




CIO&LEADER

STATE OF AI IN INDIAN ENTERPRISES 2026

*Benchmarking AI Adoption, GenAI
Production Reality, Governance
Maturity, and the Agentic Enterprise*



Contents



Foreword	03
Executive Summary	04
AI Strategy, Investment and Ownership	08
The Economics of AI	16
The Cost Surprise	18
Workforce and Talent Readiness	22
AI Talent Strategy	24
AI Training & Literacy	25
GenAI — From Pilot to Production	26
Agentic AI — The Emerging Frontier	34
AI Governance and Risk Management	40
Security and Regulatory Compliance	48
Future AI Outlook	54
Conclusion	59
About This Study	60
Appendix	61
A. Survey Instruments	61
B. Notes on Data Interpretation	61
C. About CIO&Leader	61

Conceptualized and Executed by

MANAGEMENT

Chief Executive Officer: **Rohit Gopakumar**

Chief Operating Officer: **Sunita Quadros**

Senior Director: **Sachin Nandkishor Mhashilkar**

EDITORIAL

Editorial Director: **R Giridhar**

Chief Editor: **Jatinder Singh**

Senior Correspondent: **Jagrati Rakheja**

Senior Correspondent: **Punam Singh**

DESIGN

Art Director: **Shokeen Saifi**

Assistant Manager – Design: **Manish Kumar**

SALES & MARKETING

Assistant Director - Brand & Community: **Vandana Chauhan**

Senior Manager - Marketing: **Rajiv Pathak**

Director - Projects: **Naveen Chand Singh**

Assistant Director - Projects: **Hafeez Shaikh**

Chief Manager - Projects: **Sidharth Ranga**

Senior Project Manager: **Supriya Sahoo**

Senior Project Manager: **Aanchal Gupta**

COMMUNITY ENGAGEMENT & DEVELOPMENT

Senior Manager - Community Development: **Ekta Srivastava**

Manager - Database: **Neelam Adhangale**

Manager - Community Development: **Vaishali Banerjee**

Manager - Community Development: **Reetu Pande**

Manager - Community Development: **Snehal Thosar**

Executive - Database: **Dipti Gamre**

OPERATIONS

Senior Manager - Digital Operations: **Jagdish Bhainsora**

Senior Manager- Video Content: **Sunil Kumar**

Manager - Operations: **Himanshu Kumar**

Manager - Operations: **Amit Singh**

PRODUCTION & LOGISTICS

Manager - Operations: **Mahendra Kumar Singh**

Executive - Operations : **Vishwas Gandhi**

LEGAL & COMPLIANCE

General Manager - Legal: **Anis Kazi**

OWNER

© 2026 All rights reserved

Reproduction of the contents of this document, in whole or in part, is prohibited without the written permission from

Times Strategic Solutions Ltd

7th Floor, Sunteck Icon,

BKC Junction, Kalina, Santacruz East

Mumbai 400098

Foreword

Before the Next Wave, Solve the Last One

A year ago, CIO&Leader's State of Enterprise Technology survey found that 93.5% of Indian enterprises planned to increase their AI and analytics spending. Most did. The investment has now arrived at the point of deployment. What this year's data reveals is what happens when intent meets execution at scale — and the distance between the two turns out to be considerably wider than the spending commitments would suggest.

Three findings from this year's study deserve to be read together, because in sequence they describe the central challenge facing Indian enterprise AI in 2026.

First: 52% of AI pilots never reach production. **Second:** only 12% of organizations report significant, measurable returns from their AI investments. **Third:** 46% of our Survey 3 respondents — the cohort thinking hardest about governance and risk — identify agentic AI as the capability most likely to reshape enterprise operations over the next eighteen months.

The spending arrived. The returns have not.

The 2025 survey showed the groundwork forming: 71% of enterprises had executive buy-in, IT operations led as the proving ground, and every respondent called AI/data literacy essential. This year proves that recognition hasn't turned into investment—15% still have no formal AI training despite naming productivity as their top goal. Data quality remains the biggest barrier to scaling, unchanged from last year. These are structural issues persisting across two survey cycles, not problems money alone will fix.

The 2025 survey said AI had moved past experimentation to become the enterprise's new operating layer—live in production across IT, finance,

cybersecurity, and content. That's true at the frontier. But 60% of this year's respondents are still stuck at the pilot stage, even after two-plus years of trying. The ambition is there. The execution isn't.

