



Sovereign by design

Building India's independent,
inclusive and future-ready AI



October 2026

kpmg.com/in

KPMG. Make the Difference.



Glossary

Key term

What we mean in this paper

Sovereign AI

A nation's ability to develop, deploy and govern AI on its own terms: control over critical data, models, compute and rules, not isolation from global technology

Sovereign compute

AI compute (GPUs/accelerators, data centres, cloud) located in India, under Indian jurisdiction, with assured access for national priorities

Indic LLMs

Large and small language models trained or adapted for Indian languages, scripts, speech and context

DPI

Digital Public Infrastructure: shared, interoperable public platforms such as Aadhaar, UPI, DigiLocker and Bhashini

PoC to institution

The journey from proof of concept and pilots to AI services embedded in budgets, processes, people and policy

Responsible AI

AI that is safe, fair, transparent, accountable and privacy-preserving by design

Foreword

Dear Readers,

India stands at a pivotal moment in its AI journey. Over the last decade the country has shown the world that digital public infrastructure can be built at population scale, with sovereignty and inclusion at its core; in 2023 the G20 New Delhi Leaders' Declaration endorsed a G20 Framework for Systems of Digital Public Infrastructure. Aadhaar gave more than a billion people a digital identity; UPI made real-time payments a public good. Artificial intelligence is the next and more consequential frontier, because AI is expected to increasingly shape how citizens are served, how enterprises compete and how the state makes decisions.

As adoption accelerates across sectors, one question has moved from the margins to the center of policy and boardroom conversations: who controls the intelligence that runs India? Digital sovereignty in AI is about ensuring that India has the capability, control and confidence to innovate, deploy and govern AI on its own terms. It is not about building walls. It is about deciding deliberately which layers of the AI stack India must own, which it should co-create with global partners, and which it can safely procure.

Our central argument is simple: India must be sovereign by design. Sovereignty cannot be bolted on after systems are built, any more than security or trust can. It must be engineered into data architectures, model choices, procurement contracts, governance frameworks and talent strategies from day one.

We hope this paper offers policymakers, industry leaders and innovators a practical framework and a set of actionable priorities to build an independent, inclusive and future-ready AI India by 2030. We thank all the contributors who generously shared their time and perspectives.



Brijendra Kumar

Partner and Head GovTech and
Lead AI in Government
KPMG in India



At a glance

The pilot era has to end. Now India must own its AI future.

India has built public AI capacity at speed: India AI Mission set out in March 2024 to deploy 10,000 or more GPUs, and by December 2025 more than 38,000 had been onboarded for common use.¹ Thus, the question for the next four years is no longer whether India can build AI, but whether it can control, scale and trust it. Five imperatives, each answering a critical question for India, define the path.

Indicator	Where India stands
IndiaAI Mission outlay	INR 10,371.92 crore, approved in March 2024 across seven pillars
Common compute	More than 38,000 by December 2025, at a subsidised INR 65 per GPU-hour; a further 20,000 GPUs announced in February 2026 by IT Minister Shri Ashwini Vaishnaw at Day Two of India AI Impact Summit 2026 ²
Sovereign foundation models	Sarvam AI selected in April 2025 to build one of India's first indigenous sovereign foundation model; portfolio expanded to 20 indigenous model proposals (12 LLMs and 8 SLMs) for support by July 2026. ³
Datasets	More than 15,000 datasets, 300 AI models and 30 toolkits on AIKosh as of August 2026. ⁴
Language	22 scheduled languages ⁵ ; Bhashini as national language-AI infrastructure
Clean power	Non-fossil sources passed 50% of installed capacity in June 2025; 283.46 GW of non-fossil capacity by 31 March 2026 ⁶

¹ Transforming India with AI, Press Information Bureau, Government of India, December 2025

² India's AI Moment: 20,000 additional GPUs power India's next phase of AI leadership, Press Information Bureau, Government of India, February 2026

³ Govt identifies 20 indigenous AI foundation model proposals under IndiaAI Mission: Parliament, Asian News International, July 2026


⁴ IndiaAI Mission Organises Workshop on AIKosh to Strengthen India's

AI-Ready Data Ecosystem, Press Information Bureau, Ministry of Electronics and Information Technology, August 2026

⁵ Census of India 2011, Paper 1 of 2018: Language, India, States and Union Territories, Office of the Registrar General and Census Commissioner, India, June 2018

⁶ India Ranks third globally in Renewable Energy Installed Capacity: Shri Pralhad Joshi, Press Information Bureau, Ministry of New and Renewable Energy, April 2026

Imperative	In one line
1. Make AI growth inclusive	AI that speaks every Indian language and reaches the last mile, not just the metros
2. Build trusted data and AI governance	Trust is the licence to operate: govern data and models by design, not after deployment
3. Move from PoC to institution	Fund platforms and outcomes, not pilots; build the plumbing that makes AI durable
4. Strengthen the innovation ecosystem	Make government the first and repeat customer, and back Indian builders with compute, patient capital and IP protection
5. Set out a sovereign AI future for 2030	Define what sovereign means, measure it, and hold the system accountable to milestones



01

The Strategic Imperative for Sovereign AI





Sovereignty in AI is not about doing everything ourselves. It is about never being in a position where someone else can switch off our ability to serve our citizens.



