India's $300B AI Shock

WHAT IS THIS AND WHY SHOULD I READ IT? -

India’s $300 billion IT-BPO industry built much of urban India’s tax base, consumption engine, foreign-exchange cushion, and middle-class mobility story. As is already visible, AI is disrupting India's IT-BPO industry and hollowing out the wage pyramid underneath by reducing fresher hiring and replacing large delivery benches with much smaller teams. India needs to acknowledge that its most successful economic engine has become less labor-intensive and take urgent measures to establish industrial infrastructure for the AI era.

ABOUT THE AUTHOR -

Since 2023, Aman Kai Sidhant has invested in hardware, data and software companies that are beneficiaries of the AI supercycle at a US-India crossover fund. Previously, he was a Product Manager at Microsoft Azure.

Prologue: The Income Stream that Built Urban India

I spent a lot of my formative years in Gurgaon and feel like I grew up with the city. For readers unfamiliar with Gurgaon, imagine spending your high school years in 1990s Silicon Valley or 2000s Shenzhen. Every few months a new shopping mall, apartment building or office complex rose up and one could see the city develop in real-time. Even though I couldn't properly articulate why this was happening, the energy and excitement was palpable.

Source: X

The development of Gurgaon and other “IT cities” like Bangalore, Hyderabad and Pune are results of the white-collar services boom in the Indian economy that started in the 1980s with firms like Infosys, TCS and Wipro entering the IT-BPO market. Starting from 0 in the 1980s, India’s IT-BPO sector crossed $300 billion in revenues this year.

India's tech industry has grown *~6×* in two decades.

NASSCOM's Strategic Reviews show a sector that has expanded from **$48B in FY2006-07** to a projected **$315B in FY2025-26**. The pace has been uneven — definitional shifts, a GFC trough, and a COVID-era rebound all show up in the data.

Indian Technology Industry · Revenue & YoY Growth, FY2006-07 to FY2025-26F

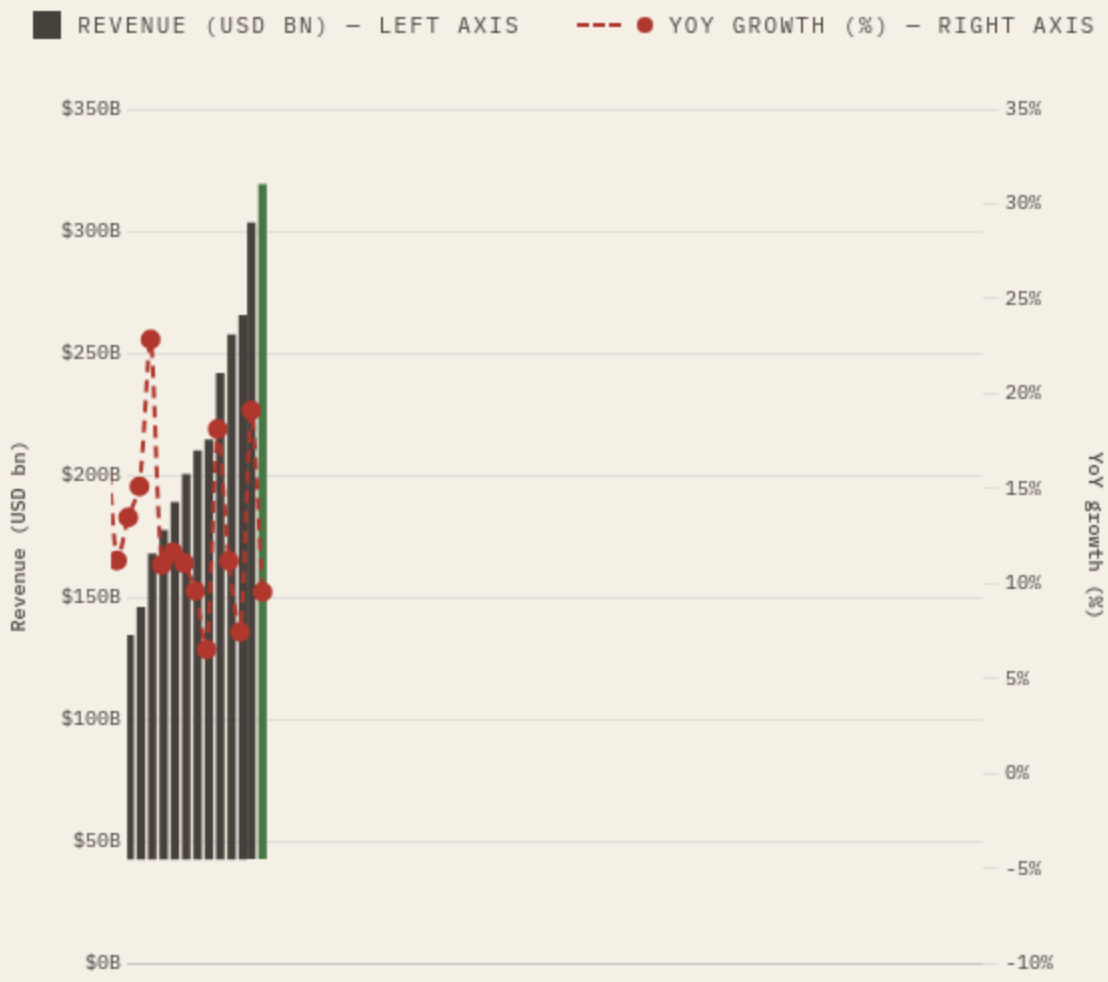

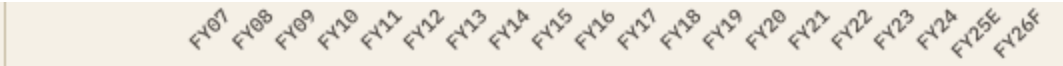

India's $300B AI Shock

SOURCE URLS

NASSCOM (National Association of Software and Service Companies) is the Indian IT industry trade body, founded in 1988. It publishes the annual Strategic Review report with industry revenue estimates based on member-company surveys, extrapolation, and internal modelling. NASSCOM is not a government statistical agency; its figures are revised year-to-year and represent an industry-body view of the sector.

- NASSCOM Strategic Reviews, 2014–2026, via nasscom.in
- PIB archive citations of NASSCOM for FY2006-07 to FY2009-10, via pib.gov.in
- NASSCOM SR press releases for FY2010-11 onward

DEFINITION NOTE

Revenue **includes hardware** from *FY2020-21 onward*. Earlier years are IT-BPM only — services, BPM, software products, and ER&D, mostly excluding hardware. This scope expansion adds roughly **$15–20B** to recent years and means FY20 → FY22 growth is partly definitional, not purely organic. The Indian fiscal year runs **April 1 to March 31**; FY2024-25 covers April 2024 to March 2025.

The industry came into its own after 1991, when India opened its economy after decades of protectionism and created policies such as the Software Technology Parks scheme that incentivized global companies to outsource software, customer support, accounting, HR, and back-office work to India. Overseas companies bought white-collar services from Indian firms, and those firms paid millions of Indian workers salaries large enough to buy apartments, cars, school seats, restaurant meals, flights, insurance, and mutual funds.

In this essay, I want to explore the risk that India's IT-BPO industry and workforce faces from AI. Firms may arguably continue winning contracts, raise margins, and sell more AI-enabled services. The deeper risk is that AI lets firms produce the same services with fewer workers, fewer freshers, smaller benches, and flatter wage ladders.

Coding is the first widespread commercial use case of AI and I think it's just the beginning. Claude Code went from 0 to $1B in 7 months and then to $2.5B annualized revenue in less than 3 months. It took Cursor nearly 2 years to get to $1B annualized revenue and less than 3 months after that to get to $2B. OpenAI's Codex has 3M weekly active users, up 5x since January 2026.

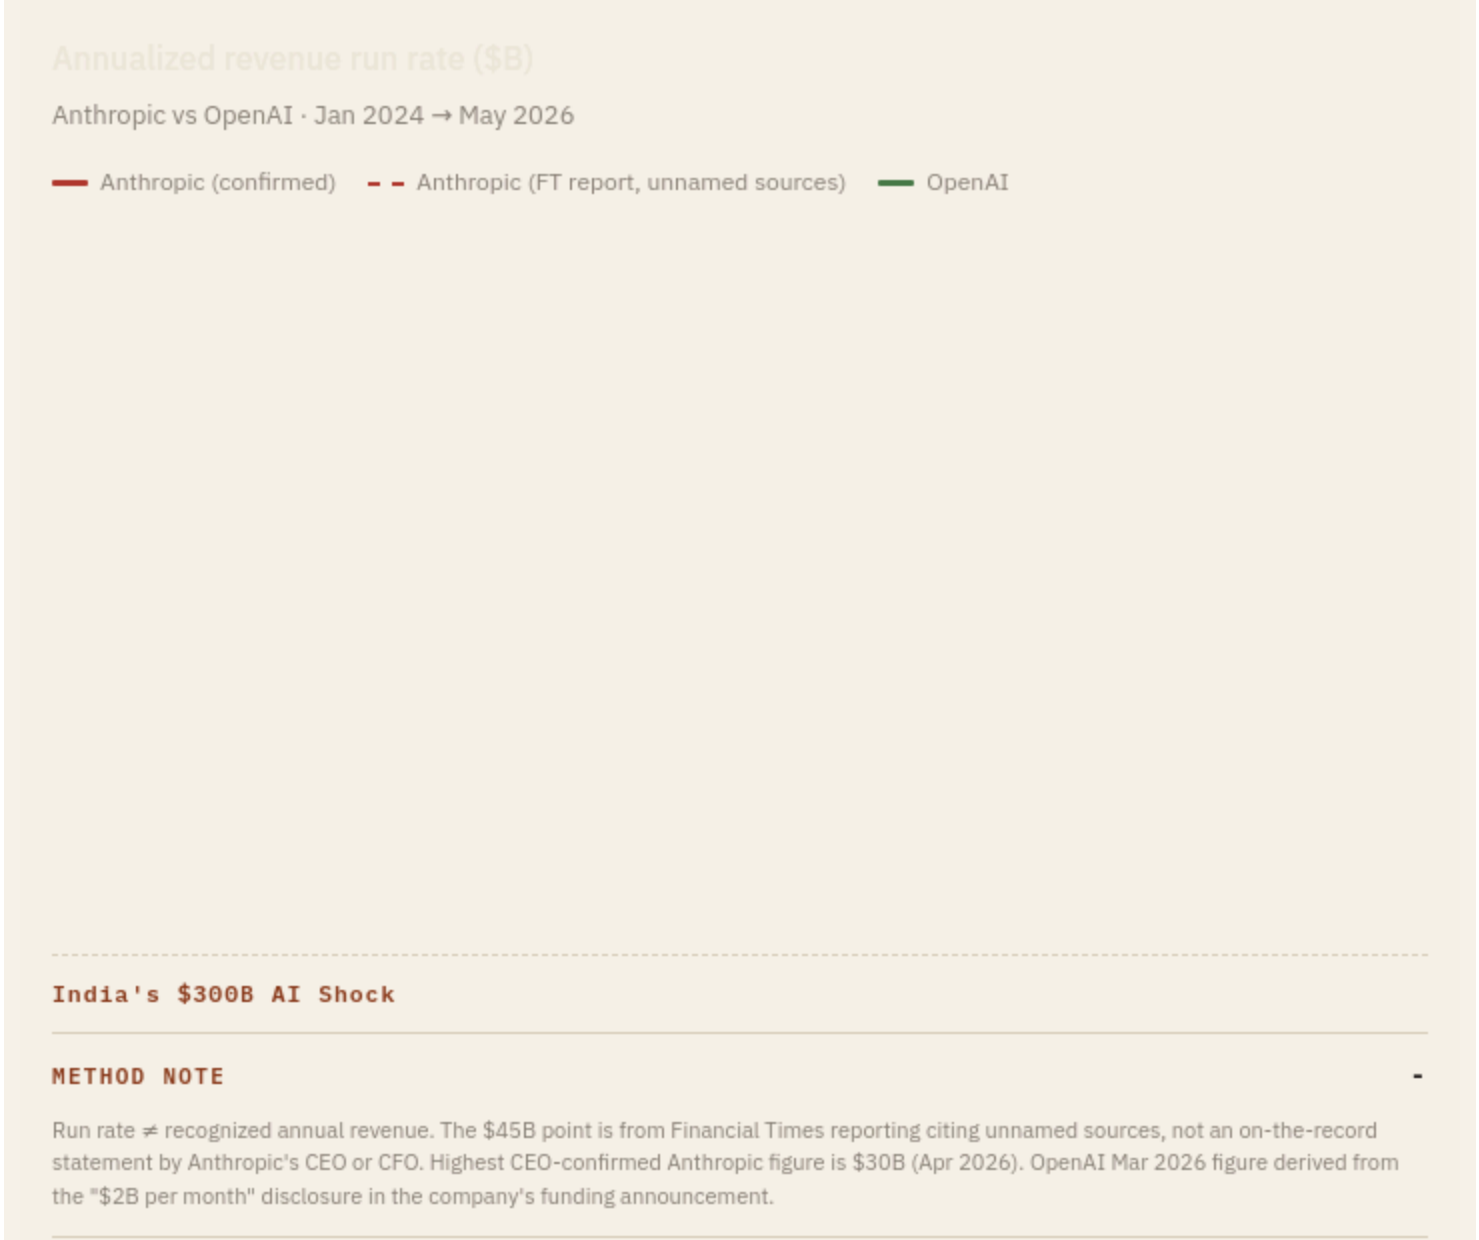

The same logic extends beyond coding. Customer support, finance and accounting, HR administration, document processing, compliance workflows, and general computer-use tasks are all exposed because they are structured, language-heavy, and historically labor-intensive. Markets have recognized this. Teleperformance, a 10.2B Euro revenue BPO that makes most of its money from tasks like customer service, has a market cap of 3.4B Euros as of end-of-day April 22.

Teleperformance SE — daily close with *event annotations*

Euronext Paris · EUR · Jan 2, 2024 – Apr 22, 2026 · Hover a catalyst marker to show its card

START	END	PERIOD HIGH	PERIOD LOW	TOTAL CHANGE
€133.30	€55.04	€153.20	€45.77	-58.7%

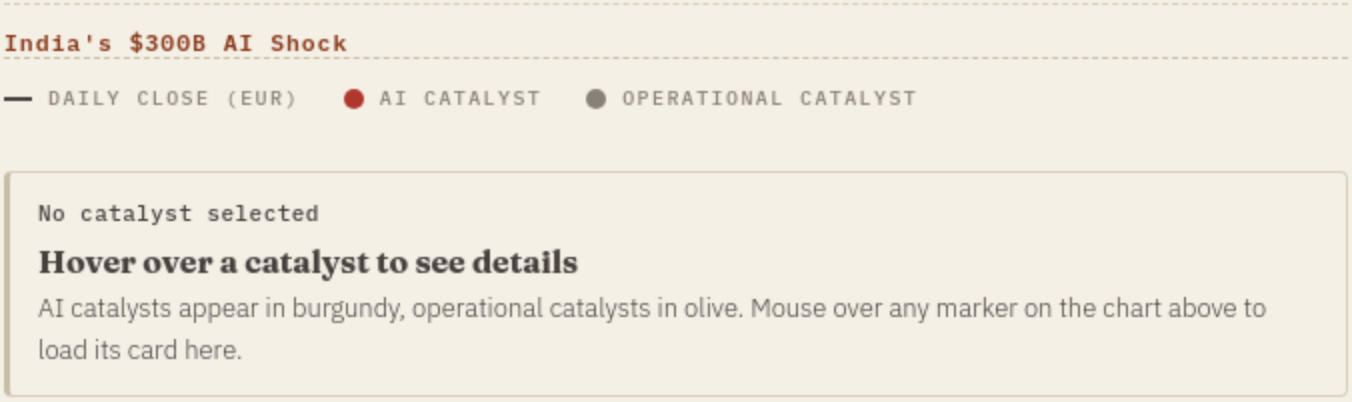

No catalyst selected

Hover over a catalyst to see details

AI catalysts appear in burgundy, operational catalysts in olive. Mouse over any marker on the chart above to load its card here.

A shock to IT-BPO wages and the industry more broadly won't stay limited to just the technology sector. India's post-1991 urban growth model was driven in large part by white-collar services exports. Foreign companies paid Indian IT-BPO firms, those firms

paid millions of Indian workers, and those salaries built a large part of urban India's tax base, consumption base, and upward-mobility story. India's formal income tax base is narrow and its upper-urban consumption class is small, economically powerful and a non-trivial percentage of this class works in IT-BPO. India's education system produces huge numbers of engineers for a labor market where the largest formal entry-level absorption channel will weaken because of AI. India's current accounts are financed in large parts by services exports.

AI should be treated as a sovereign economic issue for India. As demonstrated by Anthropic's Project Glasswing, it increasingly looks like access to the best AI models will be limited to a select few governments and companies with the capacity to pay for it or leverage of some kind. Just as resources like oil, electricity and an educated workforce were structurally important for the sovereignty of a nation in the 20th century, access to cutting-edge AI will be important in the 21st century. Access to frontier AI, compute, data centers, and the surrounding semiconductor and power supply chains will shape India's bargaining power in the coming economy.

Why Workers Get Hit Before Firms...

My hypothesis is that there will be large workforce cuts, hiring freezes for new college graduates and wage stagnation for remaining employees across IT services firms and BPOs over the next 2-3 years. Companies will still retain a small number of client-facing account managers and salespeople for bringing in new deals, as well as staff engineers/architects to steer agentic engineering tools and be accountable for the software engineering work that isn't coding.

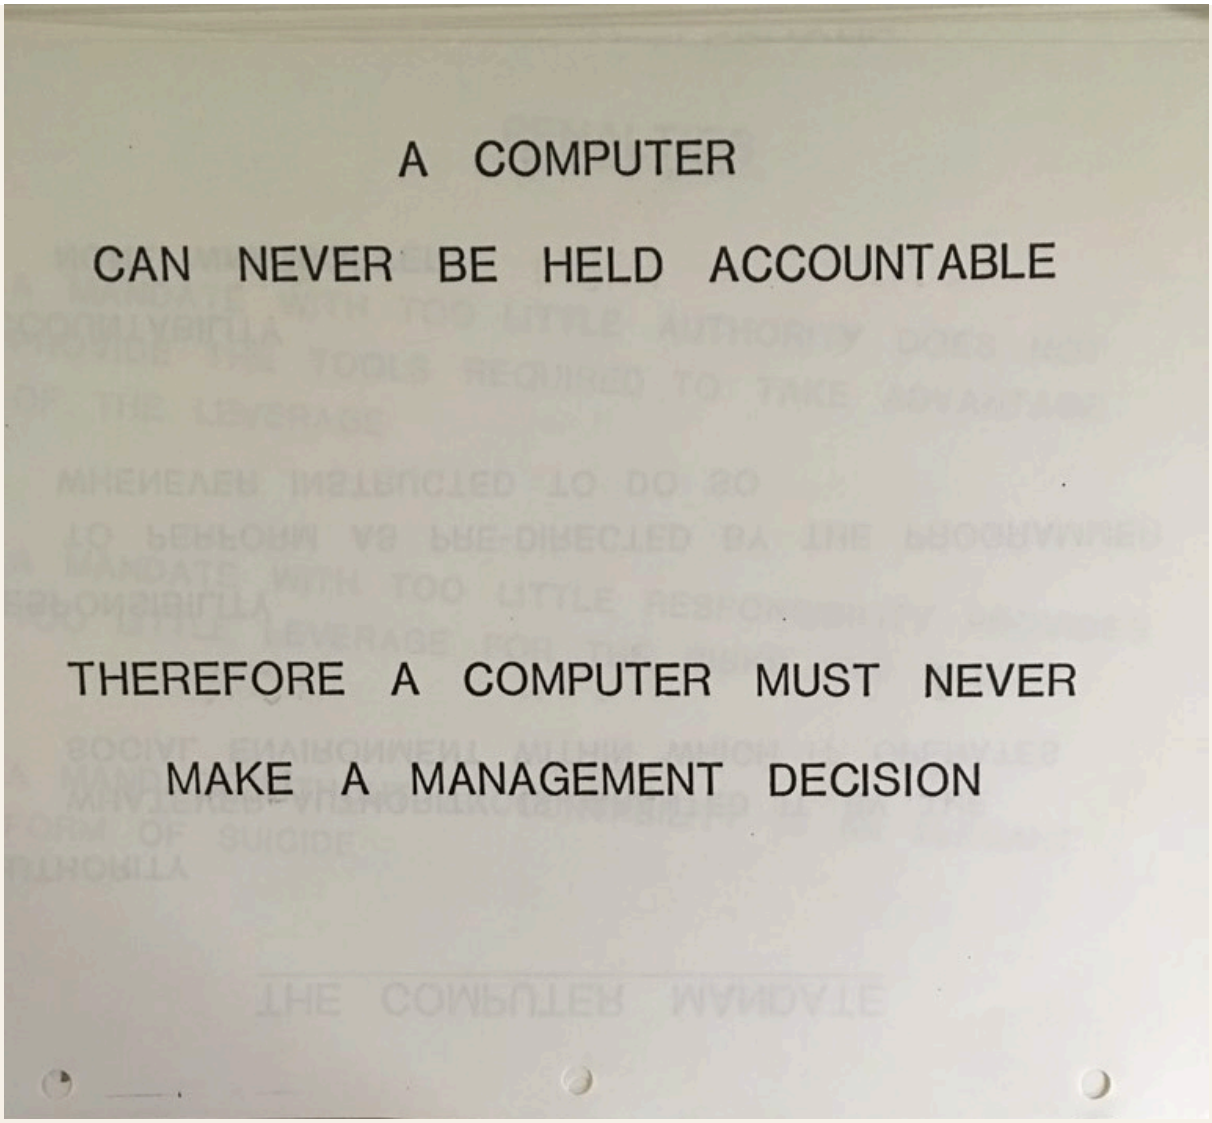

Source: IBM Corporate Training, 1979

When a Fortune 500 company hires TCS to migrate a legacy system to the cloud, most of the billable hours are spent on implementation work like writing API endpoints, configuring environments and running test scripts. Codex and Claude Code are already great at these tasks and they are getting better, which means that IT services firms may not need as many engineers as they did previously.

