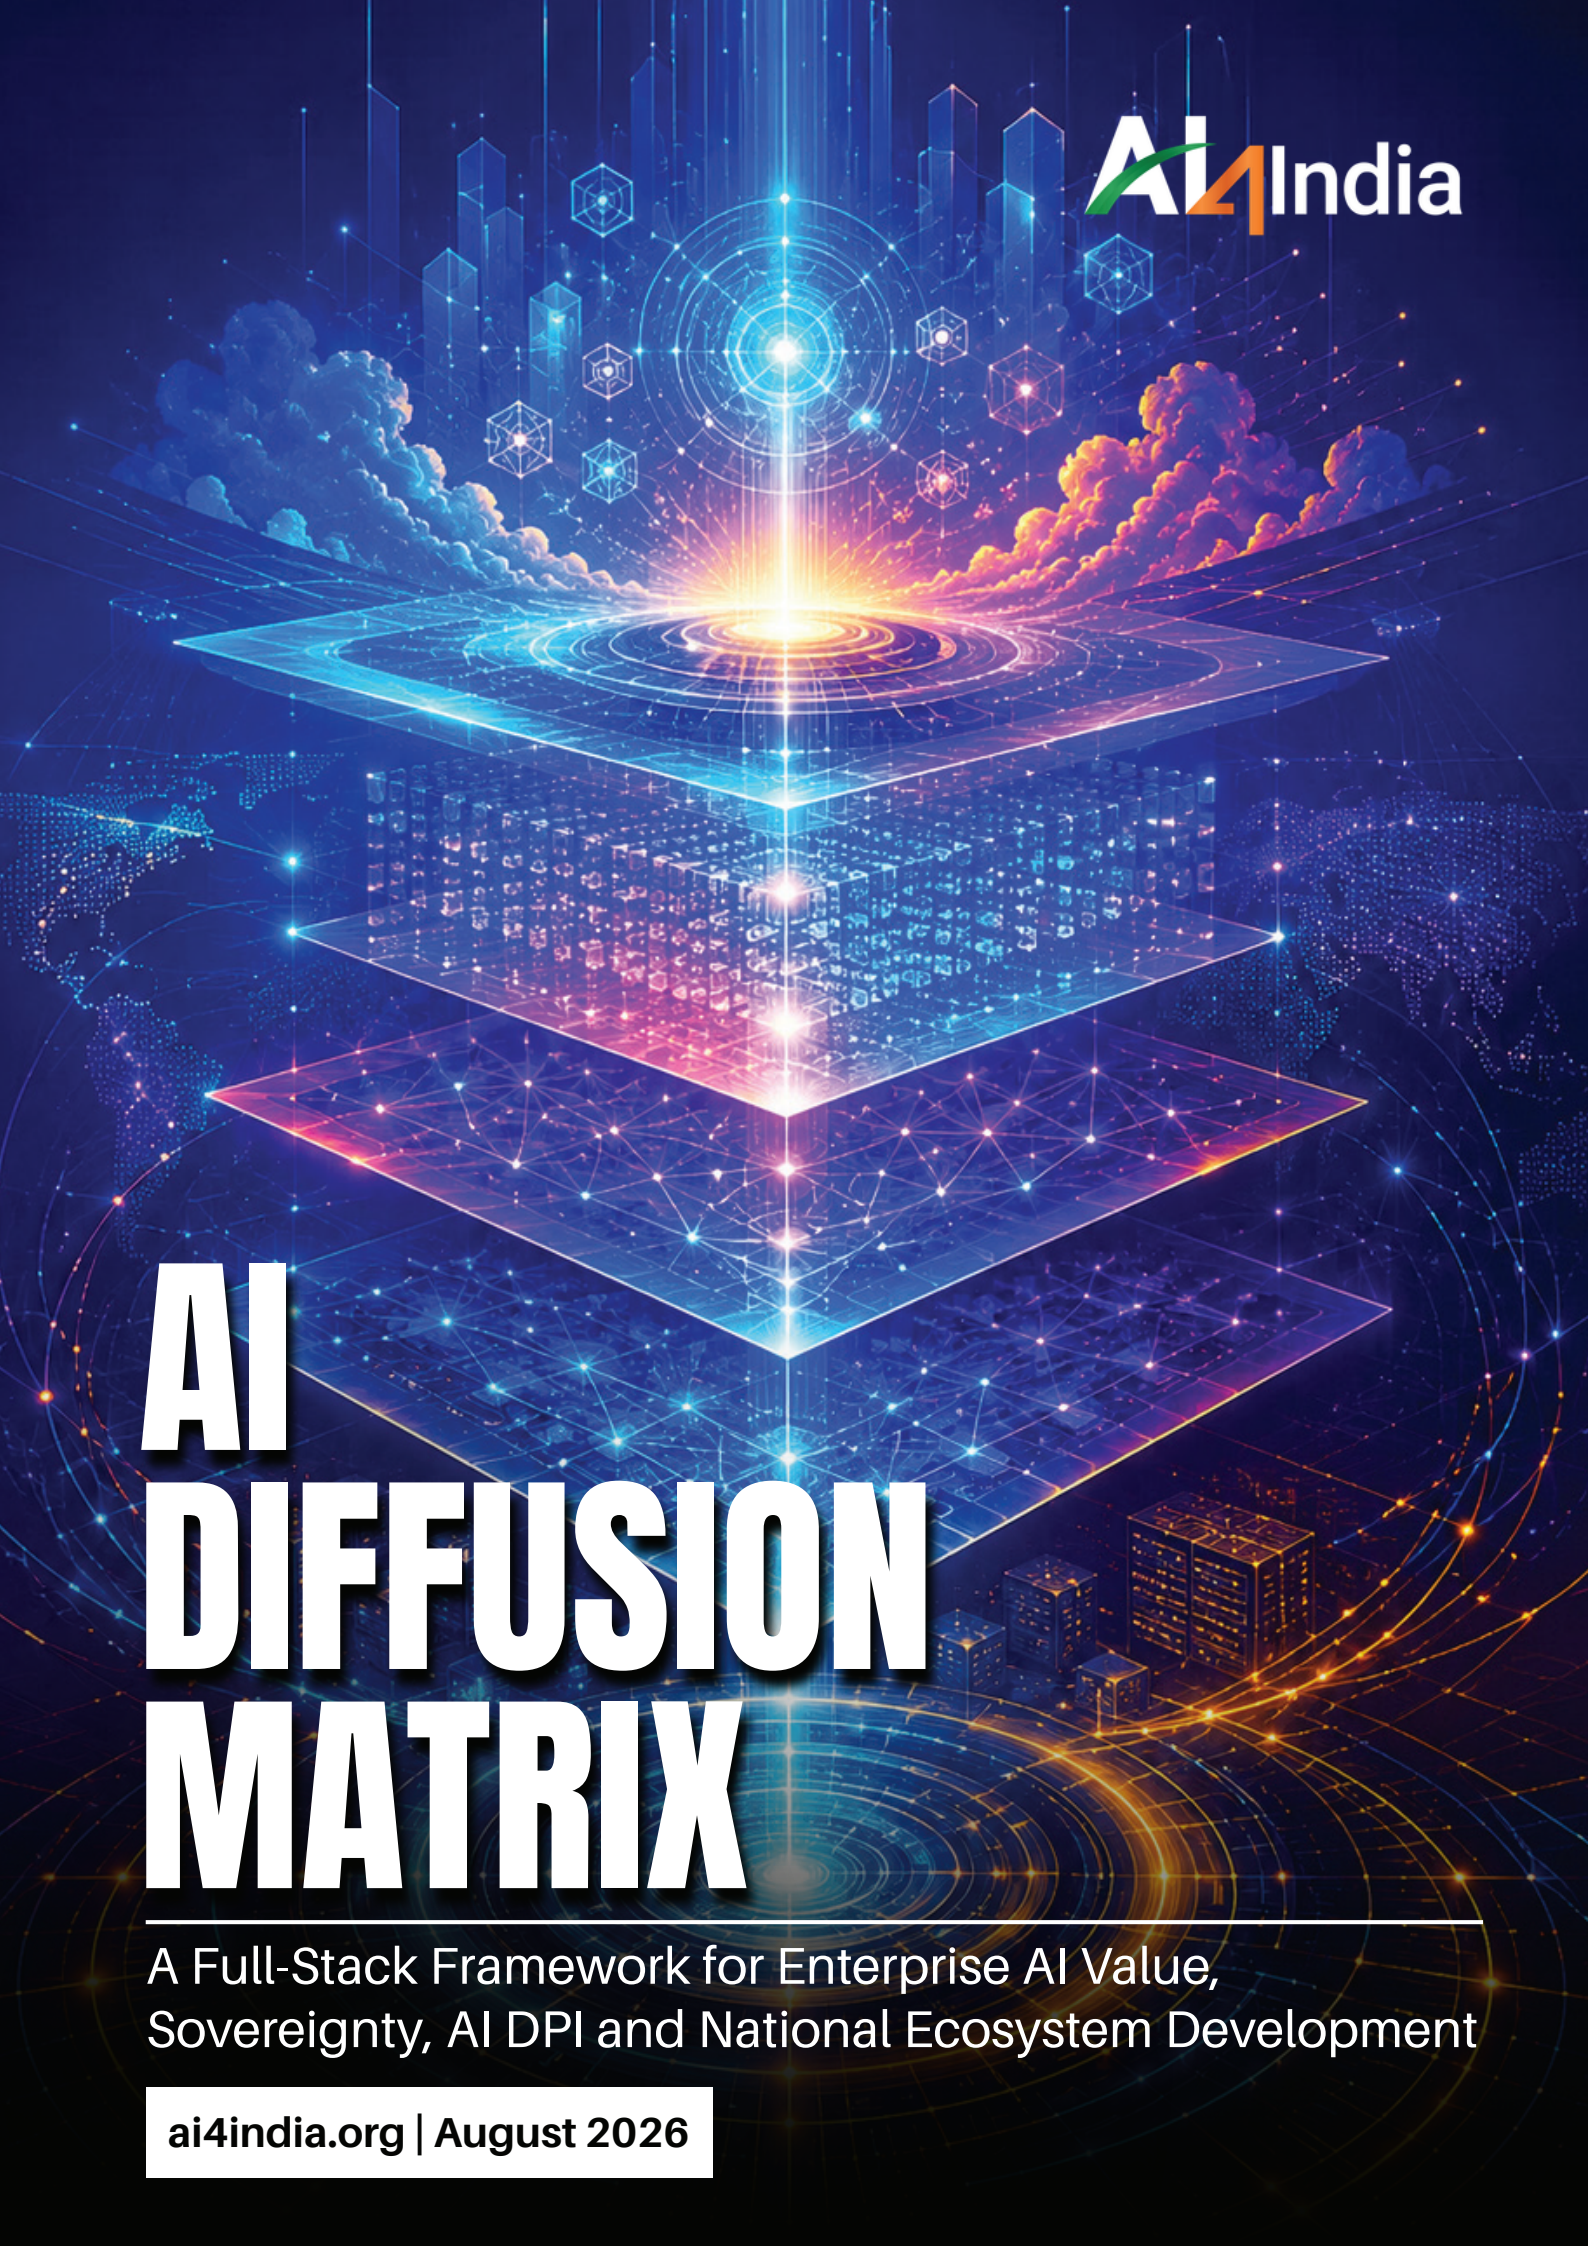

AI DIFFUSION MATRIX

A Full-Stack Framework for Enterprise AI Value, Sovereignty, AI DPI and National Ecosystem Development

ai4india.org | August 2026

Scan the QR code to access the *AI Diffusion Matrix* Assessment Tool.

aidm.ai4india.org

The AI Diffusion Matrix:

A Full-Stack Framework for Enterprise AI Value, Sovereignty, AI DPI and National Ecosystem Development

AI DIFFUSION MATRIX

The AI Diffusion Matrix: A Strategic Blueprint for Sovereign AI Adoption

The AI Diffusion Maturity Model

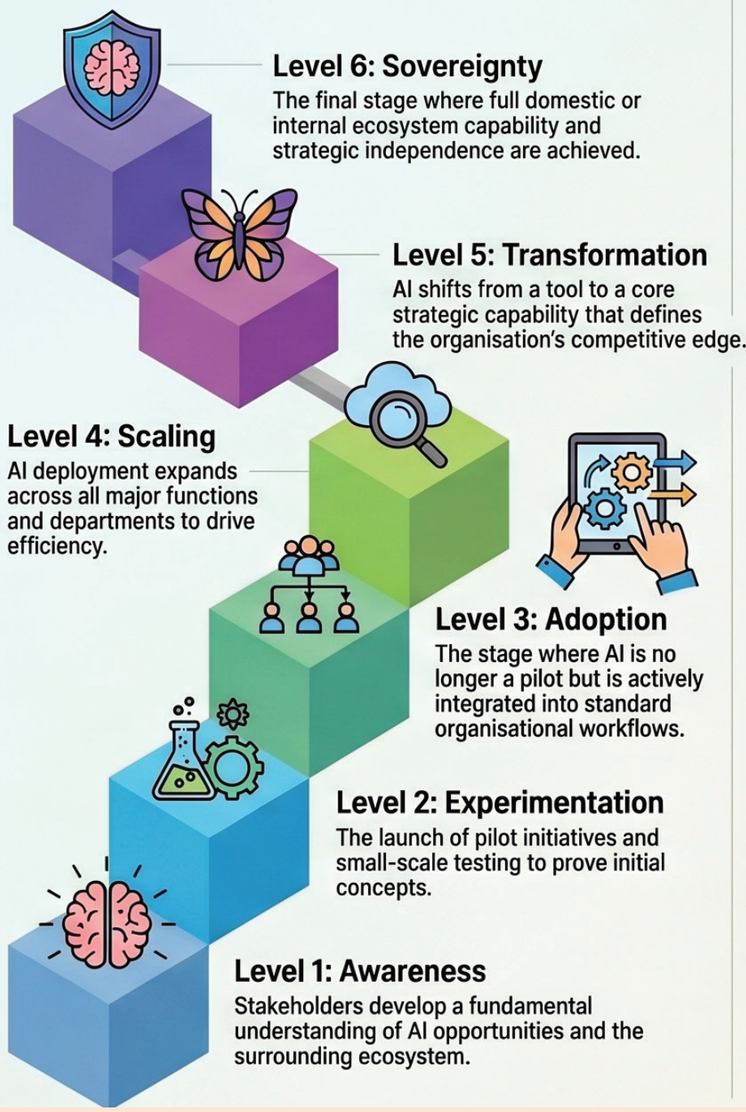

Level 6: Sovereignty
The final stage where full domestic or internal ecosystem capability and strategic independence are achieved.

Level 5: Transformation
AI shifts from a tool to a core strategic capability that defines the organisation's competitive edge.

Level 4: Scaling
AI deployment expands across all major functions and departments to drive efficiency.

Level 3: Adoption
The stage where AI is no longer a pilot but is actively integrated into standard organisational workflows.

Level 2: Experimentation
The launch of pilot initiatives and small-scale testing to prove initial concepts.

Level 1: Awareness
Stakeholders develop a fundamental understanding of AI opportunities and the surrounding ecosystem.

AI Architecture (10 layer stack)

Layer 9: Applications
The visible user interface, such as enterprise chatbots, healthcare assistants, and manufacturing tools.

Layer 8: AI Agents
Autonomous systems capable of planning, decision-making, and multi-step tool usage.

Layer 7: AI Services
Reusable capabilities like speech recognition and recommendation engines delivered via APIs.

Layer 6: Inference Infrastructure
The engine powering real-time operations, prioritising low latency, scalability, and cost-efficiency.

Layer 5: Models
The intelligence layer (LLMs, Multimodal) evaluated for 'Model Sovereignty' and support for local languages.

Layer 4: Datasets
The informational foundation, requiring high quality, diversity, and robust governance frameworks.

Layer 3: Cloud Infrastructure
Managed platforms providing storage, networking, and the compute power required for scaling AI.

Layer 2: Data Centres
Physical infrastructure assessed by energy availability, sustainability, and geographic resilience.

Layer 1: Compute Hardware
Focuses on GPUs, AI accelerators, and edge devices necessary for processing complex AI workloads.

Layer 0: Semiconductors
The absolute foundation of the stack, focusing on domestic manufacturing capabilities and supply chain resilience.

Five Strategic Questions for Diffusion

Awareness: Do we understand the dependencies?
Stakeholders must look past the interface to understand the hardware and data layers supporting the AI.

Sustainability: Is it viable long-term?
Evaluation of economic ROI, environmental energy consumption, and social workforce impact.

Ecosystem Engagement: Are we collaborating?
Measuring active participation in open-source communities, standards, and research partnerships.