— Nilachal Mishra,
Head of Government and Public Sector, KPMG in India

For the last two decades, technology choices in government were largely procurement decisions. With AI, they have become strategic decisions. Models encode language, values and assumptions; compute is concentrated in a handful of supply chains; and we believe the data that trains AI should be treated as a national asset. Three pressures make sovereignty an immediate priority for India.

Three pressures

- **Geopolitics and supply concentration** - Advanced AI chips, frontier models and hyperscale cloud are concentrated in a few countries and companies. The United States holds about 75 per cent of the world's AI supercomputer performance and China about 15 per cent⁷, while AWS, Microsoft and Google together held 63 per cent of the cloud infrastructure market in the first quarter of 2026⁸. Export controls, licensing changes or service withdrawals can affect access overnight. A nation that depends entirely on external compute and models for critical public services carries a strategic risk.
- **Data, privacy and dignity** - Citizen data, from health records to land and welfare data, is among one of the most sensitive data any state holds. The Digital Personal Data Protection Act, 2023 and the DPDP Rules, 2025, notified in November 2025, create obligations on how personal data is processed; most of these obligations take effect in May 2027, at the end of an 18-month transition. AI that is trained, fine-tuned or run on such data must operate within clear jurisdictional and governance boundaries.
- **Language, context and inclusion** - India's linguistic diversity is both its greatest AI challenge and its one of the greatest opportunity. Global models are improving rapidly in Indian languages, but performance on low-resource languages, dialects, code-mixed speech and local context remains uneven. AI that does not understand Bharat cannot serve Bharat.

The sovereignty spectrum

Sovereignty is not binary. We recommend thinking of it as a spectrum, using the KPMG in India's framework below¹⁰, and choosing a deliberate position for every layer of the AI stack.

⁷ Trends in AI Supercomputers, Epoch AI, Konstantin F. Pilz, Robi Rahman, James Sanders and Lennart Heim, April 2025

⁸ Cloud Market Share Q1 2026: AWS, Microsoft, Google Battling In AI Era (citing Synergy Research Group data), CRN, Mark Haranas, May 2026

⁹ DPDP Rules, 2025 Notified, Press Information Bureau, Government of India, November 2025

¹⁰ KPMG in India's 2026 analysis, based on the firm's work with central and state governments. This applies to all frameworks, ratings, indicators and the roadmap in this paper

Position	Description	When appropriate
Dependent	Fully reliant on foreign-owned and foreign-hosted AI, with limited control over data or continuity	Non-critical, low-sensitivity productivity use cases
Partnered	Global technology used under Indian contractual, data-residency and audit terms	Most of the enterprises and many government workloads
Controlled	Indian control over data, governance, deployment and switching, with a mix of Indian and global components	Critical citizen services, regulated sectors, strategic data
Self-reliant	Indigenous capability end to end, built and operated in India	Defence, national security, select language and public-good models

India has done this before

- Aadhaar and UPI showed that public infrastructure, open standards and private innovation can coexist, with public institutions owning the rails and the market building on them.
- The same model, public rails plus open innovation, is one of the most credible route to AI sovereignty: national compute, datasets, language infrastructure and assurance as public goods; products and services built competitively on top.
- The goal is 'controlled' sovereignty on the layers that matter, not autarky across the whole stack.




Key Points

- **Sovereignty is a strategic design choice, not a procurement afterthought.**
- **Choose a sovereignty position for every layer of the stack and every class of workload.**
- **India's DPI playbook is its competitive advantage in AI.**



02

The Sovereign AI Stack



To make sovereignty actionable, we break AI into six layers with each having a different strategic weight, a different current position for India and a different answer to the build-versus-partner question.

