cognizant

India's public AI framework

Building 80% of a globally replicable model for development

The moment

The world stands at an inflection point where artificial intelligence is emerging as the primary driver of economic growth, social development and competitive advantage. Yet the benefits of AI remain unevenly distributed, concentrated within advanced economies and a limited set of technological ecosystems, leaving nearly six billion people at risk of exclusion from AI-led productivity and value creation.

India offers a distinct pathway, one that moves beyond isolated deployment toward the industrialization of AI at population scale. By embedding AI into digital public infrastructure and service delivery systems, India demonstrates how AI can be translated into measurable, inclusive outcomes.

Over the past three to five years, India has built an AI-native system of intelligence that combines inclusive access with optimized unit economics. This approach challenges conventional assumptions about cost, complexity and scalability, proving that enterprise-grade reliability and impact can be achieved without fragmented, high-cost implementations.

This document shows how India has systematically built 80% of a seven-layer architecture for public-sector AI—from foundational digital infrastructure through applied sectoral services to emerging governance frameworks. It also defines the remaining 20%, the frontier of cross-sector integration, predictive governance and global knowledge transfer needed to translate AI investment into measurable enterprise and public value.

In that context, Cognizant is positioned not simply as an implementation partner, but as an AI builder and accountable partner for business outcomes that can help industrialize this model at scale.

The seven-layer architecture: Understanding India's approach

India's public AI framework rests on a foundational architecture that distinguishes it from adhoc technology deployment. Rather than treating AI as a series of isolated initiatives, India has built a coherent seven-layer stack adapted from international best practices (OECD, Bertelsmann Stiftung, Open Future Foundation) and tailored to India's scale, diversity and inclusion requirements. This is the difference between fragmented adoption and industrializing AI through connected, enterprise-grade systems of intelligence.

Understanding these layers is essential to grasping both what India has accomplished and what remains to unlock combinatorial innovation:

Layer 1: Applications

Sectoral services reaching citizens—healthcare triage, agricultural advisory, education personalization, grievance redressal, civic services and more

Layer 2: Models

Foundation models, domain-specific models and orchestration systems with built-in policy compliance

Layer 3: Data

Data assets, consent regimes, openness frameworks and pipelines balancing representation, innovation and sovereignty

Layer 4: Compute infrastructure

Chips, cloud platforms and data centers forming a national compute backbone with regional distribution

Layer 5: Energy

Reliable, cost-stable and sustainable power systems underpinning AI operations

Layer 6: Talent and skills

Human capability across government, academia, ecosystem partners and vendors

Layer 7: Governance

Rules, institutions, assurance mechanisms, procurement standards and accountability frameworks ensuring public value and safety

India has materially advanced all seven layers, with foundational governance, data and delivery mechanisms already in place. The remaining 20% is not basic rollout; it is the next stage of value realization, cross-sector integration, predictive governance and global knowledge transfer that can compound outcomes across the full stack.

Part one: The 80% already built

Layer 1: Applications

Real impact at scale

This layer is where India's framework moves from policy intent to operational value. Across critical sectors, AI-enabled services are being embedded into delivery systems that reach citizens at scale and demonstrate how AI-led productivity can translate into measurable outcomes.

India has already shown that application-layer value is strongest when AI is embedded into delivery systems rather than deployed as isolated pilots. Together, the following examples illustrate a broader point: The public sector opportunity is not simply AI deployment, but industrializing AI through connected workflows, trusted data exchange and outcome accountability.

For healthcare, digital public infrastructure such as ABDHM creates the conditions for AI-assisted diagnostics to reach underserved populations with lower cost and faster intervention.

Cerebra AI's TB detection deployment demonstrates how AI-enabled diagnostics can shorten diagnosis cycles and improve access in underserved districts.

In agriculture, AgriStack-linked advisory services show how context-infused, enterprise-specific capability can improve farmer decisions on a scale. Additionally, AgriStack-linked advisory and tools such as CottonAce show how AI-led productivity can improve farm outcomes through timely, localized recommendations.

With education, Shiksha Setu in Assam showed how AI could improve attendance, reduce leakages and strengthen dropout tracking, with the next step focused on personalized learning linked to DigiLocker credentials and employment pathways. Across finance, India used Aadhaar, UPI, Account Aggregator and OCEN to move credit beyond collateral, widening faster formal access for informal workers, farmers and microenterprises.

