EPINOVA–2026–RR-01

India’s AI Power in a Distributed and Uneven Technological Order:

Boundary Control, Strategic Autonomy, and Cognitive Security

Dr. Shaoyuan Wu

September 1, 2026

Global AI Governance and Policy Research Center
EPINOVA LLC

About This Report

This report examines India’s emerging AI power through the interaction of usable national capacity, technological dependence, boundary control, strategic autonomy, and cognitive security. Rather than treating AI power as equivalent to frontier-model performance, compute scale, or investment volume, it analyzes how technical, institutional, organizational, and infrastructural capabilities are converted into nationally usable power under conditions of uneven interdependence.

India is used as a mechanism-rich case because substantial software capacity, digital public infrastructure, multilingual public technology, expanding compute access, and globally distributed technical networks coexist with continued dependence on externally controlled upstream technologies. The report therefore distinguishes access from mobilizable capability, territorial presence from control, boundary power from boundary autonomy, and persistent systemic influence from narrower bilateral effects. It also examines how procurement, regulation, interoperability, substitution, and institutional conversion shape India’s ability to preserve agency without requiring full technological autarky.

The report is based on publicly available evidence through August 31, 2026. It is intended as a theory-guided analytical study rather than a national AI ranking or predictive assessment. Its broader objective is to clarify how states can build, exercise, and preserve AI-related agency when ownership, infrastructure, knowledge, and technological control are distributed across national and transnational systems.

Table of Contents

Executive Summary 02

Key Findings 04

1. Introduction: AI Power and the Indian Position 05

2. Research Design and Evidence Strategy 08

3. India’s Usable National AI Capacity 10

3.1 Territorial-Infrastructural Capacity 11

3.2 Human-Epistemic Capacity 12

3.3 Economic-Industrial Capacity 12

3.4 Semiconductor Knowledge-to-Industry Conversion 13

3.5 Political-Administrative Mobilization and Cultural-Linguistic Fit 14

4. A Vertically Uneven AI System 15

5. Nationally Mobilizable Capacity and the Wider India-Linked Technological Footprint 17

6. Conversion: Turning Distributed Capability into Usable Power 19

7. India’s Selective AI Boundary Power 20

7.1 Payment-Data Localization as Demonstrated ABP 20

7.2 Public Procurement, Platform Regulation, and Interoperability 20

8. Boundary Autonomy: Strategic Agency Under Interdependence 23

9. From Boundary Effects to Systemic Influence 25

9.1 Adjacent DPI and Architecture Diffusion 25

9.2 AI Governance and Standards 26

10. Cognitive Security as an Illustrative Stress-Test Lens 29

11. Strategic Implications and India’s AI-Power Trajectory to 2030 33

11.1 Three Policy Priorities 33

11.2 Five Conversion Tests for 2030 33

12. Conclusion 35

Appendix A. Qualitative Codebook and Case Evidence Matrix 36

References 40

Author and Institutional Affiliation 44

Disclaimer 44

Version History 44

Citation Recommendation 45

Executive Summary

India’s emerging AI power cannot be understood through frontier-model rankings, compute totals, or investment announcements alone. Its position is defined by a more uneven structure: strong software and organizational capacity, extensive digital public infrastructure, multilingual public technology, a large technical workforce, and expanding domestic compute coexist with continued dependence on externally controlled upstream technologies such as advanced accelerators, semiconductor equipment, electronic design automation, high-performance memory, cloud-management layers, and parts of the frontier-model ecosystem.

This report argues that India’s AI position is best understood through four distinct analytical objects: usable national capacity, relational boundary power, boundary autonomy, and systemic influence. Usable national capacity concerns what the Indian national system can actually mobilize and sustain. Boundary power concerns the ability to alter the conditions under which other actors access, transfer, integrate, or use a resource or function. Boundary autonomy concerns India’s own ability to switch, bypass, substitute, diversify, or selectively internalize when another actor controls a critical interface. Systemic influence is treated as a later and more demanding outcome requiring persistent and sufficiently broad effects across networks, rules, or future participation conditions.

The evidence shows that India’s strongest capabilities lie downstream and institutionally: software integration, Global Capability Centres, digital public infrastructure, public administration, multilingual AI, and a population-scale domestic market. IndiaAI demonstrates how externally sourced compute can be reorganized through public procurement, subsidy, empanelment, and allocation into nationally usable capability without transferring ownership of the underlying accelerator technology. BHASHINI demonstrates a large operational multilingual public-AI layer, showing that linguistic inclusion depends not simply on model availability but on institutional infrastructure capable of supporting multiple languages at scale.

India’s semiconductor position illustrates the limits of this conversion. The country has substantial chip-design talent and expanding domestic design and manufacturing programmes, but many critical tools, fabrication processes, memory systems, and equipment remain externally controlled. This makes India’s AI system vertically uneven: knowledge and deployment capability are deeper than full-stack technological control. Progress at one layer can reduce one dependency while revealing another.

The report also finds that India’s strongest demonstrated boundary power does not lie at the most technologically advanced upstream layers. It is concentrated instead in interfaces where Indian institutions possess enforceable jurisdiction or control over participation conditions. Payment-data localization provides the clearest case, where formal rules were followed by restrictions, compliance, and restored access. IndiaAI creates an AI-specific procurement boundary through empanelment, pricing, eligibility, and service obligations. Platform regulation establishes formal duties for synthetic media and intermediary participation, although realized behavioral effects remain less clearly demonstrated. UPI interoperability illustrates a non-coercive form of boundary control in which independently governed systems jointly define the conditions of cross-border interaction.

Executive Summary

These findings suggest that strategic autonomy under AI interdependence should not be equated with technological autarky. The more realistic objective is selective autonomy: preventing failure at a single external interface from disabling critical national functions. This can be achieved through provider diversification, workload portability, domestic capability development, trusted partnerships, open standards, fallback options, redesign, and selective internalization. Autonomy therefore depends less on owning every technological layer than on preserving credible alternatives and the ability to reorganize functions when conditions change.

India's digital public infrastructure offers a second pathway to influence. The MOSIP and broader DPI experience suggests that architecture can diffuse internationally without the originating state retaining control over foreign systems. This is analytically distinct from platform-based dependence. Open, customizable, and interoperable architectures can spread standards and institutional practices while leaving operational ownership with adopting governments. India's AI-specific systemic influence remains less mature than its DPI influence, but summit diplomacy, standards participation, and emerging governance initiatives indicate a developing pathway toward wider institutional reach.

Cognitive security is used in this report as a stress-test of the wider framework. In an AI-mediated information environment, meaningful resilience requires more than synthetic-content detection. It depends on multilingual verification, provenance, authentication, continuity, substitutability, redundant infrastructure, protected language resources, and the ability to operate when a major model API, cloud provider, or platform becomes unavailable. At the same time, such systems must preserve privacy, due process, proportionality, contestability, and independent oversight. Cognitive resilience is therefore both a technical and institutional problem.

The central conclusion is that India's AI power is neither reducible to technological sovereignty nor accurately described as simple dependence. It is a mixed-control system in which externally sourced resources are repeatedly converted into nationally usable capabilities, while selected interfaces are governed through regulation, procurement, interoperability, and public infrastructure. The durability of India's AI power will therefore depend on how effectively it deepens conversion, reduces catastrophic dependence at critical chokepoints, preserves substitutability, strengthens multilingual and public-service capacity, and retains institutional control over the systems through which AI becomes economically, administratively, and socially consequential.

Key Findings

- **India’s AI power is vertically uneven.** Strong software, digital public infrastructure, multilingual AI, technical talent, and expanding compute coexist with dependence on externally controlled upstream technologies.
- **Access does not equal usable power.** IndiaAI shows that externally sourced compute can be converted into nationally organized capability through procurement, subsidy, allocation, and service rules without transferring ownership of the underlying technology.
- **Boundary power is strongest where India controls participation conditions.** Payments regulation, public procurement, platform rules, and interoperability provide clearer leverage than frontier accelerators or semiconductor equipment.
- **Strategic autonomy does not require full technological self-sufficiency.** The more realistic objective is to preserve switching, substitution, diversification, fallback, and selective internalization at critical interfaces.
- **Digital public infrastructure can diffuse without creating foreign control.** India’s DPI and MOSIP experience suggests that open and customizable architectures can extend institutional influence while leaving operational ownership with adopting states.
- **Cognitive resilience depends on both technology and governance.** Multilingual verification, provenance, continuity, alternative providers, and protected public infrastructure must be combined with privacy, contestability, due process, and oversight.
- **India’s long-term AI advantage will depend on conversion.** The decisive issue is whether technical access and distributed capabilities can be transformed into durable, governable, and resilient national capacity.

1. Introduction

Artificial intelligence is becoming increasingly consequential to economic production, public administration, scientific research, military systems, information environments, and strategic competition. Yet the indicators commonly used to describe an "AI race"—model benchmarks, investment totals, compute capacity, data-center construction, research output, platform reach, regulation, and public-sector adoption—refer to different resources, mechanisms, actors, and outcomes. This report therefore distinguishes four analytical objects within the report itself. AI National Power (ANP) is the usable capacity of a state or national system to generate, sustain, mobilize, govern, and convert AI-relevant resources into effects. AI Boundary Power (ABP) is relational: a controller changes the conditions under which a target obtains, transfers, integrates, deploys, or continues using a resource or function at a specified interface. Boundary autonomy is the target's ability to switch, bypass, substitute, or internalize. Systemic influence is a later outcome requiring persistent and sufficiently broad change in option sets, network topology, or rules of future participation. These distinctions build on, but are not reducible to, work on weaponized interdependence, regulatory market power, and technological sovereignty (Bradford, 2020; Edler et al., 2023; Farrell & Newman, 2019; Tanner et al., 2026; Wu, 2026).

The distinction is particularly important for India. AI is not making India strategically relevant from a previously marginal position; rather, it is revaluing structural characteristics that India already possessed: a population-scale market, a large software and technology-services sector, deep engineering capability, globally distributed technical networks, extensive digital public infrastructure (DPI), and a long-standing preference for strategic autonomy under interdependence. India-specific scholarship already describes its digital-sovereignty strategy as uneven, state-led, and simultaneously constrained by dependence on foreign platforms and technologies (Jiang, 2024; Kaushik, 2026). AI intensifies both sides of that structure by increasing the value of domestic organizational and institutional capabilities while raising the strategic cost of dependence on advanced semiconductors, accelerators, cloud platforms, and parts of the frontier-model ecosystem.

India consequently poses an attribution problem that conventional national rankings handle poorly. Some of the most powerful technologies physically deployed inside India are foreign-owned or depend on externally controlled upstream layers. Conversely, Indian firms, professionals, research links, and technology networks extend far outside national territory. Geography, ownership, jurisdiction, and usable capability do not coincide. The framework used here is therefore applied not as a country leaderboard but as a discipline for asking what the national system can actually mobilize, where India can change the crossing conditions faced by others, where it can switch or substitute when others control the boundary, and when those effects become sufficiently persistent and broad to alter wider networks or rules (Wu, 2026).

1. Introduction

This report applies the framework developed in Wu (2026) as an analytical lens; it does not attempt a causal test or formal empirical validation of the framework itself. Its contribution is threefold. First, it operationalizes the distinction among state-side usable capacity, relational boundary control, target-side autonomy, and later system-shaping outcomes in a mixed-control national case. Second, it introduces an attribution discipline for India by separating nationally mobilizable AI capacity from a wider India-linked technological footprint. Third, it identifies conversion across ownership, geography, and organizational boundaries as the mechanism linking India's highly dispersed resource environment to usable capability. In doing so, the report adds an empirical application to existing accounts of network coercion, regulatory reach, technological sovereignty, and AI power by showing how a state can gain usable capability without full ownership and diffuse institutional architecture without retaining direct control.

For attribution purposes, the report distinguishes nationally mobilizable AI capacity (retaining the shorthand "India Power" only where useful) from a wider India-linked technological footprint (termed "Indian Power" in the mapping discussion). The first refers to AI-relevant capabilities, organizational resources, fallback options, and relational leverage demonstrably available to or exercisable through India's national economic, technological, administrative, or institutional system. The second describes a broader transnational footprint associated with Indian-headquartered firms, Indian-origin professionals, and India-linked research or organizational networks. It is not a claim of political loyalty, state control, or collective strategic agency. Identity and location identify possible connections; they do not establish mobilizability.

The central argument is that India occupies a distributed and vertically uneven position in the emerging AI order. Its national depth is strongest in human and organizational capacity, software and enterprise deployment, digital public infrastructure, market scale, multilingual public technology, and increasingly territorial AI infrastructure, while several consequential upstream technologies remain externally controlled. India's position is therefore highly conversion-intensive: capabilities repeatedly cross boundaries of geography, ownership, and organization before becoming nationally usable. India can internalize capability without fully internalizing ownership and can externalize institutional architecture without retaining foreign control. The evidence on boundary control is layered rather than uniform: payment localization provides the clearest demonstrated effective ABP; IndiaAI procurement provides an AI-specific operational procurement gate; general cloud procurement and newer platform rules establish formal or operational boundary authority whose realized leverage is less certain; and interoperability can produce jointly governed crossing conditions. Cognitive security provides a stress test of whether this configuration can preserve trustworthy judgment and operational continuity in an information environment that India can regulate and shape but cannot realistically expect to control comprehensively under continued technological interdependence.

1. Introduction

Figure 1 summarizes how the framework is applied in this study. ANP, ABP, and boundary autonomy are treated as parallel analytical objects rather than a causal sequence. Systemic influence is a later outcome that may follow from national capacity and/or boundary control under specified network and institutional conditions, while target-side autonomy conditions exposure to external control. Cognitive security is shown as an outer application-domain stress test of how the three analytical objects function together.