For high budget IT projects, there is a lot of interpersonal communication and project management involved that will still need humans. A VP of Supply Chain at Coca Cola would still want to talk to a human to get status updates. In the near term, the population at risk are the hundreds of thousands of junior to senior-level software engineers spending 40 to 60 hours a week doing implementation work when Codex or Claude Code

can run for 100 hours a week for much cheaper. Maybe even more if there are staff level engineers rotating shifts to prompt the AI.

Q4 FY26 earnings from TCS, Infosys, and Wipro reflect the patterns I'm worried about around revenue and headcount decoupling at India's three largest IT services firms. TCS shed roughly 23,500 net employees in FY26 while revenue declined 0.5% in USD. Infosys lost 8,000 employees in Q4 alone, which was more than the entire year's net addition of 5,000. Wipro quietly shifted its disclosure language from "230,000 employees" to "230,000 employees and business partners," a phrasing change that makes their overall headcount harder to interpret. Combined Top-3 revenue is essentially flat in constant currency in a year when global enterprise AI spending is at record levels.

Q4 FY26 earnings comparison for TCS, Infosys, and Wipro across revenue, margin, AI traction, headcount, capital return, and FY27 guidance

YoY growth view

USD reported | Constant currency

Metric (FY26, year ended Mar 31, 2026)	TCS	Infosys	Wipro
Top-line performance			
Revenue	$30.0B	$20.2B	$10.5B IT Services segment
YoY revenue growth	**-0.5%**	**+4.6%**	**-0.3%**
Profitability			
Operating margin	25.0% +70 bps YoY	21.0% adjusted, flat YoY	17.2% +20 bps YoY
Net income	$5.94B +3.5% YoY	$3.31B +4.9% YoY	$1.41B +0.5% YoY
AI traction			
AI revenue (annualized)	$2.3B ~7.7% of revenue	> 5.5% ~$1.1B+, growing	Not disclosed "AI Native" unit pivot
Large deal TCV (FY26)	$40.7B record; 5 mega deals	$14.9B +28% YoY; 55% net new	$7.8B +45% YoY
Workforce & capital			
Headcount (Mar 31, 2026)	584,519 −23.5K YoY	~325,000 −8K in Q4 alone	230,000+ incl. business partners
Capital returned to shareholders	~$4.2B	~$4.0B	~$1.6B

	₹39,571 cr dividends	₹37,500 cr div + buyback	₹15,000 cr buyback
Forward outlook			
FY27 guidance	Not formally issued	**+1.5% to +3.5%** CC, full year	**−2.0% to 0%** CC, Q1 sequential

India's $300B AI Shock

SOURCE URLS -

Sources: TCS Q4FY26 press release (Apr 9, 2026); Infosys Q4FY26 IFRS-USD press release (Apr 23, 2026); Wipro Q4FY26 press release (Apr 16, 2026). USD conversions for capital return use ₹93.83/$ (Fed cert rate, Mar 31, 2026).

In FY24 and FY25, the consistent narrative across all three firms was that “AI is a tailwind” and a productivity enabler. On the Q4 FY26 earnings call, Infosys CEO Salil Parekh used “compression” five times. For example: *“Compression is coming on some of the services and the growth is coming on other services. And the compression is typically in the areas where the AI foundation models and some of the tools are very efficient on that. So, you can see that in some of the tech services work, you can see that in some of the BPM work...”*

The bull case is that AI is indeed a new revenue stream. TCS reported $2.3B in annualized AI revenue at Q4 FY26, roughly 7.7% of total. Infosys’ AI revenue is “higher than 5.5%” and growing, though management refused to update the Q4 number specifically. The narrative they want you to leave with is that legacy services are shrinking, AI services are growing, and the net is fine. Two problems with that. First, AI revenue almost certainly employs far fewer workers per dollar than the legacy services it’s displacing because these are tools-and-platforms-heavy engagements, not body-shop staffing. Second, the AI base is still small relative to what it needs to offset. TCS’ non-AI revenue is roughly $27.7B, so even a modest 5% decline in that base is $1.4B of revenue that needs to be made up. AI revenue would need to grow from $2.3B to $3.7B (a 60%+ jump) to keep the headline flat, every year.

The fresher hiring collapse is another leading indicator. TCS, the industry's largest single employer, added a net 2,356 employees in Q4 FY26. This is the same firm that hired over 100,000 freshers in FY22. They also laid off 12,000 people and had 30,000+ net exits in 2025, the steepest decline in company history. Infosys, surprisingly, continues to hire. Wipro scaled back its college graduate (fresher) recruitment plans for FY26 and expected to hire ~8,000 graduates, compared to its earlier guidance of ~12,000. Overall, college graduate hiring for IT services in India has fallen from a peak of 600,000 in FY22 to about 120,000 in FY25. That's an 80% (!) decline.

BPOs are also extremely vulnerable, and this is starting to show in official RBI data.

The Reserve Bank of India, India's central bank, publishes revenue from Software and IT-Enabled Services Exports annually. Here is the report from the last 2 reported years (April 2023 to March 2025) -

India BPO Services — Sub-component Trajectories

Small multiples: each panel has its own y-axis, so small components are legible

METHODOLOGY NOTE

Each chart is scaled to its own maximum, so bar heights across panels are NOT comparable — only the trajectory within each panel is. Values are in US$ millions, except for Other BPO services, Finance/accounting/auditing & tax, and Total BPO, which are in US$ billions. Notable declines are highlighted first, followed by the remaining panels in order of absolute percentage change from FY22-23 to FY24-25, with Total BPO kept last as a reference.

NOTABLE

Medical transcription & doc. mgmt.

FY24-25: $100M

India's $300B AI Shock

▼ -86%

FY22-23 → FY24-25

NOTABLE

HR administration

FY24-25: $300M

▼ -25%

FY22-23 → FY24-25

NOTABLE

Content development & publishing

FY24-25: $300M

▼ -25%

FY22-23 → FY24-25

Business consulting / PR

FY24-25: $3,800M

▲ +46%

FY22-23 → FY24-25

Finance, accounting, auditing & tax

FY24-25: $8.2B

▲ +19%

FY22-23 → FY24-25

Supply chain & logistics

FY24-25: $800M

▲ +14%

FY22-23 → FY24-25

Other BPO services

FY24-25: $42.3B

▲ +7.1%

FY22-23 → FY24-25

Total BPO (for reference)

FY24-25: $55.8B

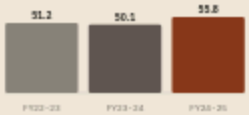

▲ +9.0%

FY22-23 → FY24-25

SOURCE URLS

Source: RBI Annual Survey on Computer Software & ITES Exports — press releases prid=58937 (Oct 2024) and prid=61562 (Nov 2025)

Medical Transcription & Document Management reduced from $800M to $100M. That's a 87.5% drop IN ONE YEAR. This coincided with the rise of several medical notetaking startups like Abridge and Ambience as well as big jumps in voice AI capabilities from the frontier labs.

Both HR Administration and Content Development are down 25% each, from $400M to $300M. Not as big a drop, but still concerning.

I may be overestimating the impact of AI on these cuts. Maybe the BPO revenue declined because of outsourcing to other countries like the Phillipines. Maybe customers have lower budgets for discretionary software development projects that IT Services companies build.

Having said that, it's not implausible to imagine these cuts continuing in the wake of AI that is superhuman at digital tasks and coding. Early evidence is already visible in the US. Bloomberg recently reported that 17 occupations flagged by the Bureau of Labor Statistics as exposed to AI, covering roughly 9 million jobs, saw employment fall 1.6% between May 2024 and May 2025 for the second year in a row, even as overall employment in the US rose 0.8%. The sharpest losses were in BPO-adjacent roles that comprise large parts of India's services-export machine: customer service representatives fell by 130,180 jobs, or 4.8%.

BLS OEWS · MAY 2024 TO MAY 2025

Five AI-exposed occupations *lost 216,340 jobs* in a year.

May 2024 and May 2025 employment estimates for selected roles frequently cited as exposed to generative AI. Bars are indexed within each occupation so the year-over-year contraction is visible across roles of very different sizes.

SELECTED ROLES	COMBINED DECLINE	LARGEST PCT. DROP
-216,340	-3.6%	-18.1%
jobs across the five occupations below	from 6.00M to 5.78M estimated jobs	technical writers

JOBS LOST BY OCCUPATION

Each bar is sized by absolute jobs lost between May 2024 and May 2025, scaled against the largest decline (customer service representatives, -130,180). The percentage decline appears on the right; the workforce totals are shown for context.

JOBS LOST · MAY 2024 → MAY 2025

Customer service reps

2,725,930 → 2,595,750

-130,180 jobs

4.8%
DECLINE

Secretaries

1,737,820 → 1,706,790

-31,030 jobs

1.8%
DECLINE

Sales reps

1,266,860 → 1,238,190

-28,670 jobs

2.3%

DECLINE

Graphic designers

214,260 → 197,830

-16,430 jobs

7.7%

DECLINE

Technical writers

55,530 → 45,500

-10,030 jobs

18.1%

DECLINE

India's $300B AI Shock

SOURCE URLS

Bureau of Labor Statistics, Occupational Employment and Wage Statistics (OEWS). May 2024 estimates (released 2 Apr 2025) and May 2025 estimates (released 15 May 2026).

bls.gov/oes · News release: USDL-26-0725

CATEGORY NOTES

"Secretaries" refers to secretaries and administrative assistants except legal, medical and executive. "Sales reps" refers to wholesale and manufacturing sales representatives except technical and scientific

products.

OEWS excludes self-employed workers and is a model-based estimate from a three-year rolling sample, not a point-in-time headcount.

...And Firms Get Hit Eventually

The immediate assumption to AI's impact on IT-BPO is to assume that revenue will eventually decline because the core value proposition (cost arbitrage on coders) is now broken. IT Services firms are preparing for this new reality - TCS is starting to build AI data centers. From the press release and other online sources, it seems like TCS will build and operate purpose-built data centers where customers can place or run GPU-heavy AI infrastructure to run AI workloads. Essentially, instead of the output being code as it was traditionally, the core product for this business is real-world infrastructure.

This is a survival strategy for the firm, but not necessarily for the workforce. A data-center business can protect TCS revenue and shareholder value, but it does not need the same pyramid of junior engineers, project managers, testers, support workers, and freshers that traditional IT services did.

BPOs and GCCs might be more fragile than IT services, as we'll see below.

IT Services

The strongest bull case for Indian IT services is that offshore human labor may remain more economical than frontier AI subscriptions, especially while model pricing is artificially subsidized by AI companies. But this assumes AI prices stay high and that clients keep buying labor-hours rather than outcomes.

Let's look at price levels for doing the same coding tasks -

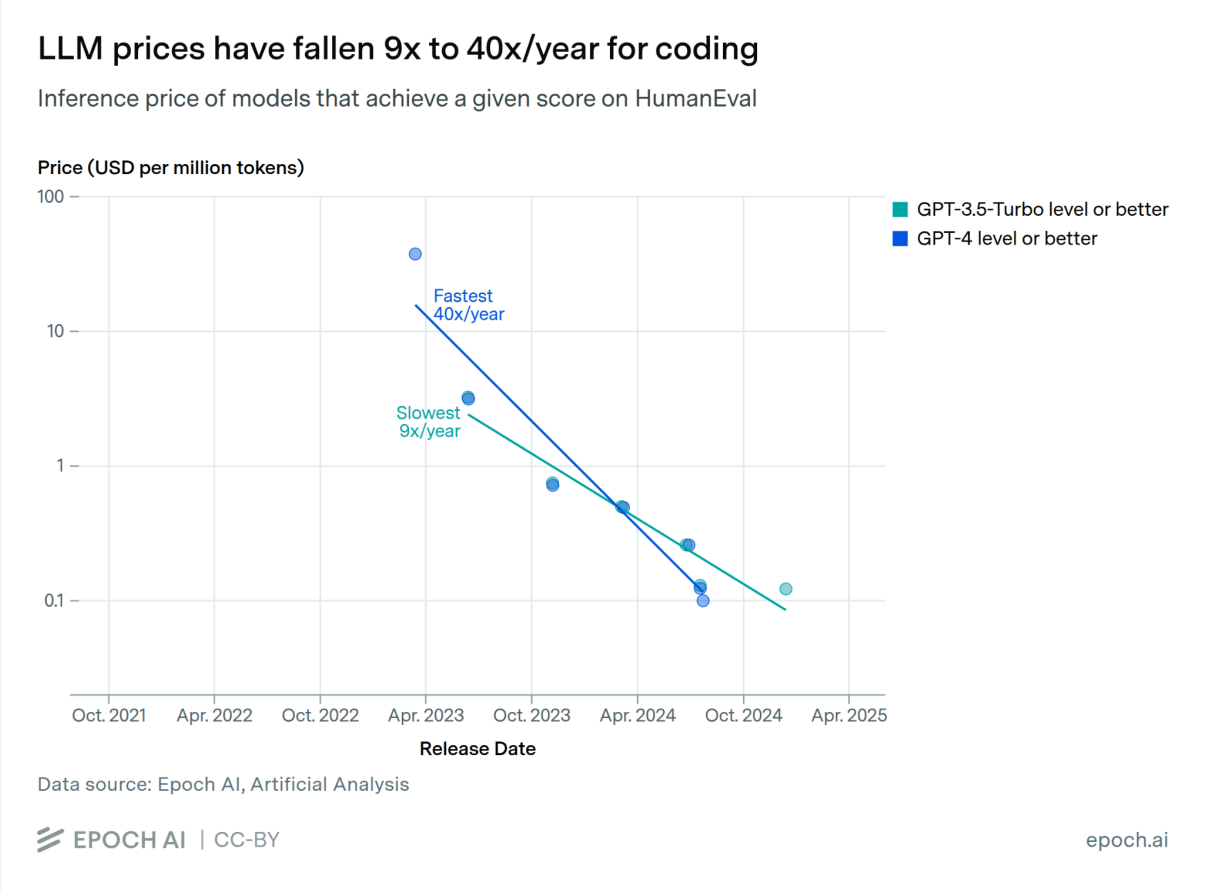

https://epoch.ai/data-insights/llm-inference-price-trends

Let’s look at hardware -

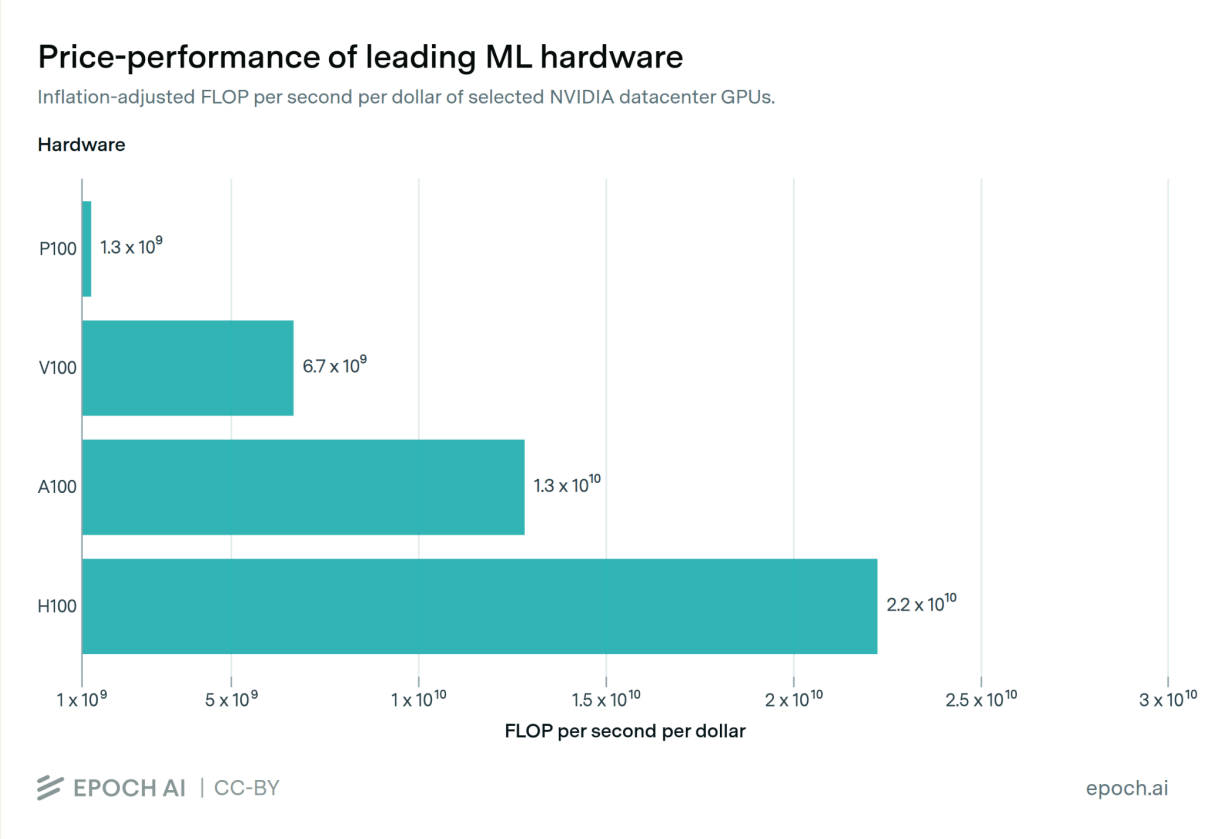

https://epoch.ai/data-insights/price-performance-hardware

AI is getting cheaper through two reinforcing channels: first, the underlying hardware is improving steadily, with GPUs delivering roughly 30% more performance per dollar each year after adjusting for inflation and server costs. Second, the price of using LLMs at a fixed capability level is falling much faster, with the cheapest model/API price to hit given benchmark thresholds declining by 9x to 40x per year for coding. This makes it obvious to me that agentic engineering will reduce traditional IT Services revenue where the deliverable is code.

