundaries between the major layers of the Ecosystem.

Open-standards & anti-lock-in principles

The Contractor shall prefer open, published, and interoperable standards and interfaces throughout. Proprietary interfaces, where used, shall be documented and accompanied by an equivalent open or standards-based access path wherever technically feasible.

Sustainability & efficiency principles

The solution shall be designed for energy and cooling efficiency appropriate to the local climate, with declared and measurable efficiency metrics.

4 SCOPE OVERVIEW & DELIVERY MODEL

i. In-scope / out-of-scope summary

In scope: all six ecosystem layers and cross-cutting capabilities; primary and DR colocation hosting; sovereign-model development; affordable and subsidised tiered inference services; Compute-as-a-Service; enterprise hosting of third-party (locally-developed) models; the national developer ecosystem; the AI marketplace including investor-connection mechanisms; the impact-measurement and feedback capability; and all integration, commissioning, operations, and capability transfer.

Out-of-scope items, if any, shall be explicitly listed by the procuring authority; absent such listing, the Successful Bidder shall assume end-to-end responsibility.

ii. End-to-end stack responsibility

The Successful Bidder shall be responsible for delivering and integrating every layer of the Ecosystem into one operational system, irrespective of how many sub-vendors or partners contribute components.

iii. Single prime accountability

Bids may be submitted by a single entity or by a consortium. Where a consortium bids, a single prime Successful Bidder shall hold end-to-end contractual accountability for delivery, integration, and all service levels across every layer. The procuring authority shall deal solely with the prime Successful Bidder for all obligations under these TORs.

iv. Parallel-track delivery model

The Successful Bidder shall deliver early-value services (notably the tiered open-weight inference service) in parallel with longer-duration tracks (notably sovereign-model training and facility build-out), in accordance with the phased delivery plan required later in these TORs.

v. Relationship between TORs and Evaluation Criteria

These TORs state the binding requirements of the programme. The Technical Evaluation Criteria score a bidder's proposed capability to meet them. Every measurable threshold in the Evaluation Criteria is reflected in these TORs.

5 TERMS OF REFERENCE

The NAIIH AI Infrastructure Project shall be implemented as a national-level AI compute and operations initiative under the National AI Innovation Hub (NAIIH). The project shall comprise of a fully integrated end-to-end ecosystem, structured into six functional layers - (1) National AI Compute Infrastructure (Foundation); (2) AI Platform & Model Stack; (3) Resource Orchestration & Access; (4) Development & Deployment; (5) AI Marketplace & Adoption; (6) Analytics, Impact & Feedback — supported by cross-cutting capabilities (security, governance & compliance, monitoring & logging, and external-systems integration) and serving the four beneficiary groups. The successful bidder shall deliver all six layers and all cross-cutting capabilities as a single integrated system. The reference architecture diagram provided in Figure 1 is indicative of function only.




Figure 1: National AI Ecosystem Reference Architecture

The technical thresholds specified in this TOR have been sized to support the minimum viable national AI programme scope: distributed pre-training and fine-tuning of a sovereign LLM at 70 billion parameters; concurrent inference serving of at least 20 open-weight models; and simultaneous access by at least four beneficiary tiers. The 64-accelerator floor reflects the minimum cluster scale at which efficient tensor- and pipeline-parallel training of a 70B-parameter model is achievable without architectural compromise. The 1 PB parallel storage floor reflects the combined working-set requirements of active fine-tuning data, model checkpoints, and the inference context store at this scale. All thresholds are binding minimums; bidders proposing configurations that exceed these figures will be scored accordingly under the tiered evaluation criteria.

The Successful Bidder shall be responsible for the end-to-end design, supply, installation, integration, commissioning, operation, maintenance, and structured transfer of a production-grade, sovereign, and scalable AI infrastructure platform in accordance with the Company's requirements and international best practices.

The scope of work shall encompass the complete AI Stack as a single integrated system, including all six layers and cross-cutting capabilities comprising AI infrastructure, colocation, AI software platform

deployment, operations, security, governance, and structured capacity building. The Successful Bidder shall ensure full interoperability, operational readiness, and seamless integration across all components.

All hardware procured under this RFP shall be owned solely by the Company and shall be free from vendor-imposed usage restrictions, remote enforcement mechanisms, proprietary lock-ins, or kill-switches, ensuring full operational sovereignty and long-term national control over the infrastructure. The scope of work under this RFP includes, but is not limited to, the following components:

26.1 LAYER 1: NATIONAL AI COMPUTE INFRASTRUCTURE (FOUNDATION)

This Part defines the foundational compute, interconnect, storage, equipment-management, orchestration, monitoring, and infrastructure-resilience requirements. All numeric thresholds in this Part are binding minimums and correspond exactly to the Technical Evaluation Criteria. All requirements are technology- and vendor-neutral; any architecture, datatype, memory technology, interconnect, or file system meeting the stated outcome, evidenced by manufacturer documentation, is acceptable.

26.1.1 AI Accelerators

The Successful Bidder shall supply brand-new, training-class AI accelerators suitable for both large-scale distributed training and high-throughput inference. The accelerators may be of any architecture and from any manufacturer, provided every threshold below is met and substantiated by manufacturer datasheets. Standalone add-in-card-only configurations that do not meet the form-factor and interconnect requirements are not acceptable. Rack-scale or super-node architectures are expressly permitted.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide training-class AI accelerators in a baseboard/module-integrated form factor (integrated baseboard, mezzanine, or proprietary integrated module, or equivalent), suitable for training and inference. PCIe add-in-card-only configurations not meeting interconnect/form-factor requirements are not acceptable.	≥ 64 accelerators
2	Provide low-precision tensor compute throughput per accelerator, at a vendor-declared datatype (8-bit, 6-bit, 4-bit floating-point, 8-bit integer, or equivalent), with dense/sparse basis disclosed.	≥ 1.5 PFLOPS per accelerator
3	Provide 16-bit floating-point (or equivalent) dense training compute per accelerator.	Range ≥ 0.4 PFLOPS to ≥ 1.2 PFLOPS per accelerator
4	Provide on-package high-bandwidth memory (or equivalent on-package memory) per accelerator; every accelerator shall meet this.	Range ≥ 96 GB to ≥ 144 GB per accelerator
5	Provide sustained on-package memory bandwidth per accelerator.	≥ 3.5 TB/s per accelerator
6	Provide aggregate on-package high-bandwidth memory across the delivered fleet (host memory excluded).	Range ≥ 6TB to ≥ 18 TB aggregate
7	Provide a high-bandwidth, low-latency intra-node (scale-up) accelerator-to-accelerator interconnect of any topology, supporting direct peer memory access without host staging. This minimum reflects the tight intra-domain coupling	≥ 700 GB/s per accelerator

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
	required for efficient large-model (tensor- and pipeline-parallel) training at national scale; architectures whose intra-node accelerator-to-accelerator bandwidth falls below this threshold cannot sustain the required collective-communication efficiency and are not acceptable.	
8	Provide high-speed, RDMA-capable inter-node (scale-out) connectivity using any standards-based or proprietary RDMA transport (or equivalent), with direct accelerator-to-network paths.	RDMA-capable; per-node injection declared
9	Provide hardware-enforced partitioning of each accelerator into multiple isolated instances with dedicated compute and memory.	Hardware partitioning supported
10	Provide a hardware-isolated trusted execution environment protecting model weights and data in use from privileged host administrators.	Confidential computing / TEE supported
11	Declare per-accelerator and per-node maximum power draw and supported cooling method (air, direct-liquid, or hybrid, or equivalent).	Declared; deployable in standard or integrated racks
12	Provide a complete, supported software toolchain (compiler, drivers, optimised libraries, integration with ≥ 2 major training frameworks). Dependence on any single closed ecosystem shall not be a precondition of operation.	Complete toolchain; portability documented
13	Supply latest-generation, in-production hardware (not end-of-life/end-of-sale) through authorised official channels.	Manufacturer authorisation provided

26.1.2 AI Interconnect & Networking Fabric

The Successful Bidder shall provide a high-performance interconnect fabric linking all accelerator nodes. The bidder may propose any topology (fat-tree/Clos, mesh, dragonfly, torus, or equivalent); the fabric shall be evaluated solely on the quantified properties below, not on its architectural form. Any RDMA-capable transport (standards-based or proprietary, or equivalent) is acceptable. No specific switch, transport, or topology is mandated. All figures apply at full delivered cluster scale.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a non-blocking fabric with no oversubscription on the AI plane, across all tiers, at full cluster scale.	1:1 (zero oversubscription)
2	Provide multiple equal-cost independent (disjoint) paths between any two accelerator endpoints.	≥ 4 disjoint paths
3	Limit the single-switch fault domain.	$\leq 12.5\%$ of accelerators per switch
4	Provide topology growth headroom without adding switch tiers or re-architecting.	$\geq 25\%$ growth headroom
5	Provide aggregate fabric injection bandwidth per accelerator node.	≥ 3.2 Tb/s per node

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
6	Guarantee lossless RDMA transport with zero packet loss under full congestion (hardware flow/congestion control, or equivalent).	Zero RDMA packet loss under congestion
7	Sustain a high proportion of theoretical bisectional bandwidth on collective operations (e.g. all-reduce), by any mechanism (in-network, host/accelerator-based, or equivalent).	≥ 90% of bisectional bandwidth
8	Deliver low end-to-end accelerator-to-accelerator latency for small messages.	≤ 3 μs for messages < 4 KB
9	Provide a non-blocking storage plane allowing all nodes to reach full aggregate storage bandwidth simultaneously.	100% simultaneous reach
10	Provide a direct accelerator-to-network data path bypassing host CPU and system memory.	Host-bypass DMA path present

26.1.3 Compute Infrastructure — Host, Control-Plane & Data-Pipeline Tiers

The Successful Bidder shall provide the server infrastructure hosting the accelerators, the management/control-plane nodes, and a dedicated data-pipeline (general-purpose CPU) tier. Processor architecture is open: any current-generation enterprise server processor (64-bit x86 or Arm-based, or equivalent) from any manufacturer is acceptable. Core-count requirements are stated as aggregate per host/node so that no socket-packing or processor architecture is excluded.

i. Accelerator Compute Host Tier

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1.1	Provide accelerator host servers with sufficient processor cores per host (aggregate across sockets).	≥ 128 cores per host (aggregate)
1.2	Provide host system memory per accelerator host with a multi-channel memory subsystem (8-channel or equivalent).	≥ 1.5 TB per host
1.3	Provide high-bandwidth host-to-accelerator I/O (PCIe Gen5 or equivalent coherent link) to every accelerator.	≥ PCIe Gen5 or equivalent
1.4	Provide aggregate local high-speed (NVMe or equivalent) data-cache storage across the accelerator-host tier, supporting direct-to-accelerator staging.	≥ 240 TB aggregate (raw)
1.5	Provide mirrored, redundant OS media per host, separate from data storage.	RAID-1 (or equivalent) OS media
1.6	Provide host data-plane network ports per host, separate from the accelerator-fabric adapters.	≥ 2 × 100 GbE or equivalent
1.7	Provide redundant, high-efficiency power supplies sized to peak node draw.	N+N or A/B redundant; Titanium-grade or equivalent

ii. Management & Control-Plane Node Tier

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
2.1	Provide cluster lifecycle-management nodes in a high-availability pair with rapid automatic failover.	≥ 2 nodes; failover ≤ 60 s
2.2	Provide orchestration control-plane nodes forming a fault-tolerant quorum.	≥ 3 nodes; tolerate ≥ 1 failure
2.3	Provide processor cores (aggregate per node) and memory per control-plane node.	≥ 96 cores/node; ≥ 512 GB/node
2.4	Provide mirrored OS media plus encryption-capable local data storage per control-plane node.	RAID-1 + ≥ 15 TB; SED-capable
2.5	Provide control-plane data and dedicated out-of-band management network ports per node.	≥ 4 × 25 GbE + ≥ 2 × 1 GbE OOB or equivalent
2.6	Distribute control-plane nodes across separate racks/cabinets and power circuits.	≥ 2 racks/cabinets; separate circuits

iii. Data-Pipeline / General-Purpose Compute Tier

This tier supports data ingestion, cleansing, ETL, tokenisation, and dataset curation. It need not contain accelerators.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
3.1	Provide a dedicated, physically separate compute tier for ingestion, cleansing, ETL, tokenisation and curation, tolerant of any single-node failure without halting pipeline processing.	Dedicated tier; single-node-failure tolerant
3.2	Provide aggregate processor cores across the data-pipeline tier.	≥ 1,500 cores (aggregate)
3.3	Provide memory per data-pipeline node for in-memory dataset transforms.	≥ 1 TB per node
3.4	Provide high-throughput working/scratch distributed block storage (flash-class, NVMe or equivalent) with sustained aggregate read and write throughput. The offered hardware and software should support deduplication, compression function, as well as two-copy, three copy and EC data protection modes.	≥ 500 TB raw; ≥ 100 GB/s read and write (aggregate)
3.5	Provide data-plane network connectivity per data-pipeline node.	≥ 2 × 25 GbE or equivalent
3.6	Provide a continuous data path from the pipeline tier to the parallel storage system and onward to the accelerators.	End-to-end data path demonstrated

26.1.4 Parallel Storage System

The Successful Bidder shall provide an external, shared, parallel storage system presenting tiered namespaces. The system is specified by capacity, performance, and data-path outcomes, not by vendor, product, controller architecture, node count, or file-system internals. Tiers are defined by performance role and capacity, not by media type; any media (NVMe, high-density flash, hybrid, or

equivalent) meeting the figures is acceptable. Storage-fabric switches are not required to share a vendor with the compute fabric. Performance floors are anchored to the minimum delivered scale (≥ 64 accelerators).