Indigenisation: Is the capability domestic?
Assessment of how much strategic capability and intellectual property is owned and controlled locally.

Actionability: What steps come next?
Identification of the specific interventions and investments required to accelerate the roadmap.

Strategic Intent of AI Diffusion Matrix

The intent is not to build everything domestically or internally. The intent is to know what matters, measure value rigorously, govern risks, preserve choice and deepen control over the layers that determine resilience and competitiveness. If you value these, read on.

Authors

Shashi Shekhar Vempati (Co-Founder, AI4India)

Alok Agrawal (Co-Founder, AI4India)

Reviewer

Adarsh S Lathika (Founder, Anatomy of Work)

The **AI Diffusion Matrix** has been developed by **AI4India** in consultation with **NITI Aayog**.

प्रो. अभय करंदीकर
Prof. Abhay Karandikar
सदस्य
Member

सत्यमेव जयते

भारत सरकार
नीति आयोग
Government of India
NATIONAL INSTITUTION FOR TRANSFORMING INDIA
NITI Aayog

MESSAGE

AI has the potential to accelerate India's rise as an economic power, strengthen its digital sovereignty, and enhance its global leadership in the coming years. For a nation of our scale and ambition, true technological sovereignty cannot be measured merely by front-end adoption. It will depend on deep, systemic AI diffusion, integrated directly into the core of our economy.

To realize our nation's full potential, India must bridge the gap between technological possibility and nationwide impact with intelligence deeply embedded across every critical domain: Governance, Manufacturing, Healthcare, Education, Agriculture, Infrastructure, Science & Technology, IT and Financial services.

The **AI Diffusion Matrix**, developed by **AI4India** in consultation with **NITI Aayog**, serves as a strategic blueprint for this mission. By evaluating the full technological depth through a robust 10-layer stack, the Matrix enables stakeholders to map dependencies, diagnose vulnerabilities, and engineer resilient, indigenous capabilities.

For the nation, this strategic AI Diffusion secures a technological self-reliance, drives rapid economic growth, and unlocks population-scale impact across public and private sectors alike. For the Indian industry, the AI Diffusion Matrix provides a practical, structured assessment tool to accelerate innovation, build world-class enterprise capabilities, and establish sustainable global competitiveness.

Through this framework, India moves beyond consuming global technologies to actively shaping them. I invite enterprise leaders, policymakers, and innovators across the nation to leverage this report to accelerate our transformation into a sovereign architect of the global AI landscape.

(Abhay Karandikar)

NITI Bhawan, Sansad Marg, New Delhi - 110001
Phone : +91-11-23096566, 23096567 | E-mail : member.ak@niti.gov.in

Alok Agrawal
Co-Founder, AI4India

"Organisations are no longer asking whether they should adopt Artificial Intelligence, but how they can adopt it meaningfully, securely, and at scale. This is not an easy task. Pace of AI development is unprecedented. There is tremendous confusion with conflicting claims. Most organisations view AI purely from the lens of user interface and applications. However long-term value will be created only when they have visibility and control over the entire AI stack, down to the foundational layers.

The AI Diffusion Matrix developed by AI4India is a practical framework to help organisations identify high-impact opportunities, determine the right value, assess readiness, prioritise investments, and build the capabilities required for successful AI adoption. Our objective is to enable enterprises to move beyond pilots and isolated use cases towards structured, responsible, and measurable AI transformation."

Adarsh S Lathika
Founder, Anatomy of Work

"The next phase of AI adoption will separate organisations that merely consume AI from those that understand and govern it. Visible applications are only the surface. Durable advantage will depend on whether enterprises can map their dependencies, redesign workflows, build human capability, measure value, and retain decision authority across critical layers of the AI stack. The AI Diffusion Matrix provides a much-needed discipline: it shifts the conversation from tool adoption to responsible, measurable, and sovereign AI capability."

FOREWORD

Prof. B. Ravindran

Professor & Head, Department of Data Science and AI, Wadhwani School of Data Science & AI (WSAI) & Centre for Responsible AI, IIT Madras

The AI Diffusion Matrix makes an argument I find compelling: that a nation's or an enterprise's real AI capability is defined not by how advanced its models are, but by how deeply and reliably AI diffuses into everyday production, governance and life. I would add one qualification to this, drawn from our work at the Centre for Responsible AI (CeRAI): **diffusion without responsibility does not scale, it stalls.**

Every layer of the stack described in this report, from datasets and models to inference, agents and applications, carries decisions about who benefits from an AI system, who bears its risks, and who is left out of it. When those decisions are made invisibly, without proportionate oversight, explainability or recourse, trust breaks down long before the technology does. And it is trust, not compute or capital, that ultimately decides whether an AI system survives contact with a hospital ward, a credit desk, or a citizen-service counter. This is precisely the "Executive Trap" this report warns against, extended one layer further: leaders can lose sight not only of their technical dependencies, but of who is affected by the systems built on top of them.

This is not a theoretical concern. In many conversations with industry leaders, from banking, retail or technology, I hear that the biggest obstacle to scaling AI is rarely technical capability. Instead it is the absence of governance calibrated to risk: explainability for a credit decision, audit trails for a dynamic pricing engine, human oversight for a coding agent. This is precisely the gap the Reserve Bank of India's FREE-AI Committee, of which I was a member, set out to address: shared infrastructure and sandboxes on one side, board-level accountability and incident reporting on the other, applied proportionately rather than uniformly. This is a discipline the India AI Governance Guidelines I chaired generalise nationally, grading liability by function, risk and due diligence across the AI value chain. Proportionate, clearly assigned accountability should run through every enterprise and government AI programme that hopes to move past pilots.

I would therefore urge readers to treat Responsible AI not as a downstream compliance layer on this Matrix, but as a dimension running through all five of its strategic questions. Awareness of dependencies should include awareness of who is affected. Sustainability should include the sustainability of public trust. Ecosystem engagement should include the people the ecosystem serves. Indigenisation of capability should go hand in hand with indigenisation of accountability. Get this right, and Responsible AI becomes the reason people, institutions and a nation are willing to depend on AI at scale.

Shashi Shekhar Vempati
Co-Founder, AI4India

From AI Innovation to AI Diffusion - Reframing India's AI Strategy

Artificial Intelligence has rapidly become the defining general-purpose technology of the twenty-first century. Governments across the world are investing unprecedented resources in frontier AI research, semiconductor ecosystems, compute infrastructure, talent development and startup ecosystems. Success is increasingly measured through the number of foundation models developed, patents filed, research publications produced, benchmark performances achieved and capital invested.

These indicators undoubtedly matter. They reveal a nation's capacity to create new technological knowledge. Yet they tell only part of the story. As political scientist Jeffrey Ding argues in his influential work on scientific and technological power, policy discourse has become overwhelmingly innovation-centric and nations are routinely judged by their ability to generate frontier innovations. Little attention is paid to an equally important capability – diffusion. The ability of a nation to diffuse those innovations throughout their economy according to Ding matters more in the long run. Ding argues that a country's long-term technological and economic strength depends not only on invention, but on its capacity to ensure that new technologies are widely adopted, continuously adapted and deeply embedded across productive sectors. When innovation significantly outpaces diffusion, countries experience what he describes as a "diffusion deficit" which according to him is a structural weakness that ultimately constrains economic transformation despite impressive scientific achievements.

This insight provides an important starting point for thinking about Artificial Intelligence in India. AI presents an even broader challenge than the diffusion of previous industrial technologies. Artificial Intelligence is not merely another technology sector. It is a foundational cognitive capability capable of augmenting decision-making, communication, reasoning, creativity and automation across virtually every economic activity. Unlike earlier technologies that transformed particular industries, AI has the potential to transform every industry simultaneously.

Consequently, the strategic question before India is no longer simply whether we can build world class AI but more importantly it is about whether India can become an AI-diffused economy. The focus of this AI4India White Paper is thus on how national AI strategy should be conceived from the perspective of population scale, economy wide diffusion.

History repeatedly demonstrates that technological revolutions are not won solely by inventors. They are won by societies capable of absorbing innovation into the routines of production, governance, education and everyday life. Electricity became economically transformative not when generators were invented, but when every factory reorganised around electric power. The internet reshaped economies not because networking protocols were invented, but because businesses, governments and citizens reorganised themselves around digital connectivity. Artificial Intelligence will follow the same historical pattern. Its greatest value will not arise from the existence of increasingly capable foundation models alone. Its value will emerge when millions of organisations redesign their workflows around AI, when governments routinely deploy AI to improve public services, when professionals across disciplines use AI as a cognitive assistant, when small businesses access capabilities previously available only to large corporations, and when citizens interact naturally with AI in their own languages.

Jeffrey Ding defines diffusion as the spread and adoption of innovations throughout a country's productive system, distinguishing it from innovation itself. His work reminds us that invention and diffusion are distinct national capabilities, often governed by different institutions, incentives and forms of capital. Countries may excel at scientific discovery while failing to achieve widespread technological adoption. Others may not pioneer frontier technologies but nevertheless derive enormous economic gains through their ability to absorb, adapt and deploy them at scale.

AI diffusion should not be understood simply as the deployment of AI software or the distribution of AI tools. Rather, AI diffusion is the process through which AI becomes embedded within the economic, institutional and social fabric of a nation. Diffusion occurs when AI ceases to be viewed as a specialised technology and becomes an ordinary capability. Seen through this lens, AI diffusion is not a single event but a cumulative national process. Only when complementary capabilities evolve together can AI become a pervasive engine of productivity rather than an isolated technological achievement.

In this sense, diffusion reflects the ability of an economy to continuously absorb technological advances and convert them into sustained improvements in productivity, competitiveness and public value. India occupies a distinctive position in this global transition. Unlike many countries beginning their AI journey from fragmented digital ecosystems, India has already demonstrated an extraordinary capacity to diffuse digital technologies at population scale. The success of Aadhaar, UPI, DigiLocker, CoWIN, FASTag, Account Aggregator and ONDC illustrates that India has developed institutional mechanisms capable of rapidly scaling

trusted digital infrastructure across hundreds of millions of users. These achievements represent more than successful technology programmes. They demonstrate an institutional capability for diffusion. India's Digital Public Infrastructure has shown that interoperability, open standards, public-private collaboration and inclusive design can dramatically reduce barriers to adoption while enabling innovation at scale.

The next stage of India's digital journey is therefore not merely to develop frontier AI models. It is to build the institutional architecture through which AI itself can diffuse across every layer of the economy. This White Paper therefore proposes a shift in perspective. Rather than evaluating India's AI ambitions solely through measures of innovation capacity, research expenditure, publications, patents, foundation models or startup valuations it argues that equal attention must be devoted to India's capacity for AI diffusion. It also argues that AI diffusion should be recognised as a distinct dimension of national capability.

The purpose of this white paper is to help Organisations (Public and Private) assess their capacity to convert frontier advances in Artificial Intelligence into broad-based productivity, institutional effectiveness, economic competitiveness and social inclusion by ensuring that AI reaches everyone, benefits everyone and transforms everything.

Content

1. Executive Summary

Artificial intelligence adoption is often treated as a software decision: select an application, buy a model/copilot license, integrate an API, launch a pilot and report productivity gains. This approach can be short-sighted and does not do justice to Enterprise AI adoption.

The AI Diffusion Matrix developed by AI4India challenges this view. Meaningful and resilient AI adoption must be evaluated across the full stack, from visible applications down to data, inference, cloud, data centres, compute hardware and semiconductors.

The central risk is the **Executive Trap**: organisations see user interface but miss deeper dependencies that determine cost, reliability, sovereignty and long-term value. The matrix converts that strategic warning into an operational assessment framework. It helps leaders ask what they know, where the gaps are, which capabilities must be domestic or internally controlled, and what actions must follow.

For an enterprise, AI value is not captured by the number of pilots or the number of employees using a tool. Value is created when AI changes a workflow thus improving a measurable business outcome. This paper expands the value realisation question into a measurable model that covers productivity, revenue, cost, risk, speed, customer experience, employee capability, innovation and sovereignty.

For government and ecosystem builders, AI diffusion is a national capability challenge. India has already taken major steps through the IndiaAI Mission, AI compute capacity, AIKosh and responsible AI policy foundations. These initiatives create the beginning of an AI DPI layer that can democratise access to compute, data and AI services.

The proposed AI Diffusion Matrix organises the AI stack into three tiers: The Interface Tier, the Intelligence Engine and the Physical Foundation. The Interface Tier creates visible adoption that creates value. The Intelligence Engine determines the quality, adaptability, safety and cost of AI. The Physical Foundation determines capacity, control, sustainability and resilience. A mature enterprise or country must know its position across all three.

Key Takeaway

AI adoption without dependency awareness creates a fragile and risk prone environment. AI adoption with stack-level assessment, value measurement and sovereignty planning creates durable strategic capability.

2. Why The AI Diffusion Matrix

The AI Diffusion Matrix is needed because unlike the previous era of Digital Technology, AI is not just an application or tool. It is becoming an operating layer for knowledge work, product development, customer interaction, risk management, public service delivery and industrial automation.

- **Pilot proliferation:** Many proofs of concept are launched, but few translate into adopted workflows. As of September 2025, an MIT study found that only 5% of AI pilots successfully scale into production.
- **Lack of measurable value:** Benefits are reported as anecdotes or usage metrics instead of financial and operational outcomes.
- **Dependency without visibility:** Leaders do not know which models, data, clouds, hardware or jurisdictions support their AI systems.
- **Strategic fragility:** Cost, access, compliance or geopolitical shocks can undermine the enterprise AI operating model.

This problem is visible globally. IBM reported in 2024 that 42% of surveyed enterprise-scale companies with over 1,000 employees had actively deployed AI, while 40% remained in the exploration or experimentation phase. McKinsey's State of AI research stresses that high performers are more likely to have practices across strategy, talent, operating model, technology, data, adoption and scaling. The implication is clear: AI value is not a tool purchase; it is a management system.

Matrix Logic

Every row in the matrix asks four practical questions: What is the current status? What are the gaps? What is the indigenisation quotient? What is the call to action? This converts AI strategy from a technology wish list into a governance and investment agenda.

Strategic Question	Why it Matters
Where is AI being used?	Identifies high-value workflows and prevents scattered, low-impact pilots.
What dependencies support it?	Maps available datasets, model(s), inference, cloud, compute, hardware and chip-level exposure.
What value is expected?	Connects AI investment to outcome metrics, financial value and capability building.
What must be controlled?	Defines the indigenisation or internal control threshold for each layer.
What interventions are needed?	Turns assessment into roadmap, investment, partnerships and governance.

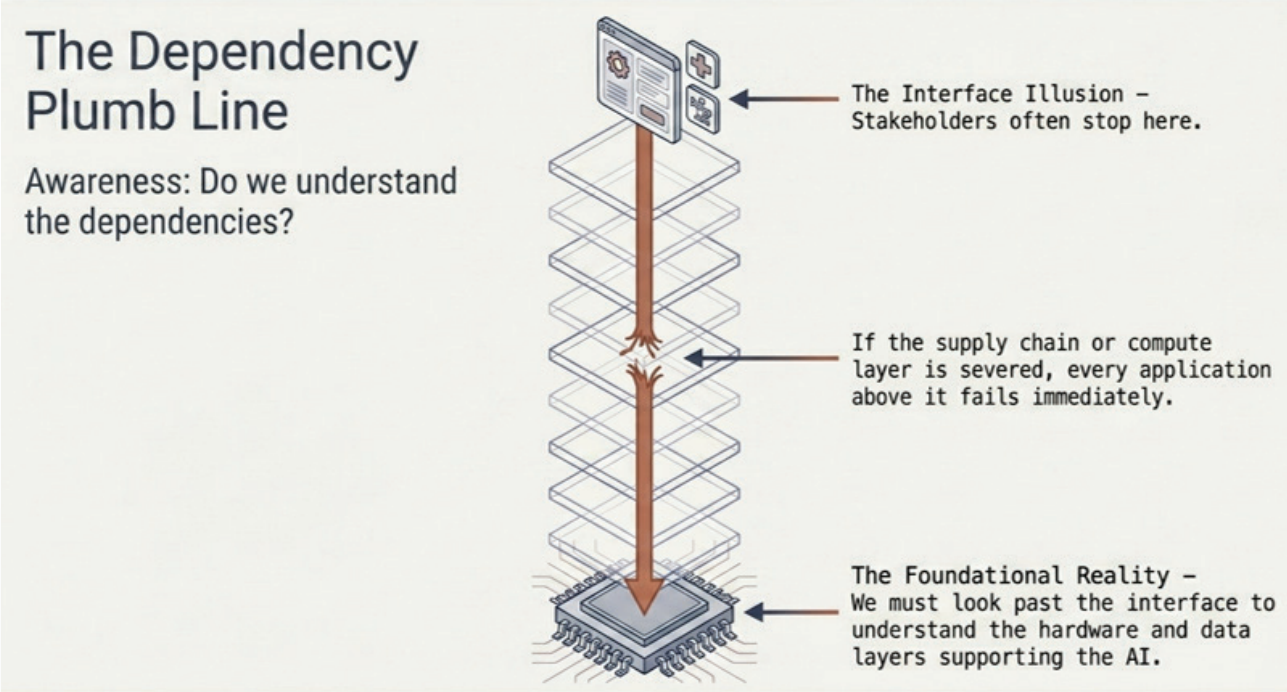

Figure 1. The dependency plumb line: the application layer depends on the full stack beneath it.