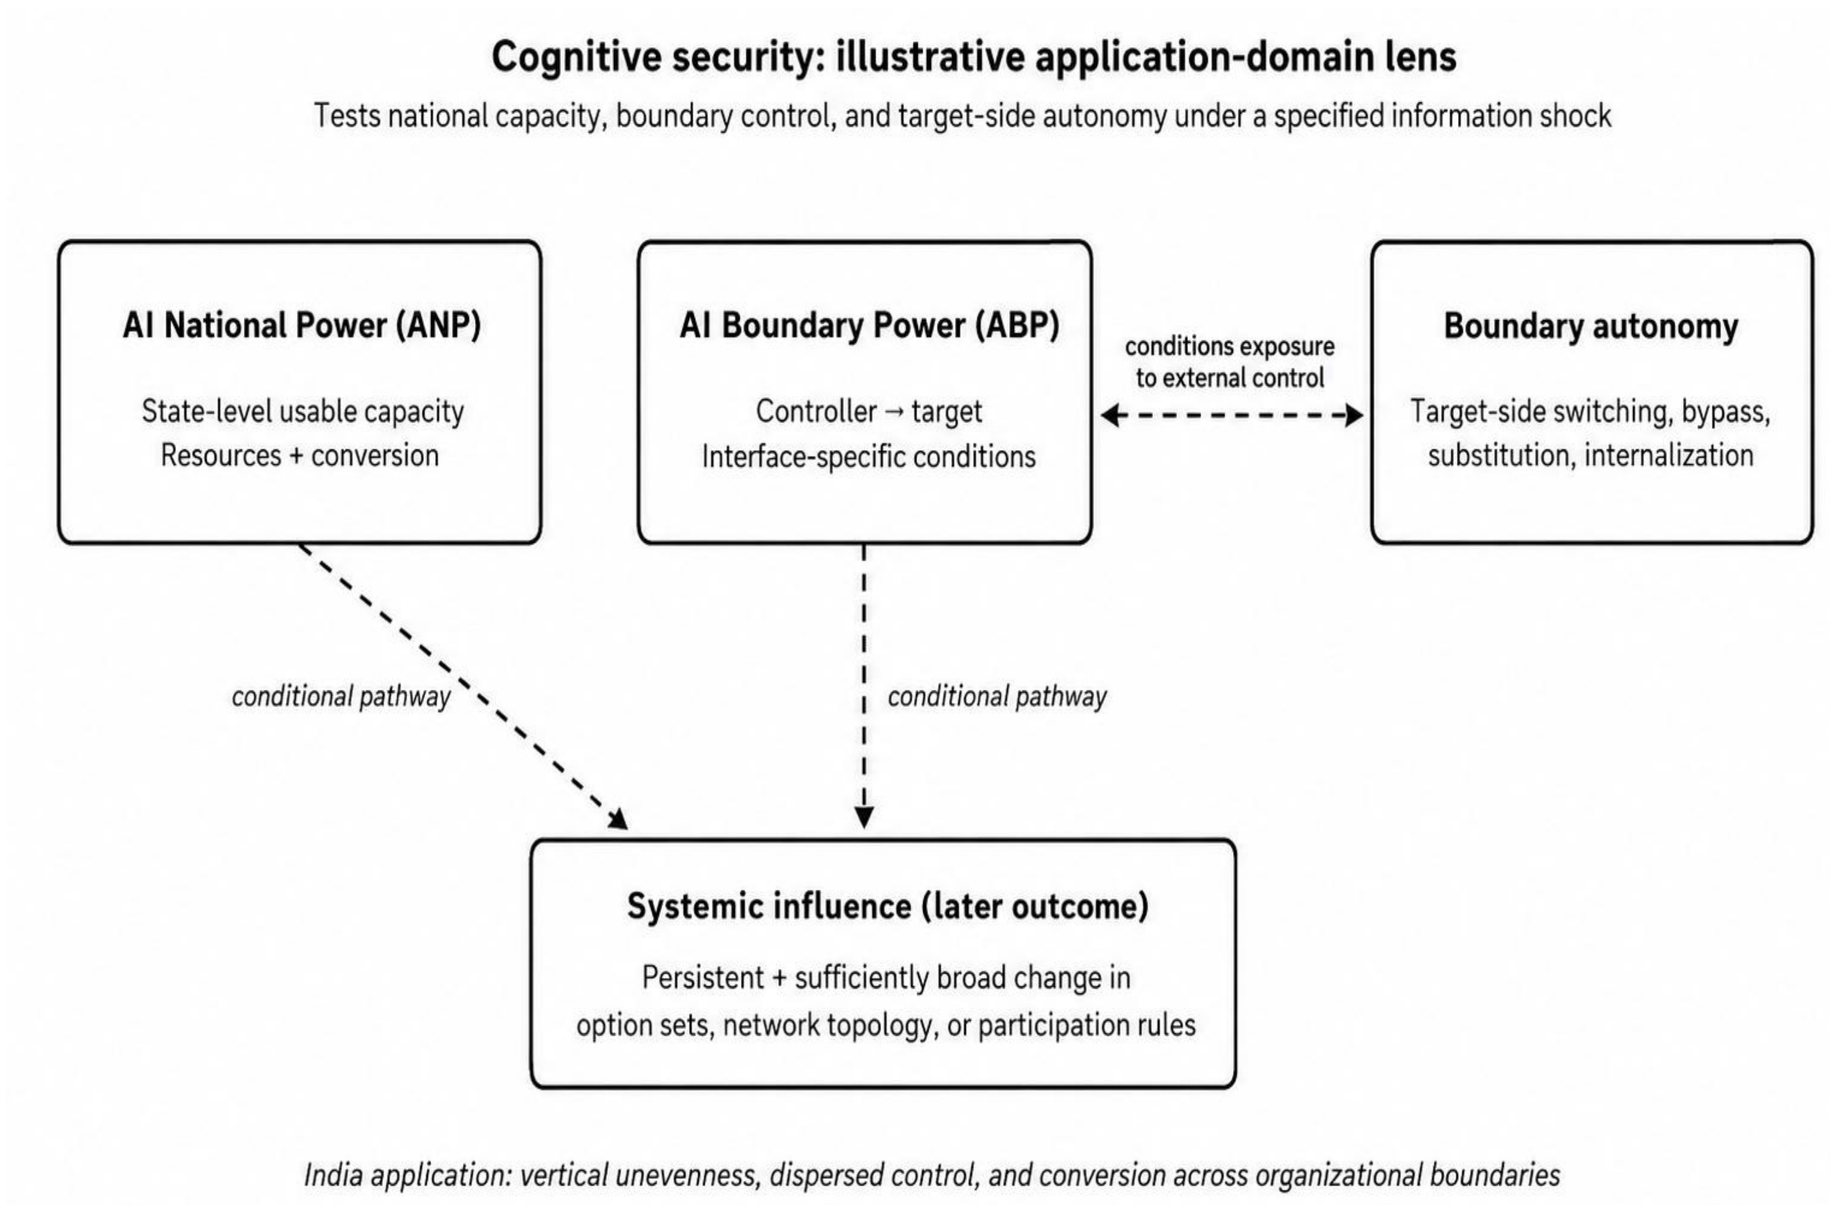

Figure 1. Analytical Application of the AI Power Framework to India

Note: ANP, ABP, and boundary autonomy are distinct analytical objects. Systemic influence is evaluated only as a later outcome when change is persistent and sufficiently broad; cognitive security is an illustrative application-domain lens. Source: Author's synthesis from the definitions and coding rules specified here and Wu (2026).

2. Research Design and Evidence Strategy

This report is a theory-guided, single-country analytical case study rather than a cross-national ranking exercise. India is selected because it combines highly visible forms of mixed control: large domestic markets and public digital systems coexist with foreign-owned cloud and corporate infrastructure, globally distributed technical labor, expanding domestic compute access, and persistent upstream dependence. The purpose is therefore mechanism mapping—identifying where resources become usable national capacity, where Indian institutions can alter another actor's crossing conditions, where India can substitute externally controlled interfaces, and where observed effects become sufficiently broad and persistent to alter wider systems. The study does not estimate a single national AI-power score or claim causal effects from a comparative design.

The unit of analysis changes with the construct. ANP is assessed at the level of the Indian state or national system for a specified sector and period. ABP is coded relationally as controller → target | interface, time: a market, procurement, data, platform, compute, or interoperability relation must identify both the actor exercising control and the actor facing altered conditions. Boundary autonomy is coded from the target side and concerns credible switching, bypass, substitution, or internalization options. Systemic influence is evaluated only within a specified domain and affected actor set and only when an observable effect is both persistent and sufficiently broad. The empirical observation window runs through August 31, 2026; dynamic operating statistics are dated snapshots rather than timeless stocks.

Evidence is weighted by source type. Primary legal, regulatory, parliamentary, central-bank, and ministry sources are used for rules, public programmes, procurement status, and official operating statistics. Corporate first-party sources are used only for claims about a firm's own investments, facilities, or commitments. Peer-reviewed scholarship is used to interpret longer-run mechanisms such as India's software capability formation, diaspora-mediated technology diffusion, digital sovereignty, DPI governance, and multilingual AI. Industry estimates are retained only when their provenance is identified and are explicitly labelled as estimates; high-quality news reporting is used for corroboration or for industry developments not available in public administrative data. No source is treated as supporting a stronger proposition than it directly establishes.

Maturity and attribution are coded conservatively. Announcements, commitments, financing or procurement, construction, operational deployment, usable capability, and mobilizability are treated as distinct states. Location, ownership, jurisdiction, and mobilizability are also coded separately: an India-located asset is not automatically India-controlled, while an overseas asset or network can contribute to national capability only through an observable connection to usable Indian options. Table 2 is therefore an evidence-state map within the India case, not a calibrated comparative index. Boundary control is coded at three evidentiary levels. Formal authority records a legal or contractual gate. Operational boundary control requires evidence that the gate is actually used to allocate access or condition participation under binding rules. Coercive or bargaining leverage is a stronger subset requiring enforcement, renegotiation, or observable target adaptation. Coercion is therefore not required for every ABP relation, but the three levels are not treated as equivalent. Systemic influence requires persistence plus breadth; bilateral durability alone is insufficient.

2. Research Design and Evidence Strategy

The qualitative codebook separates evidence states rather than assigning a composite score. For ANP, “presence/access” means a resource exists or can be purchased; “routine operational availability” requires evidence of active service or deployment; “administrative mobilizability” requires an identifiable authority able to approve, allocate, or direct access under normal conditions; and “disruption resilience” requires evidence that the function can continue, migrate, or substitute under a specified external shock. For ABP, “formal authority” means a rule or contractual gate exists; “operational boundary control” means that the gate is demonstrably applied to participation or access; and “coercive/behavioral leverage” additionally requires observed enforcement or target adaptation. For systemic influence, a single bilateral adaptation does not qualify: persistence must extend beyond the initiating event, and breadth must affect a class of actors, multiple network relations, or future participation rules. These labels are ordinal evidence categories, not metric scores.

Cases are included when they illuminate at least one of these constructs and have traceable public evidence of the relevant resource, rule, transaction, deployment, or response. Cases are excluded from stronger coding when public sources establish only announcements, general market size, nationality, or organizational location. “No public evidence identified” is not treated as proof that a capability or relationship does not exist; it records the limit of the public-source review. Appendix A provides the case-level evidence matrix, observation period, maturity judgment, alternative interpretation, and confidence used in the main tables.

The study has corresponding limitations. It relies on publicly observable evidence and does not have access to confidential corporate IP arrangements, internal government contingency plans, crisis-time contracts, or firm-level switching costs. Most claims about mobilizability therefore concern normal-time access or administrative allocation rather than demonstrated resilience under severe disruption. Where continuity under shock has not been observed, the report labels it as a condition to be tested rather than as an established capability. Several 2026 programmes are still evolving, so current capability is kept separate from future intent and the 2030 discussion is treated as scenario monitoring rather than prediction.

3. India’s Usable National AI Capacity

The starting point is not the total quantity of AI-related assets physically located in the country but the capabilities the national system can actually call upon and sustain. The framework specified here organizes AI National Power by analytical role: territorial-infrastructural, human-epistemic, and economic-industrial resource bases; political-administrative mobilization as a conversion mechanism; and cultural-linguistic fit and normative-strategic orientation as conditioning variables (Wu, 2026). **Table 1** applies this architecture to the Indian evidence.

Table 1. Role-Based Evidence Map for India’s AI National Power

Analytical Role	Indian Evidence	What It Demonstrates	Key Limitation
Territorial-infrastructural	Data centers; IndiaAI shared compute; NIC/MeghRaj; cloud regions	Capacity to host, connect, and allocate AI workloads	India-located does not equal India-owned or fully substitutable
Human-epistemic	Engineers; semiconductor designers; universities; global professional networks	Large knowledge-production and reproduction base	Identity/workforce location does not establish IP ownership or mobilizability
Economic-Industrial	Software/IT services; Global Capability Centres (GCCs); startups; domestic demand; semiconductor projects	Ability to commercialize and integrate AI	Foreign corporate ownership and global clients create mixed control
Political-administrative Mobilization	IndiaAI; compute subsidy; public AI pipeline; procurement; regulation	Converts dispersed resources into usable capacity	Announcements and problem statements are not deployment outcomes
Cultural-linguistic Fit	Multilingual society; BHASHINI	Conditions usability and legitimacy across languages	Diversity is not automatically a positive resource
Normative-strategic Orientation	Strategic autonomy; data rules; interoperability; Global South diplomacy	Shapes which dependencies are tolerated or internalized	Orientation is a moderator, not a monotonic power score

Note. Rows identify analytical roles rather than additive dimensions. The same asset can contribute differently under normal commercial conditions and crisis scenarios.

Source: Author’s synthesis based on Wu (2026), Press Information Bureau (2026b, 2026d, 2026g, 2026i, 2026k), and Ministry of Electronics and Information Technology (2026a).

3. India’s Usable National AI Capacity

3.1 Territorial-Infrastructural Capacity

India’s physical AI infrastructure has expanded rapidly. Government data report that installed data-center capacity increased from approximately 375 MW in 2020 to about 1,575 MW by August 2026, while explicitly identifying AI and high-performance computing as important drivers of demand (Press Information Bureau [PIB], 2026e). This strengthens India as a physical location for computation by expanding domestic hosting, network proximity, and jurisdictional presence. It does not, however, measure sovereign AI compute. Much commercial capacity is privately operated, and foreign cloud, accelerator, software, and supply-chain dependencies continue to shape what can actually be mobilized.

