The New Indian IT Playbook : Why the Next Winners Are Emerging Below the Mega-Caps

India as a Global "Anti-AI Play"

Global investors are increasingly casting India as the anti-AI trade, a characterisation that reflects scepticism around the country's ability to participate meaningfully in the early stages of the AI cycle, both in terms of sovereign large-language-model development and the heavy datacentre and hardware R&D spending that underpins true AI capability. This perceived absence of domestic AI capex only sharpens concerns around the resilience of India's IT services export engine at a time when global enterprises are aggressively experimenting with agentic AI.

HSBC estimates that the direct impact of AI adoption across the technology stack will translate into 8–10% revenue contraction for Indian IT services, as productivity gains reduce billable work in key service lines. Crucially, this pressure is not instantaneous but staggered: it is expected to materialise over three to four years as existing contracts come up for renewal, implying an annual revenue drag of 3–4% through FY25–FY27.

Service line	Share of revenues	Share of work impacted by AI	Productivity improvement	Contraction of total revenues
Custom App-development	25.0%	60.0%	30.0%	4.5%
Custom App-maintenance	30.0%	35.0%	20.0%	2.1%
Enterprise app related work	15.0%	20.0%	20.0%	0.6%
ER&D	10.0%	10.0%	10.0%	0.1%
IMS (Infra services)	10.0%	10.0%	20.0%	0.2%
BPO/BPM	10.0%	40.0%	30.0%	1.2%
Overall	**100.0%**	**34.5%**		**8-10%**

Source: HSBC

The AI Super-Capex Wave: A Far More Important Structural Story

The near-term revenue pressure on services tells only part of the story, because the underlying shift in enterprise technology investment is far more expansive and structural. Multiple reports from top-tier venture investors and enterprise research show that AI spending has moved well beyond experimental pilot budgets into core IT expenditure, with enterprise AI budgets growing at rates north of 70–75% year-over-year as organisations embed generative models and agent-driven workflows into mission-critical systems. AI is no longer a discretionary line item; it is now a fundamental part of how firms optimise operations, build products, and reshape customer engagement. EY's *AIdea of India 2026* survey reveals that while adoption is broadening in India with 47% of enterprises running multiple GenAI use cases live, the early labour-market effects are already showing through.

Hyperscalers are projected to deploy USD 2 trillion toward AI infrastructure between 2025 and 2030; a rebuild of enterprise architecture from linear workflows to multi-agent systems. Historically, whenever the tech stack undergoes a major investment cycle, it eventually and often with a lag, creates more work for services firms, even if the early narrative is

dominated by fears of disruption. Even now, while markets fixate on near-term slowdown in legacy workstreams, the structural demand curve for engineering-heavy IT services is steepening.

New Playbooks and the Rise of Agile Tech (IT?) Providers

The earliest and clearest beneficiaries of the AI shift are not the incumbents but the small and micro-cap IT firms, which are structurally more agile and culturally predisposed to pivot quickly. A broad scan across websites, pitch decks, domain registrations, and investor communication shows the same pattern: firms that once sold “digital transformation” or “application development” now describe themselves as AI engineering partners, focused on agentic workflows, enterprise copilots, and multi-agent orchestration.