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide total raw capacity across the external parallel storage system.	≥ 1 PB raw
2	Provide a high-performance hot tier (usable, post-data-protection) for the active training working set.	≥ 200 TB usable; non-disruptively expandable to ≥ 1 PB
3	Provide a bulk-capacity tier for the remainder of the pool with policy driven tiered namespaces shared with the hot tier.	Parallel Tiered namespace
4	Provide sustained aggregate read throughput to the accelerator fleet at full scale.	≥ 380 GB/s read
5	Provide sustained aggregate write throughput (checkpoint-class) at full scale.	≥ 190 GB/s write
6	Deliver sustained storage bandwidth to every accelerator node concurrently, with no node starved under full all-node load.	≥ 6 GB/s per node concurrently
7	Provide a direct storage-to-accelerator DMA path bypassing host CPU and system memory (any implementation, or equivalent).	Host-bypass DMA path present
8	Present a single parallel/distributed file-system namespace concurrently mountable by all accelerator and pipeline nodes.	Parallel tiered namespace
9	Scale capacity and throughput together near-linearly without re-architecting the namespace.	$\geq 80\%$ linear efficiency to $\geq 2\times$ capacity
10	Enforce per-tenant capacity quotas and bandwidth/IOPS QoS (hardware-offloaded or software-defined, or equivalent), plus role-based access control.	Multi-tenancy & QoS enforced
11	Provide data protection (erasure coding/RAID Protection or equivalent) tolerating concurrent disks failure/device failures with no data loss and no namespace outage during rebuild.	≥ 2 concurrent device failures
12	Connect storage over a standards-based RDMA-capable fabric; storage switches need not share a vendor with the compute fabric, provided an interoperable, single-support-accountable solution is delivered.	Standards-based RDMA fabric

26.1.5 Infrastructure & Equipment Management

The Successful Bidder shall provide out-of-band management, equipment-level power and cooling, and hardware security for all delivered equipment. Facility-level power and cooling capacity are addressed under the colocation requirements. Management interfaces shall be open and standards-based; any equivalent open management standard is acceptable.

i. Out-of-Band Management

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1.1	Provide a physically isolated, redundant management network carrying no compute, storage, or tenant data-plane traffic.	Dedicated; ≥ 2 switches, hot-standby
1.2	Provide dedicated management-controller access (primary + failover) for every node on the management network only.	Per-node management access
1.3	Provide an open, standards-based management API (Redfish or equivalent) on all controllers for provisioning, power control, inventory, telemetry, and firmware updates.	Open management API (Redfish or equivalent)
1.4	Provide remote console access with session recording and audit logging.	KVM-over-IP with audit
1.5	Provide per-node and per-accelerator telemetry with threshold alerting forwarded to operations promptly.	Alert forwarding ≤ 30 s of breach
1.6	Provide per-outlet metered power-distribution units, remotely manageable via an open protocol.	Metered, remotely managed PDUs
1.7	Provide automated hardware inventory/configuration management updated on hardware-change events.	Automated inventory / CMDB

ii. Equipment Power, Thermal & Cooling

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
2.1	Provide redundant, high-efficiency power delivery across all node classes, with no single supply or feed as a single point of failure.	N+N or A/B; Titanium-grade or equivalent
2.2	Maintain management-plane power independence (UPS-protected, separate circuits) during a primary-power event.	OOB power independence
2.3	Declare per-rack/per-cabinet and total peak power, reconciled with the colocation facility capacity.	Declared and reconciled
2.4	Dissipate the full declared per-rack/per-cabinet thermal load at sustained full load with no throttling, by any cooling method (air, direct-liquid, hybrid, or equivalent).	No throttling at sustained 100% load
2.5	Provide redundant, hot-swappable cooling components so a single cooling-component failure causes no throttling or shutdown.	Cooling redundancy
2.6	Ensure equipment (standard-rack or integrated-cabinet/super-node form) is physically compatible with facility floor-loading, dimensions, busbar, and cooling interfaces.	Facility-compatible; form-factor declared

iii. Physical & Hardware Security

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
3.1	Provide a hardware root of trust with secure boot and a hardware security/trusted platform module (TPM 2.0 or recognised equivalent) on all node classes.	Hardware root of trust + secure boot
3.2	Verify cryptographically signed firmware at every boot, with lockdown and alerting on attestation failure.	Signed firmware; attestation alerting
3.3	Provide chassis intrusion detection logged and alerted via the management plane.	Tamper detection

26.1.6 Container Orchestration (Foundation Level)

The Successful Bidder shall provide a container-orchestration platform as the foundation for the platform and application layers — a conformant, open container-orchestration system (a CNCF-conformant Kubernetes distribution or equivalent). No specific distribution or vendor is mandated.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a conformant, production-grade container-orchestration platform across the cluster (Kubernetes-conformant or equivalent).	Conformant orchestration platform
2	Provide accelerator device-plugin (or equivalent) integration so accelerators are schedulable resources.	Accelerator scheduling integrated
3	Provide a container runtime and an offline-capable container registry.	Runtime + registry (air-gap capable)

26.1.7 Infrastructure Monitoring & Telemetry

The Successful Bidder shall provide infrastructure-level monitoring, telemetry, and alerting spanning all delivered equipment, using standards-based interfaces (Prometheus-compatible metrics, Open Telemetry, or equivalent).

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide cluster-wide infrastructure monitoring (compute, accelerator, network, storage, power, cooling) through standards-based interfaces (or equivalent).	Standards-based monitoring
2	Provide telemetry collection, retention, and threshold-based alerting to a central operations function.	Centralised telemetry & alerting

26.1.8 Disaster Recovery & Backup (Infrastructure)

The Successful Bidder shall provide infrastructure-level backup and disaster-recovery capability. Facility-level disaster recovery (primary and secondary sites) is specified under the colocation requirements; this section addresses data and system protection operating across those sites.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a backup architecture for system state, configurations, datasets, and model artefacts, with defined retention.	Backup architecture with retention
2	Provide infrastructure-level disaster-recovery mechanisms operating with the primary and secondary facilities (see colocation requirements).	Infra-level DR integrated with sites

26.2 LAYER 2: AI PLATFORM & MODEL STACK

This Part defines the AI software stack, MLOps and data-pipeline platform, the sovereign large language model, the tiered public inference service, the enterprise hosting capability for locally-developed models, and knowledge transfer. It directly serves Objectives 1 (affordable/subsidised inference), 2 (enterprise hosting), 4 (sovereign Urdu/local-context model), and 5 (developer ecosystem). All capability requirements are evaluated by function, not by product; any vendor-native, third-party, or open-source implementation that delivers the stated function is acceptable, and reliance on any single closed ecosystem shall not be a precondition of compliance.

26.2.1 AI Software & Orchestration Stack

The Successful Bidder shall provide a complete AI software stack covering the model lifecycle from training through optimised, multi-tenant inference. Each capability may be satisfied by any implementation (vendor-native, third-party, or open-source, or equivalent). Framework compatibility may be met by native support or by a documented, supported import/conversion path.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a supported low-level accelerator runtime and optimised operator/kernel library with a documented programming interface.	Runtime + operator library
2	Provide a distributed training framework supporting multi-node data/tensor/pipeline parallelism, with native support for or a documented import path from a widely-used framework and at least one other.	Distributed training; ≥ 2 frameworks
3	Provide tooling for fine-tuning and alignment (supervised fine-tuning, preference/reinforcement methods, or equivalent) and dataset curation.	Fine-tuning & alignment tooling
4	Provide a production inference serving engine with dynamic/continuous batching, concurrent multi-model execution, multi-accelerator/multi-node parallelism, and standards-based endpoints with a metrics interface.	Production serving engine
5	Provide inference optimisation: quantisation (vendor-declared low-precision formats), key-value cache management, and throughput/latency optimisation (in-flight batching, speculative decoding, or equivalent).	Inference optimisation
6	Expose inference via an OpenAI-compatible (or documented equivalent open) API for client portability.	Open, portable inference API
7	Provide accelerator-native libraries for dataframe/ETL, classical ML, and/or graph analytics to accelerate pipelines.	Accelerated data-science libraries
8	Run the full stack on the container-orchestration platform with accelerator device-plugin (or equivalent) integration.	Orchestration-native
9	Provide automated lifecycle operators for accelerator drivers/runtime and for the high-performance network/RDMA components.	Accelerator & network operators
10	Provide serverless inference with concurrency/throughput-based autoscaling, scale-to-zero, versioned model management, and traffic-splitting/canary rollout with rollback.	Autoscaling + versioned rollout

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
11	Provide fractional accelerator allocation with hardware- or software-enforced memory isolation between co-located workloads.	≥ 4 isolated workloads per accelerator
12	Provide multi-tenant isolation (tenant namespaces/virtual clusters), resource quotas, network policy, and role-based access control.	Multi-tenancy & RBAC
13	Provide a zero-trust control plane (no inbound open ports to the cluster API; not exposed to the public internet) with a full user-level audit trail.	Zero-trust control plane
14	Operate fully in an air-gapped/on-premises environment with an offline registry and offline model/artifact distribution.	Air-gapped operation
15	Provide real-time per-tenant/project utilisation metering (accelerator/CPU/memory) and token-level inference metering with chargeback/cost-allocation APIs.	FinOps metering & chargeback
16	Provide 24x7x365 enterprise support with a defined Severity-1 response time and a long-term-support branch with a stated API-stability duration.	Enterprise SLA + LTS
17	Provide automatic checkpointing and resumption of distributed training after worker/node failure, with a stated maximum lost-work window.	Fault-tolerant training
18	Provide cluster-wide observability (metrics, logging, accelerator/job telemetry) via standards-based interfaces (Prometheus-compatible, OpenTelemetry, or equivalent).	Unified observability

26.2.2 AI Cluster Operations & Management Platform (Full-Stack)

The Successful Bidder shall provide a unified, full-stack AI cluster operations and management platform (an “AI factory control plane”, full-stack cluster lifecycle manager, or equivalent) that manages the entire cluster end-to-end through a single pane of glass. The platform shall integrate provisioning, scheduling, monitoring, fault recovery, multi-tenancy, and operations across all delivered infrastructure. It is specified by function; any vendor-native, third-party, or open-source implementation that delivers the stated capabilities is acceptable, and no specific product is mandated.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide unified, single-pane-of-glass management of the entire cluster (compute, accelerators, network, storage) from a single administrative control plane.	Single-pane full-cluster management
2	Automate end-to-end provisioning, configuration, imaging, and scaling of nodes from a few to many thousands of accelerators.	Automated provisioning at scale
3	Provide integrated workload scheduling/orchestration for both batch (HPC/training) and container (inference/services) workloads, with quota and priority across tenants/tiers.	Integrated scheduling (batch + container)
4	Provide cluster-wide telemetry, observability, and health monitoring via standards-based interfaces (or equivalent).	Cluster-wide telemetry & observability

5	Provide automated failure detection and autonomous self-recovery to minimise downtime and time-to-recovery for node/job failures.	Automated fault detection & self-recovery
6	Provide validation and diagnostics (pre-flight checks, burn-in, continuous health-checks) for the cluster and its components.	Validation & diagnostics
7	Provide multi-tenant operations: tenant isolation, role-based administration, quota/chargeback, and per-tenant operational views.	Multi-tenant operations & chargeback
8	Operate fully on-premises/air-gapped with no dependency on external cloud services for core management functions.	Air-gapped operability
9	Integrate with the orchestration (Layer 3), software stack (Sec. 26.2.1), and monitoring (Sec.26.1.7) layers as one coherent operations capability.	Integrated with platform layers

26.2.3 MLOps, Model Registry & Data Pipelines

The Successful Bidder shall provide the MLOps platform, model registry, dataset management, and AI safety controls underpinning model development and operation.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide MLOps pipelines for reproducible training, fine-tuning, evaluation, and deployment.	MLOps pipelines
2	Provide a model registry with versioning, lineage, and stage promotion.	Model registry & versioning
3	Provide dataset management and data-pipeline tooling integrated with the data-pipeline compute tier and parallel storage.	Dataset management & pipelines
4	Provide AI safety and guardrail capabilities (content controls, policy enforcement) applicable to served models.	AI safety & guardrails