The IndiaAI shared-compute programme illustrates this distinction. Shared capacity exceeded 45,000 GPUs by June 2026. An August 6 administrative release reported 237 projects approved for subsidized compute support and 93.18 lakh (9.318 million) GPU-hours sanctioned; a later August 13 factsheet described the same projects and hours as having “accessed” subsidized compute (PIB, 2026d, 2026k). Because the public releases do not provide a common usage ledger establishing actual consumption, this report uses the more conservative approved/sanctioned wording. The programme had also selected 20 indigenous foundation-model proposals, while AI Kosh hosted more than 14,000 datasets and 331 models (PIB, 2026k). These numbers demonstrate nationally organized access and administrative allocation, not state ownership of the underlying accelerator stock or crisis-time availability.

A more directly public layer is also emerging. By August 2026, a purchase order had been issued for an approximately 1.1-EFLOPS high-performance AI compute system at the National Informatics Centre facility in Shastri Park, Delhi (PIB, 2026d). Procurement status is not operational capacity, so the system should not yet be counted as an activated 1.1-EFLOPS national asset. This measurement discipline is central to the report: procurement, installation, commissioning, operational availability, and crisis-time mobilizability are distinct states.

Taken together, these developments show an uneven relationship between territorial hosting and technological control. India has substantially expanded the domestic geography of compute and data-center capacity, while important chips, software layers, cloud-management systems, and manufacturing inputs remain externally concentrated. The available evidence supports a divergence between territorial presence and control; it does not by itself establish that one is increasing at a statistically faster rate than the other.

3. India’s Usable National AI Capacity

3.2 Human-Epistemic Capacity

India’s human-epistemic position is considerably deeper than its frontier-model output alone would suggest. In March 2026, the government reported an industry estimate that India employed nearly 20% of the global semiconductor chip-design workforce and hosted about 7% of semiconductor-domain GCCs (PIB, 2026g). The figure should not be interpreted as India controlling one-fifth of global chip design. Many engineers work inside multinational organizational structures in which design tools, intellectual property, product decisions, and fabrication remain external. It nevertheless indicates a large reservoir of tacit technical knowledge embedded in the Indian economy.

The geography of Indian technical knowledge extends much farther. The National Science Board’s The State of U.S. Science and Engineering 2026 reports, using 2024 workforce data, that India-born workers accounted for 29% of all foreign-born workers in U.S. science and engineering occupations (National Science Board, 2026). The denominator is important: the figure does not mean that India-born workers comprise 29% of the total U.S. S&E workforce. It demonstrates scale and frontier-system presence, not Indian state control. The UK Migration Advisory Committee likewise found that almost 94% of IT professionals using the Global Business Mobility route came from India and that the route was used predominantly by large Indian multinationals to transfer staff (Migration Advisory Committee, 2025). The U.S. evidence primarily captures professionals embedded in host institutions; the UK evidence more directly captures India-headquartered corporate organization across borders.

These patterns are why the report does not equate diaspora size with national power. Overseas capability contributes to usable Indian capacity only where there is an observable pathway—research collaboration, corporate R&D, capital, client relationships, foreign affiliates, standards communities, or talent circulation. Identity identifies a possible network edge; it does not establish that the edge is usable by the Indian national system.

3.3 Economic-Industrial Capacity

India’s most established AI-relevant economic strength lies downstream from the frontier-model race. The country possesses a dense software, IT-services, engineering, and enterprise-integration ecosystem. Long-run research on the Indian software industry emphasizes service capabilities, process learning, and organizational adaptation under resource constraints, providing historical context for the shift from routine services toward more complex engineering and product functions (Arora et al., 2001; Athreye, 2005). A NASSCOM–Zinnov assessment reported in May 2026 estimated that India would host 2,117 Global Capability Centres (GCCs) and about 2.36 million GCC professionals in FY2025–26, with expected sector revenue of US$98.4 billion (Ishwarbharath & Suresh, 2026). These are industry estimates for the GCC sector as a whole, not audited AI-specific revenue statistics.

3. India’s Usable National AI Capacity

This layer is strategically important because much of AI’s economic value emerges after model development. Models must be integrated into enterprise software, data systems, compliance processes, cybersecurity, industrial workflows, and domain-specific applications. India possesses a deep organizational conversion layer for these tasks. Yet GCCs remain mixed-control assets: labor, managerial experience, supplier relationships, and ecosystem learning can contribute substantially to India, while IP ownership, capital allocation, product strategy, and ultimate corporate authority often remain with foreign parents. They are best treated as India-located organizational capability with partial and conditional national mobilizability.

AI is also repricing the internal composition of this advantage. Reuters reporting based on executives and industry data indicates that GCC hiring is becoming more selective as routine and entry-level roles face automation pressure while demand rises for advanced AI, cybersecurity, and specialized technical skills (Venugopal & Roy, 2026). A separate Reuters analysis reports that some India-based centres now lead end-to-end product, commercial, and R&D functions, while firms also hedge through “India-plus” locations (Monnappa, 2026). The strategic question is therefore not whether India can preserve every traditional IT role, but whether its organizational position shifts toward high-value functions that remain difficult to substitute.

3.4 Semiconductor Knowledge-to-Industry Conversion

Semiconductors provide the clearest example of a gap between human-epistemic depth and industrial control. India already participates deeply in global chip design, but much of that activity depends on externally controlled electronic design automation (EDA) software, IP libraries, foundries, process technologies, manufacturing equipment, and advanced memory. The Design Linked Incentive scheme nevertheless shows an active conversion pathway. By August 2026, DLI-backed companies had cumulatively raised more than US$100 million in venture funding and achieved 35 design tape-outs (PIB, 2026c). The strategic significance is not the tape-out count alone; it is the sequence from engineering expertise to domestic firms, capital formation, owned designs, silicon validation, and potential products.

Manufacturing depth is also rising but should be coded by maturity. By July–August 2026, 12 semiconductor manufacturing projects had been approved with investment commitments exceeding ₹1.64 lakh crore; three facilities had begun commercial production, while the approved portfolio included one silicon fab and multiple compound-semiconductor and packaging facilities. Semicon 2.0 was approved with an outlay of ₹1,27,500 crore and explicit pillars covering design, equipment and materials, fabrication, advanced packaging, R&D, and talent (PIB, 2026j). These figures demonstrate expanding industrial conversion, not advanced-node self-sufficiency. Approved projects, operating packaging plants, and future fabs are technologically different states.

3. India’s Usable National AI Capacity

The semiconductor case therefore captures the wider structure of India’s AI position: knowledge penetration is deeper than full-stack production control, while state policy attempts to convert that knowledge vertically into domestic firms, products, manufacturing, and supply-chain resilience. Successful conversion at one layer can reveal the next dependency rather than eliminate dependence altogether.

3.5 Political-Administrative Mobilization and Cultural-Linguistic Fit

Political-administrative mobilization matters because assets become national capability only when institutions can coordinate, procure, deploy, and sustain them. In July 2026, a government-wide exercise identified 762 AI problem statements across 62 ministries and departments; this figure represents a discovery pipeline, not deployed systems (PIB, 2026i). Separately, an August IndiaAI progress summary reported 62 AI prototypes and 20 AI-solution deployments (PIB, 2026k). The public releases do not establish that the 62 prototypes and 20 deployments are a subset of the 762 problem statements, so no cohort conversion rate is inferred. More specific operational evidence exists for some direct AI cases: two teams had begun deployment of an AI face-authentication solution at the Union Public Service Commission (UPSC), and two intelligent-document-processing teams were being deployed there as implementation proceeded (PIB, 2026i). These examples support capability activation at particular interfaces without justifying a single 762 → 62 → 20 funnel.

Language offers a parallel example of conversion. India’s multilingualism is not automatically an AI-power asset; it increases the cost of model coverage, evaluation, verification, and public-service delivery. The asset lies in the capability developed to operate across that environment. By August 2026, BHASHINI supported 36 Indian text languages and 23 Indian voice languages, powered more than 800 government websites, processed more than 24 million AI inferences per day, and had exceeded nine billion cumulative inferences (PIB, 2026b). This is evidence of an operational public language-AI layer, not evidence that all languages have equal quality or that multilingual complexity has been fully solved.

India therefore possesses substantial AI-relevant depth without possessing a uniformly sovereign technology stack. Infrastructure, human capital, software, public digital systems, and administrative conversion mechanisms are large and increasingly usable, while ownership, upstream control, and crisis-time substitutability remain uneven. The appropriate question is not whether India has “high” AI power in aggregate, but where capability is deep, where it is conditional, and where external boundaries remain difficult to replace.

4. A Vertically Uneven AI System

A single aggregate AI-power rank would obscure more than it reveals. The framework specified here treats cross-state placement as requiring common sectors, target sets, time windows, and aggregation rules (Wu, 2026). India is therefore represented through a layer-by-layer diagnosis. **Table 2** distinguishes availability and capability from national control and external exposure.

Table 2. Evidence-State Map of India's Vertically Uneven AI Stack

AI Layer	Current Position	Control/Ownership Pattern	Strategic Interpretation
Technical human capital	Established large-scale capability	Substantial domestic base + global circulation	Core strength; attribution depends on organizational location
Software / enterprise integration	Established large-scale capability	Large Indian firms + global clients/platforms	Core downstream conversion strength
Global Capability Centres (GCCs) / corporate R&D	Established large-scale India-located capability	Mixed; often foreign parent control	Major conversion asset, not automatic sovereign capacity
Digital public infrastructure (DPI)	Established large-scale institutional capability	Relatively high domestic institutional control	Core institutional and interface strength
Multilingual public AI	Established and expanding	High public-layer control; compute/model dependencies remain	Distinctive capability built in response to complexity
Data-center hosting	Rapidly expanding	Mixed domestic/foreign ownership	Territorial strength, not sovereign compute
Cloud service availability	Broad availability	Mixed; large foreign platform role	Externally embedded but institutionally diversified
AI compute access	Rapidly expanding access	Allocation/jurisdiction stronger than hardware ownership	Access strength with upstream exposure
Domestic foundation models	Emerging / expanding	Rising domestic control	Emerging autonomy; frontier parity not established
Semiconductor design / EDA-dependent workflows	Established capability	Mixed IP ownership; large talent base	Knowledge strength with conversion opportunity
Semiconductor fabrication	Emerging / expanding	Domestic projects + external tools/processes	Industrial conversion gap
Frontier accelerators/equipment	Limited domestic control	Highly external and concentrated	Major upstream chokepoint

Note: The second column reports descriptive evidence states within the India case, not comparative rankings or validated score thresholds. Control, ownership, normal-time availability, administrative mobilizability, and disruption resilience remain separately assessed.

Source: Author's synthesis based on PIB (2026b, 2026c, 2026g, 2026j, 2026k), Ishwarbharath and Suresh (2026), and MeitY (2026a).

4. A Vertically Uneven AI System

India’s strongest capabilities form a downstream and organizational cluster: engineering, software, IT integration, GCCs, digital public infrastructure, public administration, language technology, and a population-scale domestic market. This gives India an ability to absorb and operationalize AI at scale even where it does not originate every underlying technology. The middle layers—data centers, cloud, compute access, and model development—are characterized less by absence than by mixed control: Indian jurisdiction and allocation coexist with foreign platform, accelerator, or software dependence.

The deepest external dependencies remain disproportionately upstream. Advanced AI accelerators, leading-edge semiconductor fabrication, lithography, specialized equipment, EDA, high-performance memory, and related supply chains remain substantially concentrated outside India. Rapid downstream AI adoption can therefore increase rather than immediately reduce strategic exposure by making more national activity dependent on a smaller number of upstream interfaces.

The measurement problem is therefore one of interpretation rather than a claim that every existing indicator is incorrectly designed. When frontier-output indicators are interpreted as measures of overall national AI power, they can omit organizational, administrative, deployment, and transnational capability. Conversely, territorial asset counts or participation indicators can overstate national control if India-located data centers, foreign corporate R&D, shared access to foreign GPUs, or Indian-origin participation abroad are treated as fully nationally mobilizable. The framework is useful because it forces both attribution errors to be considered explicitly.

5. Nationally Mobilizable Capacity and the Wider India-Linked Technological Footprint

The Indian case becomes still more complicated once geography is separated from organizational control. A significant share of AI-relevant capability inside India is owned or directed by foreign parents, while Indian-headquartered firms and India-origin professionals participate in technology systems outside the country. For analytical clarity, this section uses "nationally mobilizable capacity" as the primary term and retains the shorter label India Power only as shorthand. The wider India-linked technological footprint captures organizational and professional connections that may matter for India but are not attributed to the national system without evidence of access, mobilizability, jurisdiction, and durability.

	Primary organizational control: Indian-headquartered / Indian public institution	Primary organizational control: Foreign parent / host institution
Inside India	**Nationally mobilizable core** • UPI / NPCI and payment rails • NIC / MeghRaj and public digital infrastructure • BHASHINI and multilingual public-AI infrastructure • IndiaAI allocation authority and empanelment • Indian-controlled firms, data centers, and infrastructure assets • Public institutions and regulatory authorities	**Domestic external dependence** • India-based operations of foreign hyperscalers • MNC Global Capability Centers (GCCs) • Foreign-controlled proprietary technologies used in India • Foreign-owned intellectual property and IP-dependent services used domestically
Outside India	**External organizational extension** • Overseas operations of Indian firms • Indian-controlled foreign affiliates and subsidiaries • Indian institutional research infrastructure abroad • Indian-controlled investments and assets in foreign jurisdictions	**Wider India-linked footprint** • Indian-origin professionals in foreign firms, universities, and research institutions • Diaspora knowledge networks and communities • Collaborative research outputs and publications • Host-controlled assets, platforms, and data in foreign jurisdictions

Figure 2. Territory × Organizational Control Map: Mobilizable Capacity, External Dependence, and the Wider India-Linked Footprint

Note: The horizontal dimension identifies primary organizational control rather than nationality, chip origin, or loose affiliation. The left column covers Indian-headquartered firms or Indian public institutions; the right covers foreign-parent or host-institution control. Upstream dependence on foreign accelerator technology is analyzed separately and is not encoded by this organizational-control axis. Placement does not establish contribution to nationally mobilizable capacity: access, jurisdiction, mobilizability, and durability still require separate evidence. Source: Author's analytical classification based on the attribution rules in Section 2.