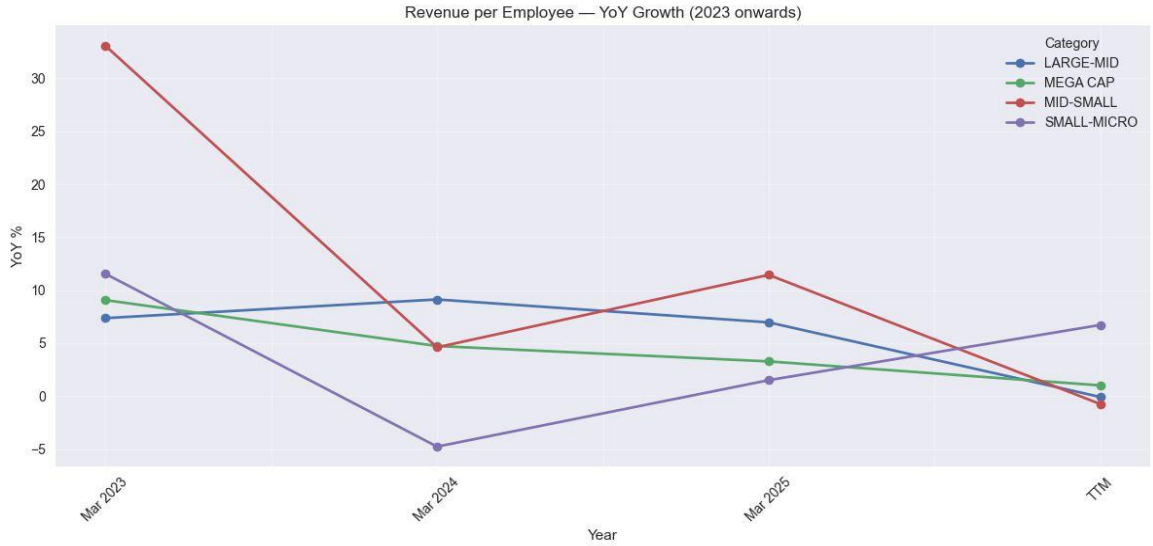

Many have adopted *.ai* domains, rewritten their service lines around model integration, RAG pipelines, autonomous agents, AI-led QA, and low-latency deployment pods. This is not branding theatre. It reflects **a** deliberate business-model shift in response to client budgets moving away from multi-year ADM/maintenance contracts toward automation sprints, AI proof-of-value cycles, and rapid prototyping engagements work that smaller vendors are structurally better positioned to execute.