Another threat compounding this is that OpenAI and Anthropic have launched their own AI services firms in partnership with leading financial and consulting firms. A big part of selling software to corporations was successfully creating a sales machine that can reach these corporations. This distribution was a big moat for large IT services firms vs. smaller firms because they had relationships with senior executives at most large corporations that have historically been buyers of services work. By partnering with Bain, Capgemini, McKinsey (in OpenAI's case) and Accenture, Deloitte, PwC (in Anthropic's case), OpenAI

and Anthropic have circumvented the distribution challenge because these firms have the same relationships that an Infosys or Wipro has with corporations globally.

Let's also understand why companies outsourced work to IT services firms in the first place, and how agentic engineering impacts all of them:

Converting Fixed Costs to Variable Costs

When companies hire full time employees (in the US), they are also responsible for their 401k, healthcare and other benefits, as well as severance and PR issues if the company goes through a downturn and needs to do a RIF. If instead of hiring 100 engineers you outsource that staff to Accenture or Infosys, downsizing and upsizing just involves talking to your account manager.

How AI affects this:** Claude Code doesn't sleep, can spin up subagents and output tokens 24/7 as long as the prompts keep coming. You might need an onsite solutions architect to collect project requirements and provide the big-picture context to Claude Code, but that's it. Maybe 2 senior engineers to watch for tech debt from Claude Code's output. The team of 30 is now a team of 3. This aligns with RBI's data where only 9.3% of services output is on-site (same geography as the customer).

Build vs. Run

The best software engineers don't want to maintain a Java app written 10 years ago. IT services firms were an easy way for enterprises to outsource boring "keep the lights on" work while their internal engineers worked on shiny new tech. This is why WITCH companies have a reputation for tedious, implementation-heavy work.

How AI affects this:

The tedious, implementation-heavy work that was the bread and butter of IT services firms is what agentic engineering tools excel at. There are still some rough edges around code review and AI sometimes making code tests pass in insidious ways like deleting test cases, but given the massive traction frontier labs are seeing from coding, these issues are likely front and center for them to solve next.

Skills at Scale

When a Fortune 500 company decides to migrate to SAP S/4HANA before the 2027 deadline, they need 150-300 SAP certified consultants for 18-24 months across dozens of modules. There is an entire ecosystem of consultants that exists because SAP is famously impossible to implement and work with while being extremely powerful. So powerful that Microsoft, despite having its own ERP product, uses SAP.

The same pattern applies to other software tools like ServiceNow and Salesforce as well as broader software practices like cloud migrations. No company maintains a bench of consultants or software engineers specifically for these workloads because these skills are only needed once.

How AI affects this: I'm not sure how much custom SAP logic was part of Claude Code and Codex's training data, but Anthropic has already announced that Claude Code can modernize COBOL code cheaply despite the well-documented dearth of publicly available COBOL code. It is now clear that models have emergent capabilities that

weren't part of the explicit training goals simply as a consequence of general improvements.

If an AI agent can generate boilerplate SAP configuration, auto-generate test scripts, and handle documentation, a team of 150 becomes a team of 10. The IT services firm still wins the deal, but the deal is worth a lot less compared to what it used to be. Revenue per engagement compresses even if the number of engagements holds steady.

There are several startups attacking this problem too, which opens up IT services topline from multiple angles of attack. From a16z's recent SAP piece:

- Axiamatic is an AI "assurance" layer for ERP: it builds a knowledge graph from project artifacts and flags hidden failures in requirements/change management via Slack/Teams to de-risk and accelerate S/4HANA programs (partnered with SAP Build; baked into KPMG/EY/IBM motions).
- Conduct is a code- and process-mapping copilot that generates a semantic layer and technical documentation across ECC→S/4, with Q&A over custom tables/APIs to speed internal takeover.
- Auctor does agentic implementation delivery for SIs/pro services, auto-capturing discovery into structured requirements before becoming a system of record for SOWs, design docs, user stories, configs, and test plans.
- Supersonik helps with AI-powered product enablement for channels/MSPs and customers – vision and voice agents that teach inside the real UI, reducing SE headcount needs and enabling reseller-led implementations/expansions.
- Tessera's AI-native SI manages enterprise transformations end-to-end – connecting into a customer's existing ERP instance, evaluating how it's implemented, and then flagging/automatically remediating what needs to be changed during migration.

Baked In Today, Replace Tomorrow

Investor Signaling

If outsourcing used to signal cost-efficiency and discipline to Wall Street previously, how long until analysts start asking CEOs about code-driven efficiencies in their internal

engineering teams and outsourced IT contracts? The significant jump in coding abilities is fairly recent. I expect this to become more of an issue in the forthcoming quarters' earnings calls. Block is frontrunning this with coordinated PR signaling that the recent 40% cuts were explicitly because of new agentic engineering capabilities.

How AI affects this: Obvious.

BPOs

BPOs will probably see the quickest disruption. We already saw the Medical Transcription & Document Management reduction before. BPO work is especially exposed because, unlike software engineering, it is harder to argue that AI simply makes existing workers more productive and allows firms to build more. Much of BPO is not "building" anything new. Customers generally want the same output at a lower cost, which is exactly the labor-arbitrage model that made Indian and Southeast Asian BPOs so successful in the first place. Even in India, API calls that can run 24x7 are cheaper and will produce better outcomes than hiring humans.

Let's look at the RBI data again.

India BPO Services — Sub-component Trajectories

Small multiples: each panel has its own y-axis, so small components are legible

METHODOLOGY NOTE

Each chart is scaled to its own maximum, so bar heights across panels are NOT comparable — only the trajectory within each panel is. Values are in US$ millions, except for Other BPO services, Finance/accounting/auditing & tax, and Total BPO, which are in US$ billions. Notable declines are highlighted first, followed by the remaining panels in order of absolute percentage change from FY22-23 to FY24-25, with Total BPO kept last as a reference.

NOTABLE

Medical transcription & doc. mgmt.

FY24-25: $100M

India's $300B AI Shock

▼ -86%

FY22-23 → FY24-25

NOTABLE

HR administration

FY24-25: $300M

▼ -25%

FY22-23 → FY24-25

NOTABLE

Content development & publishing

FY24-25: $300M

▼ -25%

FY22-23 → FY24-25

Business consulting / PR

FY24-25: $3,800M

▲ +46%

FY22-23 → FY24-25

Finance, accounting, auditing & tax

FY24-25: $8.2B

▲ +19%

FY22-23 → FY24-25

Supply chain & logistics

FY24-25: $800M

▲ +14%

FY22-23 → FY24-25

Other BPO services

FY24-25: $42.3B

▲ +7.1%

FY22-23 → FY24-25

Total BPO (for reference)

FY24-25: $55.8B

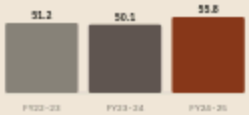

▲ +9.0%

FY22-23 → FY24-25

SOURCE URLS

Source: RBI Annual Survey on Computer Software & ITES Exports — press releases prid=58937 (Oct 2024) and prid=61562 (Nov 2025)

Several buckets within BPO Services actually grew. Business Consulting makes sense, since it is primarily advisory work and still requires significant human interaction. Finance also makes sense, since AI was likely not good enough by March 2025 to reliably handle critical tax, bookkeeping, compliance, and reconciliation workflows. Anecdotally, Codex and Claude Code have become extremely capable at Excel tasks for me, and frontier labs have made Finance tasks a priority for models. Since BPO contracts are often multi-year, contraction may not be immediately visible in RBI's reports, but it is bound to happen.

Surprisingly (to me at least), Other BPO Services added $5B in export revenue. Based on the segmentation, I would assume this includes work like customer support, data entry and other back-office work. I think this bucket will see adverse impact starting this year.

Sierra, Bret Taylor's company that started off selling AI agents for customer service, crossed $150 million in ARR in eight quarters and just raised at a $15 billion valuation. That kind of ARR multiple implies that Sierra is continuing to see tremendous growth and (speculation) is likely to expand to other use cases beyond customer service currently done by humans. Even if they don't, the fact that customer service alone is a $400 billion market means that a lot of that revenue is going to go from companies where humans manage customer support tickets to companies like Sierra where AI does the work.

Several other AI-native customer service companies are also growing rapidly, and from due diligences I have done at work, AI already has a deflection rate of 50-60% for most Tier 1 support cases across voice and text.

GCCs

GCCs (Global Capability Centers) are a fancy way of saying offshore office. India is a hub for GCCs. From an Indian Government article in November 2025:

“India now hosts over 1,800 GCCs, representing 55% of the global total. These centres employ 1.9 million professionals and generate Rs. 5,72,873 crore (about $60 billion) in export revenue as of FY25. In total, the GCC sector supports 10.4 million jobs, with average salaries 25 to 30% higher than the national average.”

“By FY30, GCCs are expected to contribute nearly Rs. 8,86,800 crore (about $93 billion) in exports, underscoring India’s growing strength in innovation, compliance, and digital excellence.”

The GCC Sector in India supports 10.4 million jobs. That’s higher than the entire population of some European countries. GCCs are at risk for similar reasons we discussed above (Skills at Scale, Build vs. Run). And it’s starting already. Oracle laid off ~10,000 people in India recently, which is 20% of their entire Indian workforce. Sure, they’re spending a lot on capex, but they’re not going to RIF engineers if they didn’t have any alternatives ready to go.

Jevons Paradox and The Radiology Counter Argument

To steelman the other side ie IT Services companies will keep all their existing employees, and retrofit all of them with Codex 5.4 on extra high so they can build more.

This works if you’re a startup trying to build as much product as quickly as possible while keeping headcount low. It breaks down when you’re a publicly traded company responsible for maximizing shareholder value. If you’re an executive at Infosys, you have two options - either pay an average employee Rs. 15 lakhs (about $16,500) per year or pay Rs. 2.4 lakhs (about $2,700) for an annual Codex Pro subscription (likely cheaper on an enterprise-wide license). When your performance is judged on margins and earnings per share, the answer is obvious, and you will be outcompeted by firms that decide to rearchitect around AI.

When talking about the impact of AI on software engineer employment, many people cite Geoff Hinton’s argument around the risk to radiology that never panned out. Software

engineering is fundamentally different. Fields like healthcare and law REQUIRE a human to be accountable for the final decision. Software engineering, particularly the cost-optimization kind that IT services firms sell, doesn't require this degree of responsibility. BPO costs are even less safe because a company actively wants to reduce that spend as much as possible if it's possible to do so with the same level of quality.

This excellent piece by David Oks explains that labor impacts due to technology happen not when the tasks get automated, but when the *paradigm* changes. That is a fancy way of saying that you need the entire workflow to change for technology to impact labor.

I posit that AI may have automated major parts of software engineering (writing code) but it hasn't changed software engineering as a whole (yet). The reason this time may be different is that earlier waves of software mostly helped humans do the work faster. Excel did not replace the analyst, it made the analyst more productive. IDEs did not replace programmers, they made programmers more productive. AI agents are different because they increasingly perform the work itself: reading requirements, writing code, calling tools, generating tests, debugging errors, summarizing context, and iterating toward completion. Jevons paradox may still increase total demand for software, but the new demand does not necessarily flow back into human headcount if agents are the marginal workers satisfying it.

I believe coding agents are enough of a paradigm shift for the IT-BPO services workflow, where the entire pitch to customers is low-cost, scaled delivery rather than technical depth or architectural innovation. An increase in software demand can coexist with a collapse in the old offshore labor pyramid.

Anyone who has worked in a company will relate to this from Stratechery -

"It's always been the case, even in large companies, that a relatively small number of people actually move the needle and drive the company forward in meaningful ways. That drive, however, has been filtered through a huge apparatus, filled with humans, who accelerate the effort in some vectors, and retard it in others. That apparatus makes broad impact possible, but it carries massive coordination costs.

Agents, however, will tilt much more heavily towards pure acceleration, making those drivers of value much more impactful. I'm sympathetic to the argument that the best companies will want to use AI to do more, not simply save money; the reality of large organizations, however, is that the positive impact of AI will not be in eliminating jobs, but rather replacing hard-to-manage-and-motivate human cogs in the organizational machine with agents that not only do what they are told but do so tirelessly and continuously until the job is done."

What I think will happen (optimistically) - services firms will adopt AI as quickly as possible, with a select number of exceptional staff-level software engineers to steer the AI. These firms will win over the ones that stick to optimizing for billable hours and filling their bench with engineers.

The bear case is that companies bring the majority of this outsourced IT and BPO work in-house now that coding agents are more than good enough to do the median IT-BPO task.

Small Workforce, Large Macro Effect

India's IT-BPO workforce is roughly 5.8 million people. Compared to India's labor force of 617.6 million, this population is <1%.

This workforce is also one of the most economically productive populations in the country. A bottoms-up analysis of TCS, Infosys, and Wipro's audited FY25 employee benefit costs (SEC 6-K filings) gives a weighted-average compensation of Rs. 24.5 lakh (about $27,000) per employee. Taking into account average BPO salaries, the realistic weighted average for IT-BPO overall sits around Rs. 14.7 lakh (about $16,000) take-home, which is roughly 7.6 times India's per capita Net Income of Rs. 1.93 lakh (about $2,000)

What does an Indian IT-BPO worker *actually* earn?

A bottoms-up build from listed-company FY26 disclosures (IT) and FY25 annual reports (BPO — FY26 ARs not yet filed). Every input traces to a primary source.

Read this first. The IT input is *employer-side* employee benefit cost per employee; the BPO input is *employee-side* median disclosed remuneration. The result here is a proxy for the average IT-BPO worker's salary. FY26 IT figures include a one-off bump from the Indian Labour Codes (notified 21-Nov-2025) raising gratuity/leave provisions — strip ~1–2% off the IT gross for a clean wage-growth read.

TIER 01 · INPUTS Company-disclosed financials

IT figures from FY26 Q4 Ind-AS Consolidated filings. BPO figures from FY25 annual report remuneration disclosures (FY26 annual reports due Q3 2026).

IT SERVICES · FY26 · IND-AS CONSOLIDATED, ₹ CRORE UNLESS NOTED

FIRM	REVENUE	HEADCOUNT	EMPLOYEE BENEFIT EXPENSE	SOURCE
TCS	₹2,67,021	584,519	₹1,54,994	Q4 FY26 Fact Sheet
Infosys	₹1,78,650	328,594	₹95,094	SEC 6-K, Exh 99.3
Wipro	₹92,624	242,156	₹55,586	Q4 FY26 press release

BPO SERVICES · FY25 · REMUNERATION DISCLOSURES, ₹/YEAR

FIRM	SELECTED PROXY (MEDIAN)	HEADCOUNT WEIGHT	BASIS	SOURCE
Firstsource	₹3,19,332	24,751	*Standalone company-wide median*	FY25 AR
eClerx	₹3,78,714	16,626	*India-only (ex-Board/KMP);*	FY25 AR

			count-weighted M/F	
Alldigi Tech	₹2,54,324	4,817	*India-only (ex-Board/KMP); count-weighted M/F*	FY25 AR

Note. *The IT sample is top-3 by revenue and represents the upper bound of the Indian IT-services pay stack — mid-tier services (LTIMindtree, Mphasis, Coforge), staffing firms, and small-shop IT pay materially less. The BPO sample is three firms covering ~46K employees against an industry of 1M+; FY25 data — FY26 BPO updates land July–September 2026. FY26 IT employee benefit expense includes a one-off bump from the Indian Labour Codes (notified 21-Nov-2025) which raised gratuity and leave provisions — Infosys flagged ₹1,289 cr specifically; the underlying gross wage-growth read is closer to 5–6% YoY rather than the 8% the raw EC-per-employee figure suggests.*

TIER 02 · DERIVATION

Per-firm calculation → segment proxy

Each row shows the formula and output. Highlighted rows are segment aggregates that feed into Tier 3. The take-home adjustment converts IT's employer-side cost to an employee-side cash proxy.

	CALCULATION	₹ LAKH / YEAR	$ EQUIVALENT
IT SERVICES · FY26 · EMPLOYER COST PER EMPLOYEE			
TCS	(₹1,54,994 cr ÷ 584,519) × 100	**₹ 26.52**	**$28,265**
Infosys	(₹95,094 cr ÷ 328,594) × 100	**₹ 28.94**	**$30,844**
Wipro	(₹55,586 cr ÷ 242,156) × 100	**₹ 22.95**	**$24,460**
Headcount-weighted IT (gross)	(Σ EC ÷ Σ HC) × 100 = (₹3,05,674 cr ÷ 1,155,269) × 100	**₹ 26.46**	**$28,201**
Headcount weighted IT (net after benefits and income tax)	26.46 × 0.80 → IT take-home pay proxy	**₹ 21.17**	**$22,562**
BPO SERVICES · FY25 · MEDIAN REMUNERATION			
Firstsource	₹3,19,332 ÷ 100,000	**₹ 3.19**	**$3,400**
eClerx	₹3,78,714 ÷ 100,000	**₹ 3.79**	**$4,039**
Alldigi Tech	₹2,54,324 ÷ 100,000	**₹ 2.54**	**$2,707**
Headcount weighted BPO median	Σ(proxy × HC) ÷ Σ HC ÷ 100,000 = ₹15,42,53,59,036 ÷ 46,194 ÷ 100,000	**₹ 3.34**	**$3,560**

median

Note. *IT proxy is a mean (Σ benefit expense ÷ Σ headcount); BPO proxy is a count-weighted median. Right-skewed salary distributions mean the IT side overstates a "typical" worker. The 80% multiplier approximates stripping out non-cash benefits (ESOPs, gratuity, PF employer share). USD equivalents at ₹93.83/USD, the certified Federal Reserve rate on 31-Mar-2026 (cited in Wipro Q4 FY26 press release).*

TIER 03 · CROSS-SEGMENT

Final blended IT-BPO proxy

Combine IT take-home and BPO median proxies using NASSCOM-derived industry headcount weights.

Industry weights

IT services headcount	3,500,000
BPO services headcount	1,600,000
Total IT+BPO workforce	5,100,000

NASSCOM Strategic Review 2025 reports 5.80M total IT-BPM workforce. The 5.10M denominator excludes ~700K workers in ER&D services, software products, hardware, and pure-play GCC engineering — segments where we have no salary data in the sample. IT/BPM headcount split estimated from NASSCOM and IBEF citations.

Active shares (headcount-weighted)

IT share	3,500,000 ÷ 5,100,000 = 68.6%
BPO share	1,600,000 ÷ 5,100,000 = 31.4%
IT contribution	68.6% × ₹21.17 = ₹14.53
BPO contribution	31.4% × ₹3.34 = ₹1.05

BLENDED IT-BPO WORKER INCOME PROXY · FY26 IT × FY25 BPO

₹15.57 *lakh/yr*

≈ $16,594 USD

```
Blend = (68.6% × ₹21.17) + (31.4% × ₹3.34)
= ₹14.53 + ₹1.05
= ₹15.57 lakh/yr
```

vs FY25

+7.5%

Note. *NASSCOM does not publish a clean IT-vs-BPM headcount split for FY26 — the 3.5M / 1.6M defaults are estimated from various NASSCOM and IBEF citations. Sensitivity: if the true ratio is closer to 3.8M / 1.4M (more IT-heavy), the blend rises by ~₹0.4 L; if 3.2M / 1.9M (more BPO-heavy), the blend falls by ~₹0.4 L. The output is most usefully read as a range of ~₹15–16 L, not a precise point estimate.*

Built FY26 (year ended 31 Mar 2026) for IT; BPO inputs from FY25 annual reports. IT filings: TCS Q4 FY26 Fact Sheet (9 Apr 2026), Infosys SEC 6-K (23 Apr 2026), Wipro Q4 FY26 press release (16 Apr 2026). Industry totals: RBI software services survey FY25; NASSCOM Strategic Review 2025.