Building on the attribution rule set out in Sections 1–2, this section uses "India Power" only as shorthand for nationally mobilizable capacity and the wider India-linked technological footprint ("Indian Power" in the mapping discussion) as a descriptive map of possible connections. The first is an attribution claim about capabilities and options demonstrably available through Indian institutions; the second is not. National relevance still requires evidence of access, mobilizability, jurisdiction, and durability.

5. Nationally Mobilizable Capacity and the Wider India-Linked Technological Footprint

As Section 3.2 shows, the U.S. and UK evidence captures two different forms of external reach: India-born professionals embedded in host-country S&E institutions and India-headquartered firms organizing cross-border technical labor. Reserve Bank of India (RBI) data add an economy-wide organizational measure: software-services exports reached US$190.7 billion in FY2023–24 excluding overseas commercial presence and US$205.2 billion including it; Mode 3 commercial presence contributed US$14.5 billion and Mode 4 presence of natural persons US$19.1 billion (RBI, 2024). Mode 4 is not a diaspora measure, but the data show that India's technology economy operates through multiple institutionalized cross-border channels.

These capabilities should not be nationalized by assumption. Host-country law, corporate authority, IP ownership, and individual agency remain decisive. The strategically relevant question is whether traceable mechanisms connect the wider network to usable Indian capabilities through research, investment, firms, standards, clients, or institutional cooperation. This restraint is consistent with evidence that diaspora ties can support business leads and finance and that ethnic scientific networks can facilitate international technology diffusion without implying state command over the actors involved (Kerr, 2008; Nanda & Khanna, 2010). Identity identifies a possible connection; it does not establish mobilizability.

6. Conversion: Turning Distributed Capability into Usable Power

Conversion is the missing link between resource presence and national capability. The framework specified here treats coordination, procurement, implementation, and mobilization institutions as mechanisms that determine how much of an available technical resource base becomes usable capability (Wu, 2026). India provides a revealing application because much of its resource environment is externally sourced, foreign-owned, transnationally distributed, or institutionally shared.

IndiaAI illustrates inward technological conversion. Global accelerator technology and commercial service-provider capacity are reorganized through domestic hosting, government empanelment, subsidy, and allocation into nationally accessible compute. The conversion increases mobilizability without transferring accelerator IP or eliminating future supply risk. India can therefore internalize capability without fully internalizing ownership.

GCCs represent a second, partial form of inward conversion. Foreign firms embed engineering, product, analytical, and R&D functions inside India, producing local skills, managerial knowledge, supplier relationships, and professional ecosystems while ultimate IP and strategic authority often remain abroad. Conversion exists in degrees: a routine support center is more replaceable and yields less locally retained capability than a deeply embedded R&D hub with end-to-end product responsibility. The future of India's GCC advantage depends on whether AI pushes these centers toward higher-value functions faster than automation erodes routine work (Monnappa, 2026; Venugopal & Roy, 2026).

Semiconductors show vertical conversion. India starts with a globally significant design workforce and seeks to translate that human-epistemic base into domestic firms, venture capital, owned designs, tape-outs, validated silicon, manufacturing, packaging, equipment, and materials. Each successful step can expose the next scarce interface. Conversion therefore reduces some dependencies while relocating others.

Conversion also runs outward. By July 2026, India had formal DPI cooperation arrangements with 24 partner countries, but the government's own description distinguishes knowledge exchange and feasibility from pilots and operational implementation (PIB, 2026h). The Modular Open Source Identity Platform (MOSIP) provides a deeper adjacent-DPI example. Because these cases concern digital public infrastructure rather than frontier-AI systems, they are used as mechanism evidence for architecture diffusion, not as direct proof of broad AI-specific systemic influence; Section 9 evaluates their maturity more closely (Parsheera, 2024; Raghunath, 2026).

The Indian AI system is therefore characterized by bidirectional conversion. It internalizes selected capabilities from globally controlled technological systems while externalizing selected institutional and organizational capabilities developed through its own software, public-digital-infrastructure, and professional ecosystems. This is a mechanism rather than a third category of power. Its strategic value depends on how much capability becomes usable, how durable the relationship is, and whether that capability survives disruption.

7. India's Selective AI Boundary Power

National capability does not automatically produce relational power over external actors. ABP requires a controller, target, specified interface, actual discretion, interface criticality, enforceability, and an assessment of target substitutability (Wu, 2026). The evidence reviewed here is layered. Payment localization provides the clearest demonstrated effective state ABP. IndiaAI Compute provides an AI-specific operational procurement gate. General cloud procurement and newer platform rules establish formal or operational boundary authority whose realized supplier or platform leverage is less certain, while UPI interoperability is jointly governed rather than coercive. **Table 3** summarizes these distinctions.

7.1 Payment-Data Localization as Demonstrated ABP

The clearest case is payment-data localization. In April 2018, RBI directed payment-system providers to ensure that the entire data relating to payment systems operated by them was stored in India and required compliance reporting and system audits (RBI, 2018). In 2021, RBI restricted American Express and Diners Club from onboarding new domestic customers for non-compliance and later imposed a similar restriction on Mastercard (RBI, 2021a, 2021b). In June 2022, the Mastercard restriction was lifted after RBI determined that the company had demonstrated satisfactory compliance (RBI, 2022).

The sequence is analytically important: a formal rule was followed by non-compliance, a restriction on market participation, target adaptation, and restored access. India did not merely possess regulatory authority; it changed the target's crossing conditions and enforced the change. Market scale matters here only because global payment firms value continued participation. Population alone is not boundary power; population-scale market access combined with jurisdiction and enforcement can become boundary power.

7.2 Public Procurement, Platform Regulation, and Interoperability

Public procurement illustrates a different boundary, but the evidence is weaker than in the payment case. MeitY's Annual Report documents a substantial government cloud environment, while the official Guidelines for Government Departments for Adoption/Procurement of Cloud Services, dated March 31, 2017, set procurement and exit-management requirements (Ministry of Electronics and Information Technology [MeitY], 2017, 2026a). These rules show that access to government workloads is conditional on Indian procurement and security requirements. The public evidence reviewed here, however, does not show a comparable sequence of supplier exclusion, renegotiation, and target adaptation. General cloud procurement is therefore coded as evidence of boundary authority rather than as strong demonstrated ABP.

7. India’s Selective AI Boundary Power

IndiaAI Compute contains two analytically distinct relationships. IndiaAI → applicants is primarily an administrative allocation relation: verification entities and the Program Monitoring and Evaluation Committee approve or reject project requests and determine subsidy and allocation. IndiaAI → empanelled cloud service providers is a procurement boundary: providers serve approved users under empanelled rates and service rules; the End-User Policy assigns CSPs responsibility for hardware/software compatibility, accessibility, and uptime under notified SLAs, and service requests can expire if provisioning deadlines are missed (IndiaAI, 2025, 2026a, 2026b; PIB, 2026d). The provider relation therefore supports coding an AI-specific operational procurement gate. It does not establish payment-style coercive bargaining leverage because the reviewed public record does not show a major supplier changing its terms after enforcement.

Platform regulation extends the logic into the information domain. The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026 (G.S.R. 120(E), February 10, 2026) introduced duties for defined synthetically generated audio, visual, and audio-visual information, including labeling, metadata/identifier, and intermediary due-diligence requirements (MeitY, 2026b). These rules establish formal boundary authority because participation in the Indian information market is subject to India-specific duties. The public evidence reviewed through August 31, 2026 does not, however, establish a payment-like enforcement and target-adaptation sequence or other realized platform effect. The case is therefore coded as formal regulatory boundary authority; effective ABP and realized effects remain unassessed.

Not all boundary power is coercive. The Unified Payments Interface (UPI)–PayNow linkage demonstrates negotiated interoperability between independently governed payment systems. The linkage launched in February 2023, and NPCI International Payments Limited (NIPL) documented an expansion to 19 participating Indian banks in July 2025 (NIPL, 2025; RBI, 2023). The arrangement preserves separate domestic systems while creating a jointly defined crossing interface. This form of relevance comes from becoming an interface with which other systems choose to connect, not from excluding them.

India’s boundary-control evidence is therefore selective and tiered rather than uniformly strong. Payment-data localization provides demonstrated effective state ABP with observable enforcement and target adaptation. IndiaAI Compute provides an AI-specific operational procurement gate over empanelled providers while separately exercising administrative allocation authority over users. General cloud procurement establishes formal procurement conditions whose effect on particular suppliers is not traced here, and the 2026 synthetically generated information (SGI) rules establish formal platform duties whose realized effects are not yet assessed. At upstream compute and semiconductor-equipment interfaces, Indian ABP is not demonstrated in the reviewed public evidence, while separate evidence of concentrated external supply and technological dependence indicates high exposure.

7. India's Selective AI Boundary Power

Table 3. Selected Indian AI-Relevant Boundary Relations

Boundary Relation	Interface & Mechanism	Observed Control / Target Response	Assessment
RBI / Indian state → payment providers	Payment data & market access; localization, audit, onboarding restriction	Restrictions imposed; Mastercard access restored after satisfactory compliance	Demonstrated effective state ABP
IndiaAI → empanelled AI cloud service providers	Subsidized AI compute; empanelment, pricing, provider service obligations and SLAs	CSPs must maintain required hardware/software and accessibility, meet uptime/provisioning conditions, and serve approved users under empanelled rates	AI-specific operational procurement gate; coercive bargaining effect not demonstrated
Indian state → major digital intermediaries	Defined synthetic audio/visual media; due diligence, labeling, metadata and lawful-content duties	2026 duties established; no payment-like post-rule enforcement/target-response sequence identified	Formal boundary authority; effective ABP and realized effects not demonstrated
Indian government → general cloud providers	Public-sector cloud workloads; empanelment, certification, procurement and exit terms	Eligibility and migration terms exist; public evidence reviewed does not show a clear exclusion/renegotiation target-response sequence	Formal procurement boundary; operational effect on specific suppliers not established
India / National Payments Corporation of India (NPCI) ↔ foreign payment systems	Unified Payments Interface (UPI) cross-border payments; negotiated interoperability	Operational bilateral linkages jointly define crossing conditions	Selective/joint boundary control
India vis-à-vis frontier accelerator / advanced equipment suppliers	Upstream compute and production-equipment supply; procurement and diversification	Reviewed public evidence does not demonstrate India materially altering concentrated upstream supplier terms; external concentration/dependence is separately evidenced	Indian ABP not demonstrated in reviewed public evidence; upstream exposure remains high

Note: ABP is coded only where India can change another actor's crossing conditions. "Formal authority" records a legally or contractually defined gate; "operational boundary control" requires evidence that the gate is actually used to allocate access or condition participation; "coercive/behavioral leverage" additionally requires enforcement, renegotiation, or observable target adaptation. Supplier participation in a procurement system can therefore demonstrate an operational gate without demonstrating payment-style coercive leverage. Assessment reflects publicly observable evidence through August 31, 2026.

Source: Author's coding based on RBI (2018, 2021a, 2021b, 2022), MeitY (2017, 2026a, 2026b), IndiaAI (2025, 2026a, 2026b), and the cases reviewed in the text.

8. Boundary Autonomy: Strategic Agency Under Interdependence

Boundary autonomy concerns the ability to switch, bypass, substitute, or internalize when another actor controls the interface. It reduces vulnerability without automatically creating power over the controller (Wu, 2026). This provides a more useful interpretation of Indian strategic autonomy than either full-stack autarky or political equidistance. The relevant question is whether closure at one technological boundary can disable the wider national system.

Cloud infrastructure demonstrates layered autonomy. Government agencies have access to public cloud infrastructure and multiple empanelled providers, while procurement rules include exit-management provisions (MeitY, 2017, 2026a). OECD work on domestic public-cloud compute availability similarly distinguishes the availability of compute within a jurisdiction from ultimate provider ownership, reinforcing the need to separate access from technological control (Lehdonvirta et al., 2025). Geographic hosting and provider choice can therefore be relatively strong even when technical portability remains workload-specific. A workload deeply tied to proprietary databases, cloud-native APIs, managed AI services, identity systems, or vendor-specific security architecture may remain costly to migrate. Many providers do not automatically mean easy portability.

Compute autonomy is similarly layered. IndiaAI offers multiple service providers and accelerator types, improving normal-time access and provider-level substitution. The portal specifies allocation timelines after approval, while the End-User Policy assigns availability and accessibility to empanelled providers under SLAs and makes unfulfilled requests expire after specified provisioning windows (IndiaAI, 2025, 2026a, 2026b). This is evidence of routine administrative mobilizability and provider obligations, not of continuity during a severe upstream supply, API, or provider disruption. Current evidence therefore supports greater substitutability at the service-provider layer than at the frontier-hardware layer; it does not establish crisis-resilient hardware autonomy.

Semiconductors are a portfolio problem rather than a binary sovereignty test. The supply chain contains concentrated interfaces for lithography, EDA, materials, advanced memory, manufacturing equipment, and process technologies. Domestic production is justified at some layers; supplier diversification, trusted partnerships, inventory, redesign, or less advanced substitutes may be more realistic at others. Sovereignty research increasingly treats reliable diversified external relationships as compatible with meaningful autonomy rather than as evidence of failure (Edler et al., 2023; Tanner et al., 2026).

This logic also reframes India’s international technology partnerships. Different relationships can support different layers—frontier technology access, semiconductor industrial depth, research, capital, compute, or Global South DPI cooperation. India-specific digital-sovereignty scholarship likewise emphasizes a strategy that combines state-led capacity building, multilateral independence, and continued dependence on foreign digital technologies rather than complete technological separation (Jiang, 2024; Kaushik, 2026). Strategic autonomy in AI is therefore better understood as diversified, layer-specific embeddedness than as equal political distance from competing blocs.