1. **Compute and semiconductors:** GPUs and accelerators, data centres, sovereign and in-country cloud, energy, and the longer-term semiconductor roadmap.
2. **Data:** Public and private datasets, consent architecture, sector data exchanges, Indic language corpora, and data quality.
3. **Models:** Foundation models, Indic LLMs and SLMs, speech and vision models, sector-specific and fine-tuned models, and open-weight options.
4. **Platforms and DPI:** Bhashini, AIKosh, model marketplaces, AI-ready government platforms, and APIs that let models plug into citizen services.
5. **Talent, research and IP:** Researchers, engineers, AI-literate public servants, academic labs, retention of top talent in India, and Indian ownership of the intellectual property they create.
6. **Governance and trust:** Policy, standards, evaluation and red-teaming, AI assurance, procurement guardrails, and grievance redressal.

Assessment: India's position and recommended approach by layer

Layer	India today	Control needed	Approach
Compute	Medium	High	Hybrid: national pool plus global partners with in-country hosting, on sites with assured clean power and water
Data	Medium	High	Build and own: national data exchanges, consent, Indic corpora
Models	Medium	Medium	Hybrid: Indic sovereign models plus open and global models via governed access
Platforms/DPI	High	High	Build and own: extend the DPI playbook to AI
Talent	Medium	High	Build: fellowships, labs, public-sector AI cadre
Governance	Medium	High	Own: Indian rules, assurance and evaluation capacity

Ratings are KPMG in India's indicative assessment

What it takes to hold a position - Choosing a 'controlled' or 'self-reliant' position for a layer is only half the decision. The position holds only if four conditions are met: Indian builders can sell beyond the first pilot, they can raise capital patient enough for deep technology, they own the IP they create, and the compute they run on has assured power and water. Miss any one, and a self-reliant position on paper becomes dependence in practice.


Power and water: the physical layer of sovereignty

Sovereign compute is only as dependable as the power and water behind it. A data centre without assured power or water can be switched off as surely as a foreign cloud. India's data centre capacity rose from 375 MW in 2020 to over 1,500 MW in 2025¹¹ and could reach 4.5 to 6.5 GW by 2030. Data centres already use about 0.5% of India's electricity and about 150 billion litres of water a year, and both are expected to more than double by 2030. Developers treat access to reliable power as the primary siting criterion, and water is a growing risk where urban demand is already high, as in Chennai and Bengaluru.¹²

India starts from strength on power. Non-fossil sources passed 50%Q of installed capacity in June 2025, five years ahead of the 2030 target, and the SHANTI Act, 2025 allows private companies, under license, to build and operate nuclear power plants. Water is the gap. In March 2026 the Government told Parliament that MeitY had received no concerns about water use by data centres. In our view that reflects an absence of measurement, not an absence of risk.


Key Points

- Own the data, platform and governance layers outright.
- Partner on frontier compute and models, but on Indian terms: residency, auditability and portability.
- Invest in Indic models where global offerings under-serve India's languages and public-good needs.
- Site compute where clean power and water are assured, and measure both.



¹¹ Water Consumption by Data Centers in the Country (Rajya Sabha Unstarred Question No. 4074), Press Information Bureau, Ministry of Jal Shakti, March 2026

¹² Scaling India's Data Centre Ecosystem: Stakeholder Perspectives on Infrastructure, Energy, and Resilience, Council on Energy, Environment and Water and SYSTEMIQ, Vishal Tripathi et al., February 2026



03

Imperative 1: Make AI growth inclusive



The central question: *How can AI create tangible impact across key sectors such as healthcare, education, agriculture, manufacturing and public services, ensuring that the benefits are widely shared?*

The challenge

- In our view, one of the most AI value today is captured by English-speaking, digitally fluent, urban users and large enterprises. Access is part of the reason: in March 2026 rural India had 48 internet subscriptions per 100 people, against 127 in urban India (TRAI).¹³ Yet India's one of the largest opportunities lie with the farmer seeking advisory in Marathi, the ASHA worker recording health data by voice, the MSME owner filing compliance in Tamil, and the citizen trying to resolve a grievance in Bhojpuri.
- If AI is designed for the top of the pyramid, it may widen digital divides rather than close them. Inclusion must therefore be a design principle, not a downstream outcome.

What good looks like

- Voice-first, language-first interfaces as the default for citizen services, built on Bhashini and Indic models.
- AI services that work on low-cost devices, low bandwidth and assisted channels such as Common Service Centres.
- Sector programmes in agriculture, health, education and MSMEs with measurable outcomes for underserved groups.
- Inclusion metrics, by language, gender, geography and disability, embedded in most of the public AI programme.

By the numbers

- India has 22 scheduled languages among 121 languages and 1,369 rationalised mother tongues. Bhashini integrates more than 350 AI models and has over 450 active customers and around 1.2 million downloads.