26.2.4 Sovereign LLM Development (Objective 4)

The Successful Bidder shall develop and deliver a fully sovereign large language model fine-tuned for the Urdu language, other local and regional languages, and the cultural context of Pakistan, together with the governed data pipelines and partnerships required to build it. The defining requirement is sovereignty: the delivered model and all artefacts shall be owned by Ignite under an unbroken permissive license chain, free of any foreign veto, usage cap, geographic, field-of-use, or output restriction.

i. Training & fine-tuning platform

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1.1	Provide a fully configured end-to-end training and fine-tuning platform on the delivered infrastructure (pre-training, continued pre-training, full and parameter-efficient fine-tuning).	Operational training platform
1.2	Provide experiment tracking, dataset/version lineage, model registry, checkpointing, and reproducible runs.	Experiment & lifecycle management

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1.3	Provide an automated evaluation harness including Urdu/local-language and domain-specific evaluation sets.	Evaluation harness
1.4	Ensure the platform is fully operable by the in-country team at contract end, with no licensing or technical lock preventing independent fine-tuning.	Handover-ready; no lock-in

ii. Governed data pipelines & local-context datasets

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
2.1	Establish a working data-partnership framework and arrangements with willing Pakistani organisations under Ignite governance and consent/legal compliance.	Partnership framework + governance
2.2	Assemble, cleanse, and curate Urdu and other local-language corpora and local-context datasets to a stated minimum scale and quality.	Urdu & local corpora (scale/quality stated)
2.3	Operate production data pipelines (ingestion, cleansing, de-duplication, PII handling, curation, training-ready) integrated with the compute and storage tiers.	End-to-end pipelines operational
2.4	Ensure all collected/curated data is resident in-country and owned by the Government of Pakistan, with full provenance, audit trail, and consent/withdrawal handling.	In-country; GoP-owned; auditable

iii. Sovereign fine-tuned model

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
3.1	Base the sovereign model on a base model under Apache 2.0, or an equivalent OSI-approved permissive license providing an explicit patent grant, with an unbroken permissive license chain from base weights through every intermediate artefact to the delivered model, such that the result is freely modifiable, redistributable, and commercialisable by the Government of Pakistan with no foreign veto, usage cap, geographic, field-of-use, or output restriction.	Apache 2.0 or equivalent permissive; unbroken chain
3.2	Fine-tune/align the model on the Urdu and local-context datasets, with demonstrated capability uplift on Urdu/local evaluation sets versus the un-tuned base.	Measured Urdu/local uplift
3.3	Transfer to Ignite the complete artefact set: model weights, training/fine-tuning code, datasets, configurations, and documentation, sufficient for independent retraining and redeployment.	Full sovereignty transfer
3.4	Apply safety alignment, content controls, and bias/appropriateness handling tuned to the Pakistani legal and cultural context.	Local-context alignment & safety
3.5	Provide a full model card (architecture, training-data summary, capabilities, limitations, evaluation results, intended-use and licensing statement).	Model card delivered

26.2.5 Tiered Public Inference Service (Objective 1)

The Successful Bidder shall stand up a public, tiered inference service offering a catalogue of open-weight models, launched in parallel with sovereign-model training so that public value is delivered early. Consistent with Objective 1, the service shall be priced below prevailing international market rates, with subsidised (concessional) access for Academia & Researchers and for Startups. The open-weight models in the catalogue need not be permissively licensed, provided each model's license permits the intended hosted-inference service.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a catalogue of open-weight models in production inference, launched in parallel with (not after) sovereign-model training, spanning a range of sizes and capabilities (general, coding, reasoning, multilingual).	≥ 20 open-weight models
2	Operate four distinct access tiers — Academia & Researchers, Startups, Public Sector, Private Sector — each with defined eligibility, quota/rate limits, data-handling/residency terms, and service levels, enforced via platform multi-tenancy/RBAC.	Four-tier access model
3	Provide subsidised (concessional) access for Academia & Researchers and Startups, and overall service rates cheaper than prevailing international market rates.	Subsidised + below-international pricing
4	Make the inference service live and publicly accessible within a stated timeline from contract start, with a committed uptime SLA.	Parallel launch; uptime SLA committed
5	Provide a standards-based API (OpenAI-compatible or documented equivalent), self-service onboarding, authentication, and per-tier metering and quota enforcement.	Open API + per-tier metering
6	Provide per-tier, per-model usage metering and reporting to Ignite (utilisation, adoption, cost allocation).	Usage metering & reporting

26.2.6 Enterprise Hosting of Locally-Developed Models (Objective 2)

The Successful Bidder shall provide the capability to host, on the national inference layer, AI models developed by Startups, Researchers, and Public- and Private-Sector organisations, enriching the inference layer beyond the open-weight catalogue. The commercial terms for any hosted third-party model are determined case-by-case once operational and are not fixed by these TORs; the Successful Bidder's obligation is to provide the hosting capability and controls.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a managed, secure, multi-tenant hosting capability enabling third-party (locally-developed) models to be deployed and served through the national inference layer.	Third-party model hosting capability
2	Provide tenant isolation and intellectual-property protection appropriate to model owners (isolation of weights, data, and endpoints).	Isolation + IP protection

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
3	Provide onboarding, deployment, versioning, metering, and service-level controls for hosted third-party models, supporting flexible commercial arrangements.	Hosting lifecycle & metering controls
4	Support configurable access and pricing arrangements per hosted model (subsidised, tiered, or owner-governed), without fixing such terms in advance.	Configurable access/pricing support

26.2.7 Knowledge Transfer, Skilling & Ecosystem Capability Building

The Successful Bidder shall transfer the knowledge, skills, and institutional capacity required for Pakistan to operate, sustain, and grow the entire Ecosystem independently — fulfilling Objective 5 (a national developer and innovation ecosystem) and reinforcing the programme’s sovereignty goals. This obligation extends beyond the operations team to the wider national developer, research, and startup community, and continues through a defined post-go-live support and mentoring period as set out in the SLA schedule. All capability-building shall be designed for local self-sufficiency, including the propagation of training capacity to Pakistani trainers and institutions ensuring the curricula is designed for heterogeneous accelerator and software-stack environments, and shall not be limited to the stack of the delivered platform.

i. Operations Team Capability Transfer

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1.1	Provide structured training for the in-country team covering platform operation, pipeline management, fine-tuning, and inference-service operation.	Operator & developer training
1.2	Demonstrate, before contract close, that the in-country team can run a full fine-tuning cycle and operate the inference service independently (acceptance gate).	Independent-operation acceptance gate

ii. National Developer Skilling & Curriculum

Ref	Requirement (the Successful Bidder shall...)	Basis
2.1	Develop a structured, tiered curriculum (foundational, advanced, and specialist levels) covering platform use, model fine-tuning, inference, data pipelines, and AI application development on the Ecosystem.	Tiered national curriculum
2.2	Develop accredited training modules and certification pathways for Pakistani developers, designed for accreditation by a Pakistani national body (e.g. HEC, NAVTTC, a university consortium, or Ignite) so that skill levels are standardised across the ecosystem; the Successful Bidder shall support the accreditation process with the relevant national institution(s).	National-accredited certification
2.3	Deliver a “training -of-trainer” programme that builds a cadre of qualified Pakistani trainers and institutions able to continue delivering the curriculum and certifications independently of the Successful Bidder.	Train-the-trainer / local propagation
2.4	Provide localised training materials and delivery, including in Urdu, to ensure accessibility across the national developer community.	Localised materials (incl. Urdu)
2.5	Transfer all curriculum content, assessment materials, and certification frameworks to Ignite for ongoing national use, with rights to maintain and extend them.	Curriculum IP transfer to GoP

iii. Community Support & Engagement

Ref	Requirement (the Successful Bidder shall...)	Basis
3.1	Establish and operate community support mechanisms — forums, discussion platforms, and knowledge bases — for Researchers, Startups, and the wider developer community, beyond static documentation.	Community forums / knowledge base
3.2	Provide a technical helpdesk / support desk for ecosystem users (Researchers, Startups, public- and private-sector developers), with published response expectations defined in the SLA schedule, distinct from the platform enterprise SLA.	Ecosystem-user helpdesk (per SLA)
3.3	Run regular developer-engagement activities — “office hours”, clinics, and developer-advocate sessions — to assist users in building and deploying solutions.	Office hours / developer advocacy
3.4	Conduct outreach and innovation programmes (e.g. hackathons, challenges, or innovation calls) to drive adoption and surface local AI use-cases across sectors.	Innovation / outreach programmes
3.5	Provide tier-specific onboarding pathways for the four beneficiary groups (Academia & Researchers, Startups, Public Sector, Private Sector), easing entry for first-time users.	Tier-specific onboarding

iv. Post-Go-Live Support & Sustained Capability

Ref	Requirement (the Successful Bidder shall...)	Basis
4.1	Continue community support, skilling, and mentoring for a defined post-go-live period as specified in the SLA schedule, to ensure the ecosystem takes root before the Successful Bidder withdraws.	Post-go-live support (per SLA)
4.2	Demonstrate, before the end of the support period, that Pakistani teams and trainers can independently operate the platform, deliver the curriculum, and sustain community support without Successful Bidder dependence.	Sustained-independence demonstration

26.3 LAYER 3: RESOURCE ORCHESTRATION & ACCESS

This Part defines how compute and platform resources are allocated, quota-managed, provisioned, access-controlled, metered, and billed across the four beneficiary tiers. It is the operational layer that makes Objectives 1 and 3 deliverable in practice: it provides the quota, credit, metering, and billing machinery through which affordable and subsidised inference (Objective 1) and Compute-as-a-Service at rates cheaper than prevailing international market rates (Objective 3) are operated. All capabilities are specified by function; any implementation (vendor-native, third-party, or open-source, or equivalent) that delivers the stated outcome is acceptable, and no proprietary platform is mandated.

26.3.1 Compute Allocation & Scheduling

The Successful Bidder shall provide a compute allocation and scheduling capability that places training, fine-tuning, and inference workloads onto accelerated and general-purpose compute efficiently and fairly across tenants and tiers.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a compute allocation engine that schedules workloads across the accelerator and general-purpose compute pools with high utilisation.	Allocation engine across pools
2	Provide job scheduling and queue management for batch, interactive, and service workloads, including gang/topology-aware scheduling for distributed training.	Scheduling + queue management
3	Enforce fair-share, priority, and pre-emption policies across the four beneficiary tiers, configurable by the governance authority.	Fair-share & priority across tiers
4	Support reservation of capacity for specific tenants, programmes, or time windows (e.g. sovereign-model training, subsidised research allocations).	Capacity reservation

26.3.2 Quota & Credit Management

The Successful Bidder shall provide a quota and credit-management capability that operationalises the tiered, subsidised access model. This capability is the mechanism by which subsidised access for Academia & Researchers and Startups (Objective 1) and differentiated CaaS entitlements (Objective 3) are administered and controlled.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a per-tier and per-tenant quota system governing consumption of inference, training, and compute resources.	Per-tier / per-tenant quotas
2	Provide a credit-allocation and top-up mechanism (including grant, subsidy, and concessional credits) administrable by the governance authority.	Credit allocation incl. subsidy/grants
3	Support resource-reservation and entitlement mechanisms tied to credits and quotas, with enforcement at allocation and runtime.	Reservation & entitlement enforcement
4	Provide configurable policies to apply subsidised/concessional rates to eligible tiers (Academia & Researchers, Startups) and standard rates to others.	Configurable subsidised vs standard policy

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
5	Provide alerting, soft/hard limits, and controlled overage handling to prevent uncontrolled consumption.	Limits & overage controls

26.3.3 Environment Provisioning

The Successful Bidder shall provide self-service provisioning of working environments for the national developer ecosystem (Objective 5), with strict tenant isolation.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide self-service provisioning of containers and managed notebook/interactive environments for eligible users.	Self-service container/notebook provisioning
2	Provide a curated template and image catalogue (frameworks, runtimes, toolchains) for rapid environment creation.	Template/image catalogue
3	Enforce tenant workspace isolation (compute, storage, network) between users and organisations.	Tenant workspace isolation
4	Integrate provisioning with quota, credit, and access-control so provisioned resources consume the correct entitlements.	Provisioning tied to quota/IAM

26.3.4 Access Control (Identity & Access Management)

The Successful Bidder shall provide identity and access management spanning all beneficiary tiers and all layers, using open standards (e.g. OIDC/SAML or equivalent). No proprietary identity system is mandated.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide identity and authentication using open, standards-based protocols (OIDC, SAML, or equivalent), with multi-factor authentication.	Standards-based auth + MFA
2	Provide role-based and attribute-based access control mapped to the four beneficiary tiers and to organisational roles.	RBAC/ABAC across tiers
3	Support federation with national and enterprise identity providers (or equivalent), enabling organisations to use their own identities.	Identity federation
4	Provide full, tamper-evident audit logging of authentication and authorisation events.	Auditable access events

26.3.5 Usage Metering, Billing & Resource Monitoring

The Successful Bidder shall provide metering, billing, and resource monitoring that support affordable and subsidised pricing (Objectives 1 and 3) and feed the impact-measurement capability (Objective 7). Pricing is operated below prevailing international market rates; the metering and billing machinery shall make such pricing administrable and transparent.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide per-tenant, per-tier, per-service, and per-model usage metering (accelerator/CPU/memory/storage and token-level inference).	Granular usage metering
2	Provide billing, invoicing, and cost-allocation supporting subsidised, concessional, and standard rate structures across tiers.	Billing & cost allocation incl. subsidy
3	Provide resource-monitoring dashboards for operators and tenants (utilisation, quota consumption, cost).	Resource-monitoring dashboards
4	Expose metering and billing data via APIs to the analytics and impact-measurement capability (Part G).	Metering/billing APIs to analytics

26.4 LAYER 4: DEVELOPMENT & DEPLOYMENT

This Part defines the developer-facing platform that enables the national developer and innovation ecosystem (Objective 5): the environments, workspaces, testing and validation frameworks, continuous integration and delivery, security scanning, and documentation/packaging through which developers across the country build, test, deploy, and scale AI solutions on national infrastructure. All capabilities are specified by function; any implementation (vendor-native, third-party, or open-source, or equivalent) that delivers the stated outcome is acceptable, and no proprietary toolchain is mandated.