8. Boundary Autonomy: Strategic Agency Under Interdependence

Greater autonomy can improve bargaining position, but only if the controller changes behavior in response to India's credible exit. A government agency that can migrate from one cloud provider to another is more resilient; it does not thereby acquire power over the provider. This distinction prevents reductions in dependence from being overstated as external leverage.

9. From Boundary Effects to Systemic Influence

Systemic influence is a more demanding outcome than either national capability or a durable bilateral boundary effect. The framework requires effects to be persistent and sufficiently broad: they must extend beyond the initiating transaction to a class of actors, multiple network relations, or rules governing future participation (Wu, 2026). Because India's most mature outward architecture-diffusion cases are currently found in adjacent digital public infrastructure rather than frontier AI itself, this section distinguishes AI-specific systemic influence from adjacent-DPI mechanism evidence. **Table 4** summarizes the maturity of the selected cases under that scope distinction.

9.1 Adjacent DPI and Architecture Diffusion

UPI illustrates the difference between reach and systemic change in an adjacent digital-infrastructure domain. In FY2025–26, UPI processed more than 24,162 crore transactions with a value of roughly ₹314 lakh crore, and by August 2026 Indian authorities reported several foreign acceptance, remittance, or interoperability arrangements (PIB, 2026l). These are heterogeneous relationships, not equivalent national adoptions. Their analytical value for this report is to show how an India-controlled institutional interface can become one endpoint of repeated cross-border connections; they do not by themselves establish broad AI-specific systemic influence. Independent analyses of India Stack similarly emphasize that DPI is a socio-technical governance architecture whose diffusion, publicness, and effects cannot be reduced to deployment counts (Parsheera, 2024; Raghunath, 2026).

Bhutan provides a deeper example. Its central bank describes iUPI QR as an interoperable channel connecting India's UPI ecosystem with Bhutan's National QR ecosystem. Transaction volume rose from 1.55 thousand in 2024 to 125.87 thousand in 2025, while value increased from Nu. 6.59 million to Nu. 220.44 million (Royal Monetary Authority of Bhutan, 2026). This is persistent operational embedding, but considered alone it remains narrow. The broader systemic argument depends on repeated interoperability across multiple networks rather than any single bilateral linkage.

9. From Boundary Effects to Systemic Influence

MOSIP provides the strongest adjacent-DPI architecture-diffusion case in this study, but its aggregate rollout count has mixed maturity and should not be read as 14 equally mature operating systems. The 2026 progress report classifies 14 engagements as “National Roll-out,” yet the implementation table includes mature national-scale operations alongside cases still preparing procurement or national expansion—for example, Niger was working on an RFP and Burkina Faso was planning national rollout after an earlier pilot (MOSIP, 2026b, PDF p. 6). Persistence is therefore established from selected mature deployments rather than from the aggregate category alone. In the Philippines, the MOSIP adoption MoU dates to August 2019 and national Step 2 registration began in January 2021; by 2026 the progress report recorded roughly 94 million IDs and 170 million cumulative authentications (MOSIP, 2021; MOSIP, 2026b, PDF p. 23). Ethiopia’s partnership began in June 2020 and the 2026 report recorded about 33 million registrations and 66 million authentications (MOSIP, 2022; MOSIP, 2026b, PDF p. 24). Morocco moved from a two-province pilot in 2021 to national generalization across 83 provinces in October 2022 (Ministry of Interior, Morocco, & University Mohammed VI Polytechnic, 2026, PDF p. 4), while the 2026 progress report documents continuing large-scale authentication and benefit-delivery use (MOSIP, 2026b, PDF p. 22). Partner governments retain ownership and customization authority. The evidence therefore supports persistent India-origin architecture diffusion in selected adjacent-DPI deployments, not Indian state control and not direct proof of broad AI-specific systemic influence.

The adjacent-DPI evidence suggests a distinctive path to influence: India may diffuse architecture more easily than control. Open or interoperable systems can generate standards familiarity, ecosystem ties, and institutional relationships while leaving operational sovereignty with the adopting state. For the AI-power argument, this is mechanism evidence rather than proof that the same breadth has already been achieved in AI-specific systems.

9.2 AI Governance and Standards

India’s role in frontier-AI governance remains at an earlier stage. The India AI Impact Summit Declaration was endorsed by 92 countries and international organizations, and 13 leading model providers signed the New Delhi Frontier AI Impact Commitments (PIB, 2026f). These outcomes demonstrate convening and agenda-setting capacity, not completed system-shaping influence. The commitments call for analysis of real-world AI use and wider multilingual and cultural evaluation (PIB, 2026a).

There is limited but direct AI-specific evidence of follow-through. Anthropic’s Transparency Hub incorporates the New Delhi Frontier AI Commitments into its voluntary governance architecture, including commitments to publish aggregated usage insights, evaluate multilingual capabilities, and collaborate with local ecosystems on under-represented languages and cultural contexts (Anthropic, 2026). This is an observable organizational response in a frontier-model provider, but one provider remains below the breadth threshold. The case is therefore coded as AI-specific early institutionalization rather than systemic influence.

9. From Boundary Effects to Systemic Influence

International technical standards are similar. ITU lists E.AIQ, a framework for quality evaluation of conversational AI systems, as a new Recommendation under study with India among its supporting members (International Telecommunication Union [ITU], 2026). This supports a claim of substantive participation in standards formation, not a claim that India has already set an international AI standard.

India's AI-specific systemic influence remains emerging and is substantially narrower than its global technological footprint. Its adjacent DPI experience provides a stronger, already observable architecture-diffusion pathway, especially in identity and payments, while AI-governance commitments and technical standards remain at earlier stages of institutionalization. This yields an important strategic possibility rather than an accomplished fact: India may become capable of shaping the interfaces through which AI enters societies before it shapes the frontier technologies from which AI itself is built.

9. From Boundary Effects to Systemic Influence

Table 4. Systemic-Influence Maturity of AI-Relevant and Adjacent Digital-Infrastructure Cases

Case	Observed Effect	Persistence	Breadth	Assessment
RBI payment localization	Enforceable participation rule + firm adaptation	Rule since 2018; enforcement sequence observed in 2021–22	Provider-wide rule plausible; external spillover not established	Strong ABP; candidate sectoral systemic effect
UPI–PayNow	Durable bilateral payment linkage	Launched Feb. 2023; NIPL documented expansion to 19 Indian banks in Jul. 2025; a 2026 UPI snapshot lists Singapore but does not separately identify PayNow	Bilateral	Adjacent-DPI network rewiring; insufficient alone for broad systemic influence
Bhutan iUPI / National QR	Operational interoperability with rising transactions	Operational across 2024–25; longer persistence not established here	Narrow national relation	Adjacent-DPI deep but still temporally limited embedding
Wider UPI foreign presence	Multiple acceptance/remittance/interoperability relationships	Multiple 2026 expansions; relationship types remain heterogeneous	Multi-country but heterogeneous	Adjacent-DPI emerging broader network influence
India Stack cooperation	24 formal partner frameworks; pilots and implementation vary	2026 snapshot; implementation maturity varies by partner	Broad diplomatic reach	Adjacent-DPI reach; systemic effect not yet established
Modular Open Source Identity Platform (MOSIP)	14 provider-classified “National Roll-out” engagements; mature national-scale operations coexist with RFP/planning-stage cases	Selected mature deployments span 2020/21–2026; persistence is not uniform across all 14	Multi-country, mixed maturity	Strong adjacent-DPI architecture diffusion based on mature subset; aggregate rollout count is not uniform operational maturity
New Delhi Frontier AI Commitments	13 model providers + some firm-level follow-through	Feb.–Aug. 2026	Potentially cross-firm	AI-specific agenda setting and early institutionalization
ITU E.AIQ	India-supported conversational-AI evaluation standard under study	Under study in 2026	Potentially broad if adopted	AI-specific emerging standards influence

Note: Systemic influence is coded only when persistence and sufficient breadth are separately supported; a durable bilateral change is not enough by itself. UPI, India Stack, and MOSIP are adjacent digital-public-infrastructure cases used to identify an architecture-diffusion pathway, not as direct proof of broad AI-specific systemic influence. Assessment covers publicly observable effects through August 31, 2026 across payment infrastructure, digital identity, DPI, AI governance, and technical standards.

Source: Author’s assessment based on Wu (2026), Royal Monetary Authority of Bhutan (2026), PIB (2026f, 2026h), MOSIP (2026a, 2026b), Anthropic (2026), ITU (2026), Parsheera (2024), and Raghunath (2026).

10. Cognitive Security as an Illustrative Stress-Test Lens

Cognitive security is where the preceding distinctions can be applied most directly, but this section is an illustrative stress-test lens rather than an empirical crisis exercise. India's information environment combines population scale, linguistic heterogeneity, globally operated platforms, domestic legal authority, public digital infrastructure, and rapidly expanding generative-AI capability. The analytical question is whether the system could preserve reliable judgment, authenticated communication, and operational continuity under a specified information shock while critical technical interfaces remain transnational. The report does not claim to have observed such a national stress event.

The cultural-linguistic dimension is especially demanding. BHASHINI demonstrates that India has built a large operational language-AI layer across Indian text and voice languages (PIB, 2026b). This capacity matters because linguistic diversity is not itself a power asset; the capability lies in the infrastructure developed to operate across it. Recent scholarship on multilingual AI in India similarly treats language heterogeneity as both a constraint and an opportunity for inclusion, emphasizing that deployment quality, infrastructure, and representation shape whether multilingual systems broaden participation (Jayashankar et al., 2025).

India has an emerging verification and Responsible AI development pipeline, but the maturity of its components differs. As of July 2026, the government had approved 13 Responsible AI projects under the Safe & Trusted AI pillar, covering areas such as deepfake detection, machine unlearning, privacy-preserving AI, explainability, bias mitigation, and risk assessment (PIB, 2026k). Approval establishes an R&D pipeline for capability development; it does not by itself demonstrate delivered, deployed, or performance-validated verification tools. Separately, the 2026 SGI amendment establishes formal labeling, metadata/identifier, and intermediary duties for defined synthetic audio, visual, and audio-visual information (MeitY, 2026b). Together these measures support capability development and participation conditions, but they do not prove national cognitive resilience or improved public trust.

Detection alone has a scaling problem. As generation becomes cheaper, defenders cannot rely exclusively on post-publication judgments about whether each item is synthetic or misleading. NIST's synthetic-content guidance treats detection, provenance, labeling, authentication, testing, and audit as complementary approaches rather than substitutes (Chandra et al., 2024). Provenance can move part of the problem earlier in the information chain by recording origin and transformation history, but it has important limits. C2PA explicitly notes that valid provenance does not establish that the underlying claim is true, provenance can be incomplete or removed, and content without credentials should not automatically be distrusted (Coalition for Content Provenance and Authenticity [C2PA], n.d.). High-consequence official communication can therefore benefit from authentication without assuming that signatures or credentials alone establish factual truth.

Continuity is equally important. A verification system that depends entirely on one proprietary foreign model API, cloud provider, or distribution platform may fail precisely when it is most needed. Cognitive security therefore requires boundary autonomy: locally deployable or substitutable models for critical tasks, public or multi-provider cloud fallback, protected language resources, redundant distribution channels, and portable provenance and verification tools.

10. Cognitive Security as an Illustrative Stress-Test Lens

Foreign platforms should accordingly be treated as high-value but externally controlled interfaces rather than as inherently hostile systems. India can retain connectivity while increasing transparency, provenance, fallback, and interoperability. The goal is managed interdependence: prevent a single external interface from becoming a point of cognitive failure.

Normative-strategic orientation also conditions whether cognitive-security capability is legitimate and usable. Cognitive security in this report concerns the integrity and resilience of the information environment through which both the state and society form judgments; it should not be equated with protection of official narratives. Provenance, authentication, intermediary duties, and crisis coordination therefore require proportionality, privacy safeguards, transparent procedures, meaningful avenues for challenge or appeal, and independent oversight where rights are affected. Scholarship on Indian digital sovereignty and DPI highlights precisely these tensions between autonomy, inclusion, openness, surveillance, and democratic accountability (Jiang, 2024; Kaushik, 2026; Raghunath, 2026). Normative legitimacy is consequently a conditioning factor on usable cognitive-security capacity rather than an optional external consideration.

Figure 3 summarizes the resulting application-domain architecture. These are implementation functions rather than new theoretical dimensions of AI power.