2.1 The Enterprise Trap

Executives often see AI through a productivity lens: documents get drafted faster, code suggestions appear faster, support tickets get summarised and employees experiment with chat interfaces. Those are useful starting points, but they do not define enterprise transformation. True diffusion occurs when AI changes the way an enterprise behaves and performs, not merely the individual task.

For example, GitHub Copilot research found that, in a controlled experiment where developers were asked to implement an HTTP server in JavaScript, participants completed the task 55.8% faster with the AI pair programmer than without it. That is a strong task-level signal. However, enterprise value still requires measurement of code quality, review time, defects, security, delivery frequency, developer satisfaction and product outcomes. The matrix therefore asks not only whether AI exists, but whether the organisation is ready, governed and measuring value.

2.2 The National Sovereignty Trap

At national level, the same trap appears when AI capability is measured by the number of applications rather than control over datasets, models, compute, cloud, data centres, hardware supply chains and semiconductor capability. India's AI strategy has long recognised the importance of AI for inclusive growth through. The next frontier is to connect adoption with sovereign access to the deeper AI stack.

2.3 Navigating AI Sovereignty

Navigating the aspects of national and enterprise AI sovereignty requires appreciation of the journey from foundational awareness to strategic independence, understanding the 10 layer interdependent architecture and looking at the compass by answering the 5 strategic questions to govern the adoption process.

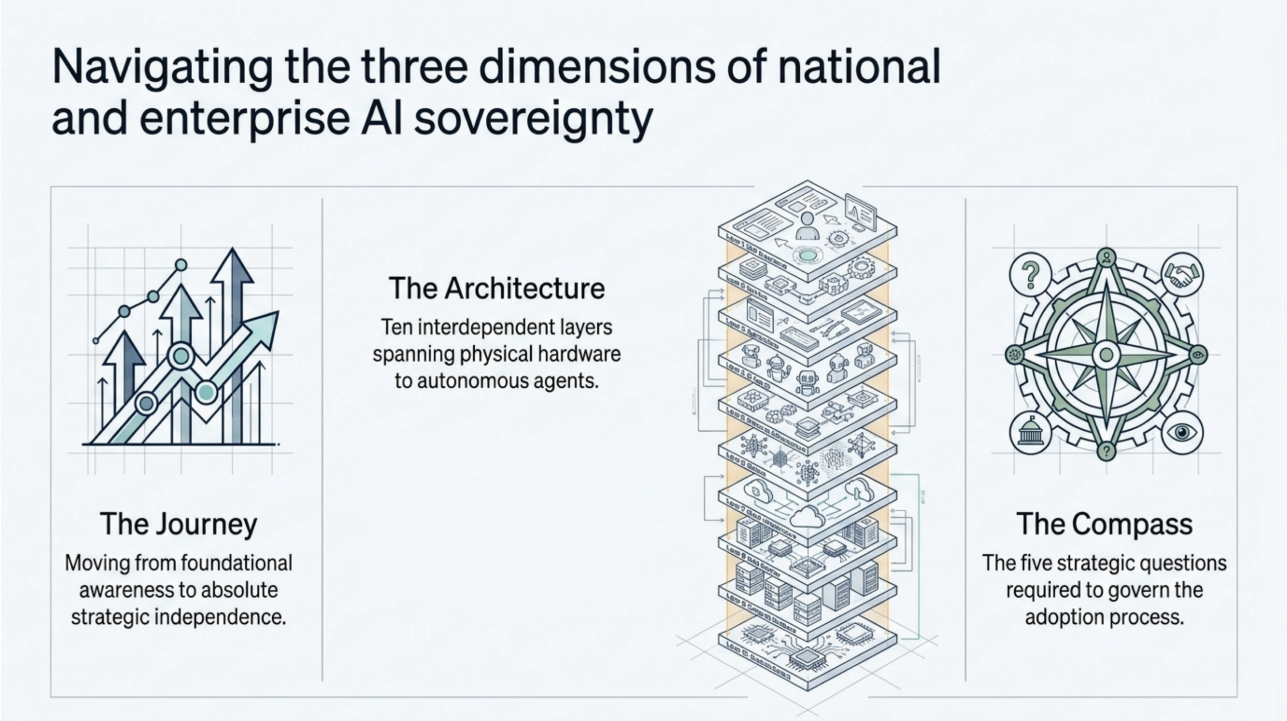

Figure 2. Navigating the 3 dimensions of national and enterprise sovereignty.

3. The Strategic Endgame: AI Sovereignty and Enterprise Control

As AI becomes foundational to economic competitiveness and national security, AI sovereignty must become a defining objective of India's national AI strategy. For a country, sovereignty means that the national ecosystem can access or build critical AI infrastructure, influence standards, protect data, support domestic innovation and maintain continuity under stress. For an enterprise, sovereignty means a clear ability to choose, switch, audit, govern and optimise AI across the full stack.

Dependent Posture

An organisation or nation relies heavily on external AI providers for models, infrastructure, data or compute, with limited visibility or control over critical dependencies. This can reduce flexibility, increase vendor lock-in and expose operations to external disruptions.

Sovereign-Controlled Posture

An organisation or nation retains strategic control over key AI capabilities while remaining connected to the global AI ecosystem. It can govern its AI systems, manage dependencies and switch providers or technologies without compromising resilience or continuity.

Dimension	Dependent Posture	Sovereign / Controlled Posture
Awareness	Views AI as a software UI or app subscription.	Maps dependencies down to data, models, inference, cloud, compute and chips.
Sustainability	Chases short-term pilot ROI and ignores energy, skills and long-term cost.	Evaluates total cost, energy, workforce, resilience and operating model sustainability.
Ecosystem	Relies on closed vendor ecosystems.	Participates in standards, open-source, research, sector communities and public-private platforms.
Indigenisation	IP, data and compute remain offshore or opaque.	Strategic capability and IP are owned, controlled or at least contractually assured.
Actionability	Buys tools without intervention plan.	Directs capital and management action into capability gaps across the stack.

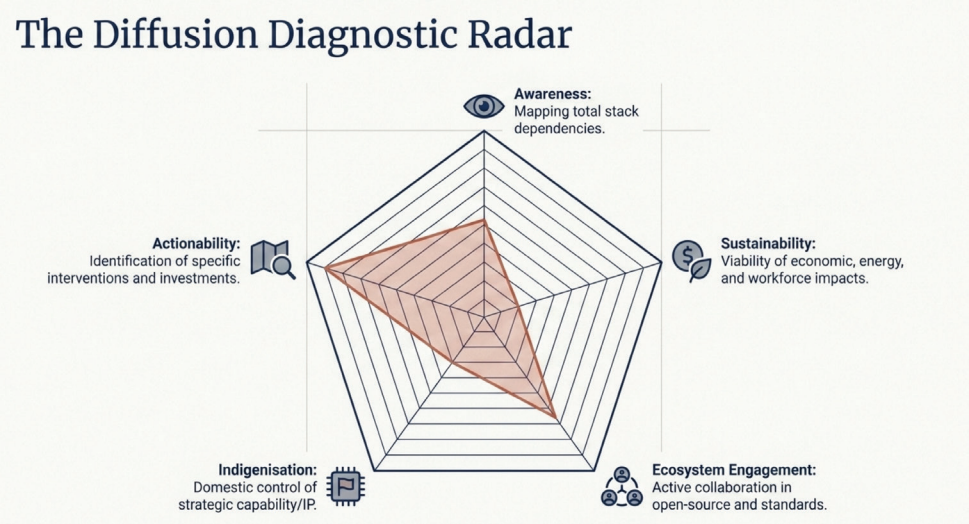

Figure 3. Diffusion diagnostic radar: an illustrative way to visualise strengths and gaps across the five strategic questions.

3.1 Sovereignty Does Not Mean Autarky

Sovereignty should not be confused with complete self-sufficiency in every layer. The practical objective is decision authority. An enterprise may choose a global cloud provider, a frontier model, an Indian data centre or an open-source model. The sovereign question is whether the organisation understands the dependency, has a fallback path, protects sensitive data, can evaluate performance, and can renegotiate or switch when needed.

In India's context, this distinction is important. The ecosystem can combine global models, domestic models, public datasets, regulated data exchange, local inference providers, India-hosted cloud capacity and sector-specific AI DPI layers. The goal is to deepen strategic control where it matters most, while remaining open to global innovation where dependence is manageable or inevitable.

3.2 Sovereignty Thresholds by Levels of Criticality

Criticality	Examples	Minimum Sovereignty Threshold
Low criticality	Internal drafting, learning assistants, market research summaries.	Basic data protection, usage policy, vendor terms, human review.
Medium criticality	Sales enablement, HR service desk, finance analytics, procurement support.	Data residency clarity, audit logs, model evaluation, workflow controls, fallback process.
High criticality	Credit decisions, medical triage, industrial control, public benefits, law enforcement support.	Human-in-the-loop, explainability, domain benchmarks, model risk review, data lineage, legal review.
Strategic / national criticality	Defence, sovereign identity, critical infrastructure, national language models, emergency response.	Domestic or contractually assured compute and data control, red-team evaluation, continuity plans, model and supply chain assurance.

4. The Full Stack AI Architecture

To fully understand dependencies, maintain visibility and extract maximum value, the AI Diffusion Matrix comprises a ten-layer architecture across three macro tiers: The Interface Tier, The Intelligence Engine and The Physical Foundation. This 10-layer architecture is significantly more nuanced than most interpretations available. This provides a granular understanding of the various aspects of eventual AI performance. At the same time an appreciation of each of these layers opens up opportunities for Indian researchers, developers and startups to build solutions and serve the entire stack.

This report expands each layer into sublayers to make the framework usable for enterprise architecture, procurement, policy design and value measurement.

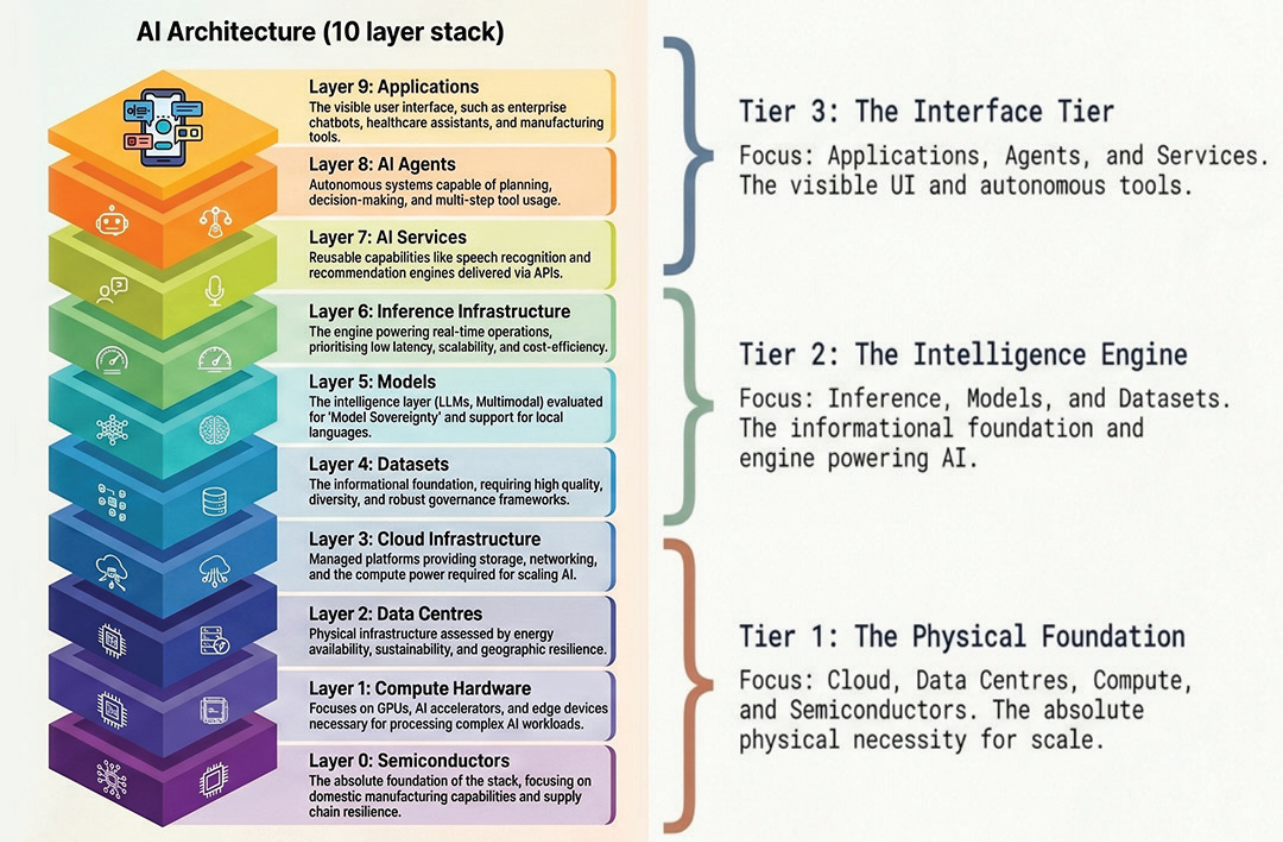

Each layer provides opportunity for Indian companies and startups to build solutions and serve the entire stack.

Figure 4. Expanded AI Diffusion Stack across the Interface Tier, Intelligence Engine and Physical Foundation.

4.1 The 10 layers: L9 to L0

Each of the 10 layers serves a Primary Role, dependent on each of the other layers of the stack. Without appreciation of these roles, end performance can vary vastly and may not produce the desired outcomes.

Layer	Name	Macro-tier	Primary Role
L9	Applications	Interface Tier	Bots, Copilots, Workflow Applications, Citizen/Customer/Employee-Facing AI Experiences.
L8	AI Agents	Interface Tier	Autonomous Systems for Planning, Tool Use, Workflow Management & Execution and Exception Handling.
L7	AI Services	Interface Tier	Reusable AI capabilities such as Voice, OCR, Translation (Multilingual Capabilities), Recommendation and Document Intelligence, Code Generation, Image and Video Creation.
L6	Inference Infrastructure	Intelligence Engine	Model serving, Retrieval, Routing, Observability, Evaluation and Runtime Cost Control.
L5	Models	Intelligence Engine	Foundation (LLMs), Domain, Multimodal and Small Language Models (SLMs) with evaluation and alignment layers.
L4	Datasets	Intelligence Engine	Enterprise, Sectoral, Public, Personal, Synthetic and Labelled Data governed for AI readiness.
L3	Cloud Infrastructure	Physical Foundation	Cloud, Location, Platform, Container, Storage, Networking, Security and AI Engineering Services.
L2	Data Centres	Physical Foundation	Physical Facilities, Location, Power (Energy), Cooling, Security, Resilience, Interconnect and Green Operations (Sustainability).
L1	Compute Hardware	Physical Foundation	GPU/Accelerator Clusters, Operating Systems, Servers, Edge Devices, Networks, Storage and Firmware.
L0	Semiconductors	Physical Foundation	Chip design, Packaging, Memory, EDA, Fabrication Ecosystem and Supply Chain Resilience.

4.2 Applications – Sub Layer Relevance: L9

The visible end-user layer where people experience AI through bots, copilots, dashboards, workflow tools, portals and mobile applications. Ease of deployment and integration, user friendly interface and extremely high reliability are the key to this layer.

Sublayer	Relevance
Enterprise copilots and knowledge assistants	Creates an always-on productivity and knowledge layer for employees by helping them search enterprise knowledge, summarise documents, draft content, analyse policies, support research, and make faster informed decisions.
Role-based workflow applications	Embeds AI into specific functions such as sales, finance, HR, legal, procurement and operations to reduce cycle time, improve process accuracy, strengthen compliance, and augment role-specific decision-making.
Customer service and citizen-service interfaces	Improves service quality and lowers operating cost by providing 24x7 multilingual assistance, faster query resolution, guided self-service, agent-assist recommendations, and consistent responses across channels.
Industry-specific applications e.g. healthcare, banking, manufacturing, agriculture, education, logistics etc.	Delivers domain-specific business value by applying AI to high-impact workflows such as clinical support, fraud and credit risk, predictive maintenance, quality inspection, crop advisory, curriculum design, AI tutors, demand forecasting and route optimisation.
Analytics and decision-support dashboards	Moves management reporting from static dashboards to proactive intelligence through forecasting, anomaly detection, scenario modelling, root-cause analysis, risk alerts, and recommended next-best actions.
Collaboration, communication and document workflows	Improves everyday enterprise throughput by automating meeting summaries, email and document drafting, translation, knowledge capture, review workflows, and cross-team collaboration across the organisation.