26.4.1 Developer Environments

The Successful Bidder shall provide managed developer environments enabling users across all four beneficiary tiers to develop AI solutions on the platform, integrated with the orchestration, quota, and access-control capabilities of layer 3.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide managed integrated development environments and interactive notebook environments accessible to eligible users.	IDEs + managed notebooks
2	Provide isolated developer sandboxes with controlled access to compute, data, and SDKs.	Isolated developer sandboxes
3	Provide SDKs and client libraries for the platform's APIs and services (in commonly-used languages, or equivalent).	SDKs / client libraries
4	Integrate environments with identity, quota, and metering so usage consumes the correct entitlements.	Integrated with IAM/quota/metering

26.4.2 Model Training & Fine-Tuning Workspaces (Developer-Facing)

The Successful Bidder shall provide self-service, developer-facing workspaces for model training and fine-tuning, distinct from the sovereign-model platform but built on the same underlying capabilities, enabling third parties to train and fine-tune their own models (supporting Objective 2 hosting and Objective 5).

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide self-service training and fine-tuning workspaces with access to scheduled accelerated compute within entitlements.	Self-service fine-tuning workspaces
2	Provide experiment tracking, dataset access, and checkpoint management for developer workloads.	Experiment tracking access
3	Provide curated base-model and dataset access (including permissively-licensed bases) for developer fine-tuning, with license clarity.	Curated bases/datasets with license clarity

26.4.3 Testing & Validation Frameworks

The Successful Bidder shall provide frameworks to test, evaluate, and benchmark both models and applications before deployment to the inference and marketplace layers.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a model evaluation harness, including Urdu/local-language and domain-specific evaluation sets, reusable by developers.	Model evaluation harness (incl. local sets)
2	Provide application testing frameworks (functional, integration, load) for AI-enabled applications.	Application testing frameworks
3	Provide benchmarking tools for performance, latency, and quality, with comparable, repeatable methodology.	Benchmarking tools
4	Provide validation checks that can be enforced prior to deployment or marketplace listing.	Pre-deployment validation checks

26.4.4 Documentation & Packaging

The Successful Bidder shall provide artifact packaging and developer documentation capabilities, ensuring solutions can be packaged for deployment and documented for the marketplace and for adopters (supporting Objective 6).

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide standardised artifact packaging for models and applications (containers, model formats, or equivalent).	Standardised packaging
2	Provide generation and hosting of developer/API documentation (standards-based, e.g. OpenAPI-style or equivalent).	Developer/API documentation
3	Support packaging of solutions for marketplace listing with their technical detail and demonstration assets.	Marketplace-ready packaging

26.5 LAYER 5: AI MARKETPLACE & ADOPTION ECOSYSTEM

This Part defines the national AI marketplace through which AI solutions are discovered, evaluated, adopted, transacted, and scaled — the home of Objective 6. The marketplace lists AI solutions with their complete technical detail, including standards-based interface documentation (e.g. OpenAPI-style specifications or equivalent) and working demonstrations, enabling discovery and adoption by public- and private-sector buyers; and it provides a mechanism to connect solution developers with investors, supporting commercialisation and growth of the national AI sector. All capabilities are specified by function; any implementation (vendor-native, third-party, or open-source, or equivalent) that delivers the stated outcome is acceptable, and no proprietary marketplace platform is mandated.

26.5.1 AI Marketplace Portal — Build & Delivery

Ref	Requirement (the Solution shall...)	Basis
1	Deliver a production-grade Marketplace web portal implementing Sections 26.5.2–26.5.8 as one integrated application.	Deployable portal application
2	Provide a solution & model catalogue with full technical detail, standards-based API documentation (OpenAPI-style or equivalent), and working demos/playgrounds per Sec.26.5.2–26.5.3.	Catalogue + docs + demos
3	Provide commercialization workflows — procurement, contracting/onboarding, subscription & usage-based billing, and revenue-sharing — per Sec.26.5.4.	Commercialization workflows
4	Provide trust & validation (ratings/reviews, performance benchmarking, certification/compliance badges) per Sec.26.5.5.	Trust & validation
5	Provide integration & adoption surfaces (SDKs/integration kits, developer documentation portal, enterprise integration APIs) per Sec.26.5.6.	Integration & adoption
6	Provide usage & adoption analytics feeding the impact-measurement capability (Layer 6) per Sec.26.5.7.	Usage & adoption analytics
7	Provide the investor-connection / matchmaking capability per Sec.26.5.8, under governance and confidentiality controls.	Investor connection
8	Integrate the portal with identity, entitlement, quota, metering, and billing (Layer 3), and operate it under Ignite governance.	Platform-integrated & governed

The Successful Bidder shall design, build, and deliver a national AI Marketplace web portal (the “Marketplace”) as an integrated, production-grade application implementing the full set of marketplace capabilities defined in Sections 26.5.2–26.5.8 below. The Marketplace is the public adoption surface of the Ecosystem (Objective 6): it lists AI solutions and models with their complete technical detail — including standards-based interface documentation (OpenAPI-style or equivalent) and working demonstrations — enabling discovery, evaluation, procurement, and adoption by public- and private-sector buyers, and connecting solution developers with investors. The portal shall be delivered as a deployable application (not merely a specification), hosted on the national infrastructure, integrated with identity, entitlement, metering, and billing (Layer 3), and operable under Ignite governance.

26.5.2 Solution & Access

The Successful Bidder shall provide the catalogue and access surfaces through which solutions and models are published and consumed, including the AI-as-a-Service surface for the inference and compute services (supporting Objectives 1 and 3).

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a solution catalogue listing AI solutions, services, and models with complete technical detail and metadata.	Solution catalogue
2	Expose the platform model registry so approved models are discoverable and consumable through the marketplace.	Model registry exposure
3	Provide an AI-as-a-Service API marketplace surface through which inference, compute, and hosted-model services are accessed.	AI-as-a-Service marketplace
4	Integrate catalogue access with identity, entitlement, quota, and metering (Layer 3).	Integrated access/entitlement

26.5.3 Discovery & Testing

The Successful Bidder shall enable buyers and developers to discover, evaluate, and trial solutions before adoption.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide sector-based discovery (e.g. health, agriculture, finance, government, education, or equivalent classification) to help buyers find relevant solutions.	Sector-based discovery
2	Provide demonstration environments in which solutions can be evaluated safely.	Demo environments
3	Provide interactive API playgrounds for hands-on evaluation of services and models.	API playgrounds

26.5.4 Commercialization

The Successful Bidder shall provide the commercial machinery enabling solutions and services to be procured, contracted, billed, and monetised across public and private sectors, consistent with the affordable/subsidised positioning of Objectives 1 and 3.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide procurement systems supporting both public-sector and private-sector buyers, including approval and compliance workflows.	Procurement (public & private)
2	Provide contracting and onboarding workflows for buyers and solution providers.	Contracting & onboarding
3	Provide billing and monetisation supporting subscription and usage-based models, and subsidised/concessional structures where applicable.	Subscription & usage-based billing

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
4	Provide revenue-sharing mechanisms between the platform and solution providers, with transparent accounting.	Revenue-sharing mechanisms

26.5.5 Trust & Validation

The Successful Bidder shall provide trust signals enabling buyers to assess solution quality, performance, and compliance.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide ratings and reviews for listed solutions and providers.	Ratings & reviews
2	Provide performance benchmarking results for listed solutions, using comparable methodology.	Performance benchmarking
3	Provide certification and compliance badging (e.g. security, data-protection, quality) administrable under governance.	Certification & compliance badges

26.5.6 Integration & Adoption

The Successful Bidder shall provide the developer- and enterprise-facing means to integrate and adopt marketplace solutions, supporting Objective 5 and Objective 6.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide SDKs and integration kits enabling rapid integration of marketplace solutions and services.	SDKs & integration kits
2	Provide a developer documentation portal with standards-based API documentation (OpenAPI-style or equivalent) and examples.	Developer documentation portal
3	Provide enterprise integration APIs enabling organisations to embed marketplace services into their own systems.	Enterprise integration APIs

26.5.7 Usage Tracking (Marketplace)

The Successful Bidder shall provide marketplace-level usage and adoption tracking, feeding the analytics and impact-measurement capability (Objective 7).

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide API and service usage analytics per solution, provider, and buyer.	API usage analytics
2	Provide adoption metrics (uptake, retention, growth) across sectors and tiers.	Adoption metrics
3	Feed marketplace usage and adoption data to the impact-measurement capability (Part G).	Data feed to impact analytics

26.5.8 Investor Connection & Commercialisation Support (Objective 6)

The Successful Bidder shall provide a mechanism connecting solution developers — particularly Startups and Researchers — with investors, to support commercialisation, funding, and growth of the national AI sector. This is a distinguishing requirement of the marketplace and shall be delivered as a managed capability, not merely a static listing.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a means for solution developers to present solutions, traction, and technical detail to a curated investor audience.	Developer-to-investor presentation
2	Provide an investor-facing discovery and matchmaking capability to identify solutions and teams of interest.	Investor discovery & matchmaking
3	Provide secure sharing of relevant technical, usage, and performance information to support investment evaluation, under appropriate consent and confidentiality controls.	Secure, consented information sharing
4	Operate the capability under governance, with eligibility, transparency, and conflict-of-interest controls.	Governed, transparent operation

26.6 LAYER 6: ANALYTICS, IMPACT & FEEDBACK

This Part defines the capability through which Ignite measures adoption, performance, and the economic and sectoral impact of the national investment — the home of Objective 7. It consolidates usage and metering data from the orchestration layer (Layer 3), the marketplace (Layer 5), and the inference and hosting services (Layer 2) into analytics, dashboards, impact measurement, and feedback loops. All capabilities are specified by function; any implementation (vendor-native, third-party, or open-source, or equivalent) that delivers the stated outcome is acceptable, and no proprietary analytics platform is mandated.

26.6.1 Usage Analytics

The Successful Bidder shall provide platform-wide usage analytics aggregating consumption across all services, tiers, and sectors.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide platform-wide usage analytics across inference, compute, hosting, development, and marketplace services.	Platform-wide usage analytics
2	Provide breakdowns by beneficiary tier, sector, organisation, service, and model.	Multi-dimensional breakdowns
3	Aggregate metering and usage data from Layers 2,3 &5 via their published interfaces.	Consolidated data ingestion

26.6.2 Impact Measurement

The Successful Bidder shall provide a capability to measure and report the economic and sectoral impact of the programme, enabling the governance authority to assess the return on the national investment.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide economic-impact measurement (e.g. value enabled, cost savings versus international alternatives, sector growth indicators, or equivalent), with a documented, repeatable methodology.	Economic-impact measurement + methodology
2	Provide sectoral-impact measurement across the sectors served (e.g. health, agriculture, finance, government, education, or equivalent).	Sectoral-impact measurement
3	Support definition and tracking of programme key performance indicators and outcomes over time.	KPI/outcome tracking over time
4	Provide reporting suitable for governance, oversight, and public accountability.	Governance/accountability reporting

26.6.3 Policy Dashboards

The Successful Bidder shall provide dashboards for Ignite to monitor the Ecosystem and inform policy.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide governance dashboards for Ignite presenting adoption, utilisation, performance, and impact.	Governance dashboards
2	Provide configurable views and drill-down by tier, sector, service, and time period.	Configurable, drill-down views
3	Provide role-based access to dashboards for authorised governance and oversight users.	Role-based dashboard access

26.6.4 Performance Monitoring

The Successful Bidder shall provide service-level performance monitoring spanning the Ecosystem's services, distinct from infrastructure monitoring (Layer 1).

