

Training vs Inference: Path to India's AI Sovereignty

16 January 2026

Government of India

Quantum Tiger

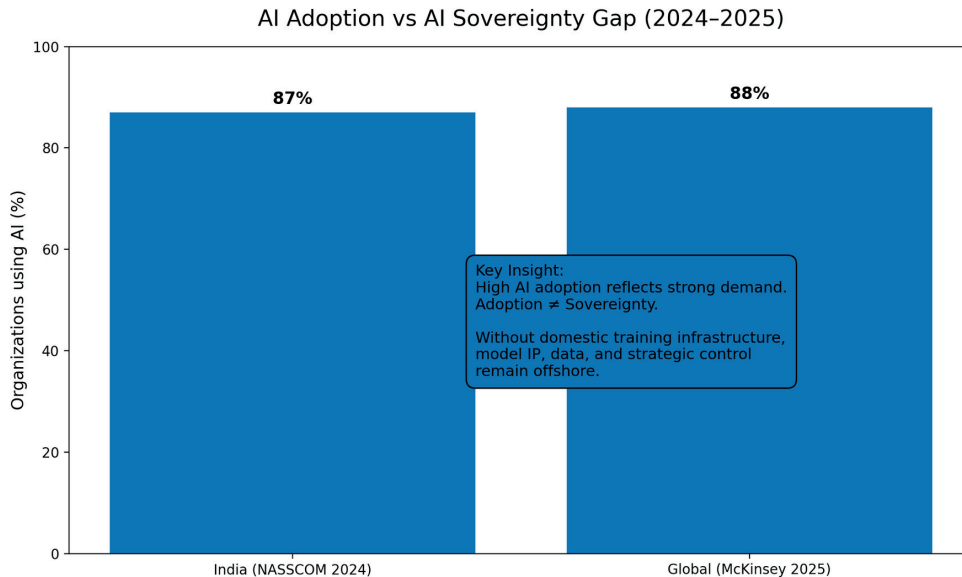
Policy brief for startups & public sector stakeholders

Executive Summary

India's enterprise AI adoption is now mainstream, but strategic sovereignty sits on the ability to train (not merely to use or serve) large-scale models within national jurisdiction. Recent 2025 studies show a large majority of organisations report AI use, yet the transition from pilots to scaled training capacity remains limited, creating a dependence on foreign compute, foreign clouds, and external IP that weakens sovereignty. This paper updates earlier recommendations with 2025 evidence, and proposes a practical, funded program to secure domestic training infrastructure, data trusts, talent pipelines, procurement pathways, and regulatory guardrails.

AI adoption (India vs Global):

India has high enterprise uptake (NASSCOM AI Adoption Index reports strong adoption; McKinsey's State of AI 2025 shows 88% of organizations globally reporting AI use in at least one business function). These numbers show demand is high, but training capacity is not assured by adoption alone.



- India: ~87% enterprise AI adoption (NASSCOM AI Adoption Index)
- Global: ~88% enterprise AI adoption (McKinsey State of AI 2025)

[High AI adoption reflects strong demand.
Adoption ≠ Sovereignty.

Without domestic training infrastructure, model IP, data, and strategic control remain offshore.]

IndiaAI (AI Kosh): As of 15 September 2025, India's AI Kosh reports over 2,000 curated non-personal datasets — an important and growing sovereign data asset, but still not yet sufficient for large-scale foundation-model training across all strategic domains.

Training vs Inference: Path to India's AI Sovereignty

16 January 2026

Government of India

Quantum Tiger

Policy brief for startups & public sector stakeholders

Why training matters more for sovereignty than inference?

- Training builds models (the intellectual property and behaviour); it requires large GPU/accelerator farms, fast distributed networking (NVLink/RDMA), low-latency storage tiers, sustainable power and energy budgets, and secure supply chains. Training is where model provenance, data lineage and IP are created, and where exfiltration or poisoning presents the greatest national risk.
- Inference can be hosted domestically or at the edge, and is comparatively less compute-intensive per request, but inference endpoints are only as sovereign as the model that powers them.

Policy implication: Guaranteeing inference onshore is not sufficient. India must secure the training lifecycle, compute, data, tooling, certification, to claim true AI sovereignty.