In urban governance, AI made service delivery more responsive through civic monitoring, UMANG and DigiYatra, improving issue detection, access and operational efficiency at scale.

In direct benefit transfers, AI helped identify duplicate and fake beneficiaries, reduce leakage and improve the precision and transparency of welfare delivery.

Additional application case studies are provided in the Appendix under Layer 1.

Layer 2: Models

From multilingual to domain-specific

India's AI applications succeed because they're built on models designed for Indian complexity. As India's AI ambitions grew, it became clear that relying on global LLMs, which serve primarily English and a few major languages, was fundamentally insufficient for inclusive public service delivery.

India invested in building sovereign, multilingual AI infrastructure. India's application layer is reinforced by models built for local complexity, multilingual access and sector-specific deployment. This is a critical AI builder principle: Generic models rarely deliver enterprise-grade reliability and safety in high-diversity environments without context engineering, orchestration and domain tuning. India's approach therefore combines multilingual capability, sovereign model development and sector-specific models to support scalable service delivery.

Bhashini demonstrates how multilingual AI can expand inclusive access at national scale, while BharatGen reflects a broader move towards sovereign, context-infused AI models built for Indian languages and public-sector use cases. Additional model examples, including BharatGPT, Sarvam AI and sector-specific models, are included in the Appendix under Layer 2.

Layer 3: Data

The fuel for inclusive AI

AI models are only as good as the data they're trained on. India faced a fundamental challenge: how to build high-quality datasets that represent India's diversity without compromising privacy or sovereignty?

Layer 3 is the data foundation that makes the rest of the stack durable. India's distinctive contribution is not simply scale of data generation, but the creation of consent-based, interoperable data flows that support AI innovation without abandoning accountability. This aligns closely with an AI builder model in which value depends on trusted context, governed access and enterprise-grade safeguards rather than unrestricted data accumulation.

Aadhaar-linked identity systems, ABDHM health records and AgriStack illustrate how digital public infrastructure creates reusable data foundations, while DEPA demonstrates how consent-based exchange can support innovation with accountability. Fuller data-layer examples are provided in the Appendix under Layer 3.

Layer 4: Compute infrastructure

Democratizing GPU access

A transformative innovation emerged in 2025–2026: India recognized that talent or models were not the bottleneck for AI development—the challenge was getting access to affordable compute.

The solution was audacious: democratize GPU access at scale.

Layer 4 addresses one of the most practical barriers to scaled AI adoption—affordable, reliable compute. India's response has been to widen access to compute infrastructure, reduce cost friction and make experimentation more viable for startups, researchers and public-interest deployments. In AI builder terms, this layer matters because optimized AI unit economics are essential to moving from pilot activity to sustained, enterprise-scale execution.

The Common Compute Portal illustrates India's push to democratize GPU access, while the Cerebra AI example shows how lower-cost compute can translate into deployable, domain-specific public-health capability. Detailed compute statistics and cost comparisons appear in the Appendix under Layer 4.

Layer 5: Energy infrastructure

Layer 5 recognizes that underpinning the entire AI infrastructure stack is reliable, sustainable energy planning. India's distributed data center strategy locates regional centers to reduce latency and enable local deployment, integrates solar and wind, deploys water-efficient AI-optimized cooling and ensures grid stability through targeted investments. Emerging economies often face intermittent power, high costs and environmental concerns. India is solving this by planning AI infrastructure around renewable availability, distributing compute regionally to reduce grid stress and optimizing power use through AI itself enabling sustainable, cost-effective deployment

Layer 6: Talent and skills

Building the workforce

An AI infrastructure without talent is inert. Layer 6 reflects the need to build a skilled workforce capable of engineering, operating and governing AI systems over time.

India has invested heavily in a world-class, cost-effective workforce: 6.6 lakh (660,000+) AI professionals trained, with more than 500,000 additional targeted by 2027 through upgraded ITIs and polytechnics with data labs, emphasizing accessibility for rural and underrepresented populations.