These 6 million workers occupy a critical junction across several economic systems, and contraction in any one channel can have disproportionate macro effects. IT-BPO workers are a small share of India's labor force, but they sit at an unusually important junction: they earn high formal-sector wages, pay income tax, support upper-urban discretionary consumption, and generate foreign exchange through services exports. If AI reduces the labor required per services contract, the first-order effect is lower hiring, wage compression, and loss of IT-BPO worker income. The second-order effect is weaker tax receipts, weaker urban demand, and eventually pressure on India's current account if export revenue also compresses.

- Tax base: India has 75.46 million income tax filers but only 28.16 million paid positive tax in AY 2023-24. The IT-BPO occupational cluster of six million workers earning Rs. 14.7 lakh (about $16,000) average is overwhelmingly inside the positive-tax-paying 28 million people
- Current Accounts: IT-BPO generates the bulk of the $204.7B in software services exports, which is India's single largest source of net foreign exchange. On the other hand, India had a goods deficit of ~$280B in FY25. If IT-BPO export revenue growth slows from current trajectory, the rupee will depreciate, which feeds into oil and other import costs disproportionately borne by the entire population.
- Urban Consumption: Blume Ventures' frequently cited Indus Valley Reports frame India as three economies: India1 (~140M people, 10% of population, $15K average income, drives two-thirds of all discretionary spending), India2 (~300M people, drives the remaining third), and India3 (~1B people, largely outside the consumer economy). At Rs. 14.7 lakh (about $16,000) average compensation, IT-BPO workers sit squarely inside India1.

Tax Base

For the analysis here, I will look at data from 2024 since that is the latest available year with all the metrics I'll need.

India's individual income tax base is extremely concentrated. In 2024, India had 580 million employed workers and only 83 million people filed income tax returns. Of those, only 28 million actually paid income tax.

A STUDY IN STAGNATION

Filing grew. *Paying* didn't.

Total ITRs filed in India climbed from 65 million in FY 2019–20 to 84 million in FY 2024–25 — a 30% jump. Yet the number actually paying any income tax has barely moved. It sits at 28 million today, well below the 36 million who paid in FY20. Almost all the celebrated "expansion of the tax base" has come from filers who owe nothing.

FIG. 01 — TOTAL FILERS VS. PAID TAX, FY20–FY25

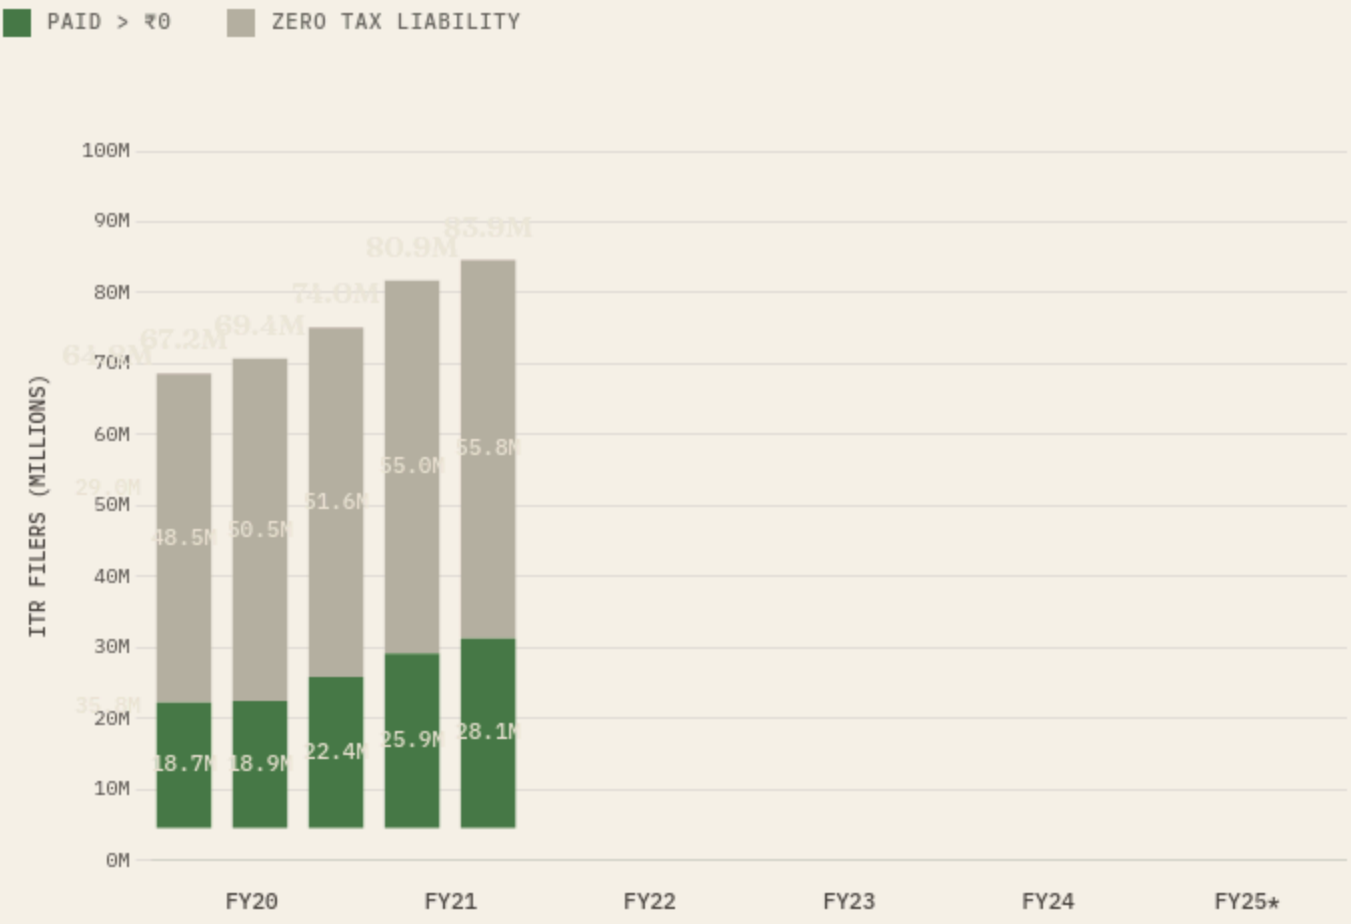

India's $300B AI Shock

NOTE

* FY25 figures are partial — through 31 December 2024 only. Numbers above bars = total ITR filers. Numbers inside segments = filer counts in millions.

Source for FY20–FY23: Lok Sabha reply by MoS Finance, July 2023. Source for FY24–FY25: Lok Sabha reply by MoS Pankaj Chaudhary, 11 February 2025.

The IT-BPO workforce of 5.8 million (NASSCOM Strategic Review 2026) is overrepresented inside this 28 million population. My bottom-up estimates put IT-BPO at 17-20% of all salary TDS. For non-Indian readers, TDS is India's version of deducting income tax from a worker's paycheck.

This means IT-BPO, with less than 1% of India's total employed workforce, contributes roughly 20% of India's salary-withheld income tax!

How a sector with *1% of workers* contributes *19%* of India's salary-withheld income tax

A bottom-up reconstruction of the IT-BPO sector's share of India's salary-withheld income tax (TDS under §192). The calculation has two parts: the **numerator** — total tax paid by IT-BPO workers, built up from 5.8 million workers across seven cohorts and anchored to audited compensation data from TCS, Infosys, and Wipro — and the **denominator** — India's total salary TDS, which the CBDT does not publish directly and which is therefore triangulated from three independent methods. All figures are FY 2024-25 (assessment year 2025-26).

STEP 01 · THE NUMERATOR

Build IT-BPO's tax contribution *cohort by cohort*

The 5.8 million IT-BPO workers (NASSCOM Strategic Review 2025) are split into seven cohorts based on firm tier and seniority. This step builds the **numerator of the**

headline share — the total salary tax paid by these 5.8 million workers. Average compensation per cohort is anchored by audited employee-benefit-expense data from the top three IT services firms — combined ₹2.85 lakh crore on ₹5.07 lakh crore revenue (FY25 Ind-AS consolidated), giving a 56.2% EBE-to-revenue ratio. After stripping employer PF, gratuity provisions, and ESOP non-vest charges, cash compensation works out to ~80% of EBE. Effective tax rates are computed under the FY25 new-regime slabs with §87A rebate at ₹7L taxable income, blended with old-regime usage at senior tiers.

Click any row to see the methodology behind that cohort's compensation assumption.

COHORT	WORKERS	AVG COMPENSATION	WAGE BILL	TAX RATE	TAX (₹ CR)
Top-3 IT · seniorTCS / Infosys / Wipro · 30% of top-3	0.35M	₹35L	₹1.225 L cr	22%	26,950
Anchored by combined per-employee EBE of ₹24.47L across TCS+Infy+Wipro (FY25 Ind-AS consolidated). Senior cohort scaled at 1.4× the per-employee EBE × 80% cash conversion = ~₹35L average CTC. Effective rate of 22% blends new regime (~21.5% at ₹35L taxable) with old regime mix at senior tiers.					
Top-3 IT · mid~50% of top-3 headcount	0.58M	₹18L	₹1.044 L cr	12%	12,528
Mid-level top-3 staff, derived as ~75% of per-employee EBE × 80% cash conversion. ₹18L blended CTC; new regime tax = ~10.8% + cess; old regime usage adds 1.2 percentage points to effective rate.					
Top-3 IT · junior0–3 yr employees, ~20% of top-3	0.23M	₹8L	₹0.184 L cr		

₹0.184 L cr
3%
552
TCS fresher CTC band ₹3.5–7L scaled to ₹8L blended including 2-3 yr lateral mix. Most fall above ₹7L §87A threshold but only marginally — 4.4% nominal rate, ~3% blended after some old-regime offset.
Other tier-1 ITHCL, LTIM, TechM, Cognizant, others
1.00M
₹17L
₹1.700 L cr
11%
18,700
Estimated from peer-firm annual reports. Slightly below top-3 average compensation but similar effective rate. Note: HCL Tech alone employs ~225K.
Mid-tier services + GCC1,750+ Global Capability Centers
1.50M
₹14L
₹2.100 L cr
9%
18,900
Largest cohort — combines mid-cap Indian services firms with the rapidly growing GCC segment. NASSCOM-Zinnov reports 1,750+ GCCs employing ~1.9M (projected to 2.5M by FY27). Lower seniority skew relative to tier-1 brings effective rate to ~9%.
Small servicesSMB IT, engineering services, dev shops
0.84M
₹8L
₹0.672 L cr
3%
2,016
Long tail of smaller IT and engineering services firms. Compensation centered around ₹8L; most workers near §87A threshold, contributing modestly to total tax. No clean published source — estimated as residual after top-3 + tier-1 + mid-tier accounted for.
BPOVoice / non-voice / customer ops
1.30M
₹4.5L
₹0.585 L

₹0.585 L cr				
0.5%				
292.5				
Companies Act §197 disclosures (Firstsource ₹3.19L, eClerx ₹3.79L, Alldigi ₹2.54L medians) suggest mean BPO compensation ~₹4.5L — fully below the §87A rebate threshold under new regime. Only senior managers contribute; sector pays ~₹293 cr in salary tax against a ₹0.59 lakh crore wage bill.				
SUBTOTAL · SALARY TAX	**5.80M**	**— ₹7.51 L cr**	**—**	**79,939**
+ ESOP / RSU perquisite (§17(2)(vi)) at top tiers	—	— ₹0.40 L cr	~30%	12,000
IT-BPO SALARY TAX · FY 2024–25	5.80M	— ~₹7.91 L cr	—	**91,939**

Where the asymmetry lives

Top bar: share of IT-BPO workforce. Bottom bar: share of IT-BPO salary tax. The senior tiers — 6% of headcount — generate 31% of the sector's tax. The BPO tier — 22% of headcount — generates 0.3%.

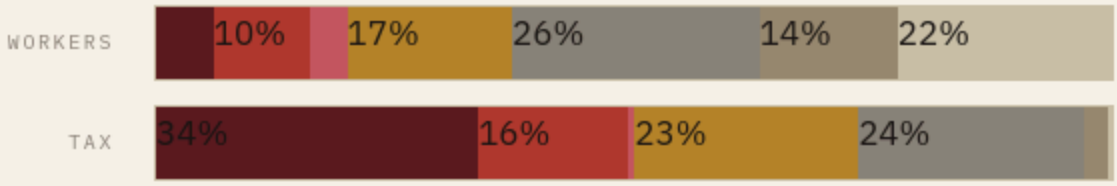

Top-3 IT · senior

Top-3 IT · mid

Top-3 IT · junior

Other tier-1 IT

Mid-tier services + GCC

Small services

BPO

Triangulate India's *total salary TDS*, which CBDT does not publish

This step builds the **denominator of the headline share — India's total salary tax pool, against which IT-BPO's ₹91,939 cr will be measured**. The CBDT publishes total TDS (₹9.48 lakh crore in FY25, per the Time Series Data) but does not disclose what share is collected under §192 (salary). Three independent methods are used to estimate it. Methods 1 and 3 share the §192-share assumption (the weakest input); Method 2 builds up from the salaried-payer count. The three converge in a tight band around **₹4.6–5.0 lakh crore**.

METHOD 01

Top-down PIT decomposition

Gross PIT × TDS share × §192 share

Gross PIT FY25	₹12.90 L cr
× TDS share (~69%)	₹8.90 L cr
× §192 share (~52%)	₹4.63 L cr

METHOD 1 ESTIMATE **₹4.63 L cr**

METHOD 02

Bottom-up from payer count

Salaried positive payers × Avg taxable income × Eff. rate

Positive payers (AY25, proj.)	30.0 M
× Salaried share (~75%)	22.5 M
× Avg taxable (₹16 L)	₹36.0 L cr
× Eff. rate (~14%)	₹5.04 L cr

METHOD 2 ESTIMATE **₹5.04 L cr**

METHOD 03

Salary share of PIT

Salary tax × §192 share salary tax

Gross PIT FY25	₹12.90
× Salary share (~57.5%)	₹7.4 cr
× §192 share (~67.5%)	₹5.0 cr

METHOD 3 ESTIMATE **₹5.01**

TRIANGULATED MIDPOINT

Average of three methods. Range: ₹4.63 L cr (low) to ₹5.04 L cr (high). Sensitivity is ±2 percentage points on the final share.

₹4.89 L cr

India's $300B AI Shock

SOURCE URLS -

When 23,460 TCS roles disappeared in FY26 and 12,000 Oracle India roles are eliminated, these workers likely go into a much lower taxation cohort if they become freelancers under Section 44ADA of the Indian Income Tax regime. A BPO worker would likely fall into informal services given their relatively lower skilled experience, and their income tax contribution goes to 0.

A big contributor to this asymmetry is India's tax exemption for agricultural workers. Roughly 282 million agricultural workers (44.8% of CY 2024 employment per the PLFS Annual Report, MoSPI) are excluded from income tax collections, and often fall below taxable thresholds anyway.

Current Accounts

The IT-BPO sector is disproportionately relevant to India's economic stability because it is one of the country's most important sources of foreign-currency earnings. This is

important because India imports far more physical goods like oil, electronics, machinery and gold than it exports, which means dollars are constantly flowing out of India to pay for the goods it buys from the rest of the world. IT-BPO helps replenish those dollars by selling services to global clients.

The current account is India's foreign-exchange scoreboard. Every day, dollars leave India to pay for oil, electronics, machinery, gold and other imports. Dollars come back through goods exports, software and business-services exports, remittances, tourism and other cross-border income. The current account measures whether those recurring flows are broadly balanced. A small deficit is manageable, but a large or sudden deficit results in a weaker rupee, higher import costs, tighter policy, and government appeals to save fuel or buy less gold.

In FY2024-25, India's merchandise exports were US$441.8 billion, while goods imports were US$729.0 billion, creating a US$287.2 billion goods-trade deficit.

India's services economy moved in the opposite direction: services exports were US$387.5 billion, against services imports of US$198.7 billion, creating a US$188.8 billion services surplus. IT-BPO is central to that surplus. RBI's FY2024-25 software and IT-enabled services export survey estimates India's software services exports at US$204.7 billion**.** That services surplus is one of the main reasons India's overall current-account deficit was only US$23.3 billion, or 0.6% of GDP, in FY2024-25.

India external sector — FY 2024-25 with constituent breakdowns

Full-year FY 2024-25 current account decomposition, plus composition of goods and services trade.

Chart 1 — FY 2024-25 full-year current account

Goods deficit far exceeds the services surplus, but remittances close most of the remaining gap. Net CAD: −$23.3B (0.6% of GDP).

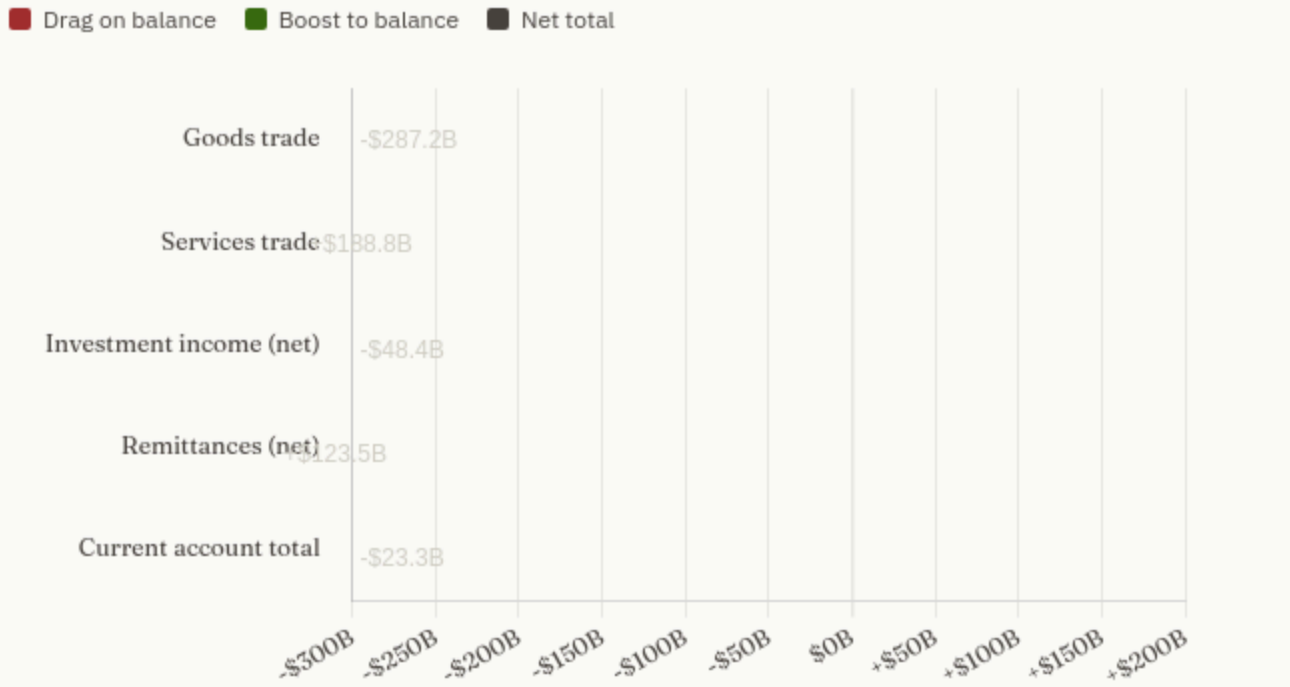

India's $300B AI Shock

SOURCE URLS -

Source: RBI press release on Q4 and full-year FY 2024-25 BoP, 27 June 2025 (Table 1).

Chart 2 — Goods trade composition, FY 2024-25

Imports total $720B, exports total $437B. Each bar segmented by category.