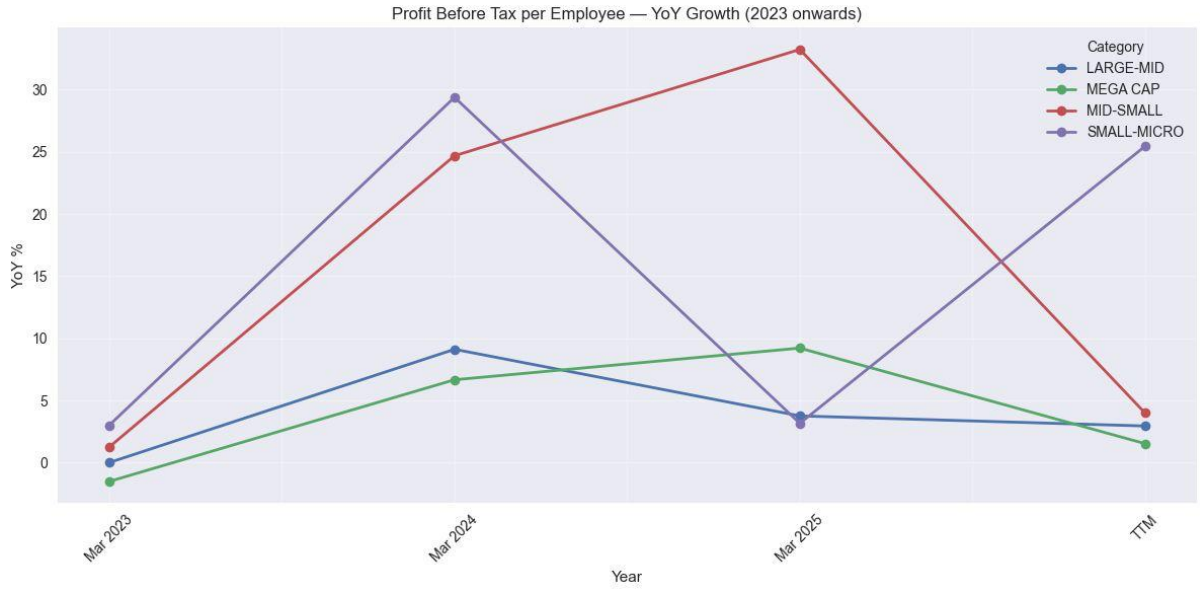

This is where the financial divergence becomes explicit. In our analysis, we saw that RPE and PBT/PE growth are materially higher for firms that have shifted their work mix toward AI-led, engineering-heavy mandates. These companies show clear gains in per-employee value creation because their revenue no longer scales linearly with headcount; automation, capability, and velocity replace the pyramid.

Conversely, firms; predominantly the mega-large cap and upper midcap incumbents which continue to rely on long-duration, effort-linked contracts, offshore-heavy pyramids, and rate-card billing show little to no movement in these metrics. Their financials are constrained precisely because this operating model insulates them from AI's upside while leaving them exposed to its deflationary pressures. The difference is unmistakable in the data: companies that rewired their business models for AI see measurable uplift; the scale players anchored to volume-driven contracting see stagnation.

Talent Flows as a Forward Indicator of AI Readiness

A core part of our thesis is that AI-led value creation in IT will follow talent, not headcount, so we built a metric to measure exactly where that talent is going. Using a multi-year panel of employee cost, employee count, and CPI-adjusted wage drift, we decomposed total payroll into two components: the inflation-adjusted cost of the existing workforce and the implied cost of new hires. Dividing this new-hire cost by the number of employees added gives an estimate of average compensation for incremental talent, and benchmarking it against prior-period average pay yields a New Hire Talent Score, effectively, a market-wide indicator of whether firms are attracting higher-skill, higher-cost, more AI-relevant talent or merely refilling the pyramid.

Table 1 : Top 10 Listed IT Companies by Talent Flows TTM			
Company	**Ticker**	**Category**	**NewHireTalentRank**
Coforge	COFORGE	MID	1
Saksoft	SAKSOFT	MICRO	2
Intellect Design	INTELLECT	SMALL	3
Sagility	SAGILITY	SMALL	4
L&T Technology	LTTS	MID	5
InfoBeans Tech.	INFOBEAN	MICRO	6
Persistent Systems	PERSISTENT	MID	7
Cyient	CYIENT	SMALL	8
Onward Technology	ONWARDTEC	MICRO	9
Affle	AFFLE	SMALL	10
Note: Coforge's score incorporates acquisition of Cigniti Technologies Ltd			

This signal is aggregated along with external sentiment data i.e. GitHub contribution patterns, job postings for LLM ops and agentic workflow roles, and the rise of AI-focused engineering footprints, all pointing to the same conclusion: the most capable, AI-native talent is flowing toward the more agile end of Indian IT, and that is where the compounding opportunity lies.

How We Interpret These Scores

The score is not a measure of scale; it's a measure of *intent*.
A high number means a company is deliberately upgrading its talent density by hiring significantly more expensive, more specialised engineers. A low number means incremental hiring is at or below the internal average, signalling continuity rather than capability-building.

When we aggregate the scores across categories, the conclusion is unmistakable:
The sharpest uplift in new-hire compensation is occurring in the small and micro-cap cohort, not among the mega caps or large mid-tier firms.

Net Value Added per Employee as a True Measure of Talent-Density Productivity

Net Value added per Employee (NVE) gives us a direct view into how much economic surplus each employee actually creates after accounting for the full cost of the workforce. It is essentially the purest signal of productivity, because it strips out headcount scale, billing-rate optics, and mix distortions and tells us whether a firm is becoming more efficient at converting talent and automation into value. Layering NVE with year-over-year NVE growth allows us to see whether that productivity is rising or stagnating, and the third layer i.e. NVE Momentum, the acceleration of that growth, acts as the most sensitive indicator of underlying change.