That is the tension this report was designed to illuminate. Indian enterprises are still working through the foundational problems of generative AI deployment — converting pilots, measuring outcomes, building governance — while the next and considerably more complex challenge is already arriving. Agentic AI does not merely generate outputs for humans to review. It acts. It executes multi-step workflows. It makes decisions within defined parameters, and sometimes beyond them. The organizations that have not yet solved the pilot-to-production problem for a text summarization tool are now being asked to think about the governance of autonomous agents acting on their behalf.

The technology leaders who responded to our three surveys represent some of the most active AI practices in Indian enterprise. They are not the laggards. If 52% of this group cannot get pilots to production and 57% cannot demonstrate or determine ROI, the numbers for the broader market are likely less encouraging.

The barriers are known. The conversion problem is measurable. The governance gap is quantified. What this year's data adds to last year's findings is urgency — because the agentic wave will not wait for the generative one to be sorted.



R. Giridhar
Editorial
Director -
Technology
ET Edge
CIO&Leader



Jatinder Singh
Chief Editor
ET Edge
CIO&Leader


Executive Summary

Two years into India's enterprise AI surge, the investment has arrived—but the execution discipline has not kept pace. This report, based on three independent surveys of 317 senior technology leaders conducted in May–June 2026, reveals a market that is broadly active yet not broadly productive. While 83% of organizations cite productivity improvement as their primary AI driver and 93% had planned to increase spending in 2025, the data exposes a widening gap between strategic ambition and operational reality.

The central finding is stark: 60% of respondents remain at the pilot or exploration stage after two or more years of investment, 52% report that fewer than one in four AI pilots reaches production, and only 12% can demonstrate significant, measurable ROI. These are not the numbers of a market failing for lack of vision; they are the numbers of a market failing for lack of process discipline—specifically, in pilot selection, business ownership, ROI measurement, and data readiness.

Compounding this challenge is the velocity of the next wave. While most organizations are still resolving the foundational problems of generative AI deployment, 64% have already moved into active piloting or production of agentic AI—autonomous systems that act, execute, and decide rather than merely generate outputs. The governance frameworks adequate for today's GenAI tools will not hold for tomorrow's agents. The organizations that build agentic-specific controls, data quality infrastructure, and DPDP compliance discipline now will be the ones that scale sustainably. Those that wait for the generative wave to be fully resolved before addressing the agentic one risk being governed by incidents rather than by design.

This report is not a prediction of what Indian enterprise AI will become. It is a calibration of where it actually stands—and a map of the four non-deferrable priorities that must run in parallel, not in sequence, for the next eighteen months.



Study Overview



AI Strategy and Investment

Enterprise AI investment is growing, productivity-driven, and structurally under-resourced relative to stated

ambitions. Productivity improvement is cited by 83% of Survey 1 respondents as the primary driver — yet 58% allocate less than 10% of the IT budget to AI, and only 14% allocate more than 20%. The CIO or central IT function holds AI strategy ownership at 40% of organizations, the largest single category by a substantial margin — an appropriate posture for the current stage, but one that will need to migrate toward shared business-unit accountability as organizations attempt to scale. Hybrid cloud plus on-premise dominates AI infrastructure at 65%, shaped by DPDP data residency requirements, inference cost economics, and legacy integration constraints.



AI Economics: Cost, ROI and Scale

The economics of enterprise AI are opaque at most organizations. Fifty-five per cent either experienced AI cost

overruns in the last twelve months or do not track AI costs separately — making rational scaling and board-level business-case construction impossible. Only 12% report significant measurable ROI. Fifty-seven per cent have no measurable ROI or cannot determine whether returns exist. Data quality and readiness (56%) and lack of measurable ROI (52%) are the top two barriers to scaling AI — and they are circular: data quality deficits impair model performance, which impairs business outcomes, which impairs ROI measurement, which impairs the investment case for addressing data quality.



Workforce and Talent Readiness

Seventy-eight per cent of organizations use a combination of external hiring and internal upskilling or rely primarily

on upskilling — the strategically sound approach. But the structural gap is in literacy investment: 15% of respondents report no formal AI training, despite citing productivity improvement as the primary AI investment driver. Shadow AI usage is widespread: 48% of respondents estimate that more than 10% of their workforce uses unsanctioned AI tools, and a further 7% lack the visibility to estimate at all. The absence of formal training does not suppress employee AI adoption; it redirects it to unmonitored channels.