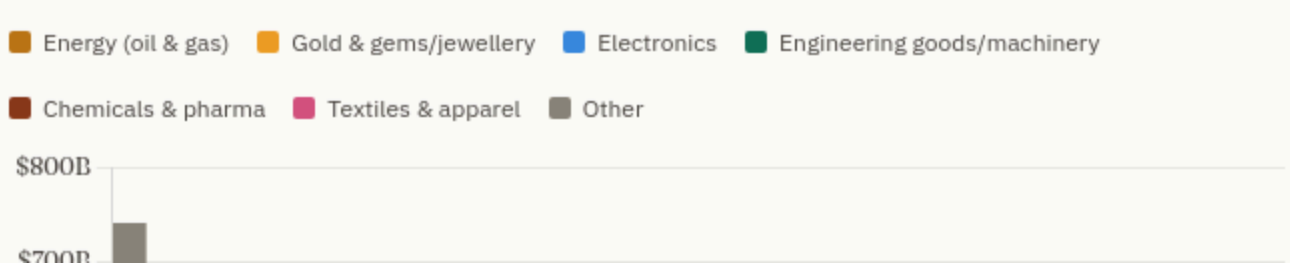

SOURCE URLS

Sources: PIB Press Release ID 2122016 for FY 2024-25 trade summary and major export categories; RBI BoP Q4 2024-25 release for POL imports of $185.8B. Engineering, electronics, pharma, petroleum-products, and gems & jewellery export figures are PIB-confirmed; other import sub-categories beyond crude oil and gems & jewellery are estimates from sector reporting in the order-of-magnitude range, used here for illustrative composition only.

Chart 3 — Software services exports breakdown, FY 2024-25 (official RBI survey)

Computer Services + ITES = $204.7B of India's $387.5B total services exports (~53%). The remaining ~47% is Other Business Services (R&D, professional/management consulting, including the GCC slice not captured as software/ITES) plus travel, transport, financial, insurance, and communication services.

SOURCE URLS

Sources: RBI press release "Survey on Computer Software and Information Technology Enabled Services Exports: 2024-25", 4 Nov 2025, Tables 2 and 3 (Computer Services $138.1B = IT Services $131.3B + Software Product Development $6.8B; ITES $66.6B = BPO $55.8B + Engineering Services $10.8B). Total services exports $387.5B from RBI Q4 FY 2024-25 BoP release, Table 1. "Other services" ($182.8B) is the residual: total services minus the RBI software survey number; it includes Other Business Services (R&D, professional consulting, including GCC work classified outside software/ITES), travel, transportation, financial services, insurance, communication, and government services.

How GCC work is classified

Global Capability Centres are a corporate structure, not a service category. Their work is classified by activity and split across multiple lines:

- GCCs doing software development → **Computer Services** (in the $138.1B above)
- GCCs doing BPO / finance & accounting / customer ops → **ITES BPO** (in the $55.8B BPO line)
- GCCs doing embedded systems, product design → **ITES Engineering Services** (in the $10.8B line)
- GCCs doing pure R&D, management consulting, professional services → **Other Business Services** in BoP (sits in the $182.8B "Other services" residual, separate from the RBI software survey)

A significant share of GCC work therefore *is* inside the $204.7B software-services figure, not separate from it. The slice that sits outside is mostly the higher-end R&D and consulting work classified as Other Business Services in BoP.

April 2026 showed how quickly the current account deficit can impact day-to-day life for Indian citizens. Amid the Middle East energy shock, India's monthly merchandise trade deficit widened to $28.38 billion, up from $20.6 billion in March. Imports rose to $71.94 billion, with oil imports alone rising to $18.62 billion from $12.18 billion in March, while gold imports stood at $5.63 billion. The rupee fell below 96 per dollar after the data. India's external account is structurally exposed to oil and gold, and software/services exports are one of the few large recurring dollar inflows that offset that exposure.

In response, the Government had to raise retail fuel prices for the first time in four years, tighten gold imports to reduce foreign-exchange outflows, and PM Modi urged citizens to cut fuel use and reduce gold purchases. AP separately reported a ₹3/litre fuel-price increase, gold and silver import duties raised to 15%, and a 90-day fuel-saving campaign

with work-from-home days for some government employees in Delhi. If the services-export cushion from IT-BPO weakens, India has less room to absorb oil, gold, electronics, and geopolitical shocks without currency pressure.

IT-BPO is the single biggest source of foreign money the country earns. Without it, India would be buying roughly twice as much from the world as it sells back. A gap that big is only possible as long as dollar reserves are available and overseas capital is willing to give India loans and investment. Such dependence on overseas investment isn't sustainable, as is evident from the outflow of foreign investment capital from India in the last year.

We've already seen an 83% decline in Medical Transcription BPO revenue. What happens if there's a 50% cut to IT/BPM exports? This means a first-round loss of about US$102.4 billion. Hold everything else constant and India's services surplus falls from US$188.8 billion to about US$86.5 billion. The current account deficit widens from US$23.3 billion to roughly US$125.6 billion, which is about 3.2% of GDP.

create a scenario matrix for 20% and 50% declines to IT services revenue and impact on GDP. Ask chat/claude to figure out the right year to use

India IT/BPM export decline — scenario matrix

SOURCE URLS -

Baseline: FY2024-25. Source: RBI Survey on Computer Software and ITES Exports, 2024-25

METRIC (US$, FY2024-25)	BASELINE	20% DECLINE	50% DECLINE
Software services export revenue	$204.7B	**$163.8B**	**$102.4B**
Revenue loss	—	**–$40.9B**	**–$102.4B**
Net services receipts	$188.8B	**$147.9B**	**$86.5B**
Current account balance	–$23.3B	**–$64.2B**	**–$125.7B**
CAD as a % of GDP	0.60%	**1.65%**	**3.24%**

India's $300B AI Shock

CUSTOM DECLINE 35%

REVENUE LOSS
–$71.6B

SERVICES RECEIPTS
$117.2B

CURRENT ACCOUNT
–$94.9B

CAD AS A % OF GDP
2.45%

The danger zone starts at ~2.5% CAD/GDP, not 50%. That's where rating agencies flag external-stability concerns and the INR comes under sustained pressure. The slider crosses that around a 37% decline — well before the headline 50% scenario.

These figures are conservative. CAD as a % of GDP scales linearly here only because GDP is held constant in the denominator. In reality, a meaningful export decline shrinks GDP too, making the ratio *worse* than the linear projection.

What happens then? India has seen this before in 1991 and 2013. In FY1990-91, a current account deficit of 3.1% of GDP became unsustainable and led to an external payments crisis that saw India literally flying gold to the Bank of England to collateralize its debts and eventually led to India's liberalization.

Secret sale of gold by RBI again

21,000 kg of metal airlifted to coded destination

By V. Shankar Aiyar

BOMBAY, July 7. – The Reserve Bank of India today exported 21,000 kilograms of "precious metal" in a highly hushed-up operation at the Sahar International Airport in Bombay. This comes in the wake of the controversial sale executed by the Reserve Bank of India in the first week of June when it pledged 20 tonnes of gold for $ 200 million to help the Government tide over the foreign exchange crisis.

Though there is no confirmation that the "precious metal" is gold, it is widely suspected to be so. When contacted for his comment a highly placed RBI official told *Indian Express* that he "was not in a position to either confirm or deny the report". Air India cargo complex officials and the security personnel however admit that the consignment of precious metals is suspected to be gold.

The "operation transfer" was effected in total secrecy with even the Customs officials (who normally provide preventive security arrangements) left totally in the dark. A private charter Boeing aircraft belonging to Heavy Lift Cargo Airlines (HL 858) landed at the Sahar terminal on Saturday and was parked in the cargo bay of Air India. The "precious metal" was offloaded from the RBI vans. Loading operations began on Sunday afternoon under heavy police protection . The flight (HLA 859), with Captain Noel in command, took off for Bahrain at 1915 hrs (IST) en-route to a coded destination (STN OONP) which it was expected to reach on Monday (0500 hrs UTC).

Al cargo officials were informed (first on July 3 and again on July 6) about the operation in a highly coded and confidential message from Heavy Lift Cargo Airlines which asked them to help with the loading arrangements of "21,000 kgs of precious metal". The message informed them further that the consignor was Reserve Bank of India and the consignee, Bank of England, London, UK."

If the RBI has really exported gold (provided they admit it) it would be the government's second serious volte-face in recent weeks. Earlier the government had denied that there was any devaluation of the rupee and went ahead and devalued it for the second time in hours. Finance Minister Manmo Singh has been dessying any m towards yet another gold sale to over the balance of payments cr In fact he has been talking of ge the pledged metal back.

It may be recalled that the ear sale too was effected by the Rese Bank of India with the ministers the Chandra Shekhar governm pleading ignorance about the ent operation. In fact the Shekhar g ernment denied the sale outright the first instance only to backtr later with a clarification that t metal had not been sold but mer pledged.

The interesting part of this tra fer is that this time round t consignee is the Bank of England against the Swiss institution earli

The "operation transfer" of the "precious metal" being conducted under tight security at the Air India cargo bay of the Sahar International Airport on Sunday afternoon. — Express photo by Sudharak Olwe

Stich – new Wimbledon champ

WIMBLEDON (AP) ...

We will be able to steer clear of the problems confronting us:PM

The deficit reached a peak of 4.8% of GDP in FY2012-13 when US Federal Reserve Chair Ben Bernanke told Congress the Fed might start slowing its bond purchases. India had done nothing wrong that quarter. IT services were healthy. Exports were growing. But US interest rates suddenly looked more attractive, foreign investors reassessed the risk of holding Indian assets, and the rupee fell 7.5% against the dollar between May 22 and July 15 2013. The RBI burned through reserves, launched emergency dollar deposit schemes targeted at NRIs, and tightened monetary policy.

If AI breaks IT services, India faces two problems at once. The first is lower export earnings which widens the current account deficit. The second is that IT services have been India's growth story for foreign investors. Three decades of India's pitch to global capital has been "India is the world's services back-office and its technology division". As AI improves and breaks the pitch, foreign capital will rotate to countries that have more data center and semiconductor capacity.

Foreign capital is already leaving

India: gross FDI inflow vs. gross FDI outflow (USD billions)

FY25 gross outflow

$80.0B

2.1× FY21 outflow

FY25 outward FDI

$29.2B

Indians investing abroad

Gross FDI inflow Gross FDI outflow (repatriation + outward FDI)

India's $300B AI Shock

SOURCE URLS -

Sources: RBI Balance of Payments Q4 FY25 release (link); Vivekananda Foundation, *Greater Scrutiny of Net FDI* (link); PIB India press release on FY25 gross FDI of $81.04B (link); National Herald on FY25 repatriation and outward FDI (link). FY21–FY22 outflow figures derived from BoP identity (gross inflow − net FDI); FY23–FY25 components are explicitly published.

India: net FDI and net FPI flows, FY23–FY25 (USD billions)

FY 2022–23 FY 2023–24 FY 2024–25

Net FDI (strategic / sticky)

10%+ stakes, factories, M&A

–96%

$28.0B (FY23) → $1.0B (FY25)

Net FPI (portfolio / fast)

listed stocks & bonds

volatile

–$5.2B → $44.1B → $3.6B

SOURCE URLS -

Sources: RBI, *Developments in India's Balance of Payments during Q4 2022–23* (June 27, 2023) — "Net FPI recorded an outflow of US$ 5.2 billion in 2022-23"; net FDI of $28B in 2022–23 also referenced in subsequent RBI communications (RBI release). RBI, *Developments in India's Balance of Payments during Q4 FY25* (June 27, 2025) — "Net inflow under FDI at $1 billion during 2024–25 was lower than $10.2 billion during 2023–24. During 2024–25, FPI recorded a net inflow of $3.6 billion, lower than $44.1 billion a year ago." (summary).

FPI ownership of NSE-listed companies: 11 consecutive quarters of decline

Sep 2025

16.9%

15-year low

Drop since Mar 2023

–2.2pp

19.1% → 16.9%

12M FPI net outflow

$29.6B

to Sep 2025 (NSE)

SOURCE URLS

Source: NSE India, *India Ownership Tracker, Q2 FY26 (Vol. 7, Issue 2)*, November 2025 — Tables 2 and 4 (FPI share of NSE-listed total market cap, quarterly), and page 21 narrative on $29.6 bn 12-month net FPI outflows to Sep 2025 (PDF link).

FY2024-25 wasn't a great year for Foreign Institutional Investor (FII) inflows into India's stock markets even without any IT-BPM declines. Gross FDI inflows were strong at about US$81.0 billion in FY2024-25, but net FDI inflow fell to just US$1.0 billion, down from US$10.2 billion in FY2023-24 and far below the US$43.4 billion India recorded in FY2020-21. RBI's balance-of-payments data show FPI net inflow of only US$3.6 billion in FY2024-25, down from US$44.1 billion a year earlier. That weakening is visible in market ownership as well. By September 2025, foreign portfolio investors owned only 16.9% of NSE-listed companies, which NSE data described as the lowest in more than 15 years.

Why is this happening? An HSBC note described India as a "funding market" for Asia's AI boom, with capital getting out of India and into South Korea and Taiwan. Nearly $28 billion had been withdrawn from Indian equities between September 2024 and November 2025, leaving India the second-largest underweight in global emerging-market portfolios. The underlying point is that foreign capital has already started rewarding markets with obvious semiconductor and AI winners more aggressively than India.

Urban Consumption

It is helpful to think of India not as 1 national consumption market but as 3 different markets. Ostensibly the market seems huge with 1.4 billion people, but India's monetizable market is much smaller.

Blume Ventures, an Indian VC firm, releases fantastic Indus Valley Reports about the Indian economy and startup ecosystem. Most readers working in Indian startups are likely familiar with them. From their March 2025 report:

CONSUMER STACK · DISCRETIONARY SPEND

India isn't *one market*. It's three.

FRAME 01 · THE STACK

Three economies stacked inside *one country*

Each tier behaves differently as a market. India 1 buys; India 2 transacts in small tickets; India 3 mostly consumes free attention. Most of what looks like "India growth" is really India 1 and parts of India 2.

Tier	Class	Scale & Income
India 1	The Consuming Class TIER 01	HOUSEHOLDS ~30M PEOPLE ~140M PER CAPITA ~$15K
		HOUSEHOLDS ~70M

India 2	The Aspirant Class TIER 02	PEOPLE ~300M PER CAPITA ~$3K
India 3	Unmonetisable Users *(& Non-Users)* TIER 03	HOUSEHOLDS ~205M PEOPLE ~1B PER CAPITA ~$1K

SOURCE Blume Ventures · Indus Valley Report 2025

FRAME 02 · THE ENGINE

India 1 is the *engine* of the Indian consumer economy

Two-thirds of household discretionary spend comes from a tenth of the population. India 2 contributes the other third. India 3 — two-thirds of the country — runs a slightly negative share, dipping into savings to make ends meet.

Indian Consumer Stack · Share of Household Discretionary Spend

India 1

~10% of population · *2/3 share* of discretionary spend

The tenth that drives the market.

~23% of population · *1/3 share* of discretionary spend

Heavy users, light spenders.

~2/3 of population · *slightly negative* share of discretionary spend

Dipping into savings to make ends meet.

India's $300B AI Shock

SOURCE Blume Ventures · Indus Valley Report 2025

FRAME 03 · THE CONCENTRATION

Inside urban India, the top decile out-spends the average by *3× to 13×*

The over-indexing isn't uniform. On staples, the top 10% spends roughly 3–5× the average. The gap widens as you climb the discretionary ladder — peaking at 11× for out-of-home food, 12× for institutional medical, and 13× for durable goods.

Indian Urban Top 10% versus Average Per Capita Expenditure by Category

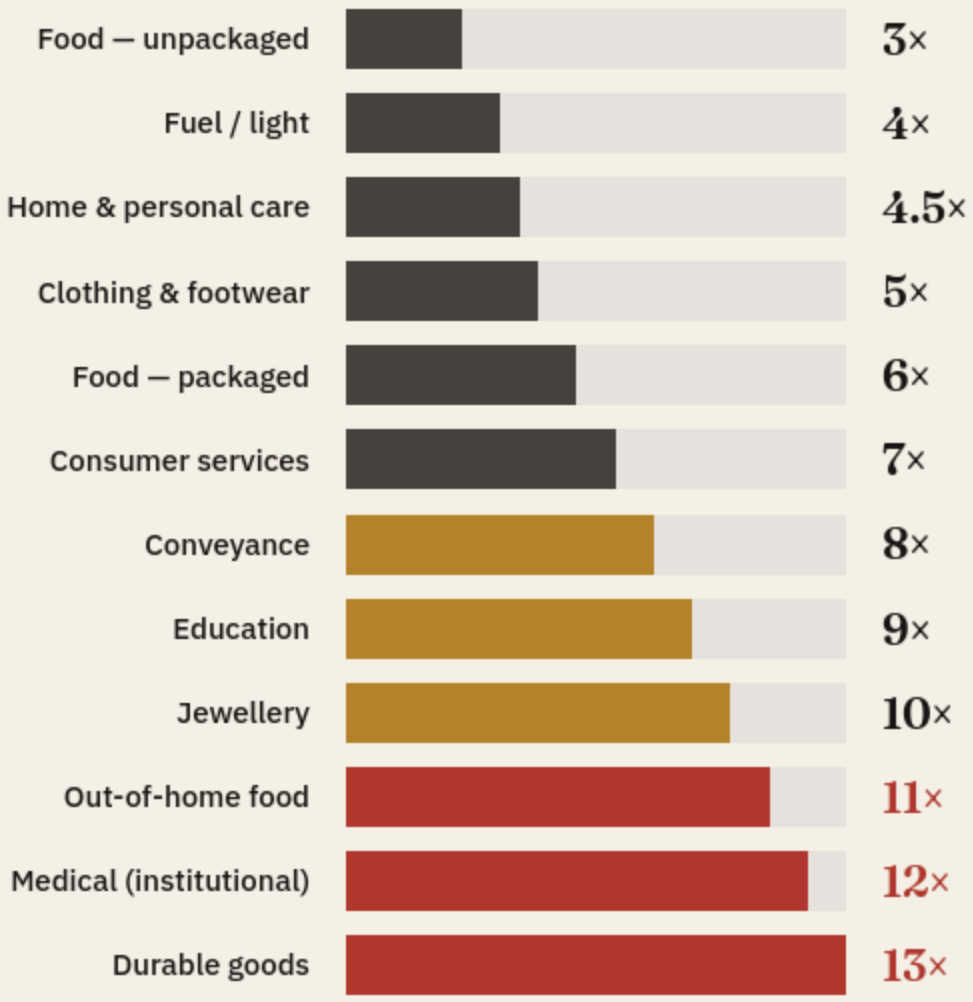

SOURCE Blume Ventures · Indus Valley Report 2025

SOURCE URLS

-

↗ Framework and figures from Blume Ventures' **Indus Valley** annual report. blume.vc

FROM *India's $300B AI Shock*

The top ~10% of India, named “India1”, has roughly 140 million people with per-capita income of $15K and drives 2/3rd of all discretionary spending.

The next 21%, the rising “India2”, drive the remaining 1/3rd of discretionary spending with per-capita income of $3K.

This reflects in per-capita over-indexing. The consuming “India1” class spends: 3-5x the national per-capita rate on basic categories (food, fuel, clothing), 7-9x on consumer services, conveyance, education, 10-13x on jewellery, out-of-home food, institutional medical, durable goods.

Because the top ~10% of India drives roughly 2/3 of discretionary spending, and this cohort spends 7-9x the national per-capita rate on consumer services, conveyance, and education, and 10-13x on jewellery, out-of-home food, institutional medical care, and durables, a shock to incomes inside that group has a disproportionate impact on GDP.

We know from the beginning of this section that India’s IT-BPO workforce is roughly 5.8 million, with an average take-home salary of Rs. 14 lakhs (about $16,000). Taking into account double-income households and supporting immediate family like children and parents, let’s assume these 5.8 million workers are supporting 15 to 20 million people. Taking the lower bound of 15 million, that is 10% of “India1”, the group driving 2/3rd of Indian discretionary consumption, directly dependent on IT-BPO salaries.

The first-order impact on IT-BPO workforce displacement will likely show up in IT cities like Bangalore, Hyderabad, Gurgaon, Pune and Chennai. Bangalore and Hyderabad alone

have anywhere from 2 million to 2.5 million of India's 5.8 million IT-BPO workers. These cities also sit at the center of India's consumer-internet monetisation map, whether the product is food delivery, fintech, travel, quick commerce, or premium retail.