Momentum captures something traditional financials never show: *when a company quietly crosses an internal inflection point*, such as consolidating low-value work, deploying automation frameworks, rolling out agentic workflows, or structurally upgrading its talent base.

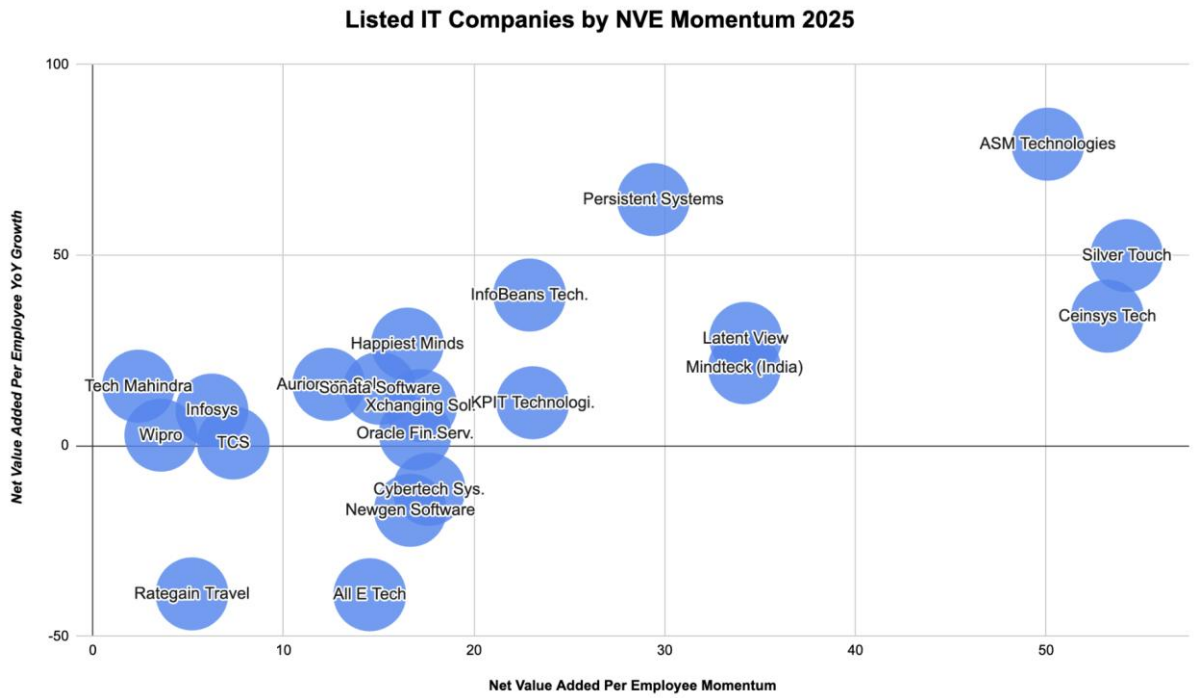

Only a small subset of firms shows positive NVE growth with rising momentum, indicating accelerating productivity and the hallmarks of real automation or AI-enabled delivery; another group shows high NVE but fading momentum, suggesting past efficiency gains that are no longer compounding.

The top of the table is dominated by smaller, more agile names such as ASM Technologies, Persistent, Silver Touch, InfoBeans, Ceinsys and Latent View — firms actively rebuilding themselves for an AI-heavy delivery model. By contrast, the bottom half is crowded with larger mid- and mega-cap incumbents like TCS, Wipro, Infosys, Tech Mahindra and OFSS, whose operating models remain tied to multi-year contracting, pyramid utilisation and rate-card billing, leaving their NVE growth flat to negative because AI is eroding the economics of scale-driven labour arbitrage faster than they can adapt. And a large share of the sector sits in the low NVE, low or negative momentum zone, where productivity is stagnant and AI rhetoric has yet to translate into operational change.