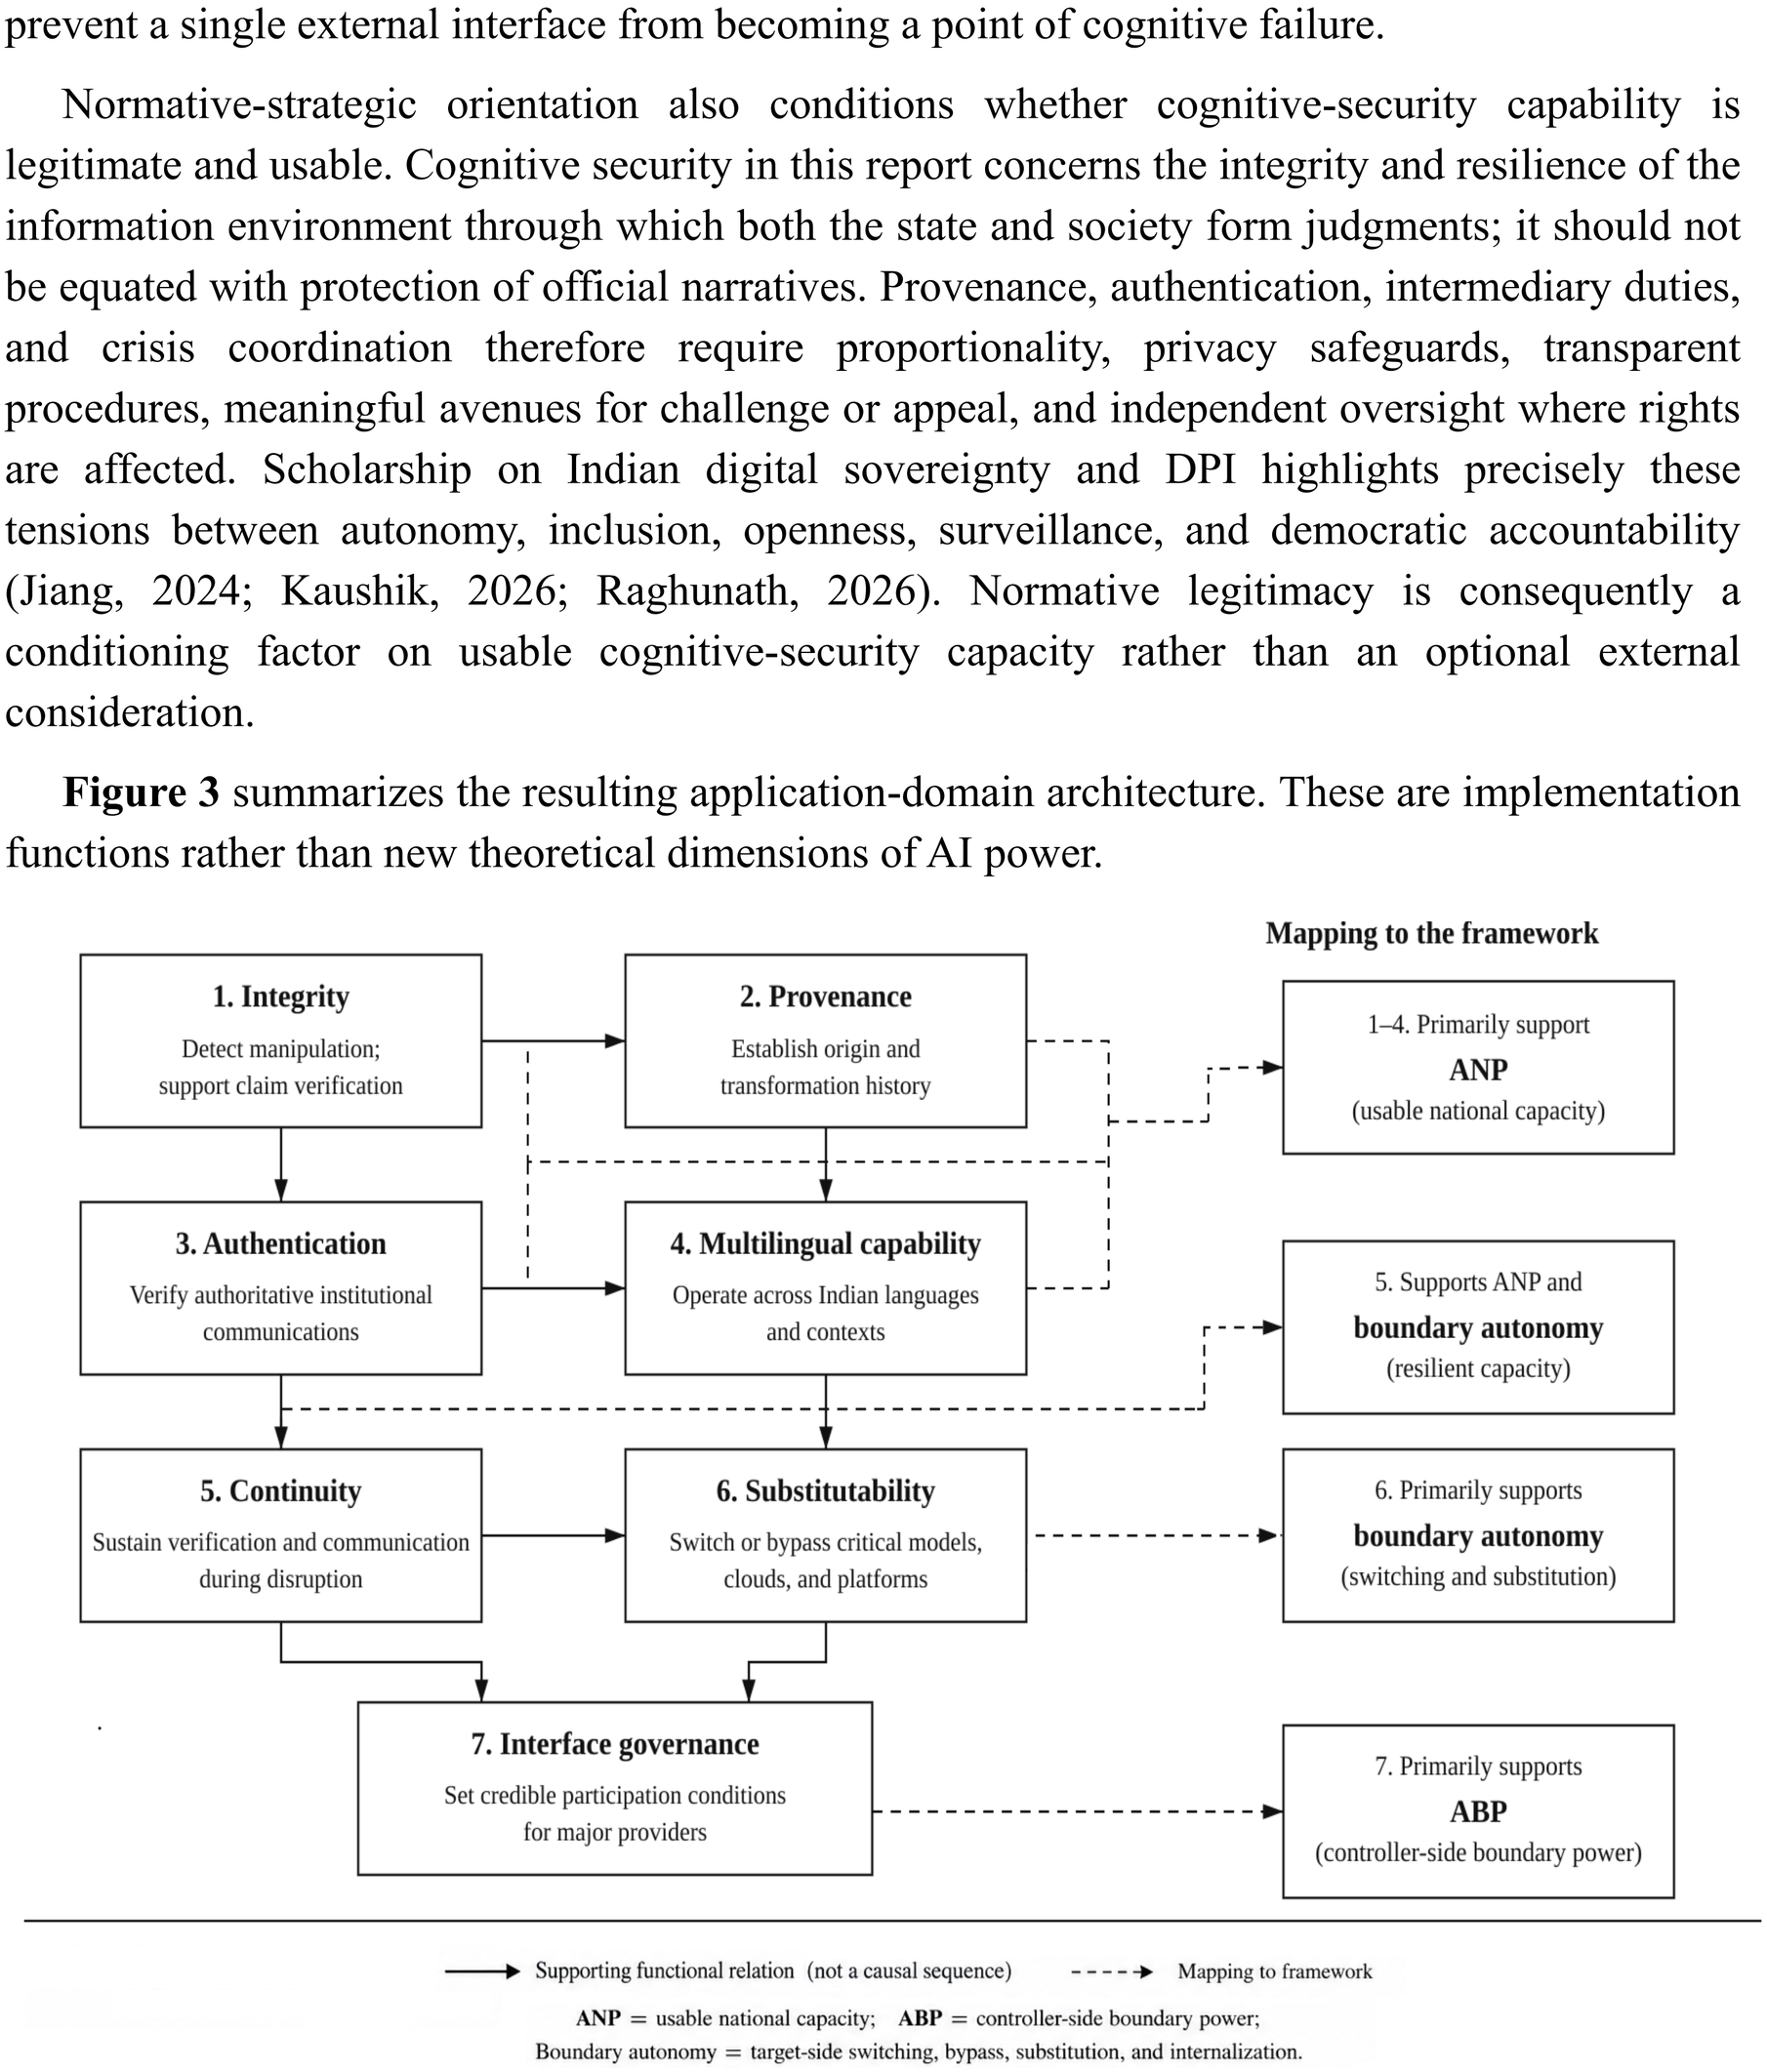

Figure 3. Cognitive-Security Architecture as an Illustrative Application-Domain Lens

Note: These implementation functions translate national capacity, boundary control, and boundary autonomy into a cognitive-security application lens; they are not additional dimensions of AI power. Provenance or valid signatures establish aspects of origin and integrity, not the truth of a claim. Source: Author's synthesis based on Chandra et al. (2024), C2PA (n.d.), Wu (2026), PIB (2026b, 2026k), and MeitY (2026b).

10. Cognitive Security as an Illustrative Stress-Test Lens

To make the lens operational without representing a hypothetical as observed fact, consider an illustrative compound scenario: during a national emergency, a multilingual synthetic audio-video message impersonates a public authority while a major foreign model API used by some verification workflows becomes unavailable for 24 hours. The exercise does not estimate real performance. It identifies what evidence would be required to claim resilience and where the present public record remains incomplete. **Table 5** specifies the evidence that would be required to treat the system as resilient under this illustrative compound shock.

Table 5. Illustrative Cognitive-Security Stress-Test Criteria

Function	Hypothetical Test	Evidence Required for a Resilience Claim	Status in This Study
Integrity	Can manipulated or misleading synthetic content be detected and subjected to reliable claim verification during the shock?	Detection performance; verification procedures; false-positive/negative rates; response latency; escalation and human-review mechanisms	Detection and Responsible AI initiatives are under development; integrated nationwide verification performance is not established
Provenance / Authentication	Can official communications be verified despite synthetic impersonation?	Signed/provenance-aware official channels; verification latency; key-compromise procedures	Regulatory and technical components exist; nationwide performance not established
Multilingual Coverage	Can verification operate across relevant languages and code-switching contexts?	Language-specific accuracy, false-positive/negative rates, coverage under emergency load	BHASHINI shows scale of language infrastructure; verification quality is not established across all languages
Continuity	Can critical verification continue during a 24-hour foreign API outage?	Locally deployable fallback, alternate cloud/model path, recovery-time objective	Normal-time compute allocation exists; disruption continuity is not publicly demonstrated
Substitutability	Can workloads migrate without unacceptable loss?	Exportability, data/model portability, migration time and performance loss	Provider diversity exists; workload-specific switching cost remains uncertain
Interface Governance	Can major platforms receive and act on authenticated emergency notices under lawful procedures?	Crisis protocol, response times, due process, appeal and audit trail	Formal intermediary duties exist; end-to-end crisis performance not established

Note: This table is an illustrative analytical exercise, not a report of an event that occurred. Integrity refers to the capacity to detect manipulation and support claim verification; it does not imply that technical detection, provenance, or authentication alone establishes the truth of a claim.

Source: Author's synthesis based on Chandra et al. (2024), C2PA (n.d.), MeitY (2026b), and the capability evidence reviewed in Section 10.

10. Cognitive Security as an Illustrative Stress-Test Lens

For this study, cognitive security is therefore the capacity to preserve trustworthy judgment, authenticated communication, operational continuity, and technological choice within an AI-mediated information environment whose critical interfaces are distributed across domestic and transnational systems. India already possesses significant components of such an architecture, but the public evidence reviewed here does not establish a fully integrated nationwide system or measured performance under the compound shock illustrated above.

11. Strategic Implications and India's AI-Power Trajectory to 2030

The evidence does not support a strategy of full-stack imitation. India is already deeply embedded in global semiconductor, cloud, software, corporate, and research systems, and many critical technologies cannot realistically be reproduced domestically in the near term. The more useful objective is selective control and resilience: mobilize what India has, reduce the cost of switching where others control the boundary, and deepen capabilities that make India itself difficult to replace.

11.1 Three Policy Priorities

First, protect critical continuity and increase substitutability at externally controlled chokepoints. MeitY, IndiaAI, NIC and sector regulators should identify a small set of critical public workloads—such as emergency communication, core administrative services, and cognitive-security verification—and specify minimum service levels, provider concentration limits, data/model exportability, fallback compute, migration time, and recovery objectives. The implementation tools are procurement requirements, multi-provider architecture, selective domestic fallback, strategic inventories, and tested transition plans. The trade-off is cost: redundancy and portability can reduce short-run efficiency, so the sovereign minimum should be limited to functions whose interruption would impose unacceptable public or strategic losses.