Generative AI: The Pilot-to-Production Reality

Function-level deployment is the most common GenAI state (36% of Survey 2 respondents); only 6% have embedded

GenAI into core enterprise workflows. Analytics and reporting (23%), customer support automation (20%), and code generation (19%) lead on measurable value — all three use cases defined by clear output metrics. The central production problem: 52% of respondents report a sub-25% pilot-to-production conversion rate, and 11% do not formally track conversion at all. The primary causes of pilot failure — data quality (42%), unclear ROI or business case (35%), and lack of business ownership (31%) — are all addressable before a pilot begins, through intake governance rather than post-mortem remediation. Most organizations (49%) use a hybrid model strategy; 57% include commercial frontier models (GPT-4, Claude, Gemini) in their portfolio, creating concentration risk that warrants active vendor dependency assessment.

Study Overview



Agentic AI: The Emerging Frontier

Sixty-four per cent of the GenAI survey population has moved beyond evaluation into active piloting or

production deployment of agentic AI — a figure higher than most market assessments of enterprise agentic readiness would predict. IT operations (66%), customer service (50%), and cybersecurity operations (48%) are the priority deployment functions, following a risk-adjusted logic: agents are directed first at domains where IT has operational ownership and where errors are recoverable. Sixty-two per cent of respondents require human approval or active human oversight of agent actions — the appropriate posture for the current state of enterprise agentic maturity, but a transitional one that must evolve through deliberate design rather than reactive escalation.



AI Governance and Risk Management

Enterprise AI adoption continues to outpace governance maturity. Only 29% of organizations have established

mature enterprise-wide AI governance frameworks, while 21% still rely on informal guidelines or no formal governance at all. Shadow AI remains widespread, with nearly half of organizations reporting unsanctioned AI usage. Data privacy (81%) and third-party AI vendor risk (49%) have emerged as the two dominant governance concerns, reflecting the growing complexity of managing AI across commercial models, cloud platforms, and SaaS ecosystems. The findings suggest that governance must evolve from policy documentation to operational discipline if AI is to scale responsibly.



Security and Regulatory Compliance

Security preparedness has not kept pace with AI adoption. Only 19% of organizations consider themselves

highly prepared for DPDP-related AI compliance despite data privacy being the industry's foremost concern. Ten per cent report confirmed AI-related security incidents in production environments, while another 14% suspect incidents without confirmation. Dedicated AI incident response capabilities remain rare, indicating that most enterprises are still adapting traditional cybersecurity frameworks to the emerging risks posed by AI systems, autonomous agents, and third-party AI services.



Future AI Outlook

Despite ongoing challenges in scaling GenAI, enterprise leaders remain firmly focused on the next wave. Agentic AI

has emerged as the capability expected to create the greatest enterprise impact over the next 18 months, with organizations increasingly moving from experimentation to operational deployment. The survey indicates that future competitive advantage will depend less on adopting new models and more on strengthening governance, improving data quality, establishing measurable business outcomes, and designing AI systems that combine autonomy with appropriate human oversight. Organizations that build these foundations today will be significantly better positioned to scale AI safely and sustainably.

FOUR STRATEGIC PRIORITIES FOR 2026–27

The data across three surveys converges on four priorities that are non-deferrable and must run in parallel, not in sequence.



Close the pilot-to-production gap. Require a named business owner committed to operationalizing the output before a pilot is approved. A two-condition intake gate — confirmed business ownership and data readiness assessment against the specific use case — eliminates the majority of pilot failure causes before resources are committed.

01

Build ROI measurement infrastructure. The 57% without measurable ROI predominantly face a measurement problem, not a performance problem. Define success metrics before deployment. One credible, rigorously measured data point breaks the ROI credibility cycle more effectively

than an abstract measurement framework built retrospectively.

02

Treat DPDP compliance as a present obligation, not a future preparation. The three immediate workstreams — consent mechanism documentation for AI-processed personal data, data

03

principal rights workflows for AI-generated outputs, and third-party AI vendor due diligence with contractual DPDP alignment — are buildable with existing governance infrastructure. None requires new technology investment.


04

Design agentic governance before production scale. The 64% already in active agentic piloting or deployment should be building agentic-specific controls now — covering autonomy thresholds,

audit trails, and escalation paths. The governance frameworks designed for generative AI tools will not hold for autonomous agents. The one sequential dependency: effective agentic governance requires data quality and consent infrastructure sufficient for agents to act on reliably. Establish that foundation first.