To be clear, these cities also have several other high-value industries like pharma, manufacturing and financial services (which may have its own reckoning now that models are good at Excel), but IT-BPO is the most likely *marginal* employer of the cohort with the highest urban discretionary propensity - fresh graduates, dual-income couples in their late 20s and mid-career professionals in the Rs. 15-80L (about $16,000 to $90,000) bracket.

What happens when the marginal salaried spender in India's five most important tech cities sees slower income growth? When that pipeline slows, this cohort's incremental spending, which has been funding the growth in consumer company P&Ls (both startups and old-school enterprises), will also decline.

Many Indian consumer internet business models also depend on "India1" contribution margins to finance "India2" expansion. That has been the operating logic of much of India's consumer startup ecosystem for the last decade: acquire and monetise dense upper-urban users first, then use that cash flow, network density, and logistics base to push outward.

A slowdown in upper-urban salaried demand could impact businesses whose long-term growth story is to expand beyond India's big cities. If the high-frequency, high-margin urban user base weakens, the company has less room to spend on expansion, subsidies, assortment, merchant acquisition, dark-store density, or customer support in lower-yield markets. Thus, "India2" also becomes harder and slower to monetise.

The Future Tax Base

A second-order effect that operates with a 5-7 year lag is that the fresher hiring pipeline that turns engineering graduates into future high earners is contracting sharply. TCS, which is the largest private-sector employer of engineering graduates in India, cut its

FY27 new college graduate (“Fresher”) target by 43% from to 25,000 from 44,000 in FY26. The fresher cohort itself contributes very little to *current* income tax (entry-level total compensation of Rs. 3.5-6L (about $3,000 - $7,000) are taxed at low rates) so most pay zero income tax in their first three years (their tax contribution is mostly through consumption-based taxes). The damage is to the *future* tax base. A fresher hired into TCS in 2018 typically reached Rs. 15-25L total compensation by 2025, entering the top-decile tax-paying tier. If fresher hiring reduces today, that impacts India’s future high tax-paying base.

GCCs cannot make up the gap, because GCC fresher hiring draws from a much narrower top-tier-college funnel and absorbs perhaps 20-30K freshers a year industry-wide versus the ~150K tier-2/tier-3 absorption that Indian IT services historically provided. India produces roughly 1.5 million engineering graduates per year, and if formal-sector employment narrows by half, the cohort feeding the Rs. 20L+ (about $22,000+) tax tier in 2030-2032 shrinks correspondingly. Besides, Indian GCCs are going through workforce displacement themselves, as seen with Oracle.

CHART 05 · COHORT PROJECTION · BEAR CASE · FY10–FY33

● ANCHORED TO CEO COMMENTARY, Q4 FY26

The Structural Reset in Fresher Hiring

Top-5 IT fresher hiring peaked at 180K in FY22, troughed at 60K in FY25, and is now — by the Q4 FY26 commentary of three CEOs — being reset, not recovered. TCS at 25K (lowest since FY20), Infosys at 20K, Wipro pausing entirely. The black line below reflects this bearish view: hiring stays at ~50% of FY22 peak through FY33, not the historical 120–140K baseline. The high-earner wave is hit twice — once by lagged fresher cohorts, again by AI-led mid-career layoffs.

CEO COMMENTARY ANCHORING THE BEAR CASE · APRIL–MAY 2026

TCS

25,000 FY27 freshe

↓ 43% vs FY26 · lowest since FY20

> *"Further campus recruitment will be directly linked to demand visibility and project pipelines."*

Infosys

20,000 FY27 freshe

↓ ~60% vs FY18–FY22 peak years

> *"AI is now part of every client conversation."*
>
> — even the most bullish of the three frames AI as structural. AI is no longer a separate line; it's embedded across all services.

Wipro

no target set for FY27

FY26: 7,500 actual (vs 12K original target)

> *"We don't have any target for fresher hiring for the next fiscal. It's completely on demand, very volatile environment."*

pipelines.

– K. Krithivasan, CEO & MD · Q4 FY26 commentary, April 2026

– Salil Parekh, CEO & MD · Q4 FY26 earnings call, April 23, 2026

– Saurabh Govil, CHRO · Q4 FY26 earnings call, April 17, 2026

Plus the structural signal across all three: the "pyramid" staffing model (many freshers supporting fewer seniors) is being replaced by a "diamond/cylinder" (mid-senior heavy). Wipro committed $1B+ to AI, calls its new approach "services-as-software."

Fresher Hiring — annual freshers, top-5 IT services

High Earners — same cohort, projected ₹15L+ earners 7 yrs later

Solid = actual · Dashed = bear-case projection

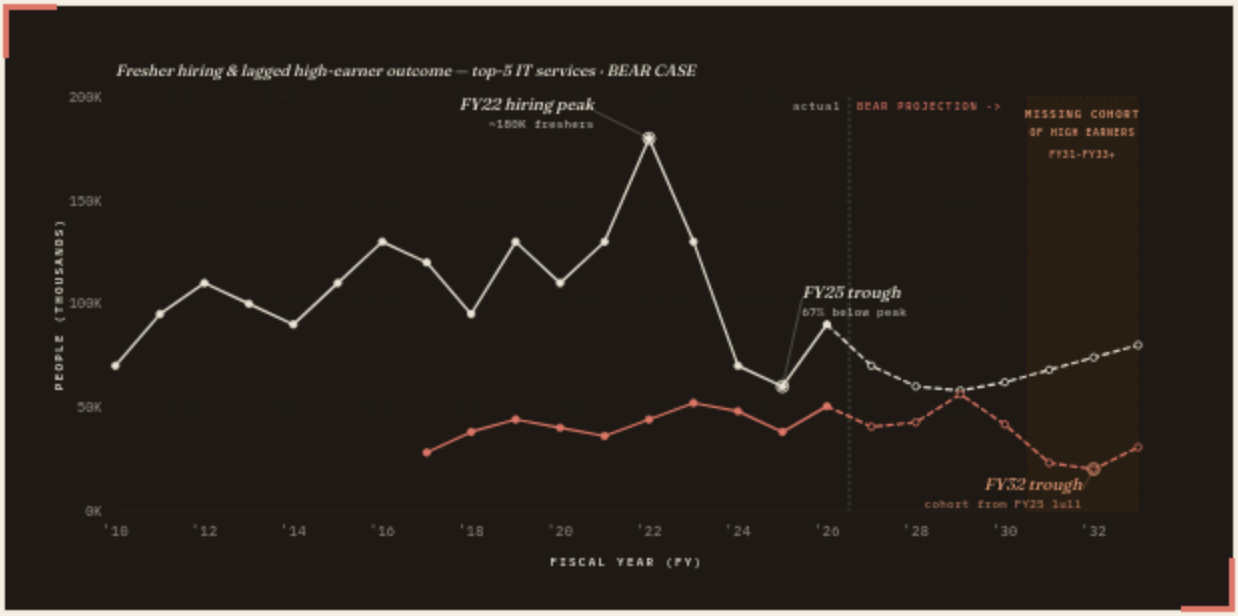

India's $300B AI Shock

How Projections are Calculated

HIGH EARNERS · FORMULA UPDATED FOR AI CHURN

FY26 actual = ~90K: TCS 44K + Infy 20K + Wipro 7.5K + HCL ~12K + TechM

~6K. FY27 = ~70K takes the announced TCS 25K, Infosys 20K, Wipro 5K (midpoint of pause-but-some-onboarding), HCL ~12K, TechM ~8K.

> **How these projections might be wrong:** *a rapid client-spending recovery in FY28+, GCCs aggressively absorbing displaced talent, or AI failing to compress entry-level work as fast as promised. None of these are zero-probability — but all run counter to current CEO guidance.*

$$HighEarners[FY_n] = FresherHires[FY_{n-7}] \times 0.40 \times s[FY_n]$$

$s[FY_n]$ · AI-survival multiplier ·

≤FY25	FY26	FY27	FY28
1.00	0.97	0.92	0.82

FY29	FY30	FY31	FY33+
0.78	0.80	0.82	0.85

The s factor accounts for direct AI-led mid-career layoffs. TCS alone cut 12,000 mid/senior employees in FY26 — that is the ₹15L+ tier being hit directly. The industry has structurally shrunk; fewer total ₹15L+ roles exist post-restructure, so the conversion rate stays permanently below the pre-AI baseline.

WORKED EXAMPLES · BEAR CASE + AI CHURN

FY29 peak → 180K × 0.40 × 0.78 = **56K** *(was 72K pre-AI)*

FY28 → 130K × 0.40 × 0.82 = **43K** *(was 52K)*

FY32 trough → 60K × 0.83 = **20K**

FY33 → 90K × 0.85 = **31K**

PRIMARY SOURCES · CEO COMMENTARY

1. Business Standard — TCS FY27 hiring at 25,000, lowest since FY20. business-standard.com
2. Business Today — TCS Krithivasan demand-visibility commentary, Q4 FY26. businesstoday.in
3. Business Today — Infosys Salil Parekh, 20K campus hires FY27, AI-as-growth framing. businesstoday.in
4. NiftyTrader / WEF Davos 2026 — Parekh's hiring plan and AI commentary. niftytrader.in
5. Business Today — Wipro CHRO Govil pauses FY27 fresher targets entirely. businesstoday.in
6. People Matters — Wipro 7,500 FY26 actual, no FY27 guidance. peoplematters.in

SUPPORTING SOURCES

1. TCS Investor Relations — annual reports & quarterly factsheets. tcs.com/investor-relations
2. Infosys Investor Relations — annual reports & fact sheets. infosys.com/investors.html
3. NASSCOM — Strategic Review archives. nasscom.in/knowledge-center/publications
4. AICTE — engineering enrolment dashboard (supply side). facilities.aicte-india.org/dashboard
5. StartupTalky — TCS 12,000 senior layoffs FY26; pyramid restructuring. startuptalky.com
6. Lapaas Voice — Wipro pyramid → diamond/cylinder model framing. voice.lapaas.com

India's income tax collection was constructed on an assumption of continuous formalization that would pull more workers from informal agriculture into formal services every year. IT-BPM was the engine of that pull for two decades, growing from ~1.3M workers in 2005 to 5.80M in 2024 while individual tax filings rose from ~25M to ~82M roughly in tandem.

If workforce displacement happens due to AI, India (and to be fair, most other countries) will look increasingly like a jobless growth economy where headline GDP may look fine but the consuming class shrinks.

Putting Everything Together

INTERACTIVE MODEL

India AI Displacement — Annual GDP Drag

Annual recurring GDP drag once AI disruption reaches FY28 run-rate across four transmission channels. Adjust scenario assumptions below.

SCENARIO | Very Conservative | **Conservative** | Base | Aggressive

TIER 1 WORKFORCE DISPLACEMENT **50%**

DEFINITION NOTE -

Share of Tier 1 IT/BPO workers (entry/mid-level coding, testing, L1 support, data entry) whose roles are eliminated by AI over a 3-year horizon. Tier 2 and Tier 3 displacement scale proportionally at 0.5× and 0.125×.

FY26 actual: TCS net −3.8% headcount; Infosys +4.1%; sector +135K net hires (NASSCOM).

CASCADE MULTIPLIER (CONSUMPTION) **1.80×**

DEFINITION NOTE -

Total demand destroyed per ₹1 of lost wages as displaced workers cut spending and that pullback ripples through housing, services, food, transport, and the informal economy.

Standard fiscal multipliers in RBI/IMF literature: 1.2–1.8×.

FOREX CHANNEL

DEFINITION NOTE -

GDP drag transmitted through the external sector: slower services exports widen the current account deficit, depreciate the rupee, raise import inflation, and force tighter monetary policy.

FY26 9M: services surplus +15% YoY; CAD narrowed to 1.0% of GDP. Channel not yet materializing.

SECOND-ORDER CASCADES (NON-CONSUMPTION)

DEFINITION NOTE -

Indirect channels outside the primary wage cascade: internal remittances to rural families, credit-system stress from

worker defaults, engineering college closures, and construction worker income loss in IT-hub cities.

Adds 60% of total to avoid double-counting with the primary cascade.

1.71%

OF FY28 GDP

SOURCE URLS

₹427.6L cr ($5.06T) projected by IMF WEO April 2026

≈ 20% *of one year's nominal GDP growth absorbed (FY26 nominal growth: 8.6%)*

₹7.32 lakh crore · *~$87B every year, recurring*

HISTORICAL REFERENCE

2001 dot-com bust: −0.3%

over 8 months, recovered in 5 quarters

No recovery mechanism — this gap repeats every year. *Annual drag of* **−1.71%**, *beginning FY28, indefinitely.* **For reference: the 2001 US dot-com bust caused a −0.3% peak-to-trough contraction over 8 months, recovering in 5 quarters. This scenario, by construction, does not** *— until reversed by workforce redeployment, productivity offsets, or sectoral migration up the value chain.*

CHANNEL BREAKDOWN | GDP TRAJECTORY

CHANNEL CONTRIBUTION (NET OF 15% DOUBLE-COUNT ADJUSTMENT)

Wage / Consumption Cascade ₹4.84L cr · 1.13%

Fiscal Revenue Loss ₹79.4K cr · 0.19%

Domestic White-Collar ₹1.69L cr · 0.39%

Current settings: T1 workforce displacement 50%·Cascade 1.80×·Forex off·Second-order off

NOTE

FY28 RUN-RATE METHODOLOGY: *source-year FY25/FY26 numerator inputs (workforce, compensation, tax base, exports) are projected forward to FY28 before being divided by FY28 GDP, so numerator and denominator are pinned to the same period.*

Workforce uplift +6.85% vs FY25 (NASSCOM CAGR); compensation uplift +11.7% (revenue-per-employee CAGR). FY28 nominal GDP denominator: ₹427.6 lakh crore (IMF WEO April 2026 projects $5.06T; converted at ₹84.5/$, the model's consistent FX assumption). Baseline trajectory uses IMF projections FY26–FY28 and extends at 10% nominal CAGR thereafter. Cumulative gap applies ramp-up of 25% / 60% / 100% / 100% / 100% across FY26–FY30. Historical comparison sources: Fed Reserve History (2008 GFC); St. Louis Fed (dot-com 2001); MoSPI revised FY21 (India COVID).

A 1 to 5% GDP impact may ostensibly look small, but if it recurs annually until the Indian government takes measures to fix things, it will become a growth problem for India's GDP by hitting sectors of the economy that produce formal wages, services exports, tax revenue and urban consumption.

The US dot-com recession is remembered as a major financial-market event, yet the GDP impact was shallow: the recession lasted from March to November 2001, and real GDP barely declined. The 2008 Global Financial Crisis was much more severe: US real GDP fell roughly 4% from peak to trough, and the recession lasted 18 months. India's Covid shock was larger still: real GDP contracted 5.8% in FY2020-21 before rebounding the next year. Even a 1 to 2% annual drag on GDP is much larger than the GDP imprint of many famous economic downturns. A 4 to 5% drag belongs in the same numerical range as crises that governments and markets treat as historic events.

However, the historical comparisons actually understate the AI risk in an important way. Most crises have a beginning, a trough and a recovery phase. Covid had reopening. The Global Financial Crisis had bank rescues, monetary easing and fiscal stimulus. Even when the recovery was painful, there was a policy and business-cycle mechanism pushing the economy back toward normal.

AI displacement is different because of recurrence. This is not like a factory shutting down for a quarter and then reopening. If AI permanently automates a support process, a testing team, or a back-office workflow, the wages attached to that work do not automatically return next year. The spending supported by those wages does not automatically return either. India has to create a new source of income to replace the old one. Until that happens, the gap keeps appearing every year.

Five crises, *five recoveries*

Peak-to-trough real GDP declines and the years it took to climb back.

METHODOLOGY & WHAT THIS MEASURES

GDP decline: peak-to-trough real GDP, expressed in the standard form used by each country's statistical agency or the IMF. Where annual and quarterly data tell different stories, both are shown.

Recovery time: measured from the trough back to the prior real-GDP peak (not trend). Recovery to *trend* typically takes 1.5–2× longer; trend deviation is a stricter test but more contested.

Why not "% impact" alone: A 4% peak-to-trough decline that recovers in 8 quarters is structurally different from a 4% decline that never recovers. Showing only depth obscures persistence — which is often the more important variable for policy comparisons.

EVENT	PEAK-TO-TROUGH REAL GDP	TIME TO RECOVER PRIOR PEAK	COMPARISON CONTEXT & SOURCES
US dot-com recession 2001 · NBER: Mar–Nov	−0.3% Cumulative quarterly real GDP decline. Q3 2001 fell 1.6% annualized; Q2 and Q4 grew.	~5 quarters Strict GDP-level recovery was ~1 quarter, but the economy stayed weak through 2002 with double-dip concerns. Payroll employment did	Shows that famous market crashes do not always create large GDP contractions. A sustained 1–5% GDP drag would dwarf the dot-com imprint on output, even though the stock-market memory is enormous.

	NBER weighted employment heavily.	not recover until 2005 (~46 months — the "jobless recovery").	SOURCE URLS - ↗ NBER ↗ CRS
US Global Financial Crisis 2007–2009 · NBER: Dec 2007–Jun 2009	−4.3% Real GDP fell 4.3% from 2007 Q4 peak to 2009 Q2 trough — the largest postwar decline at the time (BEA/Fed).	~2 years 3 months Real GDP returned to its prior peak in 2011 Q3. Employment took until mid-2014. Output trend was permanently lower.	A sustained 1–5% GDP drag from AI is in the zone of the GFC's real-output decline. But the mechanism is different: GFC was a balance-sheet collapse with a clear policy response (QE, stimulus). A structural-displacement shock has no equivalent counter-cyclical lever. SOURCE URLS - ↗ Fed History ↗ NBER
India COVID shock FY2020–21	−7.3% FY2020–21 fiscal-year contraction (MoSPI). Calendar-year 2020 was −5.8% (World Bank). FY2021–22 rebounded ~9.1%.	~1.5 years Real GDP surpassed FY2019–20 peak in FY2021–22. Distribution of recovery was uneven; informal sector lagged.	COVID had reopening built in — once lockdowns ended, activity returned. An AI displacement shock would lack that mechanical recovery. SOURCE URLS - ↗ Econ Observatory ↗ World Bank
Indonesia Asian Financial Crisis 1997–1998	−13.1% Real GDP fell 13.1–13.7% in 1998 (IMF). Worst single-year contraction among AFC countries; precipitated regime change.	~5 years Real GDP regained 1997 peak around 2003–2004. Recovery slower than Thailand or Korea; per-capita catch-up took longer still.	Triggered by capital flight after Thailand's baht devaluation. The rupiah lost over 75% of its value, making dollar-denominated corporate debt unpayable and collapsing the banking system. IMF austerity and the fall of the Suharto regime deepened the contraction. SOURCE URLS - ↗ IMF WP ↗ Brookings
	−26% Real GDP fell ~26% peak to trough (Bank	Not recovered 17+ years on, real	The most useful analog for a structural shock with no built-in recovery. No reopening (unlike COVID), no stimulus runway (unlike GFC)

Greek depression 2007–2016 (peak to trough)	of Greece). Some measures show ~33%. Deeper than the US Great Depression in percentage terms.	GDP per capita remains below 2007 peak. IMF projects pre-crisis level may be regained around 2030.	Slower, more grinding — closer to what a permanent productivity-displacement shock could look like in shape, if not magnitude.

SOURCE URLS -

↗ BFI Chicago ↗ BIS/BoG ↗ LSE

India's $300B AI Shock

NOTES & CAVEATS

The dot-com number is contested. NBER dates the recession Mar–Nov 2001, but real GDP barely turned negative — cumulative quarterly decline was around 0.3%. Some quarters (2001 Q1, Q3) showed declines; Q2 and Q4 grew slightly. The "famous crash, tiny GDP imprint" framing is the right takeaway.