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide service-level performance monitoring (availability, latency, throughput, error rates) across inference, compute, hosting, and marketplace services.	Service-level performance monitoring
2	Provide SLA tracking and reporting against committed service levels.	SLA tracking & reporting
3	Provide alerting on service-level breaches to operations and governance.	Service-level alerting

26.6.5 Feedback Collection Systems

The Successful Bidder shall provide mechanisms to collect, route, and act upon feedback from beneficiaries, closing the loop between usage and continuous improvement.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide feedback-collection mechanisms for beneficiaries across all four tiers (e.g. in-service feedback, surveys, issue reporting, or equivalent).	Beneficiary feedback collection
2	Route and track feedback to resolution and to service/product improvement.	Feedback routing & tracking
3	Feed aggregated feedback into impact measurement and governance reporting.	Feedback into impact/governance

6 CROSS-CUTTING CAPABILITIES

This Part defines capabilities that span every layer of the Ecosystem — security, governance and compliance, monitoring and logging, and integration with external systems. These requirements apply horizontally to all of Layers 1 through 6. All capabilities are specified by function; any implementation (vendor-native, third-party, or open-source, or equivalent) that delivers the stated outcome is acceptable, and no proprietary product is mandated.

27.1 SECURITY (IAM, ENCRYPTION, ZERO-TRUST)

The Successful Bidder shall apply security across all layers, consistent with the equipment-level (Layer 1) and platform-level (Layer 2) security requirements, under a coherent zero-trust posture.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Apply a zero-trust security architecture across all layers (no implicit trust; authenticate and authorise every access).	Zero-trust across all layers
2	Provide centralised identity and access management consistent with Layer 3 section 26.3.4, with multi-factor authentication and least-privilege access.	Centralised IAM; MFA; least privilege
3	Encrypt data in transit and at rest throughout the Ecosystem or recognised equivalent), with managed key management.	Encryption in transit & at rest
4	Provide a security operations capability (threat detection, incident response, SIEM or equivalent) covering all layers.	Security operations / SIEM
5	Provide vulnerability management and patching processes across infrastructure, platform, and applications.	Vulnerability & patch management

27.2 GOVERNANCE & COMPLIANCE

The Successful Bidder shall provide governance and compliance capabilities enabling Ignite to govern the Ecosystem, its data, and its models, consistent with the sovereignty and data-residency principles.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide data-governance capabilities (classification, provenance, lineage, retention, consent and withdrawal handling) across the Ecosystem.	Data governance & provenance
2	Enforce data-residency and sovereignty controls ensuring sovereign data and models remain in-country and under Pakistani jurisdiction.	Residency & sovereignty enforcement
3	Provide model governance (approval, registration, risk classification, and AI-safety policy enforcement) for served and hosted models.	Model governance & AI-safety policy
4	Provide compliance with applicable Pakistani regulation and recognised standards (e.g. ISO 27001 or equivalent), with auditability.	Regulatory & standards compliance
5	Provide comprehensive, tamper-evident audit trails across all layers for governance and oversight.	Tamper-evident audit trails

27.3 MONITORING & LOGGING

The Successful Bidder shall provide unified monitoring and logging across all layers, complementing infrastructure monitoring (Layer 1) and service-level performance monitoring (Layer 6), using standards-based interfaces.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide centralised, cluster-wide logging with retention and search across all layers.	Centralised logging & retention
2	Provide unified monitoring and observability via standards-based interfaces (Prometheus-compatible, OpenTelemetry, or equivalent).	Standards-based observability
3	Provide correlation of metrics, logs, and traces to support diagnosis across layers.	Metrics/logs/traces correlation
4	Provide alerting and notification routing to operations and governance functions.	Alerting & routing

27.4 INTEGRATION LAYER (EXTERNAL SYSTEMS & APIS)

The Successful Bidder shall provide an integration layer enabling the Ecosystem to connect securely with external government and enterprise systems through open, standards-based interfaces.

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide open, standards-based APIs (REST/gRPC or equivalent) for integration with external government and enterprise systems.	Open standards-based integration APIs
2	Provide secure connectors/gateways for external system integration, with authentication, authorisation, and rate control.	Secure connectors/gateways
3	Support standards-based identity federation with external systems (consistent with Layer 3 section 26.3.4).	External identity federation
4	Provide an API gateway with versioning, throttling, and monitoring for all externally-exposed interfaces.	API gateway with governance

27.5 INTEROPERABILITY & ACCELERATOR/HPC-AGNOSTIC AI SOFTWARE STACK

This Section requires that the AI software stack supplied under these TORs be architected for portability and heterogeneity, so that the national platform is not permanently bound to any single accelerator architecture, HPC backend, or vendor software ecosystem. The requirement is vendor- and technology-neutral: it does not mandate, favour, or exclude any accelerator architecture, any low-level programming model, or any vendor stack. It requires only that portability and heterogeneity be designed in at the layers where they are technically achievable.

Binding requirements

Ref	Requirement (the Solution shall...)	Basis / acceptance
1	Architect the AI software stack on open, standard, or documented representations at the framework, model-interchange, container, and API layers (e.g. standard	Portable layers; no single-stack precondition

	training frameworks, ONNX or equivalent model interchange, OCI-compatible containers, OpenAI-compatible or equivalent serving APIs), so that workloads are not bound to a single accelerator's proprietary programming model as a precondition of running.	
2	Ensure models and workloads can be retargeted/recompiled to a different accelerator architecture or HPC backend through a vendor-neutral compiler/abstraction path (e.g. an MLIR-based or equivalent compiler, standard kernel-abstraction layer, or documented conversion pipeline), with the supported targets and any limitations disclosed.	Retargetable via neutral compiler/abstraction
3	Support heterogeneous accelerators within the platform: the orchestration, scheduling, and management layers shall be capable of operating a fleet containing more than one accelerator architecture, or at minimum shall not preclude their coexistence and incremental addition.	Heterogeneous backend capability
4	Provide a documented migration/retargeting path enabling the Government to move models, workloads, and pipelines from the delivered accelerator/stack to a different one, stating the expected effort, tooling, and known limitations.	Documented migration path
5	Deliver all sovereign artefacts (model weights, training and fine-tuning code, datasets, and configurations) in stack-independent, open, and documented formats, so the Government can retrain or redeploy them on a different accelerator architecture (consistent with the sovereignty-transfer requirements).	Stack-independent sovereign artefacts
6	Avoid, in any layer above the unavoidable low-level kernel/driver layer, any dependency that would prevent substitution of the accelerator or HPC backend; where such a dependency is unavoidable, document it and provide an open or standards-based alternative access path wherever technically feasible.	Substitutability above kernel layer
7	Ensure the developer-facing environments, SDKs, and tooling expose portable, standards-based interfaces, so that solutions built by the national developer ecosystem are themselves portable across accelerator backends to the extent technically feasible.	Portable developer surfaces
8	Document, for each layer of the supplied stack, its portability status — portable, retargetable-with-effort, or architecture-specific — providing the Government a clear map of where lock-in does and does not exist.	Portability disclosure map

7 COLOCATION, FACILITY & PHYSICAL HOSTING

This Part defines the requirements for the primary and disaster-recovery hosting facilities. All numeric thresholds correspond exactly to the Technical Evaluation Criteria. The facility/facilities may be owned or operated by the Successful Bidder or by a named in-country colocation partner; partner facilities are acceptable provided the Successful Bidder remains contractually accountable for every service level herein. Requirements are stated as outcomes and accept any compliant facility; power and cooling are specified to match the Successful Bidder's own delivered hardware at full load plus headroom, so no fixed per-rack figure is imposed and no rack/cabinet form factor is excluded. Both a Primary and a Disaster-Recovery site are required, and both shall be located within Pakistan.

28.1 FACILITY RESILIENCY & CERTIFICATION (PRIMARY)

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a primary facility certified to Uptime Institute Tier III, or TIA-942 Rated-3, or an independently-audited equivalent, certified for the constructed/operational facility (not design-only), and concurrently maintainable (N+1 power and cooling).	Tier III / Rated-3 / equivalent (built-facility certified)
2	Commit to facility availability with defined service-credit remedies for breach.	≥ 99.982% annual availability
3	Ensure any single power-path, cooling unit, UPS, or generator is maintainable/replaceable with zero IT-load downtime.	Concurrent maintainability
4	Provide backup generation with on-site fuel autonomy at full IT load, plus a fuel-replenishment contract.	N+1 generators; ≥ 48 h fuel
5	Provide UPS sized to carry full IT load through generator start/transfer, with redundancy.	≥ N+1 UPS ride-through
6	Demonstrate operational track record to the stated tier; facilities in active certification may qualify conditionally with certification as a contract milestone.	Operational track record / roadmap
7	The co-location partner must have initiated the Cloud office accreditations under Pakistans Cloud First Policy	Evidence of process initiated / document submission.

28.2 POWER & COOLING (OUTCOME-MATCHED TO DELIVERED HARDWARE)

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Guarantee power capacity at the per-rack/per-cabinet density the delivered systems require, plus growth headroom (matched to the declared equipment envelope of Layer 1 section 26.1.5).	≥ delivered peak draw + 25% headroom
2	Cool the delivered hardware at full sustained load with no throttling, supporting the Successful Bidder's cooling method (direct-liquid, rear-door heat exchanger, high-density air, or equivalent).	No throttling at full delivered load

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
3	Accommodate the physical dimensions, weight/floor-loading (including heavy integrated super-node cabinets), busbar, and cooling interfaces of the delivered systems.	Floor-loading & form-factor accommodated
4	Operate to an efficient, climate-appropriate power-usage-effectiveness target, reported regularly.	PUE ≤ 1.5 (≤ 1.4 preferred), reported
5	Provide transparent per-circuit metered power billing with no flat-rate over-provisioning charges.	Metered power billing

28.3 SECURITY & SOVEREIGNTY (BOTH SITES)

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Locate both Primary and DR facilities within Pakistan, with all hosted data and equipment subject solely to Pakistani jurisdiction. Foreign-only siting is not acceptable.	In-country (sovereignty gate)
2	Provide multi-factor physical access control, 24x7 manned security, mantrap/airlock entry, and CCTV with retention.	Physical access control + CCTV (≥ 90 -day retention)
3	Provide a dedicated locked cage or private suite for the deployment, physically segregated from other tenants, at both sites.	Dedicated caged/suite space
4	Hold information-security certification (ISO 27001 or recognised equivalent) and comply with applicable Pakistani data-protection regulation.	ISO 27001 or equivalent
5	Provide physically diverse, redundant network entry per facility.	≥ 2 diverse fibre/carrier paths per site

28.4 DISASTER RECOVERY / SECONDARY SITE

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a secondary site at minimum Tier III-equivalent (same standard basis as the primary), located in-country.	DR site Tier III-equivalent; in-country
2	Separate the DR site from the primary by sufficient distance to survive a regional event while supporting the required replication (distance stated and justified).	Geographic separation justified
3	Provide inter-site replication with a stated recovery-point objective and declared replication architecture and bandwidth.	Replication with committed RPO
4	Provide a defined failover process with a committed recovery-time objective for resuming critical services.	Failover with committed RTO
5	Conduct scheduled DR failover testing with documented results.	DR testing $\geq$ annually

28.5 FACILITY SERVICE & COMMERCIAL TERMS

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide 24×7×365 on-site remote-hands/smart-hands support at both sites, with a defined Severity-1 response and escalation.	Remote-hands; Sev-1 ≤ 1 h
2	Commit a hosting term matching the equipment lifecycle, with defined renewal and price-escalation caps.	Committed term; escalation caps
3	Provide an installation/commissioning plan and timeline for both sites.	Onboarding & commissioning plan
4	Provide a defined exit and decommissioning plan: secure equipment removal and certified data destruction (recognised sanitisation standard), with no hostage-data or punitive exit fees.	Exit & certified data destruction

8 DELIVERY, INTEGRATION & PROGRAMME MANAGEMENT

This Part defines how the Successful Bidder shall deliver, integrate, and operate the Ecosystem as a single accountable programme across all layers. Bids may be submitted by a single entity or a consortium; in all cases a single prime Successful Bidder holds end-to-end accountability. All requirements are binding obligations of delivery.