AI Strategy, Investment and Ownership





Executive Summary

Two years into what most organizations describe as a serious AI commitment, the strategic fundamentals are becoming clear — even as the executional gaps remain pronounced. Survey 1 of the State of AI in Indian Enterprises 2026 study, completed by 117 enterprise technology leaders, maps the investment drivers, budget realities, ownership structures, and infrastructure decisions that define the current posture of Indian enterprise AI strategy.

The overarching picture is one of aspiration ahead of allocation. Productivity improvement is the dominant investment rationale, cited by 83% of respondents — yet the majority still commit less than 10% of the IT budget to the initiatives designed to deliver it. Strategy ownership sits decisively with the CIO, which is the appropriate structure for an Adopter-stage program but an increasingly limiting one as organizations attempt to migrate AI from a technology initiative to an enterprise-wide operating model change.

The 2025 CIO&Leader State of Enterprise Technology survey, conducted across a comparable enterprise readership base twelve months earlier, provides useful context for reading these numbers. That study found that only 15.8% of senior technology leaders described themselves as "highly strategic and informed" on AI, with 26.3% still in an "early-stage awareness" phase and the majority sitting between early exploration and moderate planning. The distribution is consistent with this report's finding that 60% of respondents remain at the exploration or pilot stage after two or more years of investment — a market that is broadly active but not yet broadly productive.

Two shifts between the 2025 and 2026 data are worth noting. First, leadership commitment has strengthened. The 2025 survey found 71% of organizations with executive leadership fully aligned on AI adoption, and reported that over 60% of AI strategy ownership sat with CIOs, CDOs, or business

heads. The 2026 data confirms this pattern, with CIO or central IT function holding strategy ownership at 40% — the dominant single ownership category — and executive and C-suite ownership accounting for a further 22%. Together, these figures represent the majority of survey respondents operating under recognized senior leadership accountability for AI. Second, spending intent has converted to actual investment increases, but without yet producing a step-change in proportional allocation. The 93.5% of 2025 survey respondents who planned to increase AI and analytics spending appear to have followed through — but the 2026 data shows 58% still allocating less than 10% of the IT budget to AI, suggesting that the increases are real but incremental rather than transformative.

The infrastructure data tells a similar story of pragmatic constraint. Sixty-five per cent of organizations have settled on hybrid cloud plus on-premise as their AI operating model, shaped primarily by DPDP Act data residency obligations, inference cost economics, and the integration requirements of legacy systems built for an earlier era. It is the architecturally sound choice for most Indian enterprises at this moment — and one that carries a hidden operational cost in platform engineering complexity that is consistently underestimated in technology investment planning.

The conditions for scaling AI are present in most organizations in this survey. The structural gaps are in depth of funding and breadth of ownership — and both are addressable before the next wave of capability investment arrives.

Bottom Line: The conditions for AI scaling exist at most organizations — CIO ownership, hybrid infrastructure, a clear productivity mandate. What is structurally missing is depth of commitment. Budget allocation is too thin to match strategic ambition, and ownership is too concentrated in IT to produce enterprise-wide change. The strategy is credible; the resourcing is not.


1.1

Where Do Organizations Stand?

Survey 1 asked respondents to characterize their organization's current AI journey. Thirty per cent are running pilots or POCs, 30% are exploring AI opportunities, 19% are scaling selected deployments, 17% have begun enterprise-wide transformation, and 5% describe themselves as AI-first.

The roughly equal split between pure exploration and active deployment — with less than a quarter at true enterprise scale — is consistent with the maturity distribution in Section 9. These are not organizations at the beginning of their AI journey; most have been experimenting for two or more years. The concentration at the exploration and pilot stages suggests that the barriers to moving from experiment to production are more persistent than many organizations anticipated.

Stage of AI Use



Most Indian enterprises remain in the experimentation phase, with 60% still exploring AI or running pilots despite two or more years of investment

CIO Takeaway: If you are still primarily at pilot or exploration stage after two or more years, your problem is not strategy — it is selection discipline. The organizations moving to Scaler stage share one characteristic: they stopped running pilots to discover whether AI works and started running them only when a business owner had already committed to operationalize the output. Tighten your pilot intake criteria before expanding your pilot portfolio.


1.2

What Is Driving AI Investment?