Value: Faster knowledge work, improved service response, better decision support, lower rework, personalised customer experience and improved employee productivity. As highlighted earlier, the enterprise value should be well defined and measurable before deployment.

Suggested metrics: Adoption rate, active usage, cycle-time reduction, document quality, ticket resolution time, NPS, rework rate, revenue conversion and employee satisfaction.

4.3 AI Agents: L8

Autonomous or semi-autonomous software actors that can plan, call tools, execute steps, escalate exceptions and complete workflows under policy constraints.

Sublayer	Relevance
Planning and task decomposition	Enables complex business goals to be broken into governed steps, improving process automation while requiring clear ownership, success criteria and failure controls.
Tool and API access	Determines whether agents can securely execute actions across enterprise systems instead of only generating recommendations.
Memory and context management	Preserves task history, user preferences and enterprise context so agents can deliver consistent outcomes without exposing sensitive information.
Permissions and authorisation	Ensures agents act only within approved roles, data boundaries and approval limits, reducing operational and compliance risk.
Human-in-the-loop checkpoints	Keeps critical decisions under human control and enables safe adoption in regulated, high-value or exception-prone workflows.
Agent observability and audit trails	Provides traceability of agent actions, prompts, tools used and decisions so enterprises can monitor performance and investigate failures.
Exception handling and rollback	Protects business continuity by ensuring failed or unsafe agent actions can be paused, reversed or escalated.

Value: Moves AI from advice to execution. Agents can reduce handoffs, operate across systems, complete repetitive work, and orchestrate complex workflows that span departments.

Example(s): Salesforce Agentforce describes autonomous AI agents that support employees and customers 24/7; SAP Joule Assistants coordinate AI agents across business domains.

Suggested metrics: Task completion rate, escalation rate, error rate, SLA adherence, cost per completed task, agent containment rate, compliance exceptions and human oversight time.

4.4 AI Services: L7

Reusable capabilities delivered through APIs or shared platforms, such as speech recognition, translation, OCR, search, recommendation, fraud detection and document intelligence.

Sublayer	Relevance
Speech-to-text and text-to-speech	Enables voice-first workflows, contact-centre automation, accessibility and multilingual engagement across customers and employees.
Translation and transliteration	Expands reach across languages and geographies while improving internal collaboration and customer or citizen service.
OCR and document extraction	Converts unstructured documents into searchable, process-ready data for onboarding, claims, procurement, compliance and operations.
Vector search and retrieval	Allows employees and applications to retrieve relevant enterprise knowledge from policies, documents, tickets and databases.
Recommendation engines	Personalises offers, next-best actions, content and operational decisions to improve conversion, retention and productivity.
Computer vision services	Automates inspection, verification, safety monitoring and visual quality control across manufacturing, logistics, healthcare and retail.
Fraud, risk and anomaly services	Detects unusual patterns in transactions, operations and system behaviour to reduce financial loss and operational risk.

Value: Reduces duplication by enabling reusable services. It allows multiple applications to consume common capabilities, improving consistency, cost efficiency and governance.

Example(s): Common enterprise service examples include multilingual support bots, invoice extraction services, contact-centre analytics and product recommendation APIs.

Suggested metrics: API latency, cost per call, reuse across applications, accuracy by language/domain, service availability, data leakage incidents and model drift.

4.5 Inference Infrastructure: L6

The runtime layer that serves models in production, manages prompts, retrieval, caching, routing, guardrails, observability, evaluations and cost optimisation.

Sublayer	Relevance
Model serving endpoints	Provides production-grade access to models with reliability, latency and security suitable for enterprise applications.
Prompt and context orchestration	Controls how business context, instructions and tools are assembled so outputs are consistent, relevant and auditable.
Retrieval augmented generation pipelines	Grounds model responses in enterprise data, reducing hallucination and making AI useful for knowledge-intensive workflows.
Caching and batching	Lowers inference cost and improves response time by reusing common outputs and optimising request volumes.
Model routing and fallback	Balances cost, latency, accuracy and resilience by selecting the right model and providing alternatives during failure.
Guardrails and policy filters	Enforces enterprise policies for privacy, safety, compliance and brand protection before AI outputs reach users or systems.
Evaluation harnesses	Tests model quality, safety and business performance continuously so deployments can be improved and trusted.
Observability, logging and cost telemetry	Tracks usage, quality, latency and spend to control ROI and detect failures in production.

Value: Determines reliability, latency, scalability and cost. A well-designed inference layer can lower cost, improve quality and enable model flexibility.

Example(s): Enterprises increasingly need inference routers that can select between frontier, open-source, domain-specific and small models depending on cost, risk and latency.

Suggested metrics: Latency, throughput, hallucination rate, groundedness score, cost per inference, GPU utilisation, cache hit rate, safety filter triggers and uptime.

4.6 Models: L5

The intelligence layer: foundation models, large and small language models, multimodal models, domain models, embedding and task-specific models.

Sublayer	Relevance
Foundation models	Provide broad reasoning and generation capability, but must be assessed for cost, IP, data exposure, vendor dependence and fit to enterprise use cases.
Small and domain-specific models	Deliver lower-cost, faster and more controllable performance for specialised enterprise tasks where broad generality is not required.
Multimodal models	Unlock workflows involving text, images, audio, video and documents, enabling richer automation in inspection, healthcare, design and service.
Embedding models	Convert enterprise content into searchable semantic representations, making retrieval, clustering and knowledge applications possible.
Fine-tuned and adapter models	Adapt base models to enterprise language, policies and domain data while reducing the need to build models from scratch.
Evaluation benchmarks	Establish objective measures of accuracy, safety, latency and business fit before models are approved for production.
Alignment and safety layers	Reduce harmful, biased, insecure or non-compliant outputs and help align model behaviour with enterprise policy.
Model registry and lifecycle management	Creates version control, ownership, auditability and retirement processes for models across the organisation.

Value: Models determine capability, accuracy, local language performance, reasoning depth, domain fit and cost. Model choice should be aligned with use case criticality. An enterprise may use a frontier model for complex reasoning, a small model for cost-sensitive classification, a domain model for legal/medical/financial work, and an embedding model for enterprise search.

Suggested metrics: Accuracy, true-to-original, latency, cost, bias, robustness, multilingual coverage, domain benchmark performance, explainability and model-switching flexibility.

Evaluating Model Risks - As the most important layer of the AI Stack, it is imperative to evaluate the risks associated with models, by level of criticality of use case.

Risk Tier	Example AI System	Required Control Level
Tier 1: Low	Internal drafting, summarisation of non-sensitive documents.	Basic policy, user training, disclosure and human review.
Tier 2: Moderate	Customer support assistant, HR helpdesk, sales assistant.	Data protection, evaluation, logging, role access and escalation.
Tier 3: High	Credit decision support, medical workflow support, legal review.	Formal model risk review, validation, audit trail, human-in-loop, bias testing.
Tier 4: Critical	Public benefits, critical infrastructure, safety-critical operations.	Independent assurance, red teaming, incident response, continuity and sovereignty controls.

4.7 Datasets: L4

The information foundation used to train, fine-tune, evaluate, ground and monitor AI systems.

Sublayer	Relevance
Enterprise data products	Package high-value internal data with ownership, quality SLAs and access rules so AI teams can reuse trusted data at scale.
Sector datasets	Provide industry-specific context, terminology and benchmarks needed for accurate models in healthcare, banking, manufacturing, agriculture and logistics.
Public datasets	Offer low-cost baseline data for training, evaluation and benchmarking, but require quality, licensing and bias checks.
Synthetic data	Fills gaps, protects privacy and enables testing of rare scenarios where real data is scarce, sensitive or expensive to label.
Labelled data	Supplies the ground truth required to train, fine-tune and evaluate models for accuracy, safety and domain performance.
Knowledge graphs	Captures relationships among entities, processes, policies and assets to support explainable reasoning and retrieval.
Data lineage and provenance	Shows where data came from, how it changed and whether it is reliable, enabling audit, compliance and trust.
Consent and rights management	Ensures AI uses data within permitted legal, contractual and customer consent boundaries.
Data quality and governance	Improves completeness, accuracy, timeliness and stewardship so AI outputs are dependable and defensible.

Value: Data is often the largest source of defensible enterprise value. Better data improves accuracy, personalisation, risk control and speed of deployment.

Example(s): AIKosh is positioned as an IndiaAI datasets platform to support AI research and innovation in India.

Suggested metrics: Data completeness, freshness, lineage coverage, consent coverage, labeling quality, data product reuse, retrieval precision, data access lead time and compliance incidents.

4.8 Cloud Infrastructure: L3

Managed cloud, platform and engineering services that provide compute, storage, networking, orchestration, security and developer tooling for AI.

Sublayer	Relevance
Compute instances and managed clusters	Enables scalable training, inference and experimentation without every enterprise building its own infrastructure.
Storage and data lake	Provides the common repository for structured and unstructured data required for analytics, RAG and model development.
Container orchestration	Standardises deployment, scaling and portability of AI services across cloud, on premise and hybrid environments.
Networking and identity	Secures connectivity, authentication and access between users, models, data and enterprise systems.
Secrets management	Protects API keys, credentials and tokens used by AI applications and agents to access critical systems.
MLOps and LLMOps tooling	Industrialises model deployment, monitoring, evaluation and release management across the AI lifecycle.
API gateway and service mesh	Controls traffic, authentication, rate limits and service-to-service communication across AI and enterprise applications.
Monitoring and FinOps	Tracks performance, reliability and cloud spend so AI scaling remains operationally and financially sustainable.

Value: Enables scalable deployment, faster experimentation and elastic cost management, but can create vendor lock-in if not managed deliberately.

Example(s): IndiaAI compute capacity aims to make AI compute infrastructure available through public-private partnerships and cloud services.

Suggested metrics: Deployment lead time, cost per workload, portability score, uptime, security incidents, compute availability, vendor concentration and compliance coverage.

4.9 Data Centres: L2

Physical facilities that host compute and storage, supported by power, cooling, network connectivity, physical security and operational resilience.

Sublayer	Relevance
Power availability and redundancy	Determines whether AI workloads can run reliably at scale without outages or capacity constraints.
Cooling and water management	Controls thermal performance, operating cost and environmental impact of high-density AI infrastructure.
Network connectivity and interconnect	Ensures low-latency, high-bandwidth movement of data between data centres, clouds, users and AI clusters.
Physical security	Protects servers, storage and network assets that host sensitive enterprise data and AI workloads.
Disaster recovery and geographic resilience	Maintains continuity when a site, region or network fails, reducing systemic operational risk. Factor in Geo-Political risks.
Energy sourcing	Influences AI cost, carbon footprint and long-term sustainability as compute demand grows.
Facility operations and DCIM	Provides real-time control of capacity, energy, cooling and asset performance for efficient AI infrastructure operations.

Value: Data centres determine where AI runs, the resilience of workloads, data residency, latency, energy cost and environmental impact.

Example(s): For regulated sectors, India-hosted data centres may support data residency and latency requirements. For high-scale AI, energy availability becomes a strategic constraint.

Suggested metrics: PUE, renewable energy share, power availability, downtime, latency, carbon intensity, water usage, geographic redundancy and data residency compliance.

4.10 Compute Hardware: L1

AI accelerators, servers, storage, networking, edge devices and systems that process AI workloads.

Sublayer	Relevance
GPUs and AI accelerators	Drives model training, inference throughput and cost efficiency for advanced AI workloads.
CPU and memory systems	Supports preprocessing, orchestration, traditional workloads and memory-intensive AI pipelines.
High-speed interconnect	Enables GPUs and servers to operate as a cluster, reducing bottlenecks in large-scale training and inference.
Storage systems	Determines how fast large datasets, embeddings, checkpoints and logs can be accessed and processed.
Edge devices	Brings AI inference closer to factories, branches, vehicles and field operations where latency, privacy or connectivity matter.
Firmware and drivers	Affects hardware stability, security, performance tuning and compatibility with AI software stacks.
Cluster management	Maximises utilisation, scheduling, availability and cost control across expensive AI compute resources.
Hardware security modules	Protects cryptographic keys and sensitive operations used in regulated AI and data workflows.

Value: Compute hardware drives training and inference performance, cost, energy efficiency, latency and the ability to run AI at scale.

Example(s): The IndiaAI Mission's compute portal and GPU empanelment efforts are designed to improve access to AI compute for startups, researchers, academia and public sector entities.

Suggested metrics: GPU hours, utilisation, performance per watt, cost per token, latency, memory bandwidth, queue time, cluster availability and hardware concentration.

4.11 Semiconductors: L0

The deepest physical layer: chip design, operating systems, fabrication, packaging, memory, EDA, supply chain and specialised silicon required for AI compute.

Sublayer	Relevance
Chip architecture and IP	Determines long-term control over AI performance, cost, power efficiency and strategic technology dependence.
EDA tools	Enables design and verification of chips; access constraints can limit domestic semiconductor innovation.
Fabrication ecosystem	Provides the manufacturing capability required to convert chip designs into physical silicon at scale.
Advanced packaging	Integrates memory, compute and interconnect efficiently, increasingly critical for high-performance AI chips.
Memory supply	Influences AI workload performance and cost because large models require high-bandwidth, high-capacity memory.
Testing and validation	Ensures chips meet reliability, performance and safety requirements before deployment in enterprise or national infrastructure.
Supply chain resilience	Reduces exposure to geopolitical, export-control and supplier-concentration risks across the AI hardware stack.
Domestic talent and R&D	Builds the human capital and intellectual property needed for long-term semiconductor and AI infrastructure competitiveness.

Value: Semiconductors affect long-term access, cost, energy performance and geopolitical exposure. Enterprises may not control this layer directly but must understand its implications.

Example(s): For most enterprises, action at this layer is indirect: risk assessment, supplier diversification, procurement criteria and support for domestic ecosystem initiatives.

Suggested metrics: Supplier concentration, geopolitical exposure, availability lead time, performance per watt, domestic availability, chip roadmap alignment and contingency coverage.

5. Evaluating and Governing AI Diffusion

Once the AI Architecture is understood, it is imperative to navigate each of the layers through a “Compass” with 5 Strategic Questions - Awareness, Sustainability, Ecosystem Engagement, Indigenisation and Actionability. The workbook operationalises these questions across rows such as current status, gaps, indigenisation quotient and call to action. This report expands the framework by positioning value realisation as the operating lens.

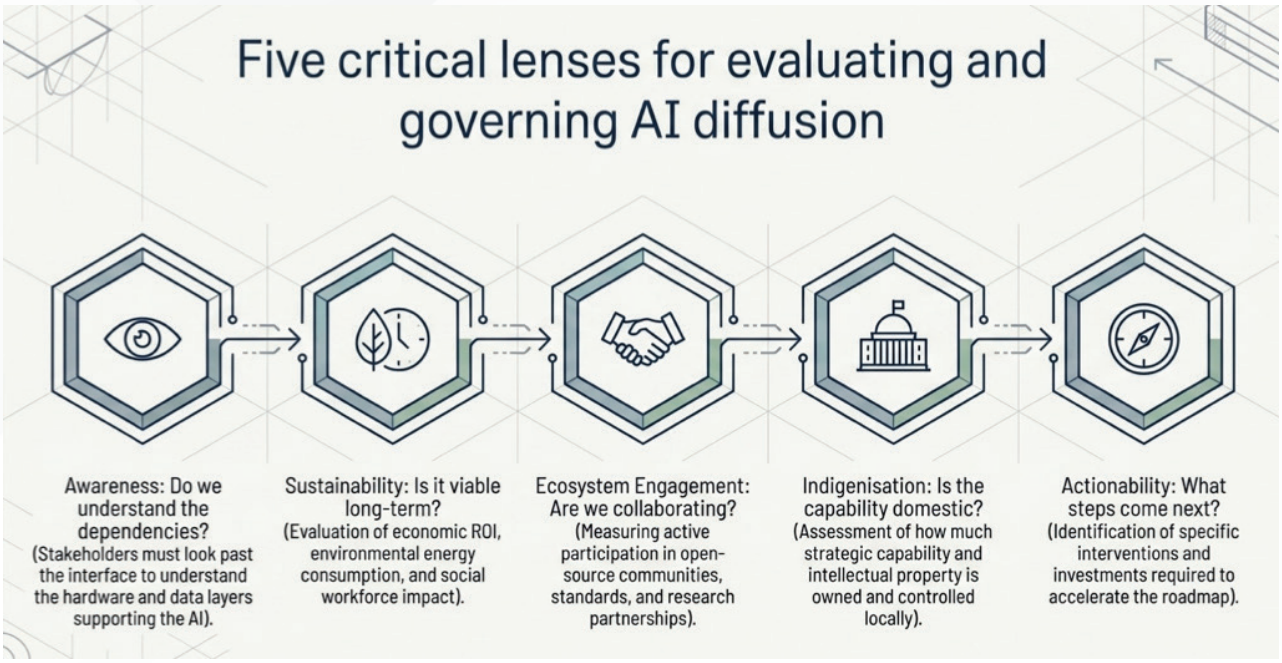

Captures all aspects related to costs, value over time, impact on human resources – hiring, training etc.