29.1 PRIME ACCOUNTABILITY & CONSORTIUM GOVERNANCE

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Designate a single prime Successful Bidder contractually accountable for all layers (compute, platform, sovereign model, orchestration, development, marketplace, analytics, facilities) and all service levels, even where partners deliver components.	Single prime accountability
2	Where a consortium is used, provide a governance structure with binding agreements ensuring each partner's commitment, dispute resolution, and continuity if a partner exits.	Consortium governance & continuity
3	Provide a responsibility/accountability matrix mapping every deliverable and interface to an accountable party.	Accountability (RACI) matrix

29.2 END-TO-END INTEGRATION METHODOLOGY

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a documented methodology integrating all layers into one operational system, with defined integration milestones, interface ownership, and acceptance gates.	Integration methodology
2	Define and own all inter-layer and inter-component interfaces, using open standards wherever feasible.	Interface ownership & standards
3	Provide integrated system testing demonstrating the layers function together at scale.	Integrated system testing

29.3 MASTER PROGRAMME PLAN & SCHEDULE

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide an integrated master schedule from contract start to full operational acceptance, with critical path, dependencies, and milestone dates.	Master schedule with critical path
2	Identify resource and supply dependencies and their lead times.	Dependency & lead-time plan

29.4 PARALLEL-TRACK & PHASED DELIVERY

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a phased delivery plan in which early-value tracks (notably the tiered open-weight inference service) proceed in parallel with longer tracks (sovereign-model training, facility build-out), with the earliest live-service date stated.	Phased, parallel-track delivery plan
2	Define value-realisation milestones aligned to the seven programme objectives.	Objective-aligned milestones

29.5 RISK MANAGEMENT & MITIGATION

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a programme risk register covering technical, supply-chain, sovereignty, and integration risks, with owners, likelihood/impact, and mitigations.	Risk register with mitigations
2	Provide supply-chain and continuity contingency, including export-control and supply-disruption scenarios and second-source or alternative-path options.	Supply/continuity contingency

29.6 ACCEPTANCE, TESTING & HANDOVER FRAMEWORK

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide an acceptance framework with measurable acceptance criteria per layer and for the integrated system.	Acceptance framework & criteria
2	Provide a capability-handover gate confirming the in-country team can operate the Ecosystem, including a full sovereign-model fine-tuning cycle and inference-service operation.	Capability-handover acceptance gate

29.7 ONGOING OPERATIONS, SUPPORT & TRANSITION TO INDEPENDENCE

Ref	Requirement (the Solution shall...)	Minimum threshold / basis
1	Provide a post-go-live operations and support model across all layers, with escalation, SLA management, and continuity.	Operations & support model
2	Provide a transition plan progressively moving operations toward in-country independence, consistent with the sovereignty objective.	Transition-to-independence plan

9 STANDARDS, COMPLIANCE, CAPABILITY & ANNEXURES

30.1 APPLICABLE STANDARDS & COMPLIANCE

The Successful Bidder shall comply with applicable Pakistani law and regulation and with recognised international standards, or equivalents, including (without limitation): information security management (ISO 27001 or equivalent); data-centre resiliency (Uptime Institute Tier classification or TIA-942 or equivalent); service management (ISO 20000 or equivalent); and quality management (ISO 9001 or equivalent). Where a specific standard is named, an equivalent recognised standard is acceptable.

30.2 PROVEN CAPABILITY REQUIREMENTS (DELIVERY PREREQUISITES)

Consistent with the Technical Evaluation Criteria (Proven Capability & Delivery Track Record), the Successful Bidder shall, as a condition of delivery, demonstrate verifiable capability — directly or through named, contractually-bound partners — across the following, with single-prime accountability retained throughout:

- Delivery and commissioning of large-scale AI/HPC infrastructure of comparable scale.
- Production training or fine-tuning and deployment of a large language model, ideally including a low-resource or non-English language.
- Operation of a multi-tenant AI platform or AI-as-a-Service offering with tiered users at scale.
- Delivery and operation of an AI/software marketplace or developer ecosystem.
- Establishment of large-scale, multi-organisation data pipelines under a governance/consent framework.
- Delivery of a sovereign, government, or regulated-sector platform with data-residency obligations.
- In-country (Pakistan) delivery, support, or partner presence; and committed, named key personnel.

10 OPERATIONALISATION, WARRANTY, SUPPORT & MAINTENANCE

This Part defines the operational lifecycle that follows supply and installation: phased acceptance, the post-deployment data-pipeline and sovereign-model operationalisation phase, warranty, and ongoing support and maintenance under defined service levels. It consolidates the lifecycle and service-level obligations of the programme; where related obligations appear elsewhere (facility service terms in Section 28.5, platform enterprise support in Layer 2 - 26.2.1 AI Software Stack sub-section 16, ongoing operations and transition in Section 29.7, and post-go-live skilling support in Section 26.2.7 sub-section iv, this Part is the authoritative reference and those sections shall be read together with it. All obligations are binding on the Successful Bidder.

31.1 PHASED ACCEPTANCE & OPERATIONALISATION

Acceptance shall be staged to reflect the reality that this is a multi-phase programme in which usable value (trained sovereign model, live services) is produced over an extended period after physical installation. The warranty and comprehensive support obligations of this Part shall commence from Operational Acceptance sub-section 3 below, not from delivery.

Ref	Requirement (the Successful Bidder shall...)	Basis
1	Achieve Delivery Acceptance upon successful installation, interconnection, and commissioning of the supplied infrastructure, evidenced against the acceptance-criteria catalogue.	Delivery Acceptance (installed & commissioned)
2	Achieve Operational Acceptance upon demonstrating that the integrated services (compute, storage, orchestration, software stack, management platform, and the parallel-track inference service) are operational and meet their stated criteria.	Operational Acceptance (services working)
3	Establish that warranty and comprehensive support obligations commence from the date of Operational Acceptance, ensuring the Successful Bidder remains fully bound through the subsequent data-pipeline and sovereign-model phase.	Warranty/support clock starts at Operational Acceptance
4	Provide a phased acceptance plan mapping each milestone (delivery, operational, sovereign-model, and final acceptance) to measurable criteria and to payment milestones (where applicable).	Milestone-mapped acceptance plan

31.2 DATA-PIPELINE & SOVEREIGN-MODEL OPERATIONALISATION

The Successful Bidder expressly acknowledges that establishing the data partnerships, ingesting and cleansing Urdu and local-language corpora, building the production data pipelines, and fine-tuning the sovereign model constitute a defined operational phase that extends for a material period beyond infrastructure Operational Acceptance. The Successful Bidder shall remain fully engaged and bound throughout this phase.

Ref	Requirement (the Successful Bidder shall...)	Basis
1	Deliver a sovereign-model operationalisation plan with milestones for data-partnership establishment, corpora ingestion and cleansing, data-pipeline production readiness, fine-tuning, and evaluation, with target durations stated.	Sovereign-model operationalisation plan
2	Provide full support, maintenance, and engineering engagement throughout the data-pipeline and fine-tuning phase, under the service levels of this Part, until the sovereign model meets its acceptance criteria.	Bound support through model phase
3	Achieve Sovereign-Model Acceptance upon demonstrating the fine-tuned model meets the agreed Urdu/local-context capability criteria and the full sovereignty-transfer obligations (Section 26.2.4).	Sovereign-Model Acceptance gate
4	Ensure the tiered open-weight inference service (Section 26.2.5) is operated and supported in parallel from its earlier launch, so public value is sustained while the sovereign model is being developed.	Parallel service sustained

31.3 WARRANTY

Ref	Requirement (the Successful Bidder shall...)	Basis
1	Warrant all supplied hardware, software, integration, and facility systems against defects and non-conformance for a defined warranty period commencing at Operational Acceptance (warranty period to be specified in the contract; see draft SLA schedule).	Warranty from Operational Acceptance

2	Repair or replace defective hardware within the response/replacement times defined in the SLA schedule, including provision of spares.	Repair/replace per SLA
3	Remediate software and integration defects, including supplying updates, patches, and fixes, throughout the warranty period at no additional cost.	Software defect remediation
4	Maintain manufacturer back-to-back warranty and authorised-channel support for all OEM components throughout the warranty period.	OEM back-to-back warranty

31.4 SUPPORT & MAINTENANCE

Ref	Requirement (the Successful Bidder shall...)	Basis
1	Provide 24x7x365 support across all delivered layers, with severity-tiered response and resolution per the SLA schedule.	24x7x365 tiered support
2	Provide preventive and scheduled maintenance (firmware, drivers, security patches, health checks) with change management and agreed maintenance windows.	Preventive maintenance
3	Maintain an in-country spare-parts holding sufficient to meet the hardware replacement SLAs, with stock levels stated.	In-country spares
4	Provide a comprehensive post-go-live support and maintenance period under the SLA, continuing community and skilling support per Section 26.2.7, and aligned with the transition-to-independence obligations of Section 29.7.	Post-go-live support period (per SLA)
5	Operate a single point of contact and an escalation matrix for all support across the integrated solution, consistent with single-prime accountability (Section 29.1).	Single support desk + escalation
6	Provide regular service-performance reporting against the SLA to Ignite .	SLA performance reporting

31.5 SERVICE LEVELS (SLA FRAMEWORK)

Support and maintenance shall be governed by the Service Level Agreement set out in the draft SLA schedule (Section 32). The SLA defines severity levels, response and resolution targets, service-availability targets per service class, service credits/penalties for breach, and escalation. The targets stated in the draft are minimum and shall be finalised by the procuring authority during contract formation.

Ref	Requirement (the Successful Bidder shall...)	Basis
1	Meet the response and resolution targets per severity level defined in the SLA schedule.	Severity response/resolution targets
2	Meet the service-availability targets per service class (core platform, inference service, marketplace, facility) defined in the SLA schedule.	Availability targets
3	Accept service credits / penalties for SLA breaches as defined in the SLA schedule.	Service credits for breach
4	Operate the escalation matrix defined in the SLA schedule for unresolved or critical issues.	Escalation matrix

31.6 MANAGED OPERATIONS OF THE SOLUTION

Beyond support and maintenance (Section 31.4), the Successful Bidder shall operate and run the entire deployed solution as a managed service for a minimum period of three (3) years commencing from complete operationalisation of the system (i.e. from the later of Operational Acceptance and Sovereign-Model Acceptance). This managed-operations obligation is a mandatory part of the scope of work and of the Service Level Agreement, and shall be separately priced by the bidder as an annual managed-operations line item (per year) so that the full-term operating cost is transparent. The Successful Bidder shall operate the solution jointly with an Ignite / in-country team from day one, under the governance and approval of Ignite , progressively transferring operational capability to the in-country team over the term.

Ref	Requirement (the Successful Bidder shall...)	Basis
1	Operate and administer all layers of the solution on a day-to-day basis — compute, accelerators, network/fabric, storage, container orchestration, software stack, and the cluster operations & management platform — to the service levels of this Part.	Day-to-day operation of all layers
2	Operate the tiered public inference service, the enterprise hosting capability, and the AI marketplace portal, including service availability, scaling, and performance management.	Operate inference, hosting & marketplace
3	Operate and continuously maintain the data pipelines on an ongoing basis — continued ingestion, cleansing, curation, and refresh of Urdu and local-context corpora — and support continued/iterative fine-tuning of the sovereign model under Ignite 's direction.	Ongoing data-pipeline & model operations
4	Perform tenant and user lifecycle administration across the four beneficiary groups — onboarding, quota/credit administration, access control, and offboarding — under the tiered-access and subsidy policies.	User & tenant operations
5	Perform capacity, performance, and availability management against the SLA, including monitoring, tuning, and proactive optimisation.	Capacity & performance management
6	Operate security operations (monitoring, patching, incident response) and execute backup and disaster-recovery procedures across the primary and secondary sites.	Security operations & DR execution
7	Perform incident, problem, and change management, with significant changes subject to Ignite 's approval through defined governance gates.	Incident/problem/change management (Ignite - approved)
8	Operate jointly with an embedded Ignite / in-country operations team from day one, progressively transferring operational responsibility so the in-country team can assume operations by end of term (aligned with Sections 26.2.7 and 29.7).	Joint operations + progressive transfer
9	Operate under the governance and approvals of Ignite , escalating decisions reserved to the governance authority and acting on its directions.	Operate under Ignite governance
10	Provide regular operational reporting and service reviews to Ignite (utilisation, availability, incidents, adoption, and impact data feeding Layer 6).	Operational reporting to Ignite
11	Deliver the managed-operations service for a minimum of three (3) years from complete operationalisation, extendable or renewable at the sole option of Ignite on the same terms.	≥ 3 years; extendable at Ignite 's option

12	Separately price the managed-operations service as an annual (per-year) line item in the financial proposal, covering the full minimum term and any optional extension years.	Separately priced per-year line item
13	Maintain, at end of term (or non-renewal), an orderly operational handover to the in-country team or a successor, with no service disruption and no hostage-data or lock-out, consistent with Section 29.7.	Orderly end-of-term operational handover

31.7 OPERATIONAL GOVERNANCE & DECISION RIGHTS

The managed operations of Section 31.6 shall be governed by a defined governance structure under the authority of the company (Ignite). The governance mechanisms below establishes how operational decisions are made, approved, and reviewed, and the respective responsibilities of the Successful Bidder and the Government. Decision rights reserved to the Government shall not be exercised by the Successful Bidder.