Complementary initiatives accelerate the pipeline:

- IndiaAI FutureSkills builds talent across government, academia and private sector
- The IndiaAI Innovation Centre (IAIC) supports applied research and commercialization
- AIKosh curates public datasets to reduce barriers for startups
- IndiaAI Startup Financing de-risks early-stage ventures

The cost advantage is substantial: ML engineers at USD 50K–70K versus USD 150K–200K globally; data scientists at USD 40K–60K versus USD 120K–180K; data annotation at Rs. 300–800 per unit versus USD 15–30.

The result: larger teams within the same budget, more comprehensive validation, faster iteration—globally competitive solutions at emerging-economy cost. India is becoming a hub for cost-effective, diverse AI talent.

Layer 7: Governance

Rules, institutions and trust

Before any AI can work at scale in the public sector, trust must be engineered into the system. India established this through a comprehensive governance framework released in January 2026.

Layer 7 establishes the rules, institutions and assurance mechanisms that make scaled AI trustworthy.

Additional details are provided in the Appendix under Layers 5, 6 and 7.

Taken together, Layers 5 to 7, complete the conditions for durable scale. They show that India is not treating AI as a narrow technology program, but as an end-to-end operating model for enterprise-grade adoption.

For energy, India is aligning compute growth with distributed infrastructure and renewable integration. For talent, programs such as India AI FutureSkills and related ecosystem investments show how capability is being built beyond a narrow elite base. For governance, the techno-legal framework, sandboxes and institutional mechanisms indicate a model focused on enterprise-grade reliability, safety and accountability.

Additional details are provided in the Appendix under Layers 5, 6 and 7.

Part two: The remaining 20%

India has built 80% of a globally replicable AI framework. Layers 1–7 are substantially operational. What remains is not more of the same, it is the frontier of combinatorial innovation where multiple AI services, data flows and decision systems work together to create compounded outcomes that single-point deployments cannot deliver.

Cross-sector integration: The multiplier layer

Today, India's AI applications work well within their sectors—healthcare AI diagnoses disease, agricultural AI optimizes irrigation, education AI tracks attendance and so on. Each is impactful. But they operate largely independently.

The frontier is cross-sector integration using digital public infrastructure, interoperable data exchange and context-aware orchestration to enable AI-native systems of intelligence that compound value across sectors.

Breaking silos

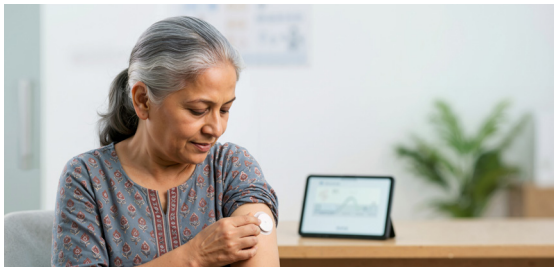

Healthcare-finance

With DEPA consent, an ABDHM diabetes diagnosis can flow securely into financial systems, prompting tailored recommendations for skill-training loans, pharmaceutical credit and chronic-disease microinsurance. Treatment adherence and financial inclusion advance together.

Timeline: 12–18 months; pilot projects underway; full integration targeted for 2027.

Education-employment

Student learning data links to verifiable DigiLocker credentials; real-time AI matches skill development to labor-market demand. Graduate unemployment can drop from 40% to single digits.

Timeline: 18–24 months; full implementation targeted for 2027.

Climate-agriculture-disaster

AI predicts crop failures from soil, weather and satellite data; on prediction, cash relief triggers via UPI, health systems prepare for outbreaks and education systems alert for dropout risk. Disaster response shifts from reactive to predictive.

Timeline: 18–24 months for data integration; 24–36 months for full automation.

Deeper sectoral integration

While cross-sector integration happens at the citizen level, sector-specific integration deepens value within each domain. AgriStack adds market intelligence price forecasting and supply-chain data, so farmers learn not just what to grow but when to sell and to which markets, with OCEN-backed credit. ABDHM records link to insurance and pharmaceutical data, enabling AI to predict treatment cost and insurance options, supporting preventive care and yielding personalized treatment pathways. Credit scoring moves beyond Aadhaar + UPI to incorporate agricultural yield, health status, education credentials and employment history, unlocking Rs. 5–10 trillion in new credit.

These deeper integrations are at advanced pilot stage; system-wide deployment is targeted for 2027.