The Survey 1 findings on investment drivers reveal a clear hierarchy. Productivity improvement is the primary cited driver, identified by 83% of respondents. Customer experience improvement follows at 71%. Innovation and new product development (62%), cost optimization (61%), and revenue growth (50%) complete the top five.

Risk reduction and compliance is the least-cited primary driver at 38%, though governance and risk concerns appear prominently as barriers to scaling rather than as drivers of investment, suggesting that the compliance case for AI is more often reactive than proactive in Indian enterprise contexts.

Primary Investment Driver for AI



Productivity and customer experience dominate AI investment priorities, while revenue growth and compliance remain secondary drivers.

CIO Takeaway: Productivity (83%) and cost optimization (61%) dominate the investment case — both defensive, both internal. Only half your peers are making a revenue growth argument for AI. If your AI investment narrative is built entirely on efficiency, you are building a cost-reduction program, not a competitive capability. Pressure-test whether your portfolio has any cases that could generate new revenue or strategic differentiation, not just cost savings.


1.3

How Much Budget Is AI Getting?

Despite strong stated commitment to AI, budget allocation tells a more measured story. Fifty-eight per cent of respondents allocate less than 10% of the IT budget to AI initiatives — 27% allocate less than 5%, and 31% allocate between 5% and 10%. Only 14% allocate more than 20% of the IT budget to AI.

For an initiative that 83% of respondents identify as a primary productivity driver, this level of funding is modest. It reflects the reality that most organizations are still in a phase where AI is one of several technology investment priorities, competing with cloud migration, cybersecurity, and core system modernization.

AI Budget (as % of IT Budget)



AI remains strategically important but financially constrained, with most organizations allocating less than 10% of their IT budgets to AI initiatives.

CIO Takeaway: Fifty-eight per cent of organizations allocate less than 10% of the IT budget to AI. If you are in this group while simultaneously describing AI as your primary productivity driver, you have a credibility gap your board will eventually surface. The investment thesis and the investment level need to be consistent. Either increase the allocation to match the ambition, or recalibrate the ambition to match the allocation — but do not carry both positions at once.


1.4

Who Owns AI Strategy?

The CIO or central IT function owns AI strategy and budget allocation in 40% of organizations — the largest single ownership category by a substantial margin. Executive and C-Suite ownership follows at 22%. AI Centers of Excellence account for 10%, shared ownership models for 9%, and direct business unit ownership for another 9%. A Chief Digital or AI Officer holds ownership in 7% of cases, and 4% report no clear ownership.

The CIO's dominant position is both a strength and a risk: it grounds AI strategy in technical reality, but may bound AI ambition within IT planning cycles and IT risk appetite.

AI Strategy Ownership



The CIO continues to lead enterprise AI strategy, though broader business ownership will be essential as AI scales.

CIO Takeaway: CIO ownership at 40% is appropriate at the Adopter stage but becomes a constraint at the Scaler stage. If AI strategy sits entirely within IT, business unit leaders have no accountability for outcomes, no authority over use case selection, and no stake in adoption rates. The transition from technology program to operating model change requires ownership migration. Start moving toward shared accountability with at least two business unit leaders before your next planning cycle, not after your next failed deployment.


1.5

What is the Infrastructure Model for AI?

Sixty-five per cent of respondents describe hybrid cloud plus on-premise as their primary AI infrastructure model. Multi-cloud follows at 13%, fully public cloud at 12%, and fully on-premise at 10%. The dominance of hybrid infrastructure reflects a practical compromise shaped by India-specific constraints: data residency obligations under the DPDP Act, the cost economics of inference at scale, legacy system integration requirements, and sector-specific regulatory mandates.

For AI workloads, hybrid architecture enables sensitive inference on-premise while using cloud for model training and burst capacity. The tradeoff is operational complexity: managing model pipelines across cloud and on-premise environments requires platform maturity that many organizations are still building.

Infrastructure Model for AI



Hybrid cloud with on-premises infrastructure has emerged as the preferred AI deployment model, balancing scalability with regulatory and data residency requirements.

CIO Takeaway: Hybrid cloud plus on-premise is the right architecture for DPDP compliance and inference cost management, but it carries a hidden operational cost: the platform engineering capacity required to manage AI model pipelines across environments is significant and is frequently underestimated. If you do not have a dedicated MLOps function or an equivalent platform capability, your hybrid infrastructure will slow deployment velocity as you scale. Budget for the operational layer, not just the infrastructure layer.