Figure 5. Five strategic questions for evaluation and diffusion.

5.1 Awareness

Evaluate whether leaders understand current and future dependencies across models, data, inference, cloud, hardware and semiconductors. Create a dependency register.

5.2 Sustainability

Evaluate long-term economics, energy impact, workforce transition, training, governance effort and environmental obligations.

5.3 Ecosystem Engagement

Evaluate whether the enterprise is a passive buyer or active ecosystem participant through partnerships, open-source contributions, sector data spaces and standards.

5.4 Indigenisation

Evaluate which layers must be Indian, India-hosted, India-compliant, locally auditable or internally controlled. Indigenisation is a risk-weighted choice, not a slogan.

5.5 Actionability

Translate insights into named interventions: procurement criteria, data readiness, model evaluation, partnerships, pilot-to-scale gates and investment priorities. Start by answering the right question. The figure and table below table serve as a guide to ask the right questions and evolve the Enterprise AI Readiness Framework.

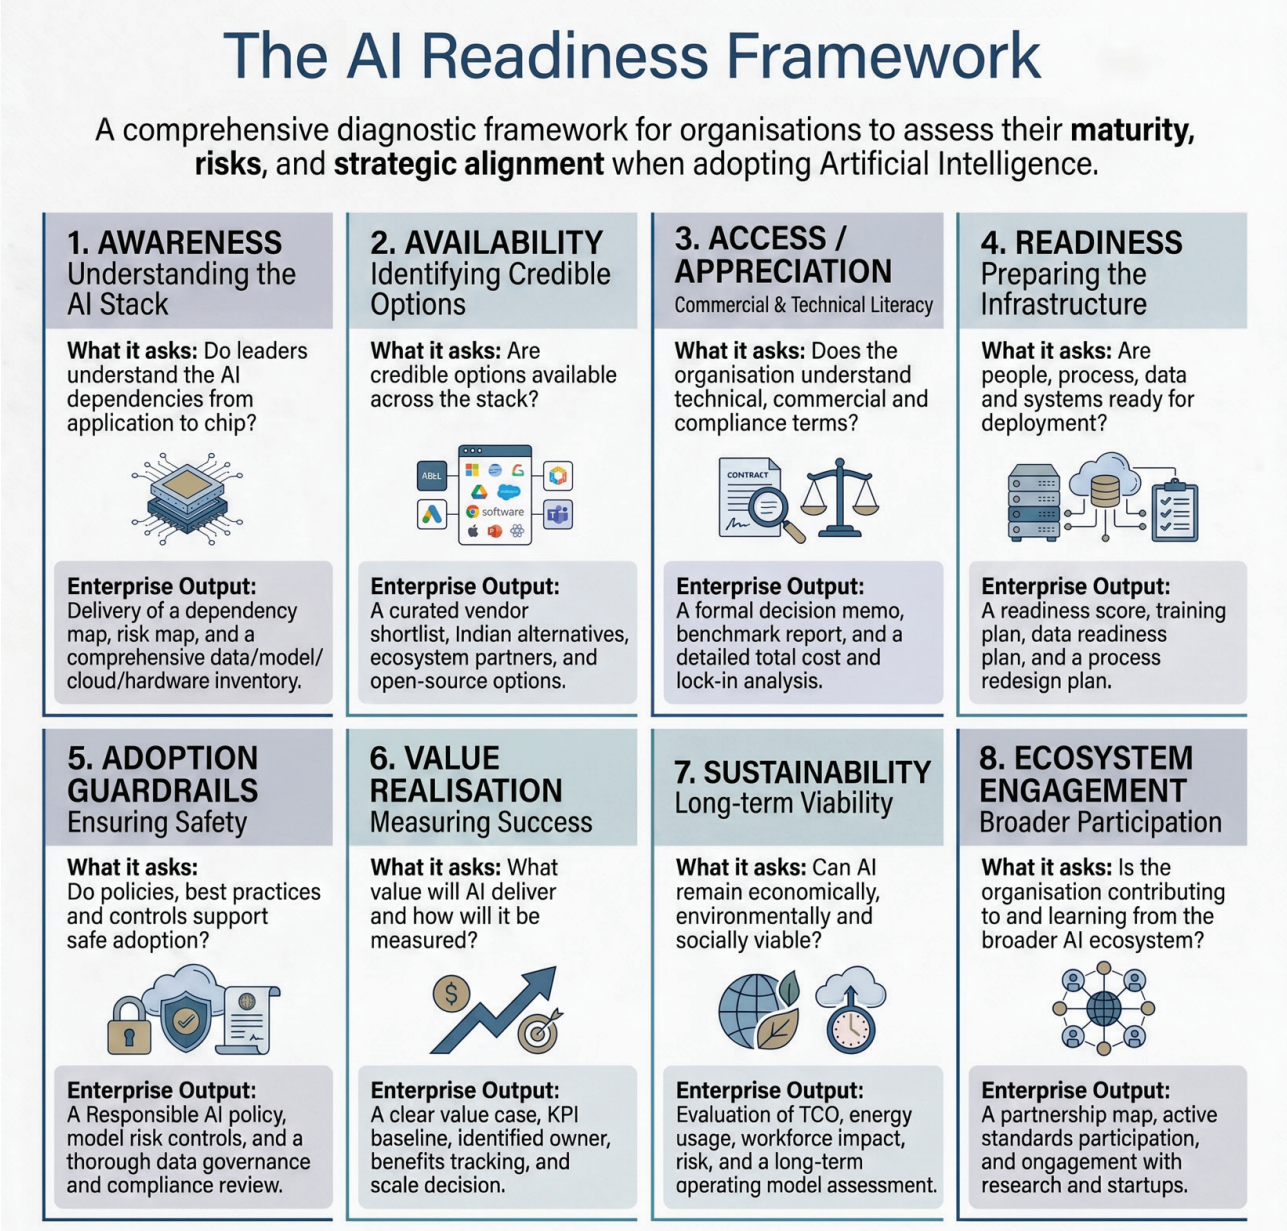

Figure 6. AI Readiness Framework.

Question	What it Asks	Output
Awareness	Do leaders understand the AI dependencies from application to chip?	Dependency map, risk map, data/ model/cloud/hardware inventory.
Availability	Are credible options available across the stack?	Vendor shortlist, Indian alternatives, ecosystem partners, open-source options.
Access / Appreciation	Does the organisation understand technical, commercial and compliance terms?	Decision memo, benchmark report, total cost and lock-in analysis.
Readiness	Are people, process, data and systems ready for deployment?	Readiness score, training plan, data readiness plan, process redesign plan.
Adoption guardrails	Do policies, best practices and controls support safe adoption?	Responsible AI policy, model risk controls, data governance and compliance review.
Value realisation	What value will AI deliver and how will it be measured?	Value case, KPI baseline, owner, benefits tracking and scale decision.
Sustainability	Can AI remain economically, environmentally and socially viable?	TCO, energy, workforce, risk and long-term operating model assessment.
Ecosystem engagement	Is the organisation contributing to and learning from the broader AI ecosystem?	Partnership map, standards participation, research and startup engagement.

6. What AI Delivers and How to Measure It

Value Realisation Is a Key Metric

Value is the economic reason to adopt AI, the management reason to scale it, and the policy reason to invest in shared infrastructure. Without value measurement, AI becomes a cost centre and a narrative exercise.

The Value Question

What value will AI deliver to the enterprise, who owns that value, how will it be measured, and what must be true across the stack for the value to be sustained?

AI delivers value when it changes a measurable business outcome. This is different from adoption. A tool can be widely used and still fail to improve P&L, risk posture or customer outcomes. The AI Diffusion Matrix therefore recommends a value architecture that links use case, workflow, data, model, operating model and financial metric.

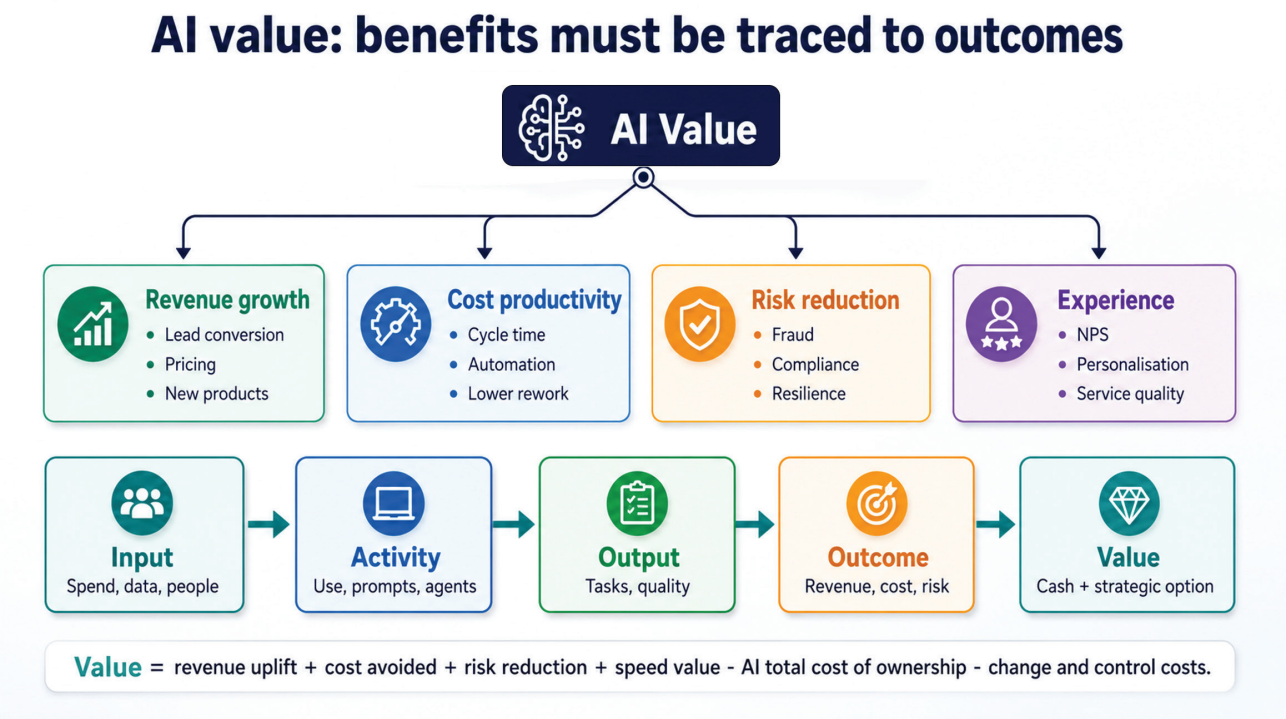

Figure 7. AI value tree and measurement chain.

External studies show both promise and caution. Google Cloud's 2024 survey reported that 74% of enterprises using generative AI were seeing ROI and that 45% of organisations reporting productivity gains estimated employee productivity had at least doubled. A Microsoft-sponsored IDC study reported an average 3.7x return per dollar invested in generative AI. These figures should be treated as directional vendor-sponsored evidence, not universal guarantees. The matrix approach prevents overgeneralisation by requiring each enterprise to baseline its own workflow and measure value over time.

6.1 The AI Value Equation

A practical value equation should combine financial benefits, operational gains, risk reduction and strategic option value. The equation below is designed for steering committee review and benefits tracking:

AI Value = Revenue Uplift + Cost Avoided + Risk Reduction + Speed Value + Strategic Option Value - AI Total Cost of Ownership - Change Cost - Control Cost

AI's total cost of ownership includes licenses, cloud, inference, data preparation, integration, evaluation, security, training, adoption support and governance. Control cost includes legal review, human oversight, audits, model risk management and incident response.

Component	How To Estimate	Evidence Required
Revenue uplift	Incremental revenue attributable to AI-enabled conversion, pricing, personalisation or new products.	Baseline vs test cohorts, sales funnel data, attribution model.
Cost reduction/ saving	Reduced effort, lower error correction, automated volume, reduced outsourcing or better capacity utilisation.	Time studies, process mining, transaction volume, cost rates.
Risk reduction	Losses avoided from fraud, compliance issues, downtime, defects or decision errors.	Historical loss data, risk probability, severity estimates, control effectiveness.
Speed value	Value of shorter cycle time in product release, service response, collections, procurement or claims.	Cycle-time baseline, bottleneck analysis, cash conversion impact.
Strategic option value	Future value created through data assets, reusable AI services, sovereign model access or platform capability.	Reuse metrics, platform adoption, model portability, ecosystem participation.
TCO and control cost	All cost to deploy, run, monitor, govern and improve AI.	License, API, GPU/inference, integration, people, training, evaluation and security costs.

6.2 Measurement: From Activity to Financial and Strategic Impact

The measurement system should be defined before the pilot begins. Every AI initiative should have a use case owner, business KPI owner, data owner, technology owner, risk owner and benefits tracker. AI pilots should not scale unless they meet the agreed outcome threshold and have a sustainable operating model.

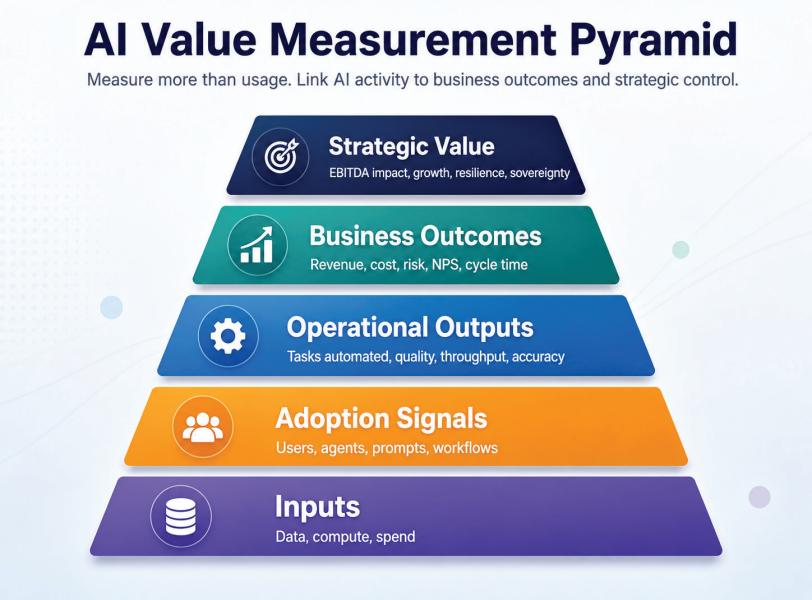

Figure 8. AI value measurement pyramid: moving beyond activity metrics to outcome and strategic value.

Metric layer	Good Metric	Weak Metric
Usage/ Inputs	Weekly active users in target workflow; workflow coverage; agent completion rate.	Number of licenses purchased or number of demos completed.
Adoption Signals	Accuracy, groundedness, defect rate, human override rate, review time.	Subjective statement that output looks good.
Operational	Cycle time, throughput, SLA, cost per transaction, escalation rate.	Number of prompts submitted.
Financial/ Business Outcomes	Revenue uplift, cost avoided, EBITDA impact, cash conversion, loss avoided.	Generic productivity claim without baseline.
Strategic Outcomes	Data asset reuse, model portability, dependency reduction, sovereign control score.	Vendor brand preference.

6.3 Examples of Value Across Enterprise Functions

Function	AI Value Example	Measurement
Sales and marketing	AI account research, proposal generation, next-best-action, personalisation.	Lead conversion, proposal cycle time, campaign ROI, win rate.
Customer service	Agent assist, automated triage, conversational self-service, sentiment routing.	First contact resolution, average handle time, containment rate, CSAT.
Software engineering	Code generation, documentation, unit tests, code review support.	Cycle time, defects, deployment frequency, review effort, developer satisfaction.
Finance	Invoice matching, variance commentary, cash forecasting, audit testing.	Close cycle time, forecast accuracy, exceptions, audit effort.
HR	Candidate screening support, onboarding assistant, HR helpdesk.	Time to hire, onboarding time, employee case resolution, training completion.
Legal and compliance	Contract review, regulatory horizon scanning, policy Q&A.	Review time, missed clauses, compliance exceptions, legal backlog.
Operations and supply chain	Demand sensing, inventory optimisation, predictive maintenance, procurement intelligence.	Forecast accuracy, downtime, inventory turns, procurement savings.