This matters because NVE Momentum is not cosmetic; it is the closest proxy for actual AI leverage inside an organisation. A turn in momentum signals that the firm is creating more value with the same workforce: automating repetitive layers, collapsing manual workflows, embedding AI tools internally, and shifting from effort-based to outcome-driven delivery structures.

How We Measured Adoption, Opportunity & Risk Using an LLM

To understand not just *whether* companies talk about AI but *how*, we built a transcript-analysis framework structured around a master LLM prompt and a classifier for categorising corporate language around AI.

For every company, all management commentary was broken into discrete sentences, and each sentence was evaluated by the LLM across three dimensions: Adoption, Opportunity, and Risk.

1. *Adoption* captures sentences where management explicitly describes AI systems that are deployed, live, implemented, or in production.
2. *Opportunity* captures forward-looking statements: plans, expectations, ambitions, or promises of future AI work.
3. *Risk* captures regulatory, execution, cost, or feasibility concerns around AI adoption.
4. Alongside this, the LLM also graded the sentiment tone of each AI-bearing sentence and the intensity of AI discussion relative to total transcript length. These individual components were then aggregated into a single AI Attention Factor that weights intensity, execution, and sentiment to show how substantively a company is engaging with AI.

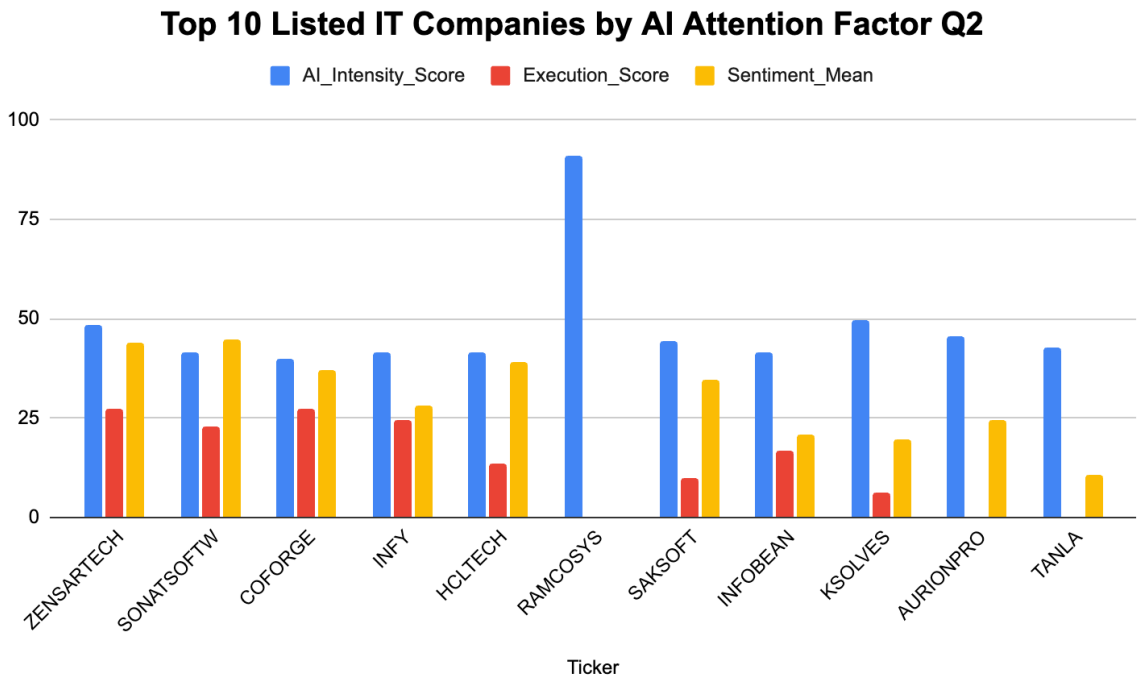

A large portion of companies show very high AI Intensity but near-zero Execution Scores, meaning they talk frequently about AI but provide almost no evidence of real deployments. Sentiment, too, is broadly muted, with many firms expressing AI optimism in abstract, non-committal language rather than in the context of concrete programs.