This image shows a full page of white paper with horizontal ruling lines. The lines are evenly spaced and run across the width of the page, typical of notebook or legal stationery. There are no margins, text, or other markings on the page.



The Economics of AI





Executive Summary

Enterprise AI has a measurement problem that is being misread as a performance problem. Survey 1 of the State of AI in Indian Enterprises 2026 study finds that the majority of Indian organizations cannot determine — and many are not systematically attempting to determine — the economic returns on their AI investments. The consequence is a strategic vulnerability: organizations that cannot demonstrate AI value to their boards are building investment cases on strategic narrative and peer pressure rather than evidence.

The scale of the measurement gap is striking. Fifty-seven per cent of respondents either have no measurable ROI from their AI investments or cannot determine whether returns exist. Fifty-five per cent either experienced cost overruns in the past twelve months or do not track AI costs separately enough to know whether they did. These are not marginal organizations; they represent the majority of a survey population that skews toward the more AI-active segment of Indian enterprise technology leadership.

The 2025 CIO&Leader State of Enterprise Technology survey provides a counterpart to these findings from the investment intent side. In 2025, 93.5% of enterprises planned to increase AI and analytics spending — with 53.2% expecting a significant increase. The dominant stated rationale was operational efficiency: 98.4% of respondents cited cost reduction and process optimization as very or somewhat important AI goals, and 100% rated better decision-making through insights as similarly important. Read alongside the 2026 data, this creates a specific tension: organizations are increasing spending on AI to deliver cost savings and decision quality improvements, while lacking the measurement infrastructure to confirm whether those outcomes are materializing.

The 2025 survey also flagged data availability and quality (90%) and demonstrating business

value (79%) as top-five barriers to AI and analytics deployment — consistent with the 2026 findings and confirming that these challenges are not new. They have also not been resolved between survey cycles. The persistence of data quality and ROI measurement as the two leading barriers across both years suggests these are structural conditions of Indian enterprise AI, not transitional friction points that will resolve with greater investment.

The barrier profile in this section reveals a circular dynamic. Data quality deficits impair model performance; impaired model performance produces weaker business outcomes; weaker business outcomes make ROI difficult to demonstrate; the absence of demonstrated ROI weakens the investment case for addressing data quality. The barriers are a feedback loop, not a list, and the loop will not break without a deliberate intervention at a specific point.

What changes in 2026 is the urgency of resolution. At the current stage of enterprise AI development, cost overruns and opaque ROI were defensible as the price of experimentation. As organizations move toward scaling, boards will expect the financial management discipline applied to AI to match that applied to any other significant capital program. The organizations that build measurement infrastructure now — imperfect, incomplete measurement, but consistent measurement — will be best positioned to sustain investment authorization when the scrutiny increases.

Bottom Line: The economics of enterprise AI are opaque at most organizations. Costs are untracked, ROI is unmeasured, and the barriers to scale are persistent precisely because they are structural rather than incidental. The risk is that investment continues on the basis of peer pressure and strategic narrative rather than evidence — until a board asks a question nobody can answer.


2.1

The Cost Surprise

Forty-five per cent reported no significant cost overruns — but the remaining 55% either experienced overruns, or had no visibility into whether overruns occurred. Twenty-two per cent experienced overruns from both inference and training costs. Twelve per cent attributed overruns primarily to inference costs, and 9% to training costs. Thirteen per cent reported that AI costs are not tracked separately at all.

Organizations that cannot distinguish AI expenditure from broader IT spend cannot assess ROI, make rational decisions about scaling specific use cases, or present a credible AI business case to the board.

AI Cost Overrun Experience



More than half of organizations either experienced AI cost overruns or lack sufficient visibility into AI spending.

CIO Takeaway: If your organization does not track AI costs separately from the broader IT budget, fix that before addressing any other AI management problem. You cannot make rational decisions about scaling, deprecating, or comparing use cases without discrete cost visibility. You cannot defend an increasing AI budget to the CFO without it. Separate AI cost tracking — by use case, by inference volume, by model — is not a nice-to-have. It is the minimum financial management standard for a program of this size.


2.2

The AI ROI Reality

Only 12% of respondents report significant measurable ROI from their AI investments. A further 13% report early signs of ROI, and 19% describe moderate ROI in selected functions. But 17% are still tracking ROI without conclusions, 14% report that AI investments have not delivered value yet, and 26% say ROI remains unclear.