7. Culture Readiness: The Human and Ecosystem

Technology alone does not determine the success of AI adoption. Organisations with similar technology investments often achieve very different outcomes because their cultures differ. A culture that encourages questioning, knowledge sharing, experimentation and responsible risk-taking enables AI to diffuse across teams and business functions. Conversely, a highly hierarchical, closed or risk-averse culture can restrict innovation, leaving promising pilots isolated and preventing them from becoming enterprise capability.

Culture Readiness is therefore included as a compact enabling dimension within the AI Diffusion Matrix. It assesses whether people, leadership and ways of working support responsible AI adoption, and whether the organisation is open to co-innovating with customers, suppliers, startups and small and medium business partners.

Figure 9. Culture Readiness enables AI and innovation to move from isolated experiments to enterprise-wide and ecosystem-wide value

7.1 Culture Readiness Assessment

The assessment uses five concise dimensions. It is not intended to create a separate culture framework, but to identify whether organisational culture is likely to accelerate or constrain AI diffusion within the enterprise and across its ecosystem.

Dimension	What To Assess	Characteristics of High Readiness
Open Culture	Can employees question assumptions, challenge respectfully and share ideas freely?	High trust, open dialogue, psychological safety and cross-functional knowledge sharing.
Innovation & Learning	Does the organisation encourage creativity, rapid experimentation and learning from failure?	Low-friction pilots, constructive feedback, learning loops and scaling of successful ideas.
AI Mindset	Is AI viewed mainly as a threat or as an opportunity to augment people and improve outcomes?	Cautious optimism, responsible experimentation and confidence in human + AI collaboration.
Leadership Enablement	Do leaders visibly sponsor innovation, empower teams and remove organisational barriers?	Clear guardrails, delegated decision rights, resources for experimentation and support for scaling.
Ecosystem Co-innovation	Is the organisation open to experimenting and co-innovating with suppliers, startups, customers and partners?	Open collaboration, partner sandboxes, governed data access and shared innovation programmes.

Ecosystem Openness in Practice

A diffusion-ready culture extends beyond internal teams. It treats the wider ecosystem as a source of ideas, capability and scale. Organisations should assess whether they:

- Experiment with ecosystem partners rather than restricting innovation to internal teams.
- Adopt an open ecosystem approach to co-innovation rather than a walled or purely transactional model.
- Provide governed sandboxes for select partners to test solutions using relevant data, APIs and environments.
- Share knowledge, training and enablement with small and medium business partners so that they can adopt technology-led innovations and collaborate effectively.

7.2 Integration With the Enterprise Assessment Tool

Culture Readiness is included as one assessment step within the AI Diffusion Matrix Assessment Tool. A short set of statements covers openness, innovation, AI mindset, leadership support and ecosystem collaboration. The result is a single Culture Readiness Score that complements the overall AI Diffusion Score.

Indicative Assessment Statements

No.	Assessment Statement
1	Employees can openly question assumptions and share ideas without fear of negative consequences.
2	The organisation supports controlled experimentation and learns constructively from unsuccessful trials.
3	AI is viewed as an opportunity to augment people and improve outcomes, while risks are responsibly governed.
4	Leaders actively sponsor innovation, empower teams and remove barriers to adoption.
5	The organisation is open to co-innovating with suppliers, startups, customers and other ecosystem partners.
6	Selected partners can access governed sandboxes, APIs, data or test environments where appropriate.
7	Knowledge, training and technology enablement are shared with small and medium business partners to improve collaboration.
8	Successful innovations can move from internal pilots or partner experiments into scaled enterprise use.

Scoring

Score	Interpretation
1–2	Restrictive
3–4	Emerging
5	Supportive
6	Diffusion Ready

Interpretation

A high Culture Readiness Score does not imply unrestricted AI use or unrestricted partner access. The target is an organisation that is open and experimentative, while maintaining clear governance, data controls, accountability, security and human oversight.

8. Mapping the AI Diffusion Journey

The maturity model is a journey from awareness to sovereignty. It should be assessed by layer. A company can be Level 4 in applications while remaining Level 1 in model evaluation, data readiness or compute dependency awareness. That gap is the illusion of progress highlighted in the deck.

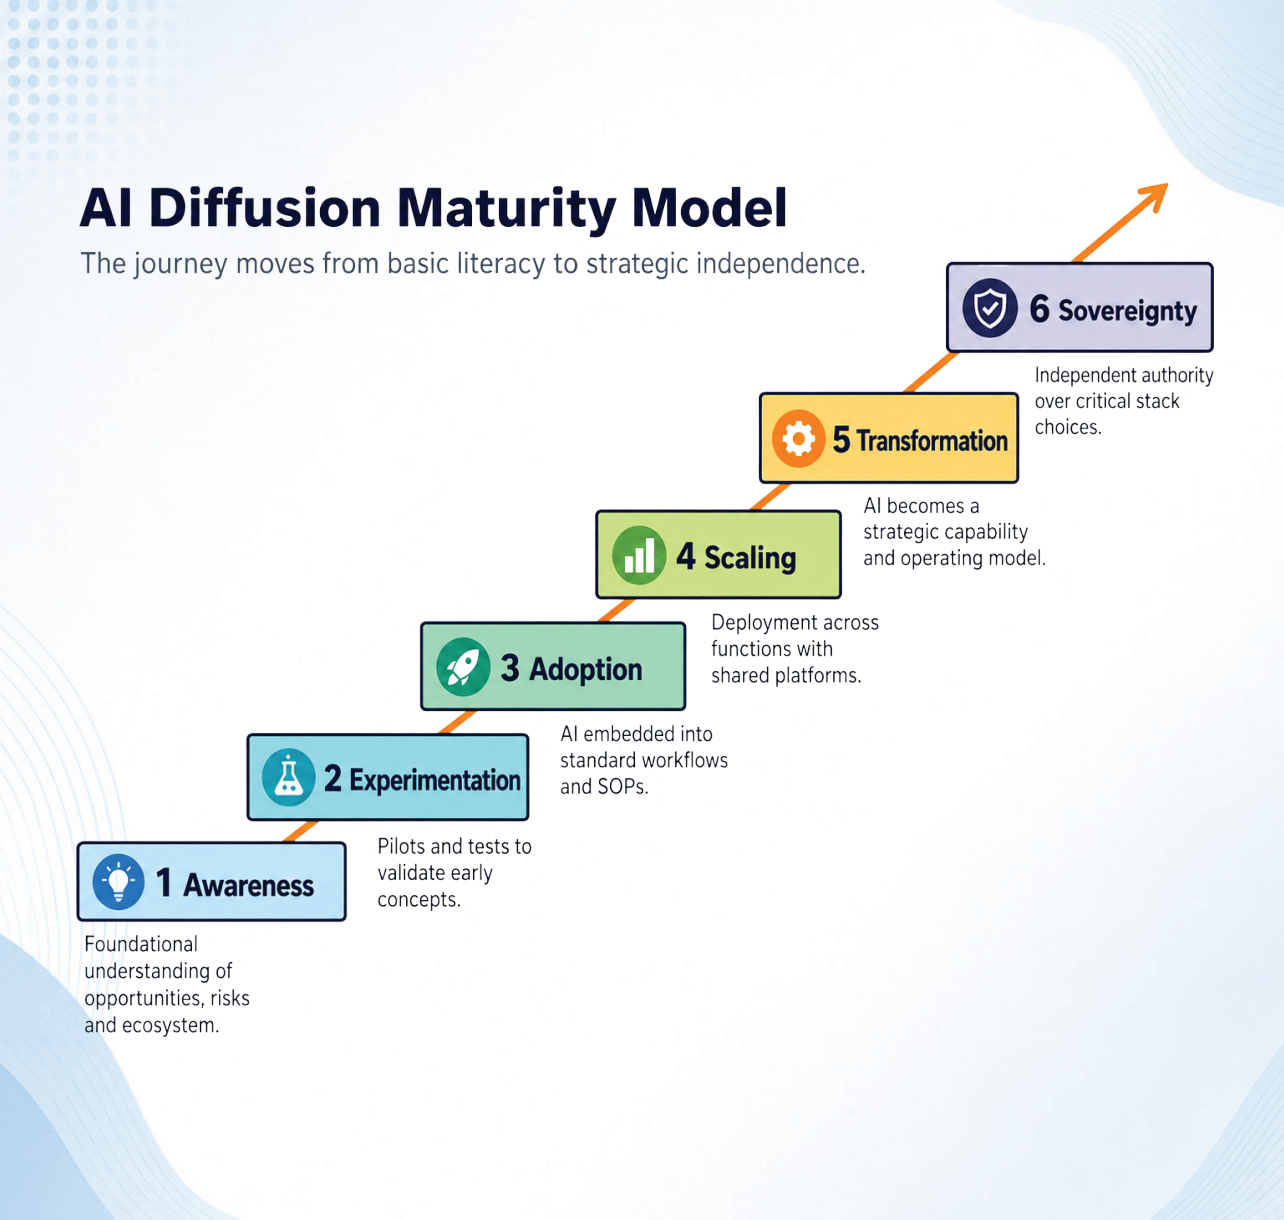

Figure 10. Six-level AI Diffusion Maturity Model

Level	Description	Evidence Required	Exit Criteria
Level 1: Awareness	Stakeholders understand AI opportunities and ecosystem basics.	Training, dependency maps, use-case inventory.	Leadership can explain where value and risk exist.
Level 2: Experimentation	Pilots and small tests prove concepts.	Prototype results, pilot metrics, user feedback.	Pilot has measurable baseline and owner.
Level 3: Adoption	AI integrated into normal workflows.	SOPs, adoption metrics, quality controls.	Workflow KPIs improve in production.
Level 4: Scaling	AI expands across functions and shared platforms.	Platform reuse, common services, governance.	Multiple business units use shared AI infrastructure.
Level 5: Transformation	AI defines operating model and competitive edge.	Process redesign, AI-native roles, value dashboard.	AI materially shifts economics or customer experience.
Level 6: Sovereignty	Strategic independence and control achieved where needed.	Dependency assurance, fallback, domestic/internal capability.	Critical AI systems remain resilient under supply, policy or vendor stress.

8.1 Scoring Methodology

The recommended scoring system combines maturity and strategic importance. A layer can score high on current adoption but still represent a major risk if it is strategically important and poorly controlled.

Score Element	Definition	Scoring Guidance
Maturity score	Current level from 1 to 6 for each layer and dimension.	1 = aware only; 6 = sovereign/ control threshold met.
Criticality score	Importance of the layer to the use case.	1 = non-critical; 5 = mission critical or regulated.
Dependency risk	Exposure to vendor, jurisdiction, cost, security or supply disruption.	1 = low; 5 = high concentration or opaque dependency.
Value potential	Expected contribution to business or public value.	1 = minor; 5 = material revenue, cost, risk or societal impact.
Action priority = (Criticality + Dependency Risk + Value Potential) - Maturity	Composite priority for intervention.	High when criticality, dependency risk and value potential are high but maturity is low.

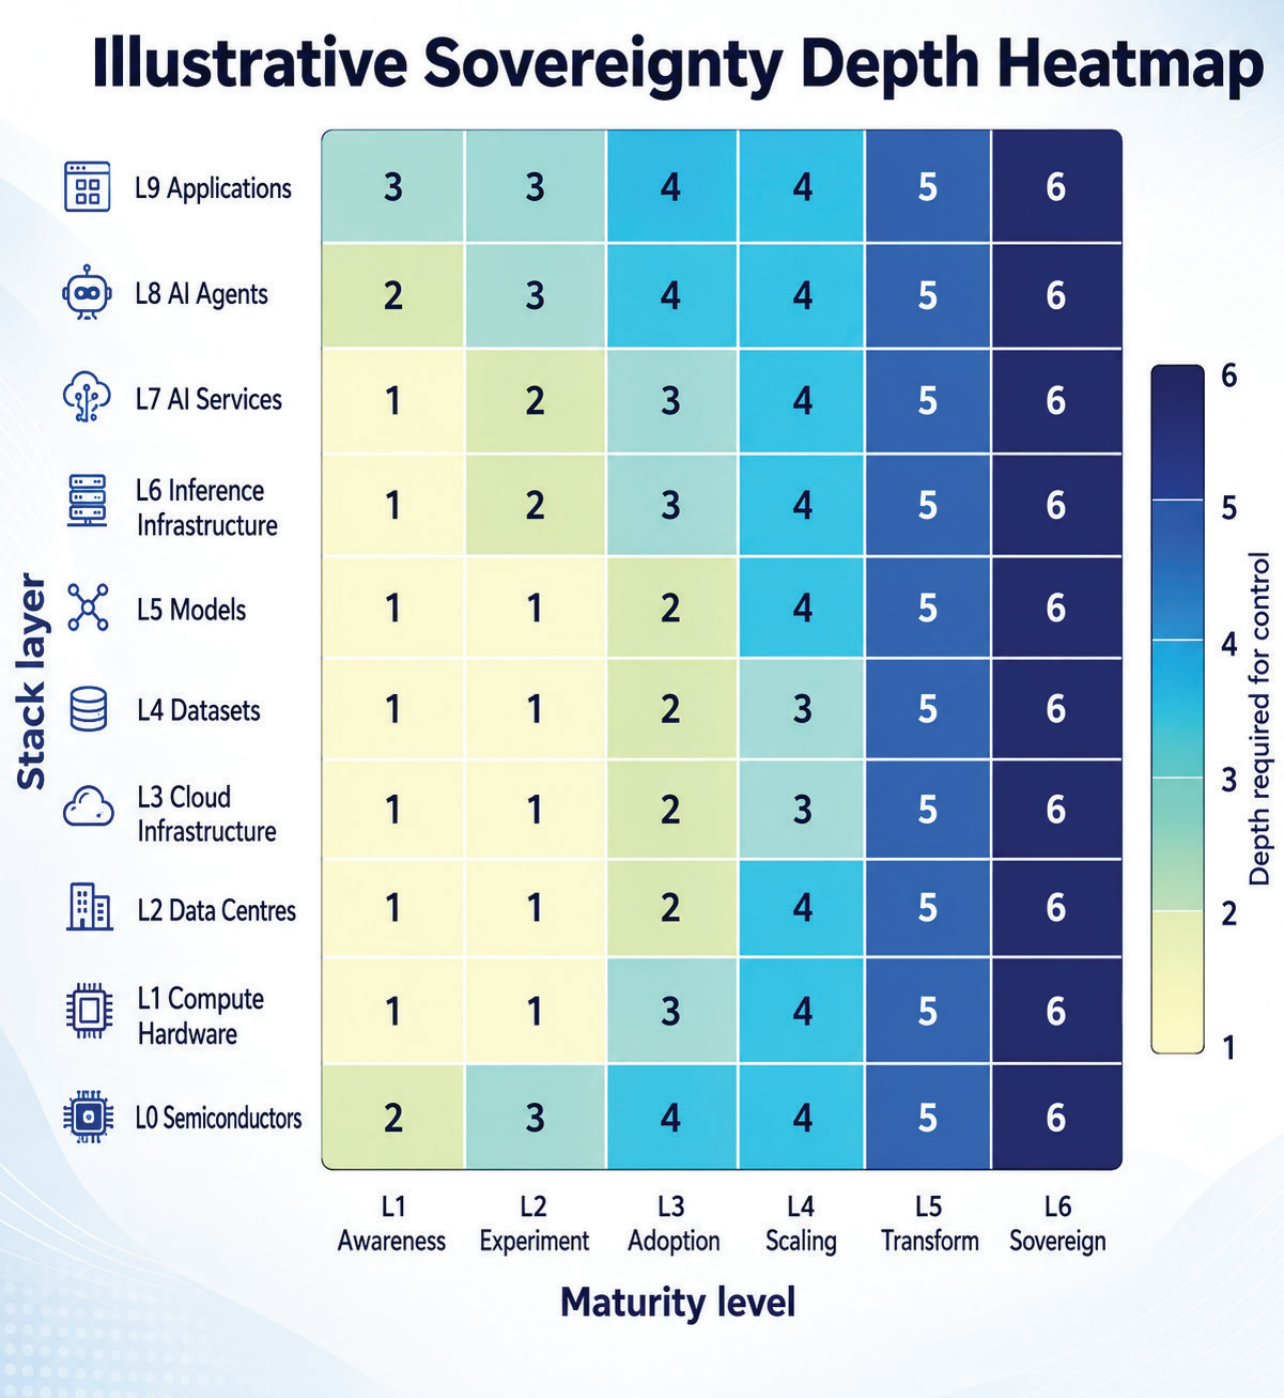

Figure 11. Illustrative maturity heatmap: sovereignty demands depth across layers.