Bhashini and BharatGen solved India's multilingual challenge. India's experience demonstrates that large-scale multilingual AI is achievable (350+ models, 1 billion tasks processed), economically viable (40%–87.5% cost advantage), and creates strategic independence. The global opportunity is vast: Sub-Saharan Africa's 2,000+ languages across 54 countries and 900M+ population; Southeast Asia's 700+ languages across 11 countries and 680M+ population; Latin America's indigenous-language preservation needs and 650M+ population; Central Asia's 200M+ population and developing digital infrastructure.

The pathway: Adapt Bhashini for non-Indian languages, replicate BharatGen infrastructure abroad, build multilingual AI consulting as a commercial export, and establish India as the global hub for minority-language AI. Partnerships are announced with African nations and Southeast Asian countries; the commercial licensing model is in development; the first European adoption (École des Ponts) demonstrates viability

India has integrated AI into defense operations. The Indian Army processed 26 years of battlefield data with AI, achieving over 94% precision in targeting analysis—demonstrating that AI can fuse decades of signals and sensor intelligence to deliver faster, more accurate strategic decisions. Civilian spillover includes predictive governance frameworks for disease outbreaks, natural disasters and economic disruptions; early-warning systems for development challenges; and strategic resource allocation based on AI-enabled forecasting. India is expanding this model to civilian governance domains, establishing AI as core to modern strategic planning.

Part three: Why India's model is globally replicable

India's approach is often assumed to be India-specific: unique demographics, governance structure, unique circumstances. But the principles are universal.

Principle 1: Digital public infrastructure first, AI second

Every successful economy—Estonia (X-Road), Singapore (digital ID), Denmark (digital credentials)—built foundational infrastructure before scaling AI.

India's contribution is not inventing this principle but demonstrating it works at massive scale (1.4 billion people) and at low cost (Rs. 50-100 per digital identity versus thousands in developed economies).

Implementation implication

DPI is foundational; cannot be rushed.

6–9 months minimum to establish, followed by sectoral AI deployment.

Principle 2: Sector-agnostic AI governance framework

Rather than building one AI system per sector, India built a single AI governance framework (NPAI) that sectors adapt.

This means developing countries don't need dozens of AI programs. They need one coherent framework that sectors implement locally. The time and cost savings are enormous.

Implementation implication

One governance framework scale across all sectors; sectoral customization happens at the application layer, not governance layer.

Principle 3: Cost advantage as structural benefit

India's cost advantage is not marginal but structural. Developing countries do not need to replicate India's full scale to benefit from it.

They can work with AI builders and accountable partners for business outcomes, or adopt operating models shaped by India's approach, to achieve faster speed to value and more optimized AI unit economics.

Principle 4: Diversity as design strength

India's linguistic, socio economic and geographic diversity initially seemed like a constraint. It became a design strength. Solutions proven for Indian complexity are inherently more robust for global deployment.

Implementation implication

Build multilingual, multi-context AI from the beginning; this ensures global applicability.

Principle 5: Government as platform builder, not service operator

India's NPAI approach treats government as the platform builder (DPI, governance frameworks, compute subsidies) rather than the sole operator.

This attracts private sector innovation. Startups and firms compete to deliver solutions on the platform. The result: faster innovation, lower government cost, broader ecosystem.

Implementation implication

Establish public-private partnership model from inception; government de-risks innovation through platform investment, not direct operation.

Conclusion

From 80% built to 20% frontier

India has built 80% of a globally replicable AI framework for development. The DPI foundation

is in place. Sectoral AI applications are deployed at scale serving 300-plus million people. Governance frameworks are established. Cost advantages are structural. The ecosystem is vibrant.

The remaining 20%—cross-sector integration, predictive governance, multilingual AI export, sovereign AI development and global knowledge transfer—is the frontier that will define the next decade of global development.

With the right global partner—one with sector expertise, engineering depth, operational discipline and a commitment to responsible AI—that frontier becomes addressable. Not over 10 years. Over three to five.

For Cognizant, this represents an opportunity to act as an AI builder that helps clients and public institutions bridge the gap between AI investment and measurable value—through AI-led productivity, industrializing AI and, where appropriate, agentifying the enterprise. The ambition is not to deploy tools in isolation, but to build context-infused systems, workflows and operating models that extend opportunity at scale.