Ref	Requirement (the Successful Bidder shall...)	Basis
1	Operate under a joint governance structure comprising an executive governance board (chaired by Ignite), a programme/operations management forum, and a technical change-advisory function, with terms of reference agreed at contract start.	Defined governance structure
2	Respect the decision-rights allocation whereby strategic, policy, sovereignty, pricing/subsidy, and security-classification decisions are reserved to Ignite, while routine operational execution rests with the Successful Bidder within agreed policies.	Reserved vs delegated decision rights
3	Submit all significant or high-risk changes (architecture, security posture, capacity, data handling, configuration baselines) to a Change Advisory Board (CAB) for Ignite approval before implementation; emergency-change procedures with retrospective approval shall be defined.	CAB change approval
4	Hold monthly operational service reviews (performance against SLA, incidents, capacity, risks, security) with the operations management forum.	Monthly operational reviews
5	Hold quarterly executive service reviews with the governance board (service outcomes, adoption and impact, financials, transition progress, forward plan).	Quarterly executive reviews
6	Hold an annual service review assessing the year's performance, value delivered, SLA compliance, and recommendations, informing any extension decision.	Annual service review
7	Maintain transparent operational records, dashboards, and audit access for the Government across all governance forums.	Transparency & audit access

The Government's reciprocal responsibilities — timely approvals, provision of the embedded in-country team, access and facilities, and decisions on reserved matters — shall be set out in the contract so that the Successful Bidder's obligations are matched by defined Government obligations.

31.8 EXIT MANAGEMENT, TRANSITION ASSISTANCE & OPERATIONAL HANDOVER

This Section consolidates the Successful Bidder's obligations on exit — whether on expiry, non-renewal, or termination — to ensure the Government, or a successor Successful Bidder or the in-country team, can assume full operation of the Ecosystem without disruption, loss of data, or dependence on the outgoing Successful Bidder. These obligations reinforce the sovereignty objective

and the transition-to-independence intent of Sections 26.2.7, 29.7 and 31.6. They are binding and survive termination.

Ref	Requirement (the Successful Bidder shall...)	Basis
1	Maintain, throughout the term, an up-to-date Exit and Transition Plan (refreshed at least annually and on request), describing how operations, knowledge, and assets will be handed over.	Living exit/transition plan
2	On exit, transfer complete, current operational documentation — architecture documentation, runbooks, operational playbooks, standard operating procedures, network and configuration diagrams, and known-issue registers.	Documentation, runbooks & playbooks transfer
3	Transfer all configurations, infrastructure-as-code, automation scripts, operational tooling configurations, and baseline definitions required to operate the solution.	Configuration & automation transfer
4	Transfer all credentials, access rights, certificates, keys, and administrative control of all systems to the Government or its nominee, and revoke the Successful Bidder's own access on completion.	Credential & control transfer
5	Transfer the source code, weights, datasets, and artefacts of all sovereign components (the sovereign model, data pipelines, and bespoke components developed for the Government), consistent with Section 26.2.4; and provide escrow or documented continuity arrangements for any proprietary third-party components necessary for continued operation, sufficient to ensure operational continuity without compromising third-party intellectual property.	Sovereign source transfer + continuity escrow
6	Export all Government and tenant data, models, and content in open, documented, non-proprietary formats, with verified integrity and completeness, and securely erase it from Successful Bidder-controlled systems thereafter.	Open-format data export + secure erasure
7	Provide a defined transition-assistance period (minimum to be specified in the contract) during which the Successful Bidder continues to operate and/or supports the incoming operator to maintain uninterrupted service.	Transition-assistance period
8	Cooperate fully with, and provide reasonable assistance and information to, a successor Successful Bidder or the in-country team, including knowledge-transfer sessions and parallel-running where required.	Successor-Successful Bidder cooperation
9	Impose no hostage-data, lock-out, or withholding of any kind; the Government's data, sovereign artefacts, and operational control shall never be contingent on payment disputes or any other leverage.	No hostage-data / no lock-out
10	Levy no punitive exit, de-installation, or data-extraction fees; exit-related charges, if any, shall be limited to fair, pre-agreed, cost-based amounts stated in the contract.	No punitive exit fees
11	Achieve a Final Operational Handover, certified against defined criteria, confirming the Government or successor can operate the Ecosystem independently with no residual Successful Bidder dependency.	Certified final handover

31.9 COMMERCIAL PROTECTION & ANTI-LOCK-IN

This Section protects the Government against cost escalation, lock-in, and loss of control over the lifecycle of the programme, consistent with the neutrality and sovereignty principles. It applies to all supply, support, managed-operations, and renewal arrangements.

Ref	Requirement (the Successful Bidder shall...)	Basis
1	Hold all recurring charges (support, managed operations, subscriptions, licences) firm for the initial term, and cap any escalation in extension years to a pre-agreed index-linked ceiling stated in the contract.	Cost-escalation cap
2	Guarantee continuity of all software subscriptions, licences, and entitlements required to operate the solution for the full term and any extension, with no mid-term withdrawal of entitlements that would impair operation.	Subscription/licence continuity
3	Vest ownership of all configurations, scripts, documentation, runbooks, and bespoke developments created for the Government in the Government, and ensure sovereign artefacts are owned by the Government per Section 26.2.4.	Configuration/documentation ownership
4	Ensure the Government is not locked into the Successful Bidder for continued operation: open or documented interfaces at layer boundaries (guiding principle 3), portable data formats, and the exit provisions of Section 31.8 shall together guarantee substitutability of the operator.	Operator substitutability / anti-lock-in
5	Offer any extension or renewal at terms no less favourable than the initial term (on a like-for-like basis), at Ignite 's sole option, with no renewal-conditional penalties.	Fair renewal protection
6	Disclose, and not unreasonably restrict, all third-party dependencies (licences, subscriptions, support contracts) whose lapse could affect operation, and assign or enable their transfer to the Government on exit where permissible.	Third-party dependency transparency
7	Ensure no foreign-controlled licence check-in, remote disablement, or external dependency can interrupt operation of the solution under Government control (consistent with program objectives).	No external kill-switch

31.10 ACCESS PRICING, REVENUE-SHARING & PLATFORM SUSTAINABILITY

The Company shall govern access pricing and revenue-sharing arrangements across the NAIH platform in accordance with the following framework, which the Successful Bidder shall develop, implement and operate:

- i. Subsidised Access Tiers: Compute-as-a-Service and inference access for the Academia & Researchers and Startups tiers shall be provided at concessional rates set below prevailing international market rates. The subsidy differential shall be funded from the approved programme budget and shall not be recovered from the Successful Bidder's managed services fee.
- ii. Cost-Recovery Tiers: Public Sector and Private Sector access shall be priced on a cost-recovery basis, with rates set by the Company to recover direct operational costs (co-location, power, OEM support) without profit extraction from public institutions. Private Sector rates may include a platform margin retained by the Company.
- iii. Marketplace Revenue-Sharing: Revenue generated through the AI Innovation Marketplace (solution licensing, API monetisation, subscription fees paid by buyers) shall be shared between the Company and solution developers on a ratio to be defined through Marketplace Governance Policy issued by the Company. The default split shall be no less than 70% to the solution developer and no more than 30% retained by the Company to fund platform operations and sustainability.

- iv. Platform Sustainability: The Successful Bidder shall operate the platform under a managed services fee structure as defined in the Payment Plan. The Company's target is for OPEX to become progressively self-financing from Year 2 onwards through cost-recovery access charges and marketplace revenue, reducing dependence on grant funding for recurring operational expenditure.
- v. Pricing Governance: All access pricing schedules shall be approved by the Company or its delegated authority prior to operationalisation and shall be reviewed annually. The Successful Bidder shall not independently set, alter, or waive any access price without prior written Company approval.

11 DRAFT SERVICE LEVEL AGREEMENT (SLA) SCHEDULE

*This SLA schedule sets out the binding minimum service levels for the programme. The targets below are mandatory minimums that bidders shall meet or exceed; bidders may propose superior levels, which become contractual once accepted. **Specific figures may be refined by the company (Ignite) during contract formation**, but shall not fall below the minimums stated here. The agreed SLA is enforceable through the service-credit and escalation mechanisms of this schedule and Section 31.*

32.1 SEVERITY LEVELS

Severity	Definition	Minimum response / resolution
Severity 1 — Critical	Complete outage or critical failure of a production service (e.g. inference service down, cluster-wide failure, data-loss risk).	Response ≤ 15 min; continuous effort; resolution/workaround ≤ 4 h
Severity 2 — High	Major degradation or partial outage; significant functionality impaired with no workaround.	Response ≤ 30 min; resolution ≤ 8 h
Severity 3 — Medium	Moderate impact; functionality impaired with an available workaround.	Response ≤ 2 h; resolution ≤ 2 business days
Severity 4 — Low	Minor issue, query, or cosmetic defect; no material service impact.	Response ≤ 1 business day; resolution ≤ 5 business days

32.2 SERVICE AVAILABILITY TARGETS (MINIMUM)

Service class	Illustrative availability	Measurement basis
Core platform & compute	≥ 99.9% monthly	Excluding agreed maintenance windows
Tiered inference service	≥ 99.9% monthly	Per-service measurement
AI marketplace portal	≥ 99.5% monthly	Per-service measurement
Hosting facility	≥ 99.982% annual	Facility-level (Tier III-class)

32.3 HARDWARE REPLACEMENT (MINIMUM)

Component class	Illustrative replacement target	Basis
Critical components (accelerator node, storage, fabric)	On-site replacement ≤ 24 h from diagnosis	From in-country spares

Standard components	On-site replacement ≤ 48 h	From in-country spares
---------------------	----------------------------	------------------------

32.4 DISASTER RECOVERY TARGETS (MINIMUM)

Service tier	Minimum RPO (data loss)	Minimum RTO (service restoration)
Critical services (sovereign model state, core platform, inference)	≤ 15 minutes	≤ 4 hours to the secondary site
Standard services (marketplace, developer, analytics)	≤ 1 hour	≤ 8 hours
Bulk/archival data	≤ 24 hours	≤ 24 hours

32.5 SERVICE CREDITS & ESCALATION (ILLUSTRATIVE)

Illustrative: service credits accrue against the relevant service charge for each SLA breach, scaled by severity and duration; repeated or sustained breaches trigger escalation through a defined matrix (support lead → programme manager → executive sponsor → Ignite governance) and, beyond defined thresholds, contractual remedies. Exact credit percentages, caps, and escalation thresholds to be finalised by the procuring authority.

Reporting (Illustrative): the Successful Bidder shall provide monthly SLA performance reports and a quarterly service review with Ignite, covering availability achieved, incidents by severity, resolution performance, and corrective actions.

12 PERFORMANCE MEASUREMENT: SUCCESSFUL BIDDER KPIs & PROGRAMME INDICATORS

This Part defines how performance is measured over the term. It deliberately separates two distinct categories, which shall not be conflated:

(a) Successful Bidder Performance KPIs (Sections 33.1–33.4) — matters within the Successful Bidder’s control: operating the solution smoothly and without technical or operational glitch, to the agreed service levels, and delivering the Successful Bidder’s contracted obligations. These carry contractual consequences (Section 33.5).

(b) Programme Indicators (Section 33.6) — programme-level outcomes such as adoption, pricing position, and economic/sectoral impact. These depend on market, policy, and demand factors beyond the Successful Bidder’s control. They are monitored and reported for the Government’s oversight only; the Successful Bidder’s sole duty is to operate the platform that enables their measurement and to supply accurate data. They do not by themselves constitute Successful Bidder underperformance and carry no penalty.

33.1 SUCCESSFUL BIDDER PERFORMANCE — MEASUREMENT FRAMEWORK

The Successful Bidder shall measure, evidence, and report the Successful Bidder Performance KPIs below, using system-derived data wherever possible (Layers 3,6), independently verifiable by Ignite, and reported through the governance forums of Section 31.7 at the stated cadence. These KPIs measure only matters within the Successful Bidder’s control.

33.2 OPERATIONAL PERFORMANCE KPIS (MONTHLY)

Ref	KPI (within Successful Bidder control)	Measurement mechanism	Cadence
1	Platform & compute availability achieved vs SLA	Monitoring/telemetry (Sec 26.1.7, Sec.27.3)	Monthly
2	Inference service availability (platform uptime)	Service-level monitoring (Sec.26.6.4)	Monthly
3	Marketplace/platform service availability	Service-level monitoring (Sec.26.6.4)	Monthly
4	Incidents resolved within SLA target (%)	Incident management records	Monthly
5	Mean time to recovery (MTTR), critical incidents	Incident records	Monthly
6	Security patch compliance & incidents handled within target	Security operations (Sec.27.1)	Monthly
7	Data-pipeline operational uptime	Pipeline monitoring (Sec.26.1.3, Sec.26.2.3)	Monthly
8	Backup success rate & DR readiness (RPO/RTO met in tests)	Backup/DR records (Sec.26.1.8, Sec.28.4)	Monthly
9	Preventive-maintenance schedule adherence	Maintenance records (Sec.31.4)	Monthly

33.3 DELIVERY & MILESTONE KPIS (MILESTONE / QUARTERLY)

Ref	KPI (within Successful Bidder control)	Measurement mechanism	Cadence
1	Milestones delivered on schedule (%)	Master schedule tracking (Sec.29.3)	Quarterly
2	Sovereign model delivered and accepted on the agreed schedule	Sovereign-Model Acceptance (Sec.31.2)	Milestone
3	Full sovereignty transfer completed (weights, code, data, rights)	Transfer acceptance (Sec.26.2.4)	Milestone
4	Urdu/local-context capability uplift delivered vs agreed target	Evaluation harness (Sec.26.2.4)	Milestone
5	Operational readiness of contracted capabilities at each phase gate	Acceptance framework (Sec.29.6, Sec.31.1)	Milestone

33.4 CAPABILITY-TRANSFER & SERVICE-MANAGEMENT KPIS (QUARTERLY / ANNUAL)

These measure transfer and service-management activities the Successful Bidder is contracted to perform (not the wider ecosystem's uptake).