Actions



Mandate language access

Make multilingual and voice access a default requirement in public AI procurement.



Fund last-mile AI

Create challenge funds for AI solutions targeting farmers, frontline workers and MSMEs.



Measure inclusion

Publish inclusion dashboards for national AI programmes.



Build assisted channels


Integrate AI into CSCs, banking correspondents and frontline worker tools.

¹³ The Indian Telecom Services Yearly Performance Indicators 2025-26, Telecom Regulatory Authority of India, October 2026

Who must act	Priority role
Government	Language mandates, inclusion metrics, last-mile delivery channels
Industry	Indic-first product design; low-cost, low-bandwidth deployment
Startups and academia	Low-resource language datasets, dialect and speech models

Key Points

- Inclusion is a design requirement, not an outcome to hope for.
- Language is India's sovereign AI moat.
- What gets measured by language and geography gets delivered.





04

Imperative 2: Build trusted data and AI governance



The central question: *What policies, standards and safeguards are required for responsible, secure and transparent use of data and AI at scale?*

The challenge

- We believe a single high-profile failure, such as a biased welfare-eligibility model, a data leak through a public chatbot or a hallucinated legal answer, can set back adoption for years.
- India has taken important steps: the DPDP Act, 2023 and the DPDP Rules, 2025, with most of the obligations in force from May 2027; the announcement of the IndiaAI Safety Institute in January 2025; and national AI governance guidelines released by MeitY in November 2025. The gap now is operational: turning principles into standards, testing, procurement clauses and accountability that departments can actually apply.

What good looks like

- Data governed as a strategic asset: classified, catalogued, consented and quality-assured.
- Sector data exchanges (health, agriculture, mobility, finance) with privacy-preserving access for AI training.
- Model evaluation before deployment, covering accuracy in Indian languages, bias, safety and security, with independent red-teaming for high-risk uses.
- Clear rules on where sensitive workloads may be hosted and processed, and contractual rights to audit, exit and port.
- Human oversight and citizen grievance redress for AI-assisted decisions.

Stage of the trust lifecycle	Key control
Collect	Consent, purpose limitation, data minimisation
Govern	Classification, cataloguing, quality, access control
Train	Dataset provenance, bias checks, documentation
Test	Indic accuracy, safety, red-teaming, security testing
Deploy	Risk-tier approval, human oversight, disclosure to citizens
Monitor	Drift, incidents, grievance redress, periodic audit

Actions



Classify AI use by risk

Adopt a tiered risk framework for public-sector AI with proportionate controls.



Standardise procurement clauses

Data residency, audit rights, model transparency, IP and exit/portability in every AI contract.



Build evaluation capacity

Indian benchmarks and test suites for Indic performance, bias and safety.



Operationalise DPDP for AI

Consent, purpose limitation and retention guidance specific to AI training and inference.

Who must act	Priority role
Government	Standards, risk tiers, assurance bodies, procurement guardrails
Industry	Responsible AI operating models, transparency, in-country controls
Startups and academia	Benchmarks, evaluation tools, privacy-enhancing technologies

Key Points

- **Govern by design:** governance added after deployment erodes trust.
- **Contracts are a sovereignty instrument:** residency, audit and exit rights matter.
- **India needs to scale and institutionalise its own evaluation benchmarks** for Indic AI, building on efforts such as MILU.



05

Imperative 3: Move from PoC to institution



The central question: *What will it take to move AI solutions from pilots and isolated use cases to enterprise and national-scale adoption with sustained value?*

The challenge

- In our work across ministries, states and PSUs, we see AI pilots in abundance and AI in production at scale only rarely. In our experience, pilots tend to stall for predictable reasons: they are funded from innovation budgets with no provision for running costs; data pipelines are built manually for the demo; there is no business owner accountable for outcomes; and procurement rules are designed for buying hardware and software, not for buying evolving models and outcomes.
- The result is ‘pilot purgatory’: impressive demonstrations that never change how citizens are served.

What good looks like

- Every PoC has a defined scale-up path, owner, budget line and success metric before it starts.
- Shared AI platforms (common models, data layers, MLOps, security) reused across departments, not rebuilt for each use case.
- Outcome-based and consumption-based procurement models for AI.
- A central AI programme office or ‘control tower’ that tracks portfolio value, risk and reuse.

Stage	Gate question	Typical owner
PoC	Does it work technically on real Indian data?	Innovation/IT team
Pilot	Does it deliver measurable value with real users?	Business department
Platform	Can it be reused securely across use cases?	State/national AI platform team
Institution	Is it embedded in budget, process, people and policy?	Department head and finance

Actions



Stage-gate the pilots

Adopt a PoC → Pilot → Platform → Institution model with clear go/no-go criteria.