A smaller subset shows balanced but low-amplitude signals of moderate intensity, some evidence of adoption, but without strong conviction or scale. Only a very narrow segment exhibits all three signals in alignment: meaningful AI discussion, credible execution cues, and constructive sentiment. The distribution in the table makes this clear—most firms cluster

in the “high talk, low execution” quadrant, a handful show steady but modest implementation, and almost none display consistent, high-quality AI discourse across all dimensions. For our thesis, this reinforces a critical insight: the industry is saturated with AI narrative, but genuine operational engagement remains scarce and only companies scoring well across all three components are actually building AI capability instead of merely performing it.

Our Investment Thesis on InfoBeans

Our philosophy has always been to catch companies *before* the market realises they are compounding engines: typically in the ₹500–5,000 crore zone, where operational shifts translate directly into valuation re-ratings. We insist on clean balance sheets, credible ROE/ROCE, disciplined capital allocation, and entry valuations that leave room for error. Against that backdrop, InfoBeans stood out immediately. At a **~₹1,300 crore market cap**, it sits squarely in the high-asymmetry corridor we prefer - small enough for meaningful multiple expansion, yet large enough to have institutional-grade governance and client depth. Its **~20× PE**, with **ROE of ~12%** and **ROCE close to ~17%**, gave us valuation comfort *and* capital-efficiency reassurance.

1. Capability Shift: Hiring Patterns Reveal Upgrading of Talent Density

Our talent framework showed that InfoBeans has been bringing in higher-skilled, higher-paid laterals, rather than refilling the bottom of the pyramid. Management’s emphasis is clearly on building capability for more complex, higher-value engineering work rather than volume hiring.

For us, this validated the “talent follows transformation” signal: when the incremental employee is meaningfully more expensive, the company is preparing for a new type of work.

2. Short-Cycle, Engineering-Led Delivery Model: Perfectly Aligned With the AI Era

One of the underappreciated advantages of InfoBeans is its short execution cycles. The firm works on projects with durations short enough to allow rapid adoption of new technologies, iteration with clients, and frequent repricing of value. There is no legacy backlog of multi-year, fixed-rate maintenance contracts that slow down AI adoption.

The concall made it clear that most wins are engineering-led, six- to nine-month cycles, with increasing wallet share from existing clients. This agility is a structural advantage: AI rewards firms that move fast and reconfigure delivery quickly.

3. Real AI Execution: Multiple Live Projects, Accelerators, and Internal Tooling

Our AI Attention Factor already highlighted InfoBeans as one of the few companies with concrete AI execution cues. The concall reinforced this: management discussed multiple active AI engagements, the development of internal accelerators, ongoing investments in retrieval-augmented systems, and their push toward context-aware AI. They also referenced using AI internally to improve development velocity.

This is not narrative or posturing; it is operational AI. Most IT firms today are still speaking in broad, high-level terms. InfoBeans is delivering units of work that reflect the actual shape of the AI services market.

4. Productivity Inflection: Evidence of Operating Leverage and Efficiency Gains

Our RPE, PBT/PE, and NVE momentum analysis already indicated a business entering a phase of productivity acceleration. The concall confirmed improved utilisation, cost structures designed around laterals rather than freshers, and operational discipline that is beginning to translate into margin expansion. Management also highlighted that the firm is crossing revenue thresholds where incremental growth now falls to the bottom line.

This is exactly what we look for: companies where internal productivity begins compounding even before revenue growth re-rates.

Against this backdrop, the valuation looked not just reasonable but misaligned with the direction of travel. And the market's reaction validated that view:

The stock re-rated over 30% within weeks of our entry, as investor attention caught up to the operating inflection we had already mapped across talent signals, AI-execution metrics, and NVE momentum. In our framework, that's exactly what a correctly timed, correctly priced small-cap position is supposed to look like.