Ref	KPI (within Successful Bidder control)	Measurement mechanism	Cadence
1	Pakistani personnel trained / certified by the Successful Bidder	Training & certification records (Sec.26.2.7)	Quarterly
2	Train-the-trainer cadre established	Trainer accreditation records (Sec.26.2.7)	Annual
3	In-country operations readiness enabled (transferable to local team)	Operations records (Sec.31.6)	Quarterly
4	Curriculum accreditation supported and achieved	Accreditation evidence (Sec 26.2.7)	Annual
5	Community support operated to agreed levels (forums, helpdesk, office hours)	Community/helpdesk records (Sec.26.2.7)	Quarterly
6	Required reporting delivered on time and complete	Governance forums (Sec.26.2.7)	Quarterly
7	Corrective-action closure within agreed time	Governance records (Sec.31.7)	Quarterly

33.5 CONSEQUENCES OF SUCCESSFUL BIDDER KPI PERFORMANCE

The Successful Bidder Performance KPIs (Sections 33.2–33.4) carry defined consequences, applied proportionately through the governance process. Programme Indicators do not carry consequences.

Ref	Provision	Mechanism	Cadence	Effect
1	Sustained or material KPI underperformance triggers a performance-improvement process	Governance escalation (Sec.31.7)	On breach	Improvement plan with deadlines
2	KPI breaches corresponding to SLA service classes attract service credits	Service-credit mechanism (Sec 32)	Per cadence	Financial credit to Government
3	Persistent failure after the improvement period escalates through the governance matrix	Escalation matrix (Sec.31.7, Sec 32)	On persistence	Executive/governance intervention
4	Cumulative Successful Bidder KPI achievement is a material input to the extension/renewal decision	Annual service review (Sec.31.6)	Annual	Informs Sec.31.6 extension at Ignite's option

Service credits and other consequences are without prejudice to the Government's other contractual remedies. Precise credit percentages, breach thresholds, and improvement-period durations shall be finalised by Ignite during contract formation, consistent with the SLA schedule (sec 32).

33.6 PROGRAMME INDICATORS (MONITORED ONLY — NO SUCCESSFUL BIDDER PENALTY)

The following indicators measure programme-level outcomes that depend on market, policy, and demand factors beyond the Successful Bidder's control. They are tracked for the Government's oversight and reported by the Successful Bidder from platform data, but they do not constitute Successful Bidder performance and carry no penalty. The Successful Bidder's only obligation in respect of them is to operate the enabling platform and to provide accurate, timely data and reporting.

Ref	Programme indicator (outcome — not Successful Bidder-penalised)	Data source / mechanism	Cadence	Tracked for
1	Inference price position vs prevailing international market rates	Pricing comparison; metering (Sec.26.3.5)	Quarterly	Objective 1
2	Compute-as-a-Service price position vs international rates	Pricing comparison; metering (Sec.26.3.5)	Quarterly	Objective 3
3	Volume of subsidised access taken up by Academia & Startups	Quota/credit data (Sec.26.3.2)	Quarterly	Objective 1
4	Active users / adoption by beneficiary tier	Usage analytics (Sec.26.6.1)	Quarterly	Objectives 1,3,5
5	Third-party models hosted; sovereign-model usage	Hosting & usage data (Sec.26.2.6, Sec.26.6.1)	Quarterly	Objectives 2,4
6	Marketplace listings, transactions, investor connections	Marketplace analytics (Sec.26.5.7, Sec.26.5.8)	Quarterly	Objective 6
7	Economic impact / value enabled; cost savings vs alternatives	Impact measurement (Sec.26.6.2)	Annual	Objective 7
8	Sectoral adoption; startups / enterprises / jobs enabled	Sectoral impact (Sec.26.6.2)	Annual	Objective 7

Where a Programme Indicator underperforms, the Company may review programme strategy, policy, pricing, or promotion; such underperformance shall not of itself be attributed to the Successful Bidder unless it is shown to result directly from a failure of a Successful Bidder Performance KPI (Sections 33.2–33.4).

13 DESIRED OUTCOMES

The following are the expected outcomes of the project:

34.1 AI INFRASTRUCTURE OUTCOMES

- i. Establishment of fully operational **sovereign AI compute infrastructure** owned by the Company under the auspices of the Ministry of IT & Telecom, Government of Pakistan
- ii. Availability of a secure, scalable, and production-grade AI platform
- iii. Reduced dependency on foreign cloud infrastructure for AI workloads
- iv. Infrastructure readiness for expansion or migration to a future national AI data center
- v. Establishment of a government-controlled AI compute and model platform capable of independent evolution beyond the contract duration

34.2 CAPACITY & ECOSYSTEM OUTCOMES

- i. Development of indigenous AI infrastructure operations capability
- ii. Availability of AI compute resources for NAIH programs and ecosystem users
- iii. Strengthened AI research, experimentation, and innovation pipelines
- iv. Improved institutional readiness for advanced AI deployment
- v. Creation of a sustainable national talent base for AI infrastructure, platform, and operations engineering

34.3 NATIONAL CAPABILITY DEVELOPMENT OUTCOMES

- i. Establishment of national expertise in AI infrastructure engineering and operations
- ii. Development of reusable national AI platform components and operational knowledge assets
- iii. Strengthened collaboration between government, academia, startups, and industry through shared AI infrastructure
- iv. Long-term reduction in dependency on external AI compute providers
- v. Creation of institutional capacity for future sovereign AI infrastructure expansion

35. GOVERNANCE STRUCTURE

The NAIH AI Infrastructure programme shall be governed the Company as part of section 31.7, meaning that all sovereignty, data governance, consent, access-tier management, monitoring, and reporting obligations defined under the TORs shall be exercised under, and remain accountable to, the Company as the designated governance authority. Regular monitoring of the programme shall be conducted by the Company through its representatives or designated quality monitors.

36. TENURE OF THE AI INFRASTRUCTURE PROJECT

The tenure of the NAIH AI Infrastructure assignment shall include:

- Procurement, Deployment, commissioning and initial operations phase (6 months); and
- **Continued operations support for a period of thirty-six (36) months.**

An extension (additional 2 years) shall be subject to Company and PPRA rules. Upon completion, operations shall transition fully to Ignite.

37. RESPONSIBILITY FRAMEWORK

37.1 PROVISION OF FUNDS

The Company shall be responsible for funding approved capital and operational expenditures in accordance with the Funding Agreement. Disbursements shall be made on a reimbursable and milestone-based basis. The Successful Bidder shall ensure the availability of funds equivalent to at least **six (06) months of operations**.

37.2 PROVISION OF CO-LOCATION FACILITY (PRIMARY +DR)

The AI infrastructure shall be hosted in an **existing Tier III or above data center facility in Pakistan** providing sites for both primary operations and DR. The co-location provider shall be responsible for power, cooling, rack density, physical security, and facility management.

37.3 PROVISION OF INFRASTRUCTURE

The Successful Bidder shall be responsible for procurement, deployment, configuration, testing, and commissioning of AI compute hardware, networking, storage, and AI software stack within agreed timelines.

37.4 OPERATIONAL HANDOVER & TRANSITION

The Successful Bidder shall ensure structured and documented handover of AI infrastructure operations to Ignite upon completion of the operations support period, subject to achievement of defined handover criteria.

37.5 INSURANCE AND RISK COVERAGE

The Successful Bidder shall ensure appropriate insurance coverage for equipment handling, transit, installation, and operational risks during the deployment and operations support period. The Successful Bidder shall remain responsible for safeguarding equipment until formal handover to the Company following commissioning acceptance.

37.6 ASSET LIFECYCLE AND DECOMMISSIONING

The Successful Bidder shall provide guidelines for environmentally responsible disposal, recycling, or refurbishment of hardware components at end of lifecycle in accordance with applicable environmental regulations and international best practices.

38. EQUITY

The Successful Bidder shall not claim any equity, ownership, or usage rights in AI infrastructure, platforms, models, datasets, or outputs developed under this project.

39. IPR

All intellectual property, configurations, customizations, documentation, and operational artifacts developed under this project shall vest exclusively with the **Company**. No rights shall accrue to the Successful Bidder beyond the contract duration

All AI models, including foundation models, fine-tuned models, and derived artifacts deployed under this project shall be free from licensing restrictions that prevent sovereign operation, modification, retraining, or independent evolution by the Company. The Successful Bidder shall disclose all applicable licenses and ensure that the Government retains full operational rights to all models deployed.

Any solution developed by the innovator/ developer/solution provider/ startup and approved on the marketplace shall be owned by the innovator/ developer/solution provider/ startup but subject to revenue sharing mechanisms in section 31.10.

40. PAYMENT / DISBURSEMENT PLAN

All payments shall be made on reimbursable model. However, a one-time mobilization advance of up to 20% of CAPEX amount may be paid to successful bidder on request, against a Bank Guarantee equivalent to the mobilization advance to be submitted by the bidder. The same will be released immediately after adjustment of mobilization advance.

Payment for required infrastructure development will be made as per the Disbursement Plan agreed with the Successful Bidder. Furthermore, payment of operating expenses including but not limited to utilities, salaries of Management Team, payment of trainings/workshops etc. will be made through the invoices duly verified by Company. The successful bidder shall submit invoices to Company as per agreed disbursement plan. The same may be verified by the Company within 30 calendar days. Payment will be made within 30 calendar days after acceptance of deliverables. In case of variance, the undisputed amount shall be recommended to finance for payment.

There shall be no expenditure incurred beyond the approved disbursement plan. However, in exceptional circumstances, for any expense beyond the approved payment plan, the Company will be the final authority to approve such payments or otherwise.

The table below defines the proposed payment plan which will be finalized in agreement with the successful bidder:

Milestone	Deliverable	% of Contract	Phase	Cost Type	Acceptance Criteria
M0	Contract Signing & Mobilization (Bank Guarantee required)	Advance: up to 20% CAPEX (recoverable vs M2/M3)	Phase 0	CAPEX (Advance)	Contract signed, Bank Guarantee submitted, mobilization completed
M1	Detailed Design & Procurement Readiness	10%	Phase 1	CAPEX	Design, architecture, BOM, and procurement plan approved
M2	Delivery of Infrastructure (verified at facility)	25% (less 10% advance recovery)	Phase 1	CAPEX	Infrastructure delivered and

					verified against approved BOM
M3	Installation & Infrastructure Integration	10% (less 10% advance recovery)	Phase 1	CAPEX	Infrastructure installed, integrated, and operational
M4	AI Platform & Software Stack Deployment	10%	Phase 1	CAPEX	AI platform, orchestration, security stack deployed and operational
M5	Portal, Marketplace & Sovereign Services Deployment	5%	Phase 1	CAPEX	Portal, APIs, marketplace, sovereign services deployed
M6	Acceptance Testing & Commissioning (SAT)	10%	Phase 1	CAPEX	Successful SAT completion and formal commissioning
M7	Operational Readiness & Initial Service Launch	5%	Phase 1	CAPEX	Service operational readiness achieved and initial launch completed
M7a	Sovereign Model Training, Acceptance & Sovereignty Transfer	5%	Phase 2	CAPEX	Sovereign models trained, accepted, and transferred to Ignite
M8	Quarterly Operations — KPI-linked Managed Services Fee	10% (quarterly pro-rata)	Phase 3	Managed Services (OPEX-linked fee)	SLA/KPI compliance and quarterly performance approval
M8 (OPEX)	Operating Expenses Reimbursement (Power, bandwidth, OEM support, etc.)	As per actuals (no cap within approved categories)	Phase 3	OPEX (Reimbursable)	Submission of valid invoices + usage proof + Ignite approval
M9	Capacity Building, Knowledge Transfer & Final Handover	10%	Phase 3	CAPEX / Close-out	Training completed, documentation delivered, final handover accepted
TOTAL FIXED CONTRACT VALUE		100% CAPEX + Managed Services Fee			
OPEX (Reimbursable, outside contract value)	Utilities, licenses, cloud costs, bandwidth, OEM AMC, etc.	Not part of 100%	Phase 3	Pure Reimbursement	Paid at actuals against verified invoices

Note: All milestone timelines stated above are indicative targets based on standard deployment lead times. Actual timelines for hardware delivery milestones (M1 and M2 in particular) may be subject to OEM and vendor hardware manufacturing and delivery schedules. Any such variance shall be formally documented by the Successful Bidder and communicated to the Company in writing within 7 calendar days of becoming known. Revised milestone dates shall require Company approval and shall not affect the percentage allocations or acceptance criteria defined above.