For India, it's consolidating its position as a model that other nations want to learn from.

For the world, it's accelerating progress on development challenges that have persisted for generations. It's not charity. It's pragmatism. It's recognizing that the next wave of economic growth, innovation and stability will come from making AI work for the 80% of the world population living in emerging and developing nations.

That moment is now. The framework exists. The technology is proven. The execution opportunity is real.

The remaining 20% awaits implementation.

Quick reference: The seven-layer architecture at a glance

Layer	Current state (80% built)	Frontier (remaining 20%)	Implementation path
Applications	Healthcare, agriculture, education, finance, governance AI deployed to 300M+	Cross-sector integration; predictive governance; automated relief systems	Sector-specific AI + integration layer
Models	350 language models via Bhashini; domain-specific models; 9–15 languages live via BharatGen	Export multilingual AI globally; enhance orchestration; 22 languages complete	Sovereign model development; language expansion
Data	DPI-enabled; DEPA-consent frameworks; high-quality datasets being built	Cross-sector dataflows; causal inference; federated learning at scale	Enhanced consent management; data integration
Compute	34,381–38,000 GPUs operational; subsidized access	50,000+ GPUs by 2027; regional distribution; cost optimization	Infrastructure expansion; regional distribution
Energy	Power planning for AI infrastructure underway	Renewable integration; sustainable AI scaling	Grid-scale renewable integration
Talent	660,000+ AI professionals trained	500,000+ by 2027; global training hubs; local capacity	Skilling programs; ecosystem development
Governance	Frameworks published; sandboxes active; principles established	Enforceable governance; automated compliance; auditability	AIGG, AISI, TPEC coordination; global alignment

References

[1] Ministry of Electronics & Information Technology (MeitY). (2026). Techno-Legal AI Governance Framework. Government of India.

[2] Government of India. (2025). National Program on Artificial Intelligence (NPAI) Annual Progress

Report. IndiaAI Mission.

[3] AI4Bharat. (2025). Bhashini: Multilingual AI for Inclusion. Annual Impact Report.

[4] BharatGen Consortium. (2025). India's Sovereign Generative AI Ecosystem. Technical Architecture and Roadmap.

[5] Cerebra AI. (2025). TB Detection AI Model: Deployment Results. Technical Case Study, Uttar Pradesh.

[6] Government of India, Ministry of Health. (2025). ABDHM Implementation and AI Integration Results. Department of Health and Family Welfare.

[7] Government of Tamil Nadu. (2024). Sovereign AI Park Chennai: Sarvam AI Partnership Agreement. Government Order.

[8] Cognizant. (2025). Healthcare AI Solutions: Deployment Framework. Technical White Paper.

[9] International Organization for Standardization. (2025). AI Governance Frameworks: Global Best Practices. ISO/IEC Guidelines.

[10] Bertelsmann Stiftung. (2024). AI Policy Framework for Developing Economies. Global AI Initiative Report.

[11] Open Future Foundation. (2025). Public AI Architecture: Global Implementation Guidelines.

[12] OECD. (2024). AI Policy Index and Global AI Governance. OECD Digital Economy Papers.

Appendix

Layer 1: Applications

Healthcare: ABDHM established a foundational health ecosystem through ABHA IDs, interoperable records and consent-based access. AI diagnostic tools such as ATMAN AI and rural screening deployments extended specialist capability into underserved areas. Cerebra AI's TB detection deployment in Uttar Pradesh's Aspirational Districts achieved 94% accuracy, reduced diagnostic time from days to minutes and lowered cost per diagnosis from approximately Rs. 300 to Rs.50.

Agriculture: AgriStack linked farmers with land, soil, weather and credit data to support farm-specific advisory. CottonAce used pheromone trap image analysis to provide early pest warnings in key cotton states and was associated with yield improvement and lower pesticide use. Predictive irrigation systems improved water efficiency, while Telangana's AI mapping used satellite imagery to monitor farmland and forecast harvest outcomes.

Education: Shiksha Setu in Assam used AI-enabled attendance monitoring across 44,000 schools to address proxy attendance, improve punctuality and identify ghost students. A future personalized learning layer aims to combine real-time learning data with verifiable credentials and employment matching.