Taken together, 57% of the respondent population — those tracking ROI without conclusions (17%), those for whom AI investments have not yet delivered value (14%), and those for whom ROI remains unclear (26%) — either have no measurable ROI or cannot determine whether it exists. As boards grow more sophisticated in their AI questions, technology leaders will face increasing pressure to move from ROI aspiration to ROI measurement.

ROI from AI Investments



Demonstrating measurable AI returns remains a major challenge, with only a small minority reporting significant ROI.

CIO Takeaway: The 57% who cannot demonstrate ROI are predominantly facing a measurement problem, not a performance problem. The models are likely producing outputs; the question is whether anyone defined success metrics before deployment. For every production AI use case you currently have, ask: what metric was defined as the ROI indicator before we went live, and what does that metric show now? If the answer is that no metric was defined, the ROI gap is a process failure at the point of deployment approval, not evidence that the use case is not working.


2.3

The AI Scaling Barriers

Data quality and readiness tops the list, cited by 56% of respondents. Lack of measurable ROI is the second barrier at 52%, creating a circular dynamic: organizations cannot justify the investment needed to address data quality, because they cannot demonstrate sufficient ROI from current AI deployments to secure that investment. Governance and compliance concerns (44%), talent shortages (44%), integration complexity (43%), and legacy infrastructure (40%) complete the barrier profile.

The simultaneous prominence of data quality and ROI measurement as the top two barriers is significant — they are not independent problems. Data quality issues directly impair model performance, which impairs business outcomes, which impairs ROI measurement.

Barriers to Scaling AI



Data quality and unclear ROI remain the biggest obstacles preventing AI from scaling across enterprises.

CIO Takeaway: Data quality (56%) and lack of measurable ROI (52%) are circular barriers — each makes the other worse. The way out of the circle is to pick one use case with relatively clean data, define the ROI metric tightly before deployment, measure it rigorously for 90 days, and use the result — whatever it is — to build your measurement methodology. One credible data point breaks the cycle more effectively than a data quality program, or an ROI framework, built in the abstract.

[illegible]



Workforce and Talent Readiness





Executive Summary

Enterprise AI does not fail because the models are wrong. It fails because people are not prepared to use them, maintain them, or extract business value from them at scale. The workforce dimension of AI strategy — talent acquisition, upskilling, and organizational literacy — consistently appears as an afterthought in technology investment planning, and consistently appears as a root cause in deployment post-mortems. Survey 1 of the State of AI in Indian Enterprises 2026 study maps the current state of workforce and talent readiness across 101 enterprise technology leaders.

The headline from the talent strategy data is broadly positive: 78% of respondents are using some combination of external hiring and internal upskilling, or are relying primarily on upskilling — the strategically appropriate posture for a period when the required AI skills are still being defined and the competitive market for specialist talent remains tight. The concern is not in the hiring strategy; it is in the literacy gap. Fifteen per cent of respondents have no formal AI training initiatives at all — the same cohort that cites productivity improvement as its primary AI investment driver. The productivity case for AI requires employees who can use AI tools effectively. Organizations investing in tools without investing in the people who will determine whether those tools are used are funding potential rather than delivery.

The 2025 CIO&Leader State of Enterprise Technology survey makes the cultural and literacy dimension of workforce readiness explicit. That study found that 100% of respondents rated AI and data literacy as very or somewhat important to their organization's AI success — the only survey item across the entire 2025 study to achieve universal agreement. Executive leadership support ranked second at 96.8%, followed by user education and change management at 96.7%. The finding that every enterprise technology leader in the 2025 cohort considers AI literacy non-

negotiable, set against the 2026 finding that 15% of organizations have made no formal provision for it, defines the implementation gap precisely.

The 2025 data also recorded that change management was cited as a major concern by 86.7% of respondents, specifically in the context of AI disrupting workflows and decision hierarchies. The 2026 workforce data illuminates this from a different angle. Organizations without formal AI training are not keeping their employees away from AI tools; they are redirecting them to unsanctioned channels. Section 6 of this report finds that 48% of respondents estimate more than 10% of their workforce uses shadow AI tools. The correlation between low training investment and high shadow AI prevalence is not coincidental: employees will seek capability whether or not the enterprise provides it.