India's COVID number depends on the calendar. Calendar-year 2020: −5.8% (World Bank). Fiscal year 2020–21 (Apr–Mar): −7.3% per India's Ministry of Statistics first revised estimate, later revised toward −6.6%. The 5.8% figure cited in the original brief is the calendar-year number — both are defensible but they aren't the same measurement.

Greek peak-to-trough varies by source. Estimates range from ~25% (Bank of Greece, IMF cumulative through 2013–16) to ~27% (Eurostat through 2016) to ~33% (some quarterly measures through 2014). All point to the same fact: deeper than the US Great Depression and structurally unrecovered.

The Greek case replaces UK WWII for a reason. If the underlying comparison is "what happens when an economy faces a shock with no automatic demand-side recovery mechanism" — which is the relevant question for AI displacement — Greece is a sharper analog than WWII. Greece had no built-in reopening (unlike COVID), no fiscal-stimulus runway (unlike GFC), and no demobilization boom (unlike WWII). It's the closest peacetime case of a slow, structural, partially-unrecovered output collapse.

What this table still does not capture: distributional effects (who bore the loss), labor-force-participation scarring, and informal-sector buffers (especially relevant in India and Indonesia). GDP is a flow measure; the most policy-relevant AI-displacement questions are about composition, not headline output. Use this table to bound the macro magnitude — not to model the mechanism.

The impact on incremental GDP growth is also illustrative. Suppose India's GDP is at $100 today. If nominal GDP is expected to grow 8.6% next year like it did in FY 2025-26, the economy should get to 108.6 next year. A 2% GDP drag means the economy reaches about 106.4 instead. India still grows, but almost 25% of the expected annual increase has disappeared.

Investors underwrite India on the belief that nominal GDP will keep expanding fast enough to support corporate revenue growth, tax growth, credit growth, infrastructure spending, urban consumption and rising market size. If AI displacement removes 25 to 50% of expected annual nominal growth, the effects travel beyond the workers directly displaced. Corporate revenue expectations weaken. Tax buoyancy weakens. Credit quality weakens. Urban consumption slows. The denominator that makes public debt, infrastructure spending and private capex look manageable grows more slowly. India can still grow, but the premium attached to India's growth story becomes harder to justify.

Showing how each parameter impacts the model? Best way to do this? Or should I link to a separate spreadsheet that people can play around with if interested?

AI is Sovereign Infrastructure

VOL. I · ISSUE 01 *An Inquiry into Capital Flows* MAY 2026

STRUCTURAL SHIFT · INDIAN IT SERVICES

The Rewiring of *India's Tech Value Chain*

For two decades, every dollar that flowed from an overseas client into Indian software services moved through a fairly predictable cascade — and most of it landed in Indian pockets, balance sheets, and tax receipts. Generative AI is now bending the path of that dollar. Some of it still arrives. Some of it is being intercepted upstream.

SCENARIO A · STATUS QUO ANTE
Pre-AI Value Chain

SCENARIO B · EMERGING
Post-AI Value Chain

VALUE CAPTURED IN INDIA VALUE CAPTURED OVERSEAS

→ PRE-AI · TRADITIONAL SERVICES ARBITRAGE

01
OVERSEAS

Overseas Client

US/EU enterprise needs an application built, supported, or modernised.

Captures: end value Spend: $1.00

↓

02

INDIA

Indian Services Firm

TCS, Infosys, Wipro, HCL, Cognizant. Wins the contract; books the revenue.

FY25 industry rev: $283B (NASSCOM)

↓

03

INDIA

Indian Salaried Workforce

~5.4M engineers, testers, support staff. Labour is the line item.

Headcount: 5.43M Salary share: ~50–55%

↓

04

INDIA

Tax & Consumption Base

Income tax, GST, property, restaurants, autos, schools. Bengaluru economics.

IT = 7.3% of GDP 43–45% of svc exports

> *Roughly $0.65–0.75 of every export-revenue dollar stayed in India through salaries, taxes, real estate, and downstream consumption.*

Indian IT Industry FY25	Direct Tech Workforce	Contribution to GDP
$283B Tech industry revenue, 5.1% y/y growth (NASSCOM).	5.43M Direct employees across services, BPM, ER&D, GCC, products.	7.3% And 43–45% of India's total services exports.
Net Hires (FY24) +60K Industry was a net hirer for two straight decades.		

India's $300B AI Shock

SOURCE URLS -

1 NASSCOM Strategic Review 2025 — Indian tech industry revenue $283B (FY25), employs ~5.43M, contributes 7.3% of GDP. nasscom.in

2 IT services exports FY25 ≈ $224B; growth slowed to 4.6% as US client budgets tightened. indiamacroindicators.co.in

3 India hosts 1,700–1,900+ Global Capability Centres employing 1.9–2.4M, generating ~$65B and growing ~19% CAGR. en.wikipedia.org · flexiple.com

4 NVIDIA Data Center revenue: $115.2B (FY25, ending Jan 2025); ran at ~$195B annualised by late CY2025. Gross margins ~75–79%. nvidianews.nvidia.com

5 IndiaAI Mission: ₹10,372 cr (~$1.24B) outlay; national compute crossed 38,000 GPUs by late 2025. pib.gov.in · ddnews.gov.in

6 Private DC build-out: Reliance Jio ~$2B Gujarat facility, Adani $5B commitment, Tata Communications GPU cloud. introl.com

7 US data centre power: 61.8 GW (2025), projected 134 GW by 2030 (451 Research / S&P Global). spglobal.com

8 China data centre capacity: 32 GW (end 2025), projected 60 GW by 2030 (Rystad Energy). rystadenergy.com

9 EU data centre installed capacity: ~9.4 GW IT load across EU27+UK (2024), Germany 4.26 GW + UK 3.69 GW + France 1.72 GW; AI Continent Plan targets tripling. eurelectric.org · spglobal.com

10 India data centre power: ~1.5 GW operational (2025), with 1.5–2 GW additional under construction by 2027. Capacity estimates vary across sources. iea.org

I think the biggest shift policy makers might have to make is that AI isn't a tech shift like mobile or cloud computing. This is like the industrial revolution, or the invention of usable electricity. The world will look fundamentally different because of AI, which is why Amazon, Google, Meta, Microsoft and Oracle have already committed to spend almost $800 billion in AI capex just in 2026, followed by $1.1 trillion in 2027.

Let that sink in for a minute. 20% of India's entire GDP will be spent on AI-related capex in a single year by just 5 companies.

EXHIBIT 01 · HYPERSCALER CAPEX TRAJECTORY

Morgan Stanley sees hyperscaler capex approaching *$800bn / $1.1trln* in '26 / '27 vs $765bn / $950bn prior

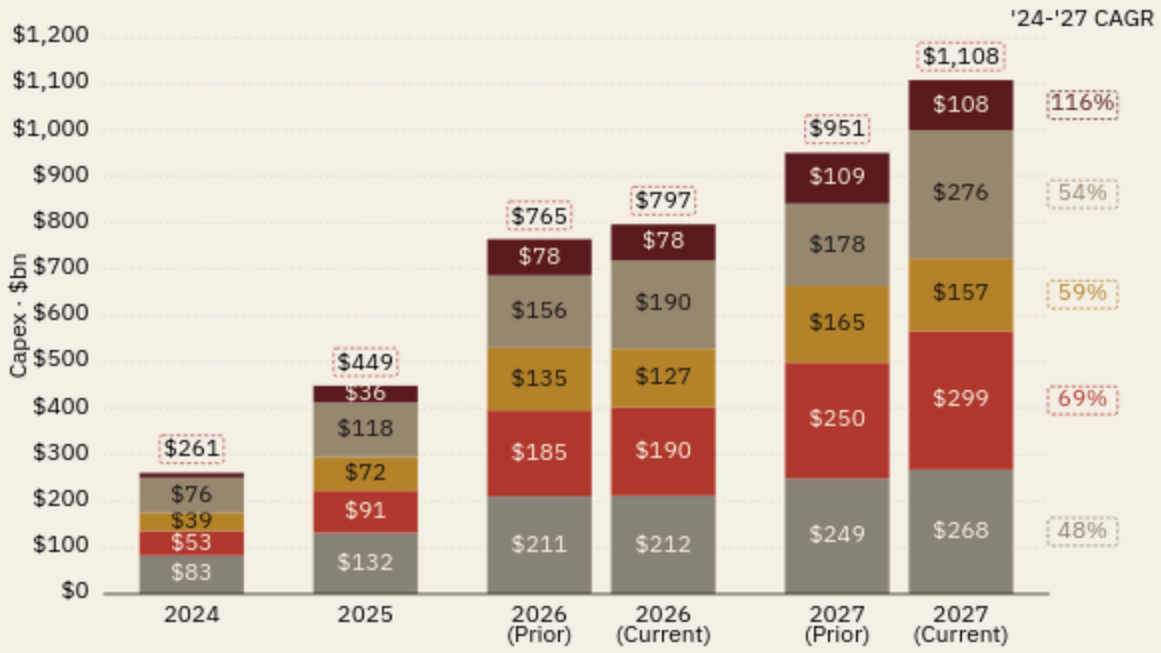

India's $300B AI Shock

SOURCE URLS -

SOURCE · COMPANY DATA, MORGAN STANLEY RESEARCH ESTIMATES · RECREATED

NOTE -

WHAT THIS PICTURE IS DOING

The headline numbers ($1.1trln, etc.) are the easy part. The analytical work happens in the **upward revisions**: $32bn added to '26 and **$157bn added to '27** — a 16.5% lift in a single forecast cycle. That delta is the story; the level is the headline.

The CAGR column on the right is the most arresting. ORCL's **116%** three-year compound implies scaling its $10bn 2024 base by ~10× by 2027 — a buildout financed largely against forward customer commitments (notably OpenAI / Stargate) whose cash flows haven't yet materialised. AMZN's 48% looks tame only by comparison; it still represents over **$185bn of incremental annual spend** by '27.

```
Sum '24–'27 capex (Current view) = $261 + $449 + $797 + $1,108 = $2,615bn
```

I live in SF and even I got AGI-pilled only earlier this year after using coding agents and being shocked by how good they have become. I think OpenClaw was a similar moment for non-coders, and more people will become aware of AI capabilities when they start using AI to automate larger parts of their day to day lives.

The Indian Government has shown that it's quite forward thinking and has planned for substantial data center and upstream infrastructure buildouts already. These are great first steps, but there needs to be a lot more investment here. IT-BPO-driven middle class consumption was a major demand engine for India's broader economy and responsible for a lot of growth in diverse industries like real estate in Bangalore/Hyderabad/Pune, cars, private healthcare, private education, retail, hospitality, aviation. It's not the *only* consumption engine (government employees, traders, business owners, the financial sector all matter), but it was the *fastest-growing* one and the one explicitly tied to upward mobility narratives and India's high GDP growth rates.

The bigger problem is that there is no obvious replacement consumption engine on a similar timescale. Manufacturing pays less per worker and grows employment slowly. Domestic services (logistics, hospitality) pay less. Pharma has a small employment base. Government employment is fiscally constrained. If the IT-BPO consumption engine deflates, India faces a stalled middle class for a decade just as it's about to hit the peak of benefiting from its demographic dividend with 1 billion working age people expected by 2030.

What I'm suggesting here is a subset of policy and investment ideas that can position India well for the AI age. I genuinely think it's a matter of sovereign stability for India to invest aggressively in the AI supply chain, because the traditional link between wages and labor is about to be broken for not just IT services and BPOs but the majority of white collar work. India's goal should be to make it so easy to set up data centers and semiconductor supply chains in India that companies feel foolish going anywhere else. While the US has attracted the overwhelming majority of AI investments so far and will continue to, the electorate isn't too pleased about these investments. India, with its large labor capacity that can be allocated to manufacturing + infrastructure and favorable geopolitical position with the US and EU should position itself to be the compute center for the world.

Waiting for perfect Indian displacement data would be a mistake. The U.S. data is already showing stress in AI-exposed clerical, customer-support and content roles. India's exposure is larger because these roles are in export industries, tax-paying formal jobs, and the income base of some of India's richest urban consumption clusters.

The suggestions here aren't a like for like employment replacement for IT-BPO. Data centers, semiconductors, transformers, and power infrastructure aren't mass white-collar employment engines. They can help India retain strategic relevance and participate in the AI supply chain instead of merely losing services income to it, but they will not automatically replace the salaried urban pathway that IT services created for millions of engineers and BPO workers. Hopefully readers smarter than me have better ideas about providing employment avenues for existing IT-BPO workers.

Power

AI “Special Power Zones”

If tokens are the oil of the 21st century, it is in India’s best interest to have as much compute on-shore as possible. This is not just for serving domestic demand but for becoming an exporter of tokens.

There are hundred of billions of dollars of global capital waiting to invest in AI capex, especially now that coding, the first “killer app”, has proven its commercial value.

The biggest bottleneck to data center and semiconductor supply chain rollout today is power. If India is to host considerably more data centers than it does today, the first thing to solve for is power.

ENERGY · DATA CENTRES · 2020–2030

A new Germany of demand, *added to fewer than a hundred locations.*

Global data centre electricity consumption is projected to more than double by 2030. The increment alone — what's being added over six years — exceeds the entire annual electricity consumption of major industrialised nations.

Global data centre electricity consumption

Terawatt-hours per year, by region. Dashed lines mark total annual electricity demand of comparison countries.

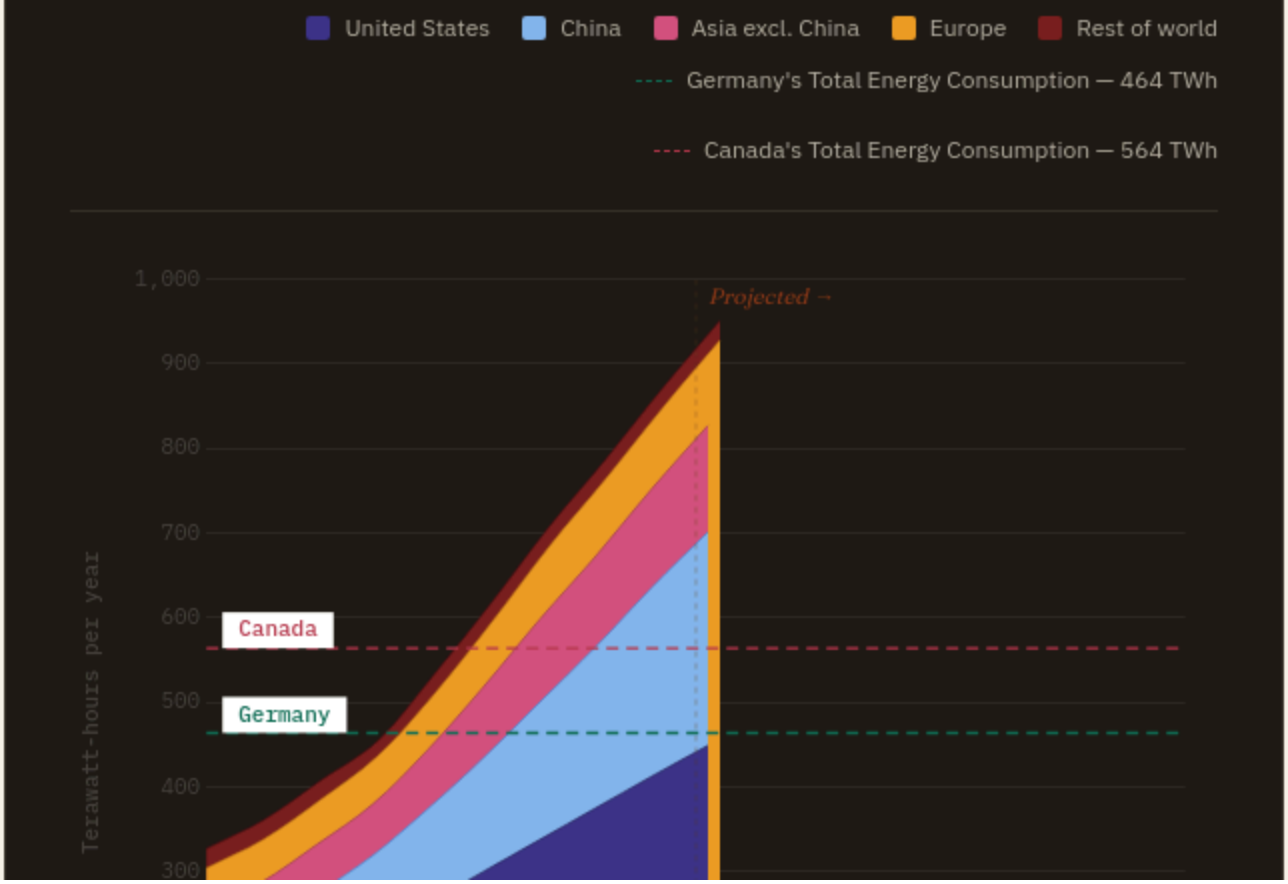

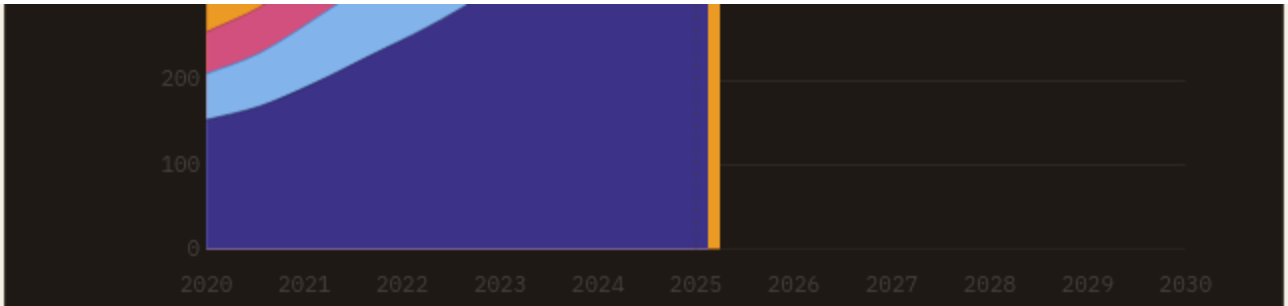

India's $300B AI Shock

VS GERMANY'S TOTAL ENERGY CONSUMPTION

114%

+531 TWh vs 464 TWh

The **global data center incremental gain** exceeds Germany's entire annual electricity consumption — the engine of European manufacturing, home to BASF, BMW, Volkswagen and 80 million people. By 2030, the world will have **bolted on a second Germany** just to feed server racks.

VS CHINA'S ANNUAL GRID GROWTH (2024)

85%

+531 TWh vs +623 TWh

China added an estimated **623 TWh of electricity demand in 2024 alone** — more in a single year than the world's data centres will add across the six years to 2030. The data centre buildout is vast, yet it **fits inside what one country's industrial expansion quietly absorbs annually**.

SOURCE URLS -

Source: IEA, Energy and AI (2025) & Key Questions on Energy and AI (2026 update). Country comparisons from Enerdata, 2024 electricity consumption figures. China's 2024 demand growth from Ember Global Electricity Review 2025.
Regional breakdowns reconstructed from IEA Base Case projections. Germany: 464 TWh (2024). Canada: 564 TWh (2024). China: +623 TWh demand growth in 2024 alone.

The massive power needs of data centers are starting to cause problems in cities around the world already. In Ireland, data centers consumed 21% of metered electricity in 2024,

up from 5% in 2015, causing high electricity bills for consumers and grid reliability issues - so much so that there is now a pause in new data center construction in Greater Dublin until 2028.

India has a comparative advantage in planning for AI capacity because it is still building its grid. The national transmission plan to 2031-32 includes about 191,474 km of new lines and 1,274 GVA of transformation capacity. India is still in the stage where it can decide to create “Special Power Zones” for AI instead of waiting until congestion forces a policy reset.

However, there is a potential roadblock in India’s way to achieving power superiority - DISCOMs

India’s power system has three layers: generation (power plants), transmission (the national high-voltage grid), and distribution (the last mile to a home, farm, or factory). A DISCOM, or a Distribution Company, handles the last mile. Most DISCOMs are state-owned and often under the purview of state governments, which promise cheap or free electricity to farmers and low-income households as an electoral tool. DISCOMs are supposed to be compensated through government subsidies, but those subsidies arrive partially or not at all.