Budget for run, not just build

Ring-fence operations and maintenance funding for AI services in production.



Reform AI procurement

Create AI-specific procurement playbooks and model framework contracts.



Build once, reuse many

Fund shared AI platforms at national and state level.

Who must act	Priority role
Government	Funding models, procurement reform, AI programme offices
Industry	Production-grade platforms, MLOps, outcome-linked commercial models
Startups and academia	Reusable components; integration with government platforms

Key Points

- Many pilots do not add up to transformation.
- Scale is an institutional problem more than a technology problem.
- Platforms and procurement are the levers that unlock scale.



06

Imperative 4: Strengthen India's AI innovation ecosystem



The central question: *How can we support AI startups, build indigenous capabilities and nurture talent to drive global competitiveness?*

The challenge

As per the Stanford AI Index 2025, India ranked second in the world for AI skill penetration between 2015 and 2024, and 74 newly funded Indian AI companies in 2024 put it fourth globally.¹⁴ It still faces four constraints: affordable access to compute, access to high-quality Indian data, markets beyond the first pilot, and capital patient enough to wait for deep technology to mature. Markets and capital, with IP ownership below, are three of the four conditions set out in Section 02 for holding a 'controlled' or 'self-reliant' position.

The IndiaAI Mission's subsidised common compute, offered at INR 65 per GPU-hour, and its support for 20 indigenous foundation model proposals are significant steps. The next step is demand: Indian AI companies need anchor customers, and public procurement accounts for roughly 20 to 22 per cent of India's GDP.¹⁵

- **Markets** - A first order is only a pilot, and a sovereign model no one buys again is not sovereign capability. The rules already let startups in: since September 2016, central ministries may relax prior turnover and prior experience conditions for startups, and by February 2025 startups had fulfilled orders worth INR 35,950 crore through GeM.¹⁶ What AI lacks is a route from a successful pilot to repeat purchase across departments. Defence shows one can be built: iDEX had signed 430 contracts by February 2025, and the Ministry of Defence had cleared the procurement of 43 items worth over INR 2,400 crore from iDEX startups and MSMEs for the Armed Forces.¹⁷ The IndiaAI Startups Global Acceleration Program has selected 10 Indian AI startups to Station F and HEC Paris to enter global markets.
- **Capital** - Self-reliant positions need Indian companies that can last for years before they earn revenue. Indian private investment in AI totalled about USD 11 billion between 2013 and 2024, against US\$471 billion in the United States. Government has begun to supply long-horizon money. The INR 1 lakh crore Research Development and Innovation (RDI) scheme, launched in November 2025, provides long-term, patient capital, and its first 19 fund managers were approved in October 2026 with INR 7,131 crore of commitments.¹⁸ The INR 10,000 crore Startup India Fund of Funds 2.0, approved in February 2026, prioritises breakthroughs that "require patient, long-term capital".¹⁹ The gap now is matching this capital to AI builders, and to the public demand that lowers their risk.

¹⁴ The AI Index 2025 Annual Report, Stanford Institute for Human-Centered Artificial Intelligence, Nestor Maslej et al., April 2025

¹⁵ Government E-Marketplace (GeM) Factsheet, Press Information Bureau, Government of India, November 2021

¹⁶ Fostering Growth and Inclusivity: The Strategic Impact of GeM on India's Economy, Press Information Bureau, Government of India, February 2025

¹⁷ Make in India Powers Defence Growth, Press Information Bureau, Ministry of Defence, April 2025

¹⁸ Research, Development and Innovation Fund Announces First Cohort of 19 Second Level Fund Managers, Department of Science and Technology, Ministry of Science and Technology, October 2026

¹⁹ Cabinet approves Startup India Fund of Funds 2.0 to mobilize venture capital for India's startup ecosystem, Press Information Bureau, Government of India, February 2026

Research and IP: invent here, own here

IP is the legal form of the control this paper argues for, and India is inventing but not yet owning at scale. It was the fifth-largest location for generative AI inventions between 2014 and 2023, with 1,350 patent families and one of the fastest growth among the top five, at 56 per cent a year; China had more than 38,000 and the United States about 6,300.²⁰ Recognition has begun: DPIIT's February 2026 startup framework created a Deep Tech Startup category, with a 20-year recognition window and a INR 300 crore turnover limit, for firms built on new scientific or engineering knowledge, significant R&D and novel IP.²¹ One open policy choice will shape the cost of building Indian models: a DPIIT committee's December 2025 working paper proposes a mandatory licence to train AI on lawfully accessed copyrighted works, with statutory royalties collected centrally.