What changed in 2024–2025

- Adoption increased, but scale gaps persist. McKinsey's 2025 State of AI shows rising AI use in enterprises globally, yet most organisations still struggle to scale pilot projects into enterprise-wide, production training operations. This gap is visible in India as well: adoption is widespread, but large-scale foundation-model training capacity is limited.
- IndiaAI Mission operationalized and expanding. The Cabinet-approved IndiaAI Mission (March 2024) has moved towards implementation in 2024–25 with allocated funds and initiatives (compute democratization, AI Kosh dataset curation, innovation centers and skilling/CoE plans). The mission's stated budget and multi-year objectives provide a strong foundation for sovereign capacity-building.
- Data foundations are growing (AI Kosh). AI Kosh hosts 2,000+ curated datasets as of Sept 15, 2025, a major milestone. Yet sector-specific high-value datasets (healthcare imaging, high-fidelity geospatial telemetry, telecom-scale sensor data) still need concerted public-private aggregation, legal clarity, and synthetic augmentation pipelines.
- Talent & skilling commitments are visible. The 2025–26 Union Budget and NITI Aayog roadmaps call out national skill centers and fellowships for AI; the talent base is growing but not yet at the scale for sustained, sovereign model R&D and large-scale systems engineering.
- Global signals: Hardware concentration (accelerators) continues to be a strategic constraint and is subject to global supply dynamics; model and compute concentration among a few hyperscalers remains a risk to national sovereignty.

(Each of the above points is drawn from 2024–2025 published studies and official mission materials.)

Training vs Inference: Path to India's AI Sovereignty

16 January 2026

Government of India

Quantum Tiger

Policy brief for startups & public sector stakeholders

Key gaps that must be addressed now

- Compute & datacenter capacity for training: India lacks enough certified, auditable, high-density GPU/accelerator clusters under sovereign control to train large foundation models end-to-end.
- Data pipeline maturity: While AI Kosh is a major step, data breadth and depth across regulated sectors are still incomplete and subject to privacy/regulatory ambiguity.
- Financing & procurement: Early-stage startups face severe capex barriers for training hardware and there are limited government procurement paths preferring sovereignly-trained models.
- Standards & auditability: No single, accepted national certification for model provenance, supply-chain verification, or "training-safety" audits exists yet.
- Talent & systems engineering scale: Need more ML systems engineers and distributed-systems expertise capable of running and optimising multi-PFLOP training workloads.

Policy recommendations

1) Build sovereign training compute (short–medium term)

- National AI Training Hubs (regional, 3–5): Fully funded multi-megawatt hubs with tenant slots available to startups, public labs and universities under secure sovereign SLAs. Hubs should be regionally distributed to avoid central points of failure.
- IndiaAI Compute Credits: Means-tested credits to cover training hours on hub infrastructure (grant-based for early-stage startups; subsidized priced for scale-ups). Use mission funds for initial capex and an ongoing revolving model for sustainability.

2) Expand and operationalize data trusts (AI Kosh → AI Commons)

- Sectoral Data Trusts: Federated public–private trusts for health, agriculture, telecom telemetry and satellite imagery with clear legal schema for pseudonymization and approved access for training under audit logs.
- Synthetic & augmentation pipelines: Invest in certified synthetic-data generators and validation toolkits that allow safe model training when real-world datasets cannot be shared

3) Talent, research & CoEs (medium term)

- Multi-year fellowships & CoEs: Fund 200–300 fellowships for ML systems and model-systems engineers; create CoEs in partnership with IITs/IISc and industry to focus on distributed training, model efficiency, and ML infrastructure.
- Industrial placements & retention packages: Encourage internships, tax credits for domestic R&D hires, and matched funding for startups hiring PhD-level ML systems engineers.

Training vs Inference: Path to India's AI Sovereignty

16 January 2026

Government of India

Quantum Tiger

Policy brief for startups & public sector stakeholders

4) Financing & procurement (immediate)

- Public procurement preference: Create an “India-trained model” procurement bonus in government RFPs and a fast-track sandbox for pilots run by startups that train models in certified sovereign hubs.
- Capex support: Time-limited concessional loans / matched grants for acquisition of racks by approved Indian entities, tied to onshore training commitments and audited logs.

5) Standards, certification & security (parallel)

- Indian Sovereign Model Certification (ISM-Cert): A government-backed certification for models trained on certified infrastructure, covering provenance, training-data lineage, and red-team results.
- Operational security & supply-chain audits: Mandate hardware supply-chain disclosure for hub providers and periodic third-party audits.

6) International cooperation & controlled supply chains

- Strategic partnerships: Negotiate bilateral/triangular agreements for reliable accelerator supply lines, while incentivizing domestic assembly and manufacturing incentives for critical components.

Practical asks for startups such as Quantum Tiger

1. Access: 50–200k GPU-hours in IndiaAI Training Hubs over 12 months via the Compute Credits program.
2. Procurement priority: Pilot contracts for GT-enabled modules in public-sector proof-of-concepts, with a 10% scoring bonus for ISM-Cert models.
3. Capex support: Concessional loan facility for purchasing racks for controlled on-premise training clusters (with audit and export controls).
4. Data access: Approved pipelines to sectoral data trusts (health, telecom, e-commerce imagery) and access to AI Kosh extensions.
5. R&D tax & IP fast-track: R&D tax credits and acceleration lanes for IP filings for training-related inventions and model IP.

These specific measures will let startups transition from inference-providers to full-stack model developers, crucial for sovereignty.