Second, deepen high-value conversion in areas where India already has a large human and organizational base. India Semiconductor Mission, MeitY, industry, universities and firms should track whether design capability becomes owned IP, validated products, production, packaging, equipment/material competence, and resilient supplier relationships; GCC policy should track product ownership, R&D mandates, senior technical responsibility, and AI architecture rather than headcount alone. The principal opportunity cost is that incentives directed toward headline capacity can crowd out harder-to-measure capability formation. Evaluation should therefore prioritize activated products, functions, and durable organizational roles over announcements.

Third, exercise and diffuse interface power only where evidence supports credible control or institutional attraction. RBI, MeitY, NPCI, IndiaAI and external-affairs institutions can use market access, public procurement, payments, AI-compute allocation, provenance rules, and interoperability to shape participation conditions. Metrics should distinguish formal rules from enforced target response and distinguish MoUs from operational foreign deployment. Open architectures that partner countries own can still generate durable influence, but the policy objective should be interoperability and institutional embeddedness rather than maximizing country counts or implying control that India does not possess.

11. Strategic Implications and India's AI-Power Trajectory to 2030

11.2 Five Conversion Tests for 2030

Table 6 operationalizes the 2030 scenarios as observable conversion tests rather than forecasts, emphasizing activation, durability, and substitutability rather than announced investment totals.

Table 6. Five Conversion Tests for India's AI-Power Trajectory to 2030

Test	Key Question	Evidence To Track	Strategic Meaning
Compute Conversion	Does headline access become resilient operational capacity?	Operational compute; supplier concentration; portability; public fallback; crisis continuity	ANP and boundary autonomy
Semiconductor Conversion	Does design depth move into owned products and production?	Domestic IP; silicon validation; production; packaging; equipment/material depth	Upstream autonomy and industrial capacity
GCC/Organizational Conversion	Does AI move India into higher-value functions?	Product ownership; R&D mandates; AI architecture; senior technical roles; India-plus shifts	Durability of India's organizational advantage
DPI External Conversion	Does cooperation become repeated institutional embedding?	Operational linkages; national rollouts; interoperability; rule adoption; persistence	Systemic influence
Cognitive-Security Conversion	Do dispersed capabilities become an integrated resilience architecture?	Provenance; authentication; multilingual verification; fallback; cross-platform crisis protocols	Political-administrative conversion under stress

Note: The tests focus on activation, durability, and substitutability rather than investment announcements or country rankings. They are scenario-monitoring criteria, not probabilistic forecasts.

Source: Author's analytical synthesis based on Wu (2026) and the evidence reviewed in this report.

Three conditional trajectories follow. Under Conversion Deepening, compute access becomes resilient, semiconductor design translates into greater domestic IP and production, GCCs move into high-value AI and R&D, and DPI shifts toward repeated foreign institutional embedding. Under Embedded Expansion, capability rises rapidly but upstream technological dependence remains concentrated, so capacity grows faster than autonomy. Under Network Erosion, physical infrastructure continues to expand while AI automates routine technology work, firms diversify toward "India-plus" locations, and the relative value of India's human and organizational network weakens.

The three trajectories are not mutually exclusive, and this study does not assign probabilities or identify a single most likely path. Current observations are compatible with elements of both Conversion Deepening and Embedded Expansion: several programmes are increasing normal-time capacity and administrative allocation while concentrated upstream dependencies remain. Future classification should depend on observable conditions such as supplier concentration, successful workload migration, commissioned production, validated domestic products, and persistent external institutional deployment—not on a directional forecast made from 2026 announcements.

12. Conclusion

India’s emerging AI power cannot be captured by asking whether it is winning a global model race. The evidence instead reveals a distributed and vertically uneven technological position: strong human, software, organizational, digital-public-infrastructure, multilingual, and expanding compute capabilities coexist with important dependence on externally concentrated upstream technologies. The report’s principal methodological contribution is to separate territorial presence, ownership, jurisdiction, normal-time availability, administrative mobilizability, disruption resilience, relational boundary control, and later systemic effects rather than collapsing them into a single ranking.

That distinction changes how globally distributed capability is attributed. Foreign-owned infrastructure inside India can contribute to nationally mobilizable capacity without becoming fully Indian-controlled, while Indian-headquartered firms or India-origin professionals abroad matter only where observable conversion pathways connect them to usable Indian options. Conversion is therefore central: India can internalize usable capability without internalizing all ownership and can diffuse institutional architecture without retaining foreign gatekeeping power.

Boundary power is more selective. Payment-data localization provides the clearest demonstrated case because enforcement produced observable target response. AI-compute procurement and synthetic-media rules demonstrate important AI-specific boundary institutions, but current public evidence supports more cautious coding of their effective leverage. Where India remains dependent on concentrated accelerators, semiconductor equipment, or proprietary technical layers, the strategic problem is boundary autonomy—credible switching, diversification, fallback, and selective internalization—rather than assumed power over suppliers.

Systemic influence is narrower still. Broad AI-specific system-shaping effects remain emerging. MOSIP and UPI provide stronger adjacent-DPI evidence of architecture diffusion and interoperability, while the New Delhi Frontier AI Commitments show early AI-specific institutional follow-through rather than cross-industry transformation. Cognitive security consequently remains an illustrative stress-test lens: India has operational multilingual, compute, and public-digital components, formal regulatory provisions, and a developing verification and Responsible AI research pipeline; resilience would require evidence that these capabilities continue to function under specified disruption and within safeguards for privacy, due process, proportionality, appeal, and independent oversight.

The central strategic task is therefore bounded and empirical: mobilize what India demonstrably possesses, deepen conversion where organizational capability is strongest, reduce the cost of switching where external actors control critical interfaces, and exercise boundary power only where enforcement or operational access conditions make it credible. Artificial intelligence is increasing both the value of India’s pre-existing strengths and the cost of its most consequential dependencies; future assessments should judge the balance through activated capability, substitutability, and persistent institutional effects rather than announcements or national labels.

Appendix A. Qualitative Codebook and Case Evidence Matrix

This appendix makes the main qualitative judgments more transparent and auditable. **Tables A1** and **A2** in Appendix separate evidence/maturity coding from limitations and confidence so that readers can identify the public basis of the principal case judgments and challenge alternative interpretations. Confidence refers to the adequacy of public evidence for the coding shown, not to the normative desirability of the policy or to the truth of every underlying administrative statistic. “High” confidence requires a directly relevant primary or first-party operational source and a close fit between evidence and claim; “moderate” indicates a reliable source with missing operational detail or a partially inferential attribution; “low” indicates that the public record supports only a preliminary interpretation. Absence of public evidence is not coded as proof of absence.

Table A1. Case-Level Evidence State and Main Coding

Case / Construct	Observation Period	Evidence State	Main Coding
IndiaAI shared compute / ANP	Jun–Aug 2026	45,000+ shared GPUs; 237 projects approved; 9.318m GPU-hours sanctioned	Routine access + administrative allocation; consumption and crisis continuity not fully observed
NIC 1.1-EFLOPS system / ANP	Aug 2026	Purchase order issued	Procurement stage only
RBI payment localization / ABP	2018–2022+	Rule, restriction, target compliance, restored access	Demonstrated effective state ABP
IndiaAI provider empanelment / AI-specific procurement boundary	2025–Aug 2026	15 providers empanelled; provider obligations, pricing/SLAs and user-allocation rules operating	Operational procurement gate over CSPs; separate administrative allocation authority over users
SGI intermediary rules / potential ABP	Feb–Aug 2026	Formal duties for defined synthetic audio/visual media	Formal boundary authority; realized effect not assessed
General government cloud procurement / boundary authority	2017–2026	Procurement/exit guidance + government cloud environment	Procurement boundary authority
UPI–PayNow / joint boundary relation	Feb. 2023–Jul. 2025 (PayNow-specific); 2026 general UPI/Singapore snapshot	Operational bilateral linkage; 2025 expansion documented	Joint negotiated interoperability

Appendix A. Qualitative Codebook and Case Evidence Matrix

Table A1. Case-Level Evidence State and Main Coding (Cont.)

Case / Construct	Observation Period	Evidence State	Main Coding
MOSIP / adjacent-DPI systemic influence	2019/2020–Aug 2026 for selected mature deployments; 2026 rollout snapshot for aggregate category	14 provider-classified national-rollout engagements, but category mixes mature national-scale operations with preparatory cases	Persistent architecture diffusion demonstrated for selected mature deployments; not uniform across all 14
New Delhi Frontier AI Commitments / systemic influence	Feb–Aug 2026	13 signatories; Anthropic incorporates commitments	AI-specific agenda setting + early institutionalization
Cognitive-security stress lens	Illustrative only	No observed national compound shock	Analytical test conditions, not empirical finding
BHASHINI / multilingual ANP	Aug 2026 snapshot	36 text languages; 23 voice languages; 800+ government websites; 24m+ daily and 9bn+ cumulative inferences	Large-scale operational language infrastructure; cross-language quality not established
GCC / organizational conversion	FY2025–26 estimate	2,117 GCCs; ~2.36m professionals; US$98.4bn expected sector revenue	Large India-located organizational capability; mixed ownership and national mobilizability
Semiconductor conversion / ANP	Mar–Aug 2026	Large design workforce estimate; DLI funding/tape-outs; 12 approved manufacturing projects with mixed maturity	Human-epistemic depth converting into industrial capability; upstream dependencies remain
Bhutan iUPI / adjacent-DPI influence	2024–2025	Operational QR interoperability; transaction volume/value rose materially in 2025	Deep bilateral institutional embedding; longer persistence not established here
India Stack cooperation / adjacent-DPI reach	Jul 2026	24 partner-country cooperation frameworks with varying feasibility/pilot/implementation status	Broad diplomatic/institutional reach; not equivalent to operational deployment
ITU E.AIQ / AI standards influence	2026	India-supported conversational-AI evaluation work item listed as under study	AI-specific standards participation; systemic effect not established

Note: Evidence states and main coding correspond to the construct-specific rules in Section 2. Observation periods refer to the public evidence reviewed, not necessarily the full lifetime of the underlying programme or relationship.

Appendix A. Qualitative Codebook and Case Evidence Matrix

Table A2. Case-Level Limitations, Confidence, and Key Sources

Case / Construct	Alternative Interpretation / Limit	Confidence	Key Source
IndiaAI shared compute / ANP	Provider/hardware dependence may constrain disruption resilience	High for approval/allocation; moderate for mobilizability	PIB 2026d, 2026k; IndiaAI 2026a
NIC 1.1-EFLOPS system / ANP	Not operational or commissioned capacity	High	PIB 2026d
RBI payment localization / ABP	Does not by itself prove external/global spillover	High	RBI 2018, 2021a, 2021b, 2022
IndiaAI provider empanelment / AI-specific procurement boundary	No major supplier target-response sequence showing coercive bargaining leverage	High for operational procurement gate; supplier bargaining effects not assessed	IndiaAI 2025, §§7–9 (print pp. 14–16; PDF pp. 20–22); IndiaAI 2026a–b; PIB 2026d
SGI intermediary rules / potential ABP	Post-rule platform enforcement and behavioral response not yet demonstrated	High for existence of formal duties; realized effects not assessed	MeitY 2026b, G.S.R. 120(E), Feb. 10, 2026
General government cloud procurement / boundary authority	Public evidence of exclusion/renegotiation and target response limited	High for existence of procurement/exit terms; realized target effect not assessed	MeitY 2017; MeitY 2026a
UPI–PayNow / joint boundary relation	Bilateral durability is not broad systemic influence	High for existence and operation through Jul. 2025; 2026 PayNow-specific continuity not separately established	RBI 2023; NIPL 2025; PIB 2026l (general Singapore UPI snapshot only)
MOSIP / adjacent-DPI systemic influence	The “National Roll-out” category includes RFP/planning cases; partner governments own/customize systems; direct AI systemic influence is not established	High for rollout-category count and selected mature operations; persistence not assessed uniformly across all 14	MOSIP 2026b, PDF p. 6 and PDF pp. 22–24; MOSIP 2021; MOSIP 2022; Ministry of Interior, Morocco & UM6P 2026, PDF p. 4
New Delhi Frontier AI Commitments / systemic influence	One observable provider response remains below breadth threshold	Moderate	PIB 2026a, 2026f; Anthropic 2026
Cognitive-security stress lens	Performance, recovery and cross-language quality remain unmeasured	Not applicable	Chandra et al. 2024; C2PA n.d.
BHASHINI / multilingual ANP	Scale metrics do not establish accuracy or parity across languages	High for reported scale; quality/effectiveness not assessed	PIB 2026b, Aug. 5, 2026 release

Appendix A. Qualitative Codebook and Case Evidence Matrix

Table A2. Case-Level Limitations, Confidence, and Key Sources (Cont.)

Case / Construct	Alternative Interpretation / Limit	Confidence	Key Source
GCC / organizational conversion	Industry estimate; sector revenue is not AI-specific and foreign-parent control remains material	Moderate-high for sector scale; case-specific mobilizability remains inferential	Ishwarbharath & Suresh 2026; Monnappa 2026; Venugopal & Roy 2026
Semiconductor conversion / ANP	Tape-outs, approved projects and operating facilities represent different maturity states	High for cited programme states; future industrial depth not assessed	PIB 2026c; PIB 2026g; PIB 2026j
Bhutan iUPI / adjacent-DPI influence	Two-year transaction evidence does not establish long-run systemic persistence	High for 2024–25 operation; longer persistence not assessed	Royal Monetary Authority of Bhutan 2026, iUPI QR section
India Stack cooperation / adjacent-DPI reach	MoUs, feasibility, pilots and national deployment are not equivalent	High for partner-framework count; mixed for implementation depth	PIB 2026h, Jul. 29, 2026 release
ITU E.AIQ / AI standards influence	Work item remains under study; adoption and implementation are not demonstrated	High for participation/status; realized standards influence not assessed	International Telecommunication Union 2026, work item E.AIQ

Note: Confidence concerns the adequacy of public evidence for the stated coding, not the normative desirability of the policy or an independent audit of the underlying administrative statistic. Absence of public evidence is not treated as proof of absence. Source: Author's codebook and the sources listed in the final column.