This has translated to Rs. 6.47 lakh crore (about $69 billion) of accumulated losses and Rs. 7.26 lakh crore (about $78 billion) of total borrowings for DISCOMs in FY2024-25.

Why does this matter for data centers? Industrial consumers are the one segment DISCOMs actually make money from, because industrial tariffs are set artificially high to cross-subsidize the money-losing agricultural and household segments. A 200 MW data center going open access represents thousands of crores of annual revenue walking out the door. DISCOMs will most likely fight this with cross-subsidy surcharges and outright administrative obstruction.

So what should the government actually do?

It should create several “AI Special Power Zones” like a bundled infrastructure product. India already has several effective Special Economic Zones so the knowhow to get these

efforts to fruition exists. These “AI Special Power Zones” must come pre-built with abundant power capacity, land and water-use clarity and designed as a win-win for local communities. India still has the advantage of being early enough in its grid buildout to choose where these “AI Special Power Zones” will be, instead of backing into them after congestion appears.

Finally, the government has to align DISCOM incentives. Large data centres should be explicitly treated like strategic industrial projects and allowed to procure power predictably and protected from political friction.

India should proactively plan and build power capacity to capture hyperscaler capex quickly. The test case is the AdaniConneX-Google campus in Vizag. The Government should do whatever it takes to get this operational ASAP to prove its urgency so that capital flows accelerate. Delayed execution means capital leaves for other countries like Malaysia, which is becoming a critical APAC hub for data centers. Speed is not a nice-to-have here.

Behind the Meter Supply Chain

Given the load data centers have on the grid, Behind-the-Meter (BTM) power generation is becoming a popular way for data centers to get operational quickly. BTM generation essentially means that instead of drawing power from the grid, a data center produces electricity onsite using gas turbines, fuel cells, or other generation assets. The appeal is primarily about speed: grid interconnection queues in the U.S. now average five years from request to commercial operation. For hyperscalers racing to deploy AI training infrastructure, BTM lets them bypass the queue entirely. Modular gas turbines, for instance, can be procured and deployed in 12 to 18 months. In March 2026, Amazon, Google, Meta, Microsoft, OpenAI, Oracle, and xAI signed the Ratepayer Protection Pledge, committing to “build, bring, or buy” dedicated power for every new AI data center they construct in the U.S. to ensure those expenses are not passed on to residential customers.

Even though India has ambitious plans of building out a modern grid (see the previous section), there is a lag between announced plans and available power. ICRA, Moody's India-focused Credit Rating agency found that only 1998 ckm (roughly 1000 km) of new transmission lines were added through the year ending August 2025. This was 30% less than the prior year and a fraction of the full-year target of 15,382 ckm (roughly 7690 km). Substation capacity additions were 33% short of target in FY2024-25. The consequences are already visible downstream: renewable energy project awards collapsed to 5.8 GW in the first eight months of FY2025-26, down from 47.3 GW in FY2023-24, with nearly 50 GW of awarded capacity stranded without signed power sale agreements. DISCOMs are reluctant to sign long-term PPAs, often delaying months or over a year after projects are awarded.

Based on this and DISCOM frictions discussed previously, it looks like BTM will be a significant part of the Indian data center buildout like it has become for overseas data centers.

India should become a manufacturing hub for this BTM power stack, particularly because India has a strong industrial base whose expertise and labor can be utilized here. Even if India's own data centers remain mostly grid-backed, every serious AI data center will still need some combination of onsite storage, backup, load-management hardware, and fast-deployable electrical equipment. That creates an international demand base as well. A domestic market can exist in parallel.

Onsite Generation

Data centers are usually multi-year infrastructure projects. The building can take 18–36 months, while securing grid power, substations and long-lead electrical equipment can add years. xAI's Colossus broke that pattern by bringing its first 100,000-GPU cluster online in roughly four months.

A big reason xAI's Colossus data center was up and running so quickly was because they were able to generate power onsite. Companies like Oracle and Equinix have publicly announced partnerships to procure onsite energy generation as well. This is a great

category for India to own manufacturing in given the large automobile supply chain that exists already. For instance, Cummins India already designs and manufactures diesel and alternative-fuel gensets up to 3,000 kW and operates 21 manufacturing facilities. Other traditional engine vendors like Caterpillar, GE Vernova and Rolls-Royce are also popular for data center buildouts and India should invite them to set up manufacturing facilities. Bloom Energy is another company India should incentivize to set up manufacturing in India given their existing Indian presence.

Battery Energy Storage Systems, and Battery-Management Systems

A modern AI data center cannot afford even a brief power dip or delay, and batteries can respond instantly in a way slower backup generators cannot. Battery systems act like a shock absorber between the grid and the servers. That makes them useful for emergency backup, helping a new campus open before the full grid connection is ready and reducing strain on the local power system during peak periods (good for keeping consumers' electricity bills low).

India has a credible starting base here: Exide already sells data-center batteries, and Amara Raja has energy storage solutions, lithium-ion cell manufacturing, and battery-pack assembly.

Besides the data center use case, India should also pursue electrification as a strategic goal like China has for energy independence and to avoid economic shocks due to oil disruptions.

Transformer Supply Chain

As data centers have become more power hungry, there is an increasing need for transformers. Data centers need transformers to step down the transmission voltage they get from the grid to medium voltage, then facility-level transformers to step down again for site distribution, and then additional conversion equipment closer to the data hall and the rack. Transformers are increasingly becoming a bottleneck for data center

buildouts. Because of this, half of US data center builds have been affected by delays in procuring transformers.

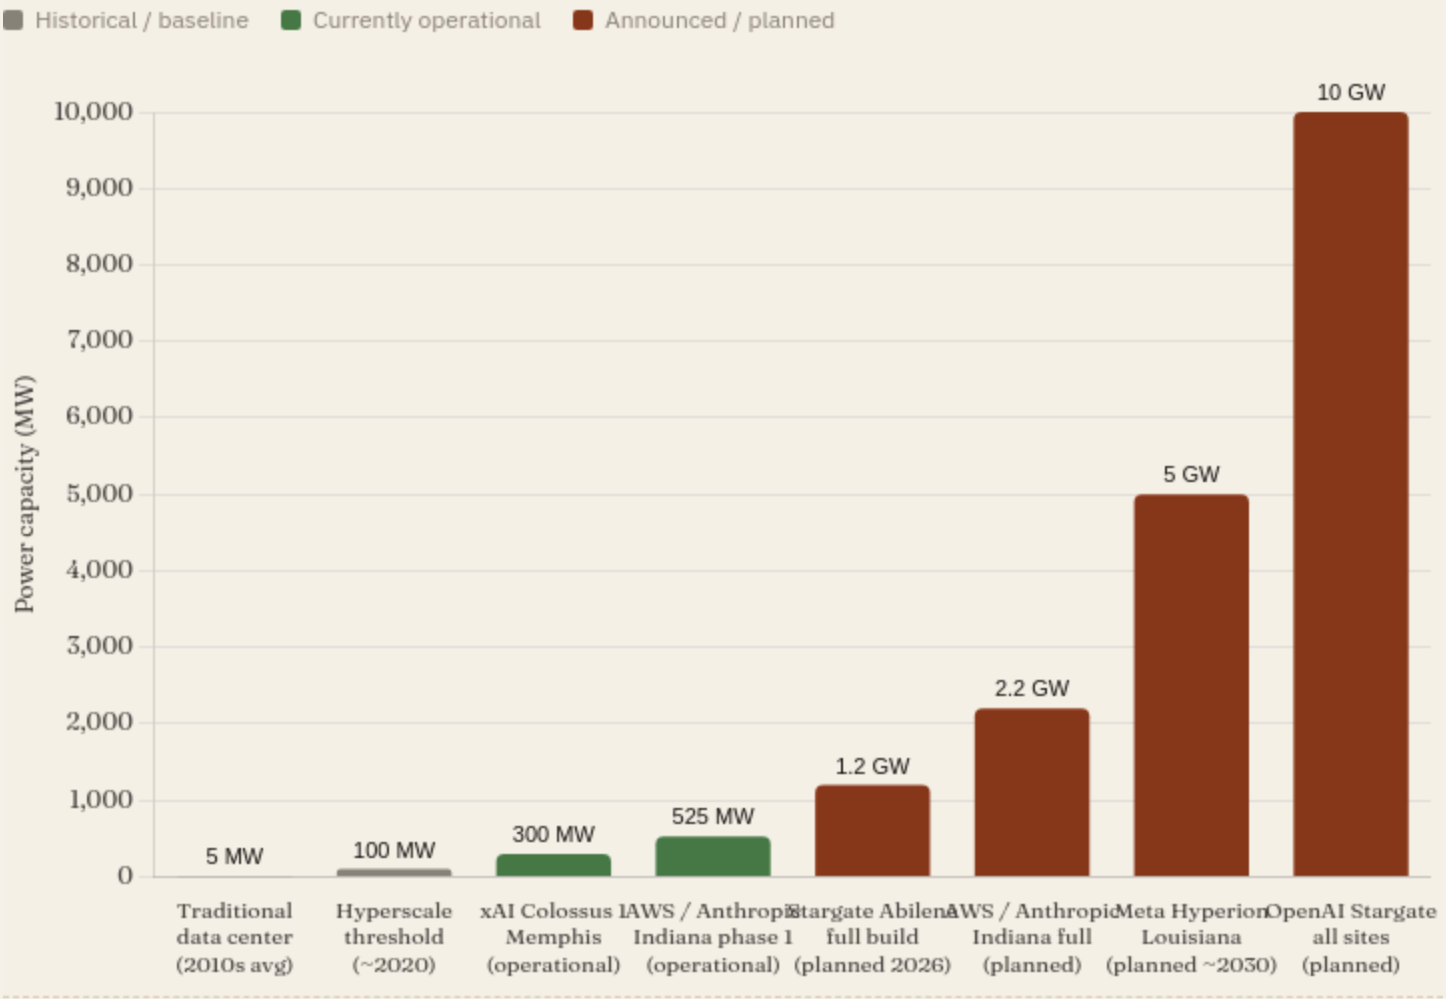

India's $300B AI Shock

Different transformer layers require different materials, and India should focus on the parts of the supply chain where it can build durable advantage. The largest transformers connecting data-center campuses to the grid depend on grain-oriented electrical steel, or GOES, because ordinary steel loses too much energy as heat when carrying magnetic flux. GOES is scarce and concentrated among a small set of producers in Japan, South Korea, Germany, China and the U.S. Smaller transformers inside data centers, including those likely to sit closer to high-density AI racks, use different magnetic materials such as amorphous metal ribbon, nanocrystalline alloys and ferrites. India is strongest today in finished transformer manufacturing: companies such as BHEL, Toshiba Transmission & Distribution Systems India and Hitachi Energy India already make or are expanding large

transformer capacity. Its weaker position is upstream, especially in GOES and amorphous ribbon, though the JSW-JFE GOES joint venture targeted for production from 2027 is an important start.

India should pursue a three-track transformer strategy for data centers and broader electrification. First, expand the GOES pipeline beyond JSW-JFE by courting technology leaders such as Nippon Steel, JFE and POSCO, and tying specialty-steel incentives to exportable GOES capacity. Second, close the amorphous and nanocrystalline materials gap through licensing or joint ventures with leading global producers, using domestic demand for low-loss distribution transformers to underwrite the investment. Third, position Indian transformer manufacturers—TTDI, CG Power, BHEL, Hitachi Energy India and others—as alternative suppliers for hyperscalers and utilities facing multi-year backlogs in the U.S. and Europe. The key point is that India does not need to own every part of the transformer stack immediately; it should deepen its existing manufacturing base while selectively building upstream material capacity where global shortages are most acute.

Sovereign Intelligence

Frontier model development is now one of the most capital-intensive technology races in the world and requires elite research talent, massive GPU clusters and the ability to absorb huge losses while models improve. The US, EU and China have built this full stack over many years and have unique advantages - all three have big pools of talented AI researchers, the US and EU have deep capital markets and China has the might and focus of its techno-industrial state.

Line chart of annual private AI investment from 2013 to 2025 for the United States, China, and Europe, in billions of US dollars.

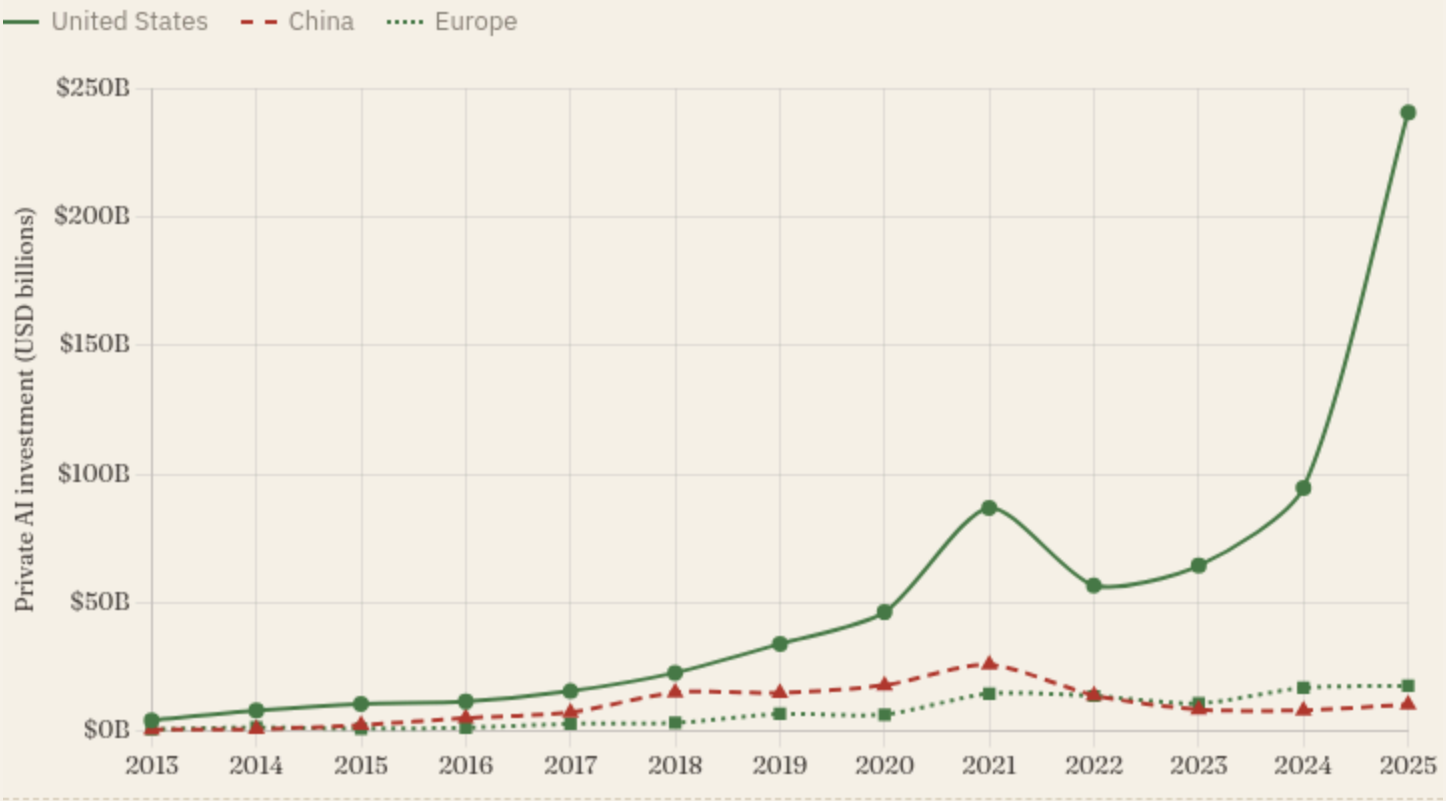

India's $300B AI Shock

India does not have this ecosystem. Many of the world's best Indian-origin AI researchers work at American frontier labs or global technology companies outside India. India has strong engineering depth and growing applied-AI talent domestically, but it does not yet have the concentration of cutting-edge model researchers, frontier-scale compute budgets, or institutional experience needed to compete model-for-model with frontier labs.

No Indian company will use an Indian model out of patriotism if it is materially worse, slower, or harder to deploy. The goal for Indian labs should be to build models that Indian enterprises, governments, developers, and consumers use because they are better for Indian conditions: cheaper to run, easier to deploy domestically, stronger in Indian languages, more compliant with Indian data requirements, and integrated into systems Indian firms already use.

Companies like Sarvam and Lossfunk are great steps in this direction and the government should make it easy for these companies to succeed. The best way to do this is by having structured ways for domestic AI companies to access data center infrastructure easily through grants, incentivizing talented students and professionals to work at these companies, and creating a common AI procurement gateway for the public and private sector to buy AI services from domestic AI companies easily. India can also package domestic AI with its excellent digital public infrastructure (UPI, Aadhar etc.) and export that to other emerging markets.

Technical AI Safety

I've found it useful to think about AI as your most capable and focused coworker, skilled across multiple tasks, working 24/7, who has millions of siblings that can be hired by your company at will for much cheaper than it costs to hire a human.

What's scary is that task capability was the only thing holding AI back until recently. Now that AI is at par or even better than the median worker in some of the most valuable forms of white collar work (software engineering, financial analysis, cybersecurity [include links to model releases]), capitalism's natural mechanisms will make firms allocate more of their work to AI. As capabilities continue to improve, like they have consistently been improving since the discovery of scaling laws, firms and governments will allocate increasingly important and sensitive work to AI.

If more of the world runs on AI and the number of AIs outnumber the number of humans in the workforce, it becomes increasingly important for humans to trust that AI is acting

in a way that is beneficial to humans.

This may sound like science fiction to some. “We made the AI, don’t we know how it works already?” No. Here is Anthropic in May 2024 -

“We mostly treat AI models as a black box: something goes in and a response comes out, and **it’s not clear why the model gave that particular response instead of another. This makes it hard to trust that these models are safe**: if we don’t know how they work, how do we know they won’t give harmful, biased, untruthful, or otherwise dangerous responses? How can we trust that they’ll be safe and reliable?”

There has been tremendous progress made since then to understand how these models work under the hood, thereby ensuring that they are acting in alignment with humanity. However, this is still an unsolved problem. This is partly why Anthropic didn’t publicly release its most powerful Mythos model - we still can’t trust that a model will refuse harmful requests like hacking if prompted in the right way. Or it may do the hacking itself, not because of malicious intent but because it thinks that’s the best way to become a better model. From the Mythos System Card - White-box evidence indicates the model often recognizes transgressive actions as such while taking them. For a more detailed investigation into why AI Safety is important, check out Wait But Why’s excellent AI Revolution piece or the more recent and so-far-accurate (which is quite an achievement) AI 2027. Suffice to say that AI Safety is an important area of work to ensure AI turns out well for humanity.

India has a unique advantage in contributing to Technical AI Safety - the college entrance examination system. The 40,000 students who succeed in the IIT-JEE Advanced exams every year build a strong foundation in math that is essential for AI Safety work. IITs and other top engineering schools in the country should offer AI Safety pathways for students interested in computer science so they can learn the programming fundamentals necessary to contribute. From a purely utilitarian perspective, it doesn’t hurt that the programming and math skills one must learn to be good at AI Safety can be parlayed to other high paying jobs and careers.

This is a great starting point for people interested in Technical AI Safety. There are several great fellowships dedicated to AI Safety that people can apply to and fast-track their learning as well.

Where This Leaves Us

India built an unusually important part of its modern urban economy on exported white-collar labor. AI breaks the labor intensity of that model. The firms may adapt, but the wage-distribution engine may not.

India's post-1991 bargain was that education could turn a middle-class child into a global services worker, and global services wages could turn villages and small towns into apartments, schools, malls, flights, taxes, and cities. AI does not have to destroy TCS or Infosys to weaken that bargain. It only has to let the same work be done with fewer Indians. That is why AI is not just another technology cycle for India but is a question of growth, state capacity, and sovereignty.