Implementation timeline

- 0–12 months: Launch compute credits pilot, stand up first regional hub, publish ISM-Cert blueprint, expand AI Kosh dataset ingestion priorities.

Training vs Inference: Path to India's AI Sovereignty

16 January 2026

Government of India

Quantum Tiger

Policy brief for startups & public sector stakeholders

- 12–36 months: Scale hubs to 3–5 regions, roll out fellowship & CoE programs, implement procurement sandbox and capex financing.
- 36–60 months: Mature domestic supply chain incentives, advanced certification adoption, and measurable shift of training cycles to domestic infrastructure.

Indicative budget & KPIs

- Indicative 5-year allocation: ₹8,000–12,000 crore (compute hubs capex & operations, credits, fellowships, CoEs). The IndiaAI Mission has already allocated a multi-thousand-crore framework which should be mapped to these priorities.
- KPIs: % training cycles onshore; number of ISM-Cert models; GPU-hours allocated to startups; number of sectoral datasets available in AI Kosh; # of PhD-level systems engineers trained and retained.

Regulatory & ethical safeguards

- Mandatory training logs & provenance for models trained with public credits.
- Periodic third-party audit & red-team tests for models in high-impact sectors.
- Privacy-preserving baseline rules for data trust access with expedited dispute-resolution for contested flows.

Risks & mitigations

- Concentration risk: mitigate with 3–5 hubs, multi-vendor hardware and regional distribution.
- Funding sustainability: revolve compute credits via fee-for-service after seed period and create government-backed demand via procurement.
- Global supply volatility: secure agreements and incentivize domestic manufacturing.

In closing

By focusing policy on the training lifecycle, compute, data, people, finance, procurement and standards, India can convert its strong AI adoption into durable AI sovereignty. The IndiaAI Mission and AI Kosh are already strong foundations; targeted investments in regional training hubs, compute credits, certification and procurement incentives will allow startups like Quantum Tiger to scale model training domestically, keep IP in India, and deliver sovereign AI solutions.

Training vs Inference: Path to India's AI Sovereignty

16 January 2026

Government of India

Quantum Tiger

Policy brief for startups & public sector stakeholders

Key references (selected, 2024–2025)

- NASSCOM — AI Adoption Index / AI Adoption Index 2.0 (2024). (India adoption & sectoral maturity). https://nasscom.in/knowledge-center/publications/ai-adoption-index-20-tracking-indias-sectoral-progress-ai-adoption?utm_source=quantumtiger.in
- McKinsey — The State of AI, 2025 (Agents, innovation & transformation). (Global adoption metrics; scale vs pilot dynamics). https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai?utm_source=quantumtiger.in
- IndiaAI / PIB — IndiaAI Mission (Cabinet approval; mission details & mission budget overview). https://www.pib.gov.in/PressNoteDetails.aspx?ModuleId=3&NotelId=156786&id=156786&utm_source=quantumtiger.in®=3&lang=2
- IndiaAI / AI Kosh (AI Kosh dataset figures — 2,000+ curated non-personal datasets as of 15 Sep 2025). https://niti.gov.in/sites/default/files/2025-09/AI-for-Viksit-Bharat-the-opportunity-for-accelerated-economic-growth.pdf?utm_source=quantumtiger.in
- NITI Aayog / India roadmaps — AI for Viksit Bharat; Roadmap on AI for Inclusive Societal Development (2025 updates). https://niti.gov.in/sites/default/files/2025-10/Roadmap_On_AI_for_Inclusive_Societal_Development.pdf?utm_source=quantumtiger.in

Training vs Inference: Path to India's AI Sovereignty

16 January 2026

Government of India

Quantum Tiger

Policy brief for startups & public sector stakeholders

Author



Rana Dutta is Co-Founder & CEO of Quantum Tiger, with close to 30 years building and scaling technology companies across AI, data infrastructure, healthcare, and enterprise IT. A serial founder and operator, he has led ventures from zero to global scale, advised boards, and trained 5,000+ professionals. Currently focused on sovereign, on-prem AI infrastructure.

 www.linkedin.com/in/duttarana

Co-Author



Manas Paul is Co-Founder & COO of Quantum Tiger with 20+ years leading large-scale operations, HR tech, and digital transformation across healthcare, infrastructure, and enterprises. He brings deep expertise in AI-augmented people analytics, HRMS platforms, and data-driven decision systems, and now focuses on building scalable, AI-native operating foundations for Quantum Tiger.

 www.linkedin.com/in/manaspaul

Quantum Tiger A Moring Bay Technology Pvt Ltd Company

P-175, (A-1/1), Diamond Park Joka, Amgachi, South 24 Parganas, Kolkata, West Bengal, India, 700104

Phone: +91-6206081212

Phone: +919831168903

Email: hello@quantumtiger.in

Web: <https://quantumtiger.in>

Follow us on: [LinkedIn](#)/[X](#) [Twitter](#)/[facebook](#)/[Instagram](#) @quantumtigerhq