9. Applying the AI Diffusion Matrix at Multiple Levels

The AI Diffusion Matrix is designed to be applied at different levels depending on the respondent, the decision to be supported and the breadth of the system being assessed. The same core framework can therefore be used to examine a single department, an entire organisation, the industry in which it operates, a state or region, or a country. What changes is not the underlying logic of the Matrix, but the unit of analysis, the evidence base and the stakeholders whose actions shape AI diffusion.

A scalable assessment lens: At the departmental and organisational levels, the Matrix focuses on use cases, value, readiness, governance and control. At the industry level, it examines shared datasets, standards, infrastructure and ecosystem dependencies. At state or country level, it expands further to public infrastructure, capability development, policy, investment, resilience and sovereignty.

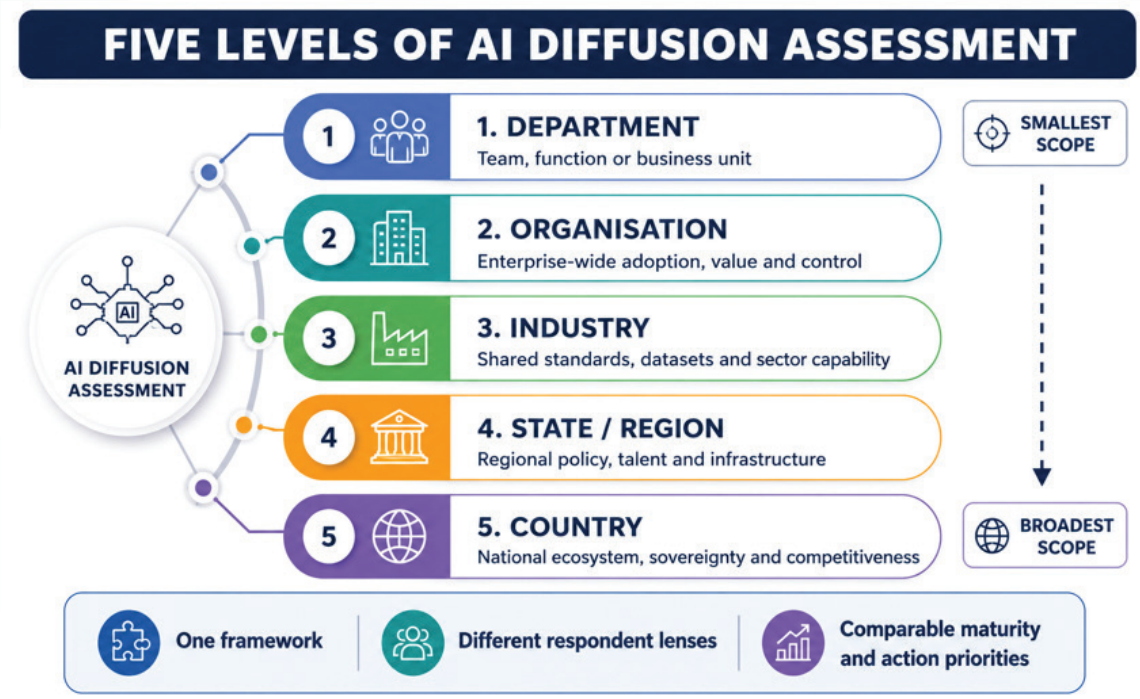

Figure 12. The same AI Diffusion Matrix can be applied across five progressively broader units of analysis.

9.1 Defining the Scope

The assessment should begin by selecting the level that best matches the respondent requirement. A department may assess a defined portfolio of AI use cases, while a country-level assessment will consider the complete ecosystem across enterprises, government, research, infrastructure and policy. Assessments can also be linked: departmental results can inform an organisational view, and organisational assessments can contribute evidence to an industry or regional analysis.

9.2 Five Assessment Segments

Each segment uses the same strategic questions and stack logic, but adapts the respondent group, evidence sources and expected output to the scale of the assessment.

Assessment Segment	Primary Focus	Typical Respondents	Indicative Output
1. Department	AI use cases, workflow impact, adoption barriers, data access and measurable value within a function or business unit.	Department head, functional leaders, process owners, technology and data teams.	Department AI Diffusion profile, priority use cases, gaps and action plan.
2. Organisation	Enterprise-wide portfolio, governance, culture, procurement, shared platforms, value realisation and strategic control.	Board, CXOs, business units, CIO/CTO/CDO teams, risk, HR and procurement.	Enterprise AI Diffusion Score, heatmap, investment priorities and roadmap.
3. Industry	Sector use cases, common standards, shared datasets, supply-chain dependencies, regulation and ecosystem capability.	Industry associations, regulators, major enterprises, startups and research bodies.	Industry maturity benchmark, shared capability gaps and sector interventions.
4. State / Region	Regional adoption, public services, talent, compute, data infrastructure, investment environment and local ecosystem readiness.	State government, regional agencies, universities, industry clusters and civic institutions.	Regional readiness profile, infrastructure priorities and public-private action agenda.
5. Country	National AI capability, diffusion across sectors, sovereignty, infrastructure, research, talent, policy and international competitiveness.	National government, industry, academia, public institutions and ecosystem leaders.	National AI Diffusion assessment, strategic dependencies and long-term capability roadmap.

9.3 Using the Segments in Practice

- Select one primary segment before beginning the assessment so that questions, respondents and evidence remain consistent.
- Use lower-level assessments as inputs to broader assessments, but avoid simply averaging scores; context and strategic importance must be considered.
- Compare results over time or across similar units to identify where AI diffusion is accelerating, where dependencies remain and which interventions should be prioritised.

Core Principle

The unit of analysis may change, but every assessment should continue to connect AI adoption with enterprise or public value, readiness, governance, ecosystem capability and the required depth of control across the AI stack.

10. Sector Playbooks and Examples

The AI Diffusion Matrix becomes most useful when applied to sectors. Each sector has different value pools, risk thresholds, datasets, compliance requirements and sovereignty needs. The following playbooks illustrate how the matrix can guide use-case selection and stack assessment.

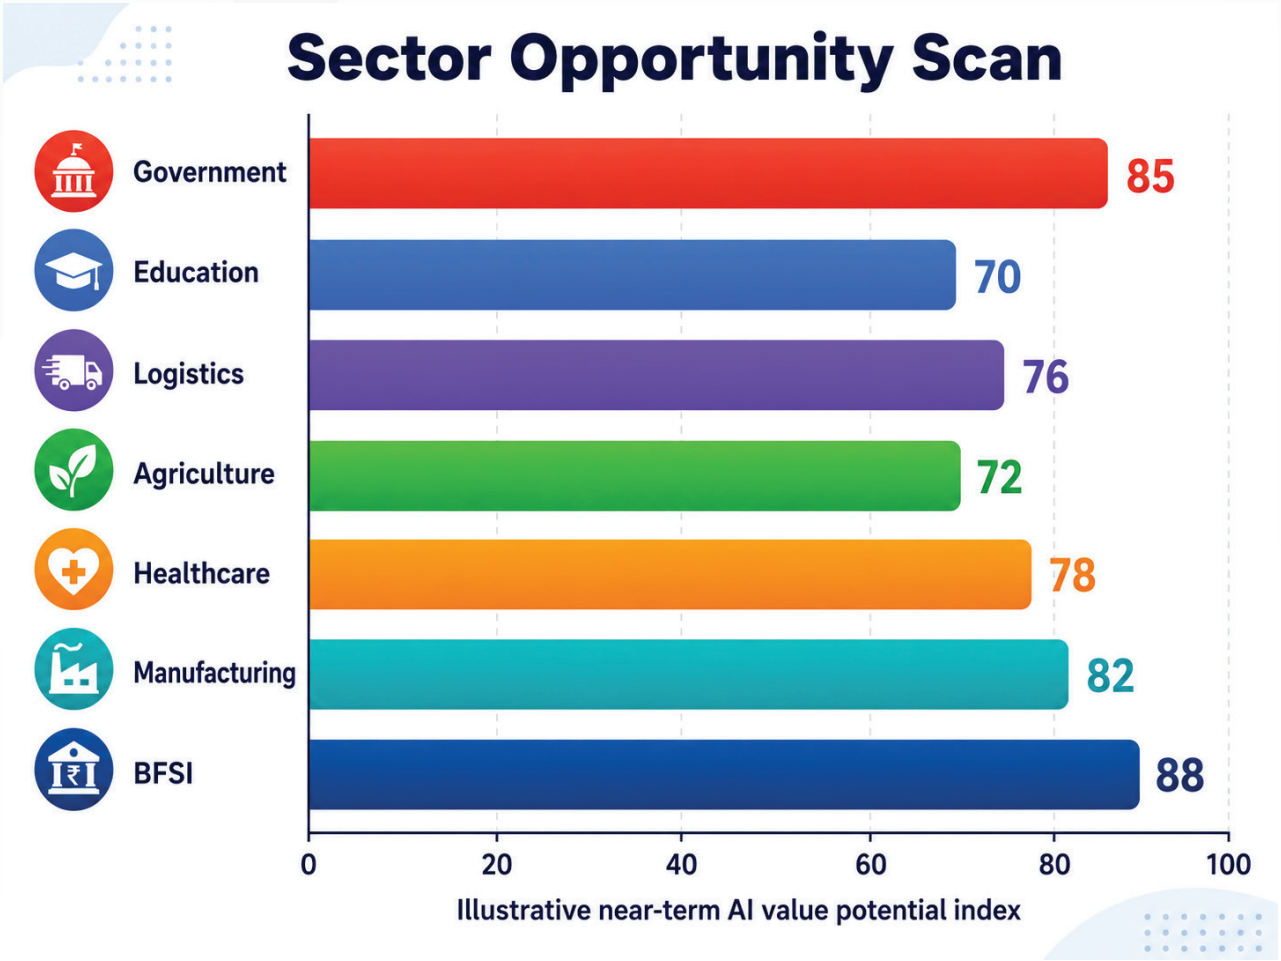

Figure 13. Illustrative sector opportunity scan. Scores are illustrative, not measured national values.

10.1 BFSI

Area	Guidance
Use cases	Credit underwriting, fraud detection, customer service, collections, risk monitoring, wealth advisory, investment analysis.
Value metrics	Fraud loss reduction, loan approval TAT, NPA prediction accuracy, customer resolution time, compliance cost, loan disbursal and deposits growth.
Sovereignty concerns	Data residency, explainability, auditability, model risk management, regulatory reporting.
Stack focus	Models, datasets, inference observability, data governance, cloud/data centre compliance.

10.2 Healthcare

Area	Guidance
Use cases	Clinical documentation, triage support, imaging assistance, patient engagement, hospital operations.
Value metrics	Clinician time saved, diagnostic turnaround, patient wait time, readmission risk, claim accuracy.
Sovereignty concerns	Patient privacy, safety, bias, auditability, clinical oversight, local language and context.
Stack focus	Datasets, model evaluation, safety guardrails, human-in-loop and sector sandboxes.

10.3 Manufacturing

Area	Guidance
Use cases	Predictive maintenance, quality inspection, production planning, safety monitoring, engineering design.
Value metrics	Downtime avoided, defect rate, OEE, energy efficiency, inventory turns, maintenance cost.
Sovereignty concerns	Industrial IP, edge compute, OT security, supplier resilience, model reliability.
Stack focus	Edge hardware, inference latency, computer vision models, data centres and secure cloud.

10.4 Agriculture

Area	Guidance
Use cases	Farmer advisory, crop disease detection, weather intelligence, input optimisation, market linkage.
Value metrics	Yield improvement, input cost reduction, advisory adoption, crop loss avoided, farmer income.
Sovereignty concerns	Indian language and local context, public datasets, smallholder inclusion, data rights.
Stack focus	AI services, mobile applications, public/sector datasets, model localisation and DPI integration.

10.5 Logistics and Supply Chain

Area	Guidance
Use cases	Route optimisation, demand sensing, warehouse automation, exception management, procurement intelligence.
Value metrics	On-time delivery, fuel cost, logistics cost, stock-outs, working capital, SLA adherence.
Sovereignty concerns	Operational continuity, supply chain data sharing, vendor concentration, cyber security.
Stack focus	Agents, AI services, data exchange, inference reliability, cloud and edge infrastructure.

10.6 Education and Skilling

Area	Guidance
Use cases	Personalised tutors, teacher assistants, assessment, content translation, career guidance.
Value metrics	Learning outcomes, teacher time saved, completion rates, language access, employability outcomes.
Sovereignty concerns	Child safety, bias, curriculum alignment, Indian languages, data privacy.
Stack focus	Applications, AI services, datasets, evaluation and governance.

10.7 Government Services

Area	Guidance
Use cases	Citizen service assistants, document processing, grievance redressal, benefit eligibility, multilingual service navigation.
Value metrics	Resolution time, service completion rate, citizen satisfaction, appeal rate, cost per service.
Sovereignty concerns	Public accountability, data protection, explainability, accessibility, service continuity.
Stack focus	AI DPI, identity/consent, data exchange, model evaluation, open standards and India-hosted infrastructure.

11. AI Governance and Operating Model

AI governance must be integrated into business governance, technology governance, data governance, risk management and procurement. NIST's AI Risk Management Framework and Generative AI Profile emphasise trustworthiness considerations for design, development, use and evaluation of AI systems.

The AI Diffusion Matrix complements such frameworks by adding stack-level dependency awareness and enterprise value measurement.

Metric Layer	Good Metric	Weak Metric
Board / CXO AI Council	Set strategic ambition, risk appetite and investment priorities.	Which value pools matter? Which domains are critical? What sovereignty threshold is required?
AI Value Office	Prioritise use cases, baseline metrics, track value and manage scale decisions.	Which pilots scale, stop or redesign? Are benefits real?
AI Architecture Board	Define stack standards, model selection, inference patterns, data architecture and platform reuse.	Which models/clouds/services are approved? What is the portability strategy?
Responsible AI and Risk Committee	Review risk tiers, data protection, safety, bias, human oversight and compliance.	What controls are mandatory before deployment?
Data Governance Council	Manage data products, lineage, access, consent and quality.	Which data can be used and under what terms?
Procurement and Legal	Contract AI vendors with portability, audit, data and continuity clauses.	What must be in contracts for model, data, IP and exit rights?
Business Process Owners	Own workflow redesign and adoption.	How will AI change SOPs, roles and KPIs?

11.1 RACI Framework for AI Governance (Responsible, Accountable, Consulted and Informed)

Roles and responsibilities for AI adoption and diffusion can be understood and mapped using the RACI Model. The below table provides responsibility, accountability, consultative and informed clarifications for each role. This would enable frictionless deployment. This paper offers this framework only as a helpful guide.

Activity	Business Owner	CIO / CTO	Data Office	Risk / Legal	AI Value Office	Ecosystem / Govt
Use-case prioritisation	A	C	C	C	R	I
Value baseline and benefits tracking	R	C	C	I	A	I
Data readiness and access	C	C	A/R	C	C	I
Model selection and evaluation	C	A/R	C	C	C	I
AI risk assessment	C	C	C	A/R	C	I
Vendor and contract terms	C	C	I	A/R	C	I
AI DPI / ecosystem participation	I	C	C	C	C	A/R
Scale decision	A	C	C	C	R	I

Legend: A = Accountable, R = Responsible, C = Consulted, I = Informed.

11.2 Procurement Practice for AI Sovereignty

A number of factors need to be part of the decision-making process during procurement of AI Solutions and Services once the Diffusion Framework has been understood and evaluated. This will ensure long term visibility and control over all aspects of the AI Stack.

- Data rights: specify data ownership, training rights, retention, deletion and secondary use.
- Model transparency: require model card, evaluation results, known limitations and update policy.
- Portability and exit: ensure export of data, prompts, embeddings, logs, customisations and evaluation records.
- Security and audit: require logs, access controls, incident reporting, penetration testing and audit rights.
- Jurisdiction and residency: specify where data and inference occur for regulated workloads.
- Performance and cost: set latency, uptime, cost-per-transaction and scaling terms.
- Human oversight: define decisions that require human validation and escalation.
- Sustainability: require energy transparency, data centre standards and efficiency metrics where relevant.