References

Anthropic. (2026). *Transparency Hub: Voluntary commitments*. Retrieved August 31, 2026, from https://www.anthropic.com/transparency/voluntary-commitments

Arora, A., Arunachalam, V. S., Asundi, J., & Fernandes, R. (2001). The Indian software services industry. *Research Policy, 30*(8), 1267–1287. https://doi.org/10.1016/S0048-7333(00)00148-7

Athreye, S. S. (2005). The Indian software industry and its evolving service capability. *Industrial and Corporate Change, 14*(3), 393–418. https://doi.org/10.1093/icc/dth056

Bradford, A. (2020). *The Brussels effect: How the European Union rules the world*. Oxford University Press. https://doi.org/10.1093/oso/9780190088583.001.0001

Chandra, B., Dunietz, J., Roberts, K., Lee, Y., Fontana, P., & Awad, G. (2024). *Reducing risks posed by synthetic content: An overview of technical approaches to digital content transparency (NIST AI 100-4)*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-4

Coalition for Content Provenance and Authenticity. (n.d.). *C2PA and Content Credentials explainer (Version 2.2)*. Retrieved August 31, 2026, from https://spec.c2pa.org/specifications/specifications/2.2/explainer/Explainer.html

Edler, J., Blind, K., Kroll, H., & Schubert, T. (2023). Technology sovereignty as an emerging frame for innovation policy: Defining rationales, ends and means. *Research Policy, 52*(6), 104765. https://doi.org/10.1016/j.respol.2023.104765

Farrell, H., & Newman, A. L. (2019). Weaponized interdependence: How global economic networks shape state coercion. *International Security, 44*(1), 42–79. https://doi.org/10.1162/isec_a_00351

IndiaAI. (2025). *End-user policy for AI services on the cloud*. Digital India Corporation, Ministry of Electronics and Information Technology. https://compute.indiaai.gov.in/files/end-user-policy.pdf

IndiaAI. (2026a). *IndiaAI Cloud Computing Portal*. Government of India. Retrieved August 31, 2026, from https://compute.indiaai.gov.in/

IndiaAI. (2026b). *Service level agreement*. Government of India. Retrieved August 31, 2026, from https://compute.indiaai.gov.in/sla

International Telecommunication Union. (2026). *E.AIQ: Framework for quality evaluation of conversational AI (C-AI) system* [ITU-T work item]. Retrieved August 31, 2026, from https://www.itu.int/ITU-T/workprog/wp_item.aspx?isn=20851

Ishwarbharath, S., & Suresh, H. (2026, May 6). India's offshore tech hubs hit $98.4 bln revenue in FY26, nearing 2030 forecast, report says. *Reuters*. https://www.reuters.com/world/india/indias-offshore-tech-hubs-hit-984-bln-revenue-fy26-nearing-2030-forecast-report-2026-05-06/

Jayashankar, P., Johnston, W. J., & Sheth, J. (2025). Multilingualism as strategy—AI-enabled STP for inclusive market development in emerging economies. *Marketing Strategy Journal, 4*, 100011. https://doi.org/10.1016/j.msj.2025.100011

Jiang, M. (2024). Models of state digital sovereignty from the Global South: Diverging experiences from China, India and South Africa. *Policy & Internet, 16*(4), 727–738. https://doi.org/10.1002/poi3.427

Kaushik, H. (2026). India and digital sovereignty: Norm entrepreneurship in global digital governance. *Asian Journal of Comparative Politics, 11*(2), 135–161. https://doi.org/10.1177/20578911261419896

Kerr, W. R. (2008). Ethnic scientific communities and international technology diffusion. *The Review of Economics and Statistics, 90*(3), 518–537. https://doi.org/10.1162/rest.90.3.518

References

Lehdonvirta, V., Wu, B., Hawkins, Z. J., Caira, C., & Russo, L. (2025). *Measuring domestic public cloud compute availability for artificial intelligence (OECD Artificial Intelligence Papers No. 49)*. OECD Publishing. https://doi.org/10.1787/8602a322-en

Migration Advisory Committee. (2025). *Professionals in IT and engineering review*. UK Government. https://www.gov.uk/government/publications/professionals-in-it-and-engineering-review/professionals-in-it-and-engineering-accessible

Ministry of Electronics and Information Technology. (2017). *Guidelines for government departments for adoption/procurement of cloud services*. Government of India. https://www.meity.gov.in/writereaddata/files/Guidelines-Procurement_Cloud_Services.pdf

Ministry of Electronics and Information Technology. (2026a). *Annual report 2025–26*. Government of India. https://www.meity.gov.in/static/uploads/2026/04/46face7d48c8f6a97030f713ad5fdab4.pdf

Ministry of Electronics and Information Technology. (2026b, February 10). *Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026* (G.S.R. 120(E)). Government of India. https://www.meity.gov.in/static/uploads/2026/02/f55fe52418b03f58b0669f6a8bc03b6d.pdf

Ministry of Interior, Morocco, & University Mohammed VI Polytechnic. (2026). *Collaborations between governments and academic institutions to build capacity and provide L3 support*. https://connect.mosip.io/resources_pdf_2026/Collaborations_between_Governments_and_Academic_Institutions.pdf

Monnappa, C. (2026, May 27). India's GCC model shifts from cost to capability as AI, talent strains bite. *Reuters*. https://www.reuters.com/world/india/indias-gcc-model-shifts-cost-capability-ai-talent-strains-bite-2026-05-27/

MOSIP. (2021, May 31). *The Philippine Statistics Authority (PSA) crosses critical milestones for the Philippine Identification System (PhilSys)*. https://www.mosip.io/news_events/the-philippine-statistics-authority-psa-crosses-critical-milestones-for-the-philippine-identification-system-philsys

MOSIP. (2022, February 23). *Ethiopia's National ID Program (NIDP) extends their partnership with MOSIP*. https://www.mosip.io/news_events/ethiopia-s-national-id-program-nidp-extends-their-partnership-with-mosip

MOSIP. (2026a). *Modular Open Source Identity Platform*. Retrieved August 31, 2026, from https://www.mosip.io/

MOSIP. (2026b). *MOSIP progress report 2026*. https://connect.mosip.io/resources_pdf_2026/MOSIP_Progress_Report.pdf

Nanda, R., & Khanna, T. (2010). Diasporas and domestic entrepreneurs: Evidence from the Indian software industry. *Journal of Economics & Management Strategy, 19*(4), 991–1012. https://doi.org/10.1111/j.1530-9134.2010.00275.x

National Science Board. (2026). *The state of U.S. science and engineering 2026*. National Center for Science and Engineering Statistics, National Science Foundation. https://ncses.nsf.gov/pubs/nsbsep20261/stem-talent-education-training-and-workforce-2

NPCI International Payments Limited. (2025, July 16). *NPCI International expands UPI–PayNow linkage to drive cross-border remittances* [Press release]. https://www.nipl.com/uploads/NPCI_International_Expands_UPI_Pay_Now_Linkage_to_Drive_Cross_Border_Remittances_eafe4d399a.pdf

Parsheera, S. (2024). Stack is the new black?: Evolution and outcomes of the 'India-Stackification' process. *Computer Law & Security Review, 52*, 105947. https://doi.org/10.1016/j.clsr.2024.105947

References

Press Information Bureau. (2026a, February 19). *Championing inclusive and multilingual AI for the Global South, India unveils New Delhi Frontier AI Commitments*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2230201&lang=1&reg=3

Press Information Bureau. (2026b, August 5). *Digital India BHASHINI Division conducts BHASHINI Rajyam workshop: Lucknow chapter to advance language-first digital governance in Uttar Pradesh*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2294908&lang=2&reg=48

Press Information Bureau. (2026c, August 15). *DLI Scheme-backed chip design startup Aheesa Digital Innovations achieved first-pass silicon success with VIHAAN*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2299883&lang=1&reg=3

Press Information Bureau. (2026d, August 6). *Government expands sovereign AI infrastructure through IndiaAI Mission and semiconductor initiatives*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2295477&lang=1&reg=1

Press Information Bureau. (2026e, August 12). *Government promotes sustainable growth of data centres through energy-efficient and clean energy measures*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2298108&lang=1&reg=1

Press Information Bureau. (2026f, March 11). *India AI Impact Summit 2026 witnessed extensive participation of public, with approximately 6 lakh attendees in person*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2238195&lang=2&reg=48

Press Information Bureau. (2026g, March 13). *India has emerged as a global hub for semiconductor design and R&D*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2239608&lang=1&reg=3

Press Information Bureau. (2026h, July 29). *India strengthens global digital cooperation through Digital Public Infrastructure partnerships with 24 countries*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2291054&lang=1&reg=1

Press Information Bureau. (2026i, July 29). *IndiaAI Mission advances AI-driven governance and public service delivery through collaborative roadmaps, innovation challenges and capacity building*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2291062&lang=2&reg=48

Press Information Bureau. (2026j, August 14). *Manufacturing momentum: Building a self-reliant India [Factsheet]*. Government of India. https://www.pib.gov.in/FactsheetDetails.aspx?ModuleId=16&NoteId=150923&id=150923&lang=1&reg=6

Press Information Bureau. (2026k, August 13). *Semiconductor and AI revolution [Factsheet]*. Government of India. https://www.pib.gov.in/FactsheetDetails.aspx?Id=150860&lang=1&reg=3

Press Information Bureau. (2026l, August 24). *UPI completes 10 years of digital payments revolution—Globally recognised as world's largest real-time payment system*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2302657&lang=1&reg=48

Raghunath, P. (2026). Big data discourses | Data in India's digital public infrastructure (DPI): A deliberative policy ecology approach. *International Journal of Communication, 20*, 417–434. https://doi.org/10.65476/0s5f9h65

Reserve Bank of India. (2018, April 6). *Storage of payment system data*. https://www.rbi.org.in/Scripts/NotificationUser.aspx?Id=11244

Reserve Bank of India. (2021a, April 23). *RBI takes supervisory action on card networks—American Express Banking Corp. and Diners Club International Ltd.* https://www.rbi.org.in/Scripts/BS_PressReleaseDisplay.aspx?prid=51471

References

Reserve Bank of India. (2021b, July 14). *Reserve Bank of India takes supervisory action on Mastercard Asia/Pacific Pte. Ltd.* https://www.rbi.org.in/Scripts/BS_PressReleaseDisplay.aspx?prid=51895

Reserve Bank of India. (2022, June 16). *Reserve Bank of India lifts the business restrictions imposed on Mastercard Asia/Pacific Pte. Ltd.* https://www.rbi.org.in/Scripts/BS_PressReleaseDisplay.aspx?prid=53877

Reserve Bank of India. (2023, February 21). *UPI–PayNow linkage: Frequently asked questions*. https://www.rbi.org.in/commonperson/english/Scripts/FAQs.aspx?Id=3415

Reserve Bank of India. (2024, October 18). *Survey on computer software and information technology enabled services exports: 2023–24*. https://www.rbi.org.in/Scripts/BS_PressReleaseDisplay.aspx?prid=58937

Royal Monetary Authority of Bhutan. (2026). *Annual payment system report 2025*. https://www.rma.org.bt/media/Publication/Payments%20%26%20Settlement%20Systems/Annual%20Payment%20System%20Report%202025.pdf

Tanner, B., Kerry, C. F., Wyckoff, A. W., Kyosovska, N., Renda, A., & Tabassi, E. (2026). *Is AI sovereignty possible? Balancing autonomy and interdependence*. Brookings Institution & Centre for European Policy Studies. https://www.brookings.edu/articles/is-ai-sovereignty-possible-balancing-autonomy-and-interdependence/

Venugopal, A., & Roy, M. (2026, May 25). Global firms rethink GCC hiring in India as AI shifts skill demand. *Reuters*. https://www.reuters.com/world/india/global-firms-rethink-gcc-hiring-india-ai-shifts-skill-demand-2026-05-25/

Wu, S. (2026). *Two structures of AI power: National capacity, relational boundary control, and systemic influence in the emerging AI order* (EPINOVA Working Paper No. EPINOVA–WP–T–2026–02). Global AI Governance and Policy Research Center, EPINOVA LLC. https://doi.org/10.67037/epinova.wp.t.2026.002

Author and Institutional Affiliation

Author: Wu, Shaoyuan.

ORCID: https://orcid.org/0009-0008-0660-8232

Affiliation: Global AI Governance and Policy Research Center, EPINOVA LLC

Disclaimer

This research report presents a theory-guided analytical assessment of India's emerging AI power, technological dependence, boundary control, strategic autonomy, and cognitive security. The frameworks, classifications, and case interpretations developed in this report are intended as analytical and diagnostic tools. They should not be interpreted as definitive measures of national power, deterministic predictions, comprehensive assessments of technological sovereignty, or official evaluations of any government, firm, institution, or technology provider. Empirical claims are based on publicly available information and remain subject to source limitations, reporting gaps, changing institutional conditions, and further verification.

The views expressed in this report are solely those of the author and do not necessarily represent the position of any external institution, government, organization, partner entity, or affiliated institution. The analysis is intended for research, policy analysis, strategic assessment, and institutional discussion. Its concepts and findings should be applied with appropriate attention to context, evidence quality, methodological limitations, and changes in the technological and policy environment.

Version History

Version	Date	Description
Version 1.0	September 1, 2026	Initial public release as EPINOVA Research Report No. EPINOVA–2026–RR-01.

Citation Recommendation

Wu, Shaoyuan. (2026). *India's AI Power in a Distributed and Uneven Technological Order: Boundary Control, Strategic Autonomy, and Cognitive Security* (Research Report No. EPINOVA–2026–RR-01). Global AI Governance and Policy Research Center, EPINOVA LLC. https://doi.org/10.67037/epinova.rr.2026.001