What good looks like

- Compute treated as a public utility: predictable, affordable and easy to access for startups and researchers.
- Government as 'first customer' through innovation procurement, sandboxes and challenge programmes.
- Open Indic datasets and benchmarks that lower the cost of building for India.
- Research-to-startup pipelines linking IITs, IISc, AI4Bharat-type labs and industry.
- Reverse brain drain: research fellowships and frontier problems that keep top talent in India.
- Markets beyond the pilot: an AI solution proven in one department can be bought by others through GeM without repeating the pilot, with startup relaxations applied by default.
- Patient capital matched to public demand: a share of RDI and Fund of Funds 2.0 commitments reaching AI models, chips and tooling, backed by multi-year government purchase commitments.
- Indian AI IP recognised and retained: public AI contracts that leave ownership of what is built with Indian builders while government keeps use rights, and an annual national register of Indian AI models, datasets and patents in deployment.

By the numbers

- Sarvam AI was selected under the IndiaAI Mission in April 2025 to build India's first indigenous foundation model, followed by Soket AI, Gnani.ai and Gan.AI.
- By July 2026 the IndiaAI Mission had empanelled 15 compute providers, sanctioned 93.18 lakh GPU-hours, supported 237 projects and identified 20 indigenous model proposals (12 LLMs, 8 SLMs) for support.

²¹ Patent Landscape Report: Generative Artificial Intelligence (GenAI), World Intellectual Property Organization, July 2024

²¹ Notification G.S.R. 108(E) on recognition of startups and deep tech startups, Department for Promotion of Industry and Internal Trade, Ministry of Commerce and Industry, February 2026

Actions



Become the first customer

Earmark a share of public AI spend for Indian startups through innovation procurement.



Simplify compute access

Self-service, transparent allocation of national compute with predictable pricing.



Open up data Expand AIKosh with high-value, well-documented Indic datasets.



Invest in people

Public-sector AI cadre, fellowships and joint industry-academia labs.



Open markets beyond the pilot Create a repeat-purchase route on GeM for AI solutions proven in one department, and grow IndiaAI Startups Global into export partnerships tied to India's DPI programmes.



Make capital patient

Direct a share of RDI and Fund of Funds 2.0 commitments to AI and compute companies, and pair them with multi-year public purchase commitments.



Recognise and protect Indian AI IP

Write IP terms into public AI contracts that leave ownership with Indian builders, settle the AI training and copyright framework, and publish an annual register of Indian AI models and patents in deployment.

Who must act	Priority role
Government	Demand creation, compute access, open data, talent programmes
Industry	Partnerships with Indian startups, co-development, skilling at scale
Startups and academia	Indic models, applied research, commercialisation

Key Points

- Demand, not just supply, builds an ecosystem.
- Government procurement is one of India's most powerful industrial-policy levers for AI.
- Talent stays where the hardest, most meaningful problems are.
- Patient capital and public demand must arrive together.
- IP created in India should stay with Indian builders.

Imperative 5: Set out a sovereign AI future for 2030



The central question: *What does a digitally sovereign and inclusive AI India look like, and what are the key milestones to get there?*

The challenge

- Ambition without measurement becomes rhetoric. For sovereignty to move from slogan to strategy, India needs a shared definition of what ‘sovereign enough’ means for every layer, a small set of indicators, and a published scorecard that holds the system accountable.
- We propose a Sovereign AI Readiness Index, built with government, industry and academia, to track progress annually.

What good looks like

- A clear national definition of sovereignty positions by layer and workload class.
- An annual, published readiness scorecard.
- Alignment of central and state AI strategies to common milestones.
- India's AI DPI can be exported as a model for the Global South, building on the G20's welcome, in 2023, of India's plan to build a Global DPI Repository.²²

Pillar	Indicative indicator
Compute	Share of critical public AI workloads on in-country compute
Data	Number and use of high-value datasets on national platforms
Models	Share of citizen-facing AI services using Indic or sovereign-controlled models
Adoption	Number of AI services in production at national and state level; citizens reached
Trust	Share of high-risk AI systems independently evaluated before deployment
Talent	AI-skilled public servants; researchers retained or returned
Markets and capital	Share of public AI spend with Indian startups; long-tenor capital committed to Indian AI companies
Research and IP	Indian generative AI patent families; Indian-owned AI models and datasets in production
Power and water	Share of national AI compute on clean power; water used per unit of national AI compute

²² G20 New Delhi Leaders' Declaration, G20 (India Presidency), September 2023

Actions



Define sovereignty Publish sovereignty positions by layer and workload class.