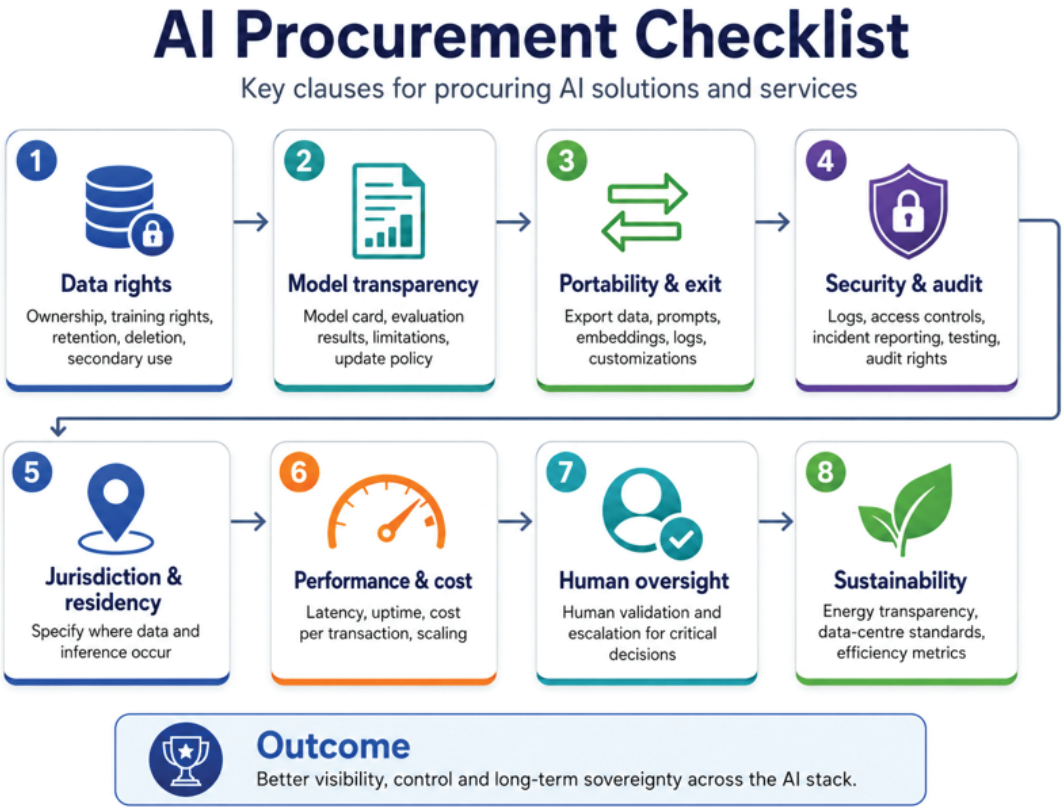

Figure 14. AI procurement Checklist.

12. Recommendations

12.1 Recommendations for Enterprises

Recommendation	Action	Expected Outcome
Treat AI as an operating model program	Create AI Value Office; assign business owners; link use cases to KPIs.	Fewer scattered pilots and higher probability of measurable value.
Map dependencies for every critical AI use case	Document data, model, inference, cloud, data centre, hardware and jurisdiction dependencies.	Improved resilience and better procurement decisions.
Build a value dashboard	Track usage, quality, operations, financial impact, risk and TCO.	Benefits are visible and scale decisions become evidence-based.
Create an internal AI DPI	Shared data products, model registry, inference layer, agent controls and governance.	Reusable infrastructure and reduced duplication.
Invest in AI literacy by role	Board, CXO, product, risk, legal, procurement and frontline training.	Better decisions and responsible adoption.
Prioritise model flexibility	Use model evaluation, routing, open-source options and fallback paths.	Lower lock-in and better cost/ performance.
Redesign workflows	Do not simply add AI to old processes; redesign roles, SOPs and handoffs.	Greater productivity and transformation value.

12.2 Recommendations for Government and Public Institutions

Recommendation	Action	Expected Outcome
Build AI DPI as shared rails	Connect compute, datasets, model registry, evaluation, shared AI services and sector sandboxes.	Lower cost and faster diffusion across startups, academia, MSMEs and public services.
Use public procurement to shape the stack	Define AI procurement clauses around data rights, audit, model transparency and India-hosted options.	Improved market incentives for responsible and sovereign AI.
Expand AIKosh and sector datasets	Prioritise Indian languages, agriculture, healthcare, education, law, climate and public services.	Better local relevance and reduced dependence on foreign context.
Support compute access	Continue subsidised compute and transparent allocation mechanisms.	Democratised innovation and research capacity.
Create evaluation benchmarks	Develop India-specific benchmarks for language, safety, fairness, legal compliance and sector performance.	More reliable adoption and model selection.
Enable sandboxes	Allow controlled experimentation in healthcare, agriculture, finance, skilling and government service delivery.	Faster learning with guardrails.
Strengthen responsible AI governance	Align with responsible AI principles and risk-based controls.	Trust, inclusion and public accountability.

12.3 Foster a Collaborative Ecosystem

AI Diffusion is a large and complex task. Only widespread Diffusion would benefit all Indians – Enterprises, Governments and Individuals alike. The right path for AI leadership and Diffusion is to foster a healthy collaboration between AI Policy Makers, Developers and Users. An illustration is offered to foster such a collaborative framework.

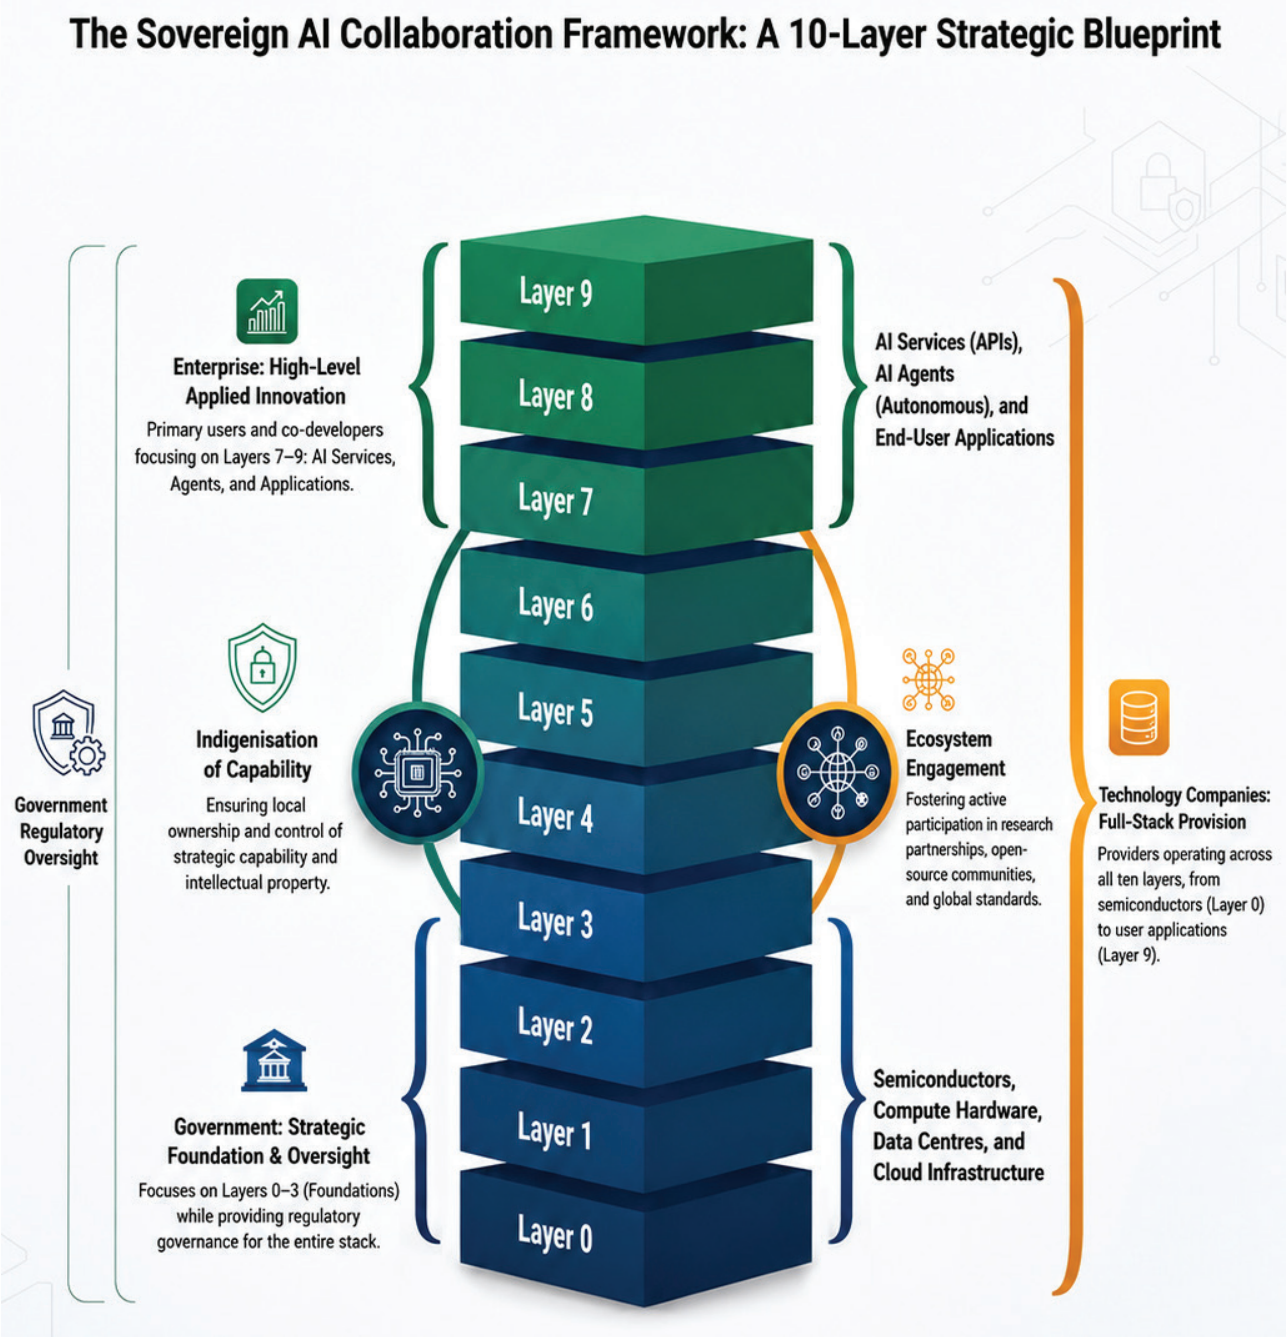

Figure 15. Diffusion as a collaborative approach.

Appendix A

AI Diffusion Matrix Scoring Template

Use this template for each priority use case. The scoring is intentionally simple so that business and technology leaders can complete it together.

Layer	Maturity	Criticality	Dependency Risk	Value Potential	Action Horizon	Owner
L9 Applications	1-6	1-5	1-5	1-5	Immediate / 6 months / 12 months	Named owner
L8 AI Agents	1-6	1-5	1-5	1-5	Immediate / 6 months / 12 months	Named owner
L7 AI Services	1-6	1-5	1-5	1-5	Immediate / 6 months / 12 months	Named owner
L6 Inference Infrastructure	1-6	1-5	1-5	1-5	Immediate / 6 months / 12 months	Named owner
L5 Models	1-6	1-5	1-5	1-5	Immediate / 6 months / 12 months	Named owner
L4 Datasets	1-6	1-5	1-5	1-5	Immediate / 6 months / 12 months	Named owner
L3 Cloud Infrastructure	1-6	1-5	1-5	1-5	Immediate / 6 months / 12 months	Named owner
L2 Data Centres	1-6	1-5	1-5	1-5	Immediate / 6 months / 12 months	Named owner
L1 Compute Hardware	1-6	1-5	1-5	1-5	Immediate / 6 months / 12 months	Named owner
L0 Semiconductors	1-6	1-5	1-5	1-5	Immediate / 6 months / 12 months	Named owner

Appendix B

Value Realisation Worksheet

Worksheet Field	Guidance
Use case name	Describe the exact workflow or decision to be improved.
Business owner	Name the person accountable for the outcome.
Baseline metric	Record the current cycle time, cost, quality, revenue or risk level.
Target improvement	Define an evidence-based improvement target and timeframe.
AI intervention	Describe the application, service, model, data and workflow change.
Measurement design	Use A/B test, pre/post comparison, matched cohort, process mining or expert validation.
Control requirements	Specify human oversight, privacy, audit and model risk controls.
TCO	Include licenses, cloud/inference, data preparation, integration, training and governance.
Scale gate	Define the minimum outcome required before wider deployment.
Sovereignty threshold	Define which layers must be India-hosted, internally controlled, auditable or portable.

Appendix C

Detailed Assessment Questions

The original workbook asks questions across current status, gaps, indigenisation quotient and call to action. The following condensed version adapts the structure for executive workshops.

Awareness

Assessment Area	Workshop Question
Current status	What is our current understanding and maturity for awareness across the use case and stack layers?
Gaps	What information, capability, data, policy, partner or process gaps prevent progress on awareness?
Indigenisation quotient	Which parts of this dimension require domestic, India-hosted, internal or contractually assured control?
Call to action	What precise intervention, owner, budget and timeline should follow?

Availability

Assessment Area	Workshop Question
Current status	What is our current understanding and maturity for availability across the use case and stack layers?
Gaps	What information, capability, data, policy, partner or process gaps prevent progress on availability?
Indigenisation quotient	Which parts of this dimension require domestic, India-hosted, internal or contractually assured control?
Call to action	What precise intervention, owner, budget and timeline should follow?

Access / Appreciation

Assessment Area	Workshop Question
Current status	What is our current understanding and maturity for access / appreciation across the use case and stack layers?
Gaps	What information, capability, data, policy, partner or process gaps prevent progress on access / appreciation?
Indigenisation quotient	Which parts of this dimension require domestic, India-hosted, internal or contractually assured control?
Call to action	What precise intervention, owner, budget and timeline should follow?

Readiness

Assessment Area	Workshop Question
Current status	What is our current understanding and maturity for readiness across the use case and stack layers?
Gaps	What information, capability, data, policy, partner or process gaps prevent progress on readiness?
Indigenisation quotient	Which parts of this dimension require domestic, India-hosted, internal or contractually assured control?
Call to action	What precise intervention, owner, budget and timeline should follow?

Adoption Guardrails

Assessment Area	Workshop Question
Current status	What is our current understanding and maturity for adoption guardrails across the use case and stack layers?
Gaps	What information, capability, data, policy, partner or process gaps prevent progress on adoption guardrails?
Indigenisation quotient	Which parts of this dimension require domestic, India-hosted, internal or contractually assured control?
Call to action	What precise intervention, owner, budget and timeline should follow?

Value Realisation

Assessment Area	Workshop Question
Current status	What is our current understanding and maturity for value realisation across the use case and stack layers?
Gaps	What information, capability, data, policy, partner or process gaps prevent progress on value realisation?
Indigenisation quotient	Which parts of this dimension require domestic, India-hosted, internal or contractually assured control?
Call to action	What precise intervention, owner, budget and timeline should follow?

Sustainability

Assessment Area	Workshop Question
Current status	What is our current understanding and maturity for sustainability across the use case and stack layers?
Gaps	What information, capability, data, policy, partner or process gaps prevent progress on sustainability?
Indigenisation quotient	Which parts of this dimension require domestic, India-hosted, internal or contractually assured control?
Call to action	What precise intervention, owner, budget and timeline should follow?

Ecosystem Engagement

Assessment Area	Workshop Question
Current status	What is our current understanding and maturity for ecosystem engagement across the use case and stack layers?
Gaps	What information, capability, data, policy, partner or process gaps prevent progress on ecosystem engagement?
Indigenisation quotient	Which parts of this dimension require domestic, India-hosted, internal or contractually assured control?
Call to action	What precise intervention, owner, budget and timeline should follow?

Appendix D

Glossary

Term	Meaning
AI Diffusion	The process by which AI capabilities move from awareness and experimentation into routine use, scaled operations and strategic transformation.
AI Sovereignty	Independent authority over critical AI development, deployment and dependency decisions.
AI DPI	Digital public infrastructure for AI: shared, interoperable rails for data, compute, models, services, governance and trust.
Agent	An AI-enabled software system that can plan and execute tasks using tools under policy constraints.
Inference	The runtime process of using a trained model to produce outputs for users or systems.
RAG	Retrieval augmented generation: a pattern that grounds model outputs in retrieved enterprise or external knowledge.
Model Router	A runtime component that chooses between models based on cost, latency, quality and risk.
Indigenisation Quotient	A measure of how much strategic capability, control or IP should be domestic or internal at a given layer.
LLMOps	Operational practices for deploying, monitoring, evaluating and governing large language model applications.
Value Realisation	The discipline of linking AI activity to measurable business outcomes, financial impact and strategic capability.

Notes

Notes

Notes