Finance: India's AI-enabled credit approach combined Aadhaar, UPI, Account Aggregator and alternative data to improve credit access for informal workers, farmers and microenterprises through OCEN-linked lending models.

Urban governance and transfers: AI-supported civic monitoring, UMANG and DigiYatra show how service access and operational efficiency can improve through AI-enabled public-service workflows. AI has also been used to reduce leakages in Direct Benefit Transfers, with reported savings from eliminating duplicate or fake beneficiaries.

Layer 2: Models

Bhashini: Built to expand inclusive AI access across Indian languages, Bhashini supports multilingual and voice-first use cases at scale and lowers the cost of language technology.

BharatGen: A government-backed consortium developing text models, speech models, datasets and infrastructure across Indian languages, with phased language expansion and sectoral applications such as e-VikrAI and Krishi Saathi.

BharatGPT and Sarvam AI: These examples illustrate the wider sovereign-model movement, including multilingual interaction, public-service accessibility and local model control.

Domain-specific models: Examples include ATMAN AI for diagnostics, agriculture-specific models for pest and irrigation decisions, education-specific models for learning and dropout prediction and finance-specific models for credit and fraud use cases.

Layer 3: Data

India's data layer is anchored in digital public infrastructure such as Aadhaar, ABDHM and AgriStack, each creating structured, interoperable data foundations for service delivery and model improvement.

DEPA provides a consent-based model for data portability and use, helping align innovation with privacy, purpose limitation and auditability. Model cards and dataset datasheets further support transparency and more accountable AI deployment.

Layer 4: Compute infrastructure

The Common Compute Portal was designed to widen access to GPUs through subsidized or lower-cost allocation, helping reduce a major barrier to AI development. The model combines public-private access and merit-based allocation.

Cerebra AI is a representative example of how improved access to compute can support practical, deployable solutions in public health. The wider compute layer also includes regional data centers, efficiency improvements and infrastructure expansion targets.

Layer 5: Energy

The energy layer focuses on distributed data-center strategy, renewable integration, cooling efficiency and grid-level reliability. These are essential to sustainable AI scaling, particularly in emerging markets where power stability and cost remain binding constraints.

Layer 6: Talent and skills

India's talent layer includes large-scale skilling through ITIs, polytechnics and national initiatives, alongside programs such as IndiaAI FutureSkills, the IndiaAI Innovation Centre, AIKosh and startup financing. The aim is to build a broad, cost-effective and inclusive workforce capable of supporting AI development, deployment and governance over time.

Layer 7: Governance

India's governance layer includes the techno-legal AI governance framework, sandboxes, watermarking and safety measures, independent testing bodies, the AI Safety Institute, AIGG and TPEC. Together, these mechanisms show how AI can be scaled through enabling governance that balances innovation, safety, transparency and accountability.

Authors

Amir Banifatemi

Amir Banifatemi is Cognizant's Chief Responsible AI Officer, leading global responsible AI strategy and governance. He is a renowned thought leader and speaker at various internal forums. He previously held key leadership roles, co founded AI Commons and helped establish the AI for Good Global Summit. His work spans AI policy, standards, innovation and enterprise deployment across international institutions. He is based in the US.

Ishan Pandey

Ishan Pandey is a Senior Manager on Cognizant's Corporate Strategy team and co-author of "India's public AI framework". He focuses on corporate strategy, market intelligence, M&A advisory and AI-led growth opportunities.

Apeksha Sagar

Apeksha Sagar is a Manager on Cognizant's Corporate Strategy team and co-author of this article. She has more than a decade of experience across sales, corporate strategy, solution positioning, portfolio expansion, due diligence and client opportunity shaping.

World Headquarters
300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233

European Headquarters
280 Bishopsgate
London
EC2M 4AG
England
Tel: +44 (01) 020 7297 7600

India Corporate Office
Siruseri-Software Technology Park of India (STPI)
SDB Block—Ground Floor North Wing
Plot No H4, SIPCOT IT Park
Chengalpattu District
Chennai 603103, Tamil Nadu
Tel: 1800 208 6999

APAC Headquarters
1 Fusionopolis Link,
Level 5 NEXUS@One-North,
North Tower Singapore 138542
Phone: + 65 6812 4000

4966050