Measure annually Launch a Sovereign AI Readiness Index with a public scorecard.



Align states Use common milestones for central and state AI missions.




Export the model Package India's AI DPI for Global South partners.

Who must act	Priority role
Government	Definition, index, scorecard, state alignment
Industry	Contribute data and benchmarks; align investment to milestones
Startups and academia	Independent measurement and research

Key Points

- What is not measured will not be sovereign.
- Milestones must be shared across centre and states.
- India's AI DPI can be exported as a model for the Global South.



08

Five hard truths from the field



Drawing on KPMG in India's work with governments and conversations with industry leaders, we see five recurring mistakes that undermines sovereign AI ambitions. Each has a practical alternative.

01

Buying GPUs without anchor workloads.

Don't do this: Capacity without demand sits idle, while costs and obsolescence accumulate.

Do this instead: Link compute investment to anchor workloads, model-training programmes and public AI services.

02

Running pilots for their own sake.

Don't do this: Dozens of PoCs create activity, not transformation.

Do this instead: Fund only pilots that have an owner, a scale-up path and a budget line for production.

03

Building models in isolation.

Don't do this: A sovereign model no one adopts is not sovereign capability.

Do this instead: Design for distribution: APIs, integration with DPI and government platforms, and developer ecosystems.

04

Treating governance as an afterthought.

Don't do this: Retrofitted controls are expensive and erode public trust.

Do this instead: Govern by design: risk tiers, evaluation and procurement clauses from day one.

05

Training talent that then leaves.

Don't do this: Skilling without career paths feeds other ecosystems. India already shows a net outflow of AI talent, at minus 1.55 per 10,000 LinkedIn members

Do this instead: Create public-sector AI cadres, research fellowships and frontier problems worth staying for.

09

Sector spotlights



Sovereignty requirements and AI value vary by sector. Ratings below are KPMG's indicative assessment (High/Medium/Low)

Sector	Investment	Adoption	Value potential	Sovereignty need
Healthcare	Medium	Medium	High	High
Agriculture	Medium	Low	High	Medium
Education	Medium	Medium	High	Medium
Manufacturing	High	Medium	High	Medium
Public services and governance	Medium	Medium	High	High

Healthcare

AI can extend scarce clinical capacity through diagnostic support, triage, clinical documentation in Indian languages and population-health analytics on the Ayushman Bharat Digital Mission rails. Health data is among the most sensitive data a state holds, so sovereignty needs are high.

Illustrative use cases

- AI-assisted screening and diagnostics at primary health centres
- Voice-based clinical documentation for frontline workers in local languages

Key Points

- Keep health data processing within Indian jurisdiction with strict consent.
- Validate models on Indian clinical populations before deployment.



Agriculture

Personalised advisory on crops, weather, pests, markets and schemes, delivered by voice in local languages, can reach smallholders at scale. Value is high but adoption depends on trust, accuracy and last-mile channels.

Illustrative use cases

- Multilingual voice advisory integrated with soil health, weather and market data
- AI-assisted crop insurance assessment

Key Points

- **Language and voice are the adoption unlock.**
- **Agri data exchanges with farmer consent are critical.**

Education

AI tutors and teacher assistants in Indian languages can personalise learning and reduce teacher workload, if designed with pedagogy and safety in mind.

Illustrative use cases

- Multilingual tutoring and assessment support
- Teacher co-pilots for lesson planning and content in local languages

Key Points

- **Child safety and data protection need stronger safeguards.**
- **Indic content quality determines outcomes.**

Manufacturing

Quality inspection, predictive maintenance, supply-chain optimisation and design acceleration can raise competitiveness, especially for MSMEs in industrial clusters.

Illustrative use cases

- Computer-vision quality control on production lines
- Predictive maintenance for MSME clusters through shared platforms

Key Points

- **Industrial data is a competitive asset; control matters.**
- **Shared cluster platforms lower MSME adoption costs.**

Public services and governance

This is where sovereignty matters most and where India's DPI advantage is strongest. AI can transform grievance redress, welfare targeting, revenue and public finance, policing and justice, and citizen communication.

Illustrative use cases

- Multilingual citizen assistants for scheme discovery and grievance redress
- AI-assisted analytics for public finance, revenue and welfare leakage detection
- Decision support for policing and justice under strong human oversight

Key Points

- **Citizen-facing and decision-making AI must sit in the 'controlled' sovereignty position.**
- **Risk-tiered governance and human oversight are non-negotiable.**

10

India 2030: What sovereign AI could look like



A day in 2030

Sunita, a farmer in Vidarbha, asks her phone in Marathi whether her crop is eligible for insurance after an unseasonal hailstorm. An AI assistant built on an Indian model, running on in-country compute, checks satellite imagery, her land record and the scheme rules, and starts her claim. She never fills in a form.

Dr. Imran, at a primary health centre in Bihar, dictates his consultation notes in Hindi. The AI structures them into his patient's health record, flags a possible drug interaction, and keeps the data within India under the patient's consent.

A district collector in Odisha reviews an AI-generated early warning on a likely spike in water-borne disease, drawing on water-quality, health and weather data. Every recommendation is explainable, logged and open to audit.

None of this requires India to build everything alone. It requires India to control the data, the rules, the critical platforms and the switch.

KPMG in India's proposed milestone roadmap

Horizon	Milestones
2026–27: Foundations	National compute scaled and easier to access; sovereignty positions defined by layer; AI procurement playbook and model contract clauses adopted; first sector data exchanges live; risk-tiered governance framework operational; power and water criteria set for national AI compute sites; a repeat-purchase route for a strong AI solutions live on GeM; first RDI and Fund of Funds 2.0 commitments reaching AI companies
2028: Scale	Indic and sovereign-controlled models powering priority citizen services; shared AI platforms in leading states; independent AI assurance and evaluation at scale; Sovereign AI Readiness Index published annually; energy and water use of national AI compute published annually; a national register of Indian AI models and patents in deployment
2030: Leadership	AI-enabled public services at population scale in all scheduled languages; a globally competitive Indian AI industry; India exporting AI DPI to Global South partners

11

Top 12 recommendations



For Government

1. **Define sovereignty by layer** - Publish which AI layers and workload classes require 'controlled' or 'self-reliant' positions.
2. **Create national and sector data exchanges** - Make high-value, consented data available for Indian AI under strong governance.
3. **Reform AI procurement** - Adopt outcome-based models and standard clauses for residency, audit, IP and exit.
4. **Fund the run, not just the build** - Ring-fence operations budgets and set up AI programme offices at centre and state level.
5. **Open markets and make capital patient** - Create repeat-purchase routes for proven AI solutions, and direct a share of RDI and Fund of Funds 2.0 capital to AI and compute companies.
6. **Plan power and water with compute** - Site national AI compute where clean power and water are assured, and require large AI data centres to disclose their energy and water use.

For Industry and enterprises

7. **Architect for sovereignty** - Design for data residency, portability and multi-model flexibility from day one.
8. **Build Indic-first** - Treat Indian languages and voice as core product requirements.
9. **Operationalise responsible AI** - Embed governance, evaluation and human oversight into the AI operating model.

For Startups and academia

10. **Open benchmarks and datasets** - Build shared Indic evaluation benchmarks and high-quality open datasets.
11. **Partner for scale** - Use government first-customer programmes and enterprise partnerships to reach production.
12. **Bridge research and industry** - Create joint labs and fellowships that turn research into deployable products and keep talent in India.

Acknowledgements and methodology



This paper draws on KPMG in India's experience with central and state governments, and on secondary research from publicly available sources. Each source is given as a footnote where its data first appears. Frameworks, ratings and the roadmap developed by the authors are KPMG in India's 2026 analysis.

KPMG in India

Aabhas Bhardwaj

Design & compliance

Arun Choudhary

Prajakta Talpade

Darshini Shah

Nidhi Agarwal

KPMG in India contacts:

Akhilesh Tuteja

Partner and Head
Clients and Markets
E: Atuteja@kpmg.com

Brijendra Kumar

Partner and Head GovTech and
Lead AI in Government
E: brijendrakumar@kpmg.com

Nilachal Mishra

Partner and Head
Government and Public Services (G&PS)
E: nilachalmishra@kpmg.com

Udit Sethi

Partner
Digital Solutions – Strategy and Insights
E: uditsethi@kpmg.com

Purushothaman KG

Partner and Head
Artificial Intelligence
E: purushothaman@kpmg.com

kpmg.com/in



Access our latest insights
on KPMG Insights Edge

Follow us on:

kpmg.com/in/socialmedia



The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act on such information without appropriate professional advice after a thorough examination of the particular situation.

The views and opinions expressed herein are those of the quoted third parties and do not necessarily represent the views and opinions of KPMG in India.

KPMG Assurance and Consulting Services LLP, Lodha Excelus, Apollo Mills Compound, NM Joshi Marg, Mahalaxmi, Mumbai – 400 011 Phone: +91 22 3989 6000, Fax: +91 22 3983 6000.

© 2026 KPMG Assurance and Consulting Services LLP, an Indian Limited Liability Partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

This document is for e-communication only. (TL_1026_AC_PT)