Two structural realities define this section's findings. First, talent strategies that rely primarily on external partners and vendors (14% of respondents) are viable at the Adopter stage but carry accumulating vendor dependency risk as organizations scale and as AI capabilities become operationally embedded. Second, the 7% of organizations reporting no formal talent strategy for AI are maintaining a position that is inconsistent with any serious ambition to move AI from departmental deployment to enterprise-wide operating model change. Workforce preparation and technology investment must grow at the same rate. Currently, for a significant minority of Indian enterprises, they do not.


Bottom Line: Most AI deployment failures are adoption failures, not model failures. The workforce readiness gap is structurally underfunded relative to the technology investment — and the organizations with no formal training are not protected by that absence. Their employees are using AI regardless; they are just doing it ungoverned.

3.1

AI Talent Strategy

Forty-five per cent describe their organization's AI talent strategy as a combination of external hiring and internal upskilling. A further 33% focus primarily on upskilling existing employees. Fourteen per cent rely primarily on external partners and vendors, which is viable at the Adopter stage but carries vendor dependency risk at scale. Only two per cent are pursuing aggressive external hiring. Seven per cent have no formal talent strategy for AI — a figure disproportionate for organizations that simultaneously report AI as a strategic investment priority.

AI Talent Strategy



Most organizations are combining external hiring with internal upskilling to build sustainable AI capabilities.

CIO Takeaway: A combination of hiring and upskilling is the right answer, but combination strategies fail without explicit allocation decisions. Which capabilities are you building internally and which are you sourcing externally? If you cannot name specific capability categories — say, prompt engineering and fine-tuning built internally, foundation model training sourced externally — your talent strategy is a sentiment rather than a plan. Build the skills taxonomy first; the hiring and training decisions follow from it.


3.2

AI Training & Literacy

Role-based training programs are in place at 37% of organizations. A further 20% are planning to implement training. Mandatory all-employee training exists at 19%, and voluntary learning programs at 10%. Fifteen per cent report no formal AI training initiatives — the same cohort that identifies productivity improvement as the primary AI investment driver.

Deploying AI tools without systematic workforce preparation creates two risks: underutilization by employees who lack confidence to use them effectively, and shadow AI usage as employees seek their own solutions outside sanctioned channels (see Section 6).

Enterprise AI Training & Literacy



Formal AI literacy programs are expanding, but a significant minority still lacks structured workforce training.

CIO Takeaway: The 15% with no formal AI training are not making a conservative choice — they are making an expensive one. Their employees are using consumer AI tools to do work-related tasks regardless, as the shadow AI data in Section 6 demonstrates. The practical question is not whether to train; it is what to train for. The most immediately effective programs focus on two things: what the organization’s sanctioned AI tools can and cannot do, and what employees must never send to an AI system. Everything else can come later.



GenAI — From Pilot to Production





Executive Summary

Generative AI has moved from strategic aspiration to operational management challenge. The question facing Indian enterprise technology leaders is no longer whether to deploy GenAI, but how to make GenAI deployments reliable, measurable, and compliant at enterprise scale. Survey 2 of the State of AI in Indian Enterprises 2026 study, completed by 88 enterprise technology leaders, captures the production reality of this transition — and the picture it reveals is one of uneven progress, persistent structural barriers, and a pilot-to-production conversion problem that is more structural than technical.

The deployment distribution in this section frames everything that follows. Thirty-six per cent of respondents are at function-level deployment; 23% report enterprise-wide adoption across multiple functions. But 22% remain at the stage of employee experimentation without structured deployment, and only 6% have embedded GenAI into core enterprise workflows. The distance between unstructured experimentation and workflow-embedded production is the operational gap that defines the current phase of Indian enterprise GenAI.

The 2025 CIO&Leader State of Enterprise Technology survey provides useful context for the function-level deployment picture. That study found IT operations leading AI adoption maturity across business functions, with 42% of respondents reporting broad deployment — followed by finance and accounts (31% broad deployment) and customer service and engagement (28%). Content creation, driven by GenAI and natural language processing tools, was registering significant uptake in marketing, communications, and CX functions. The 2026 data on measurable GenAI value tracks this pattern

coherently: analytics and reporting leads at 23% of respondents, followed by customer support automation (20%) and code generation (19%). The progression from 2025 function-level deployment depth to 2026 value realization reflects a consistent accumulation of productive use rather than a sharp break from the prior year.

The 2025 survey also identified data security and privacy (91.7% of respondents) and change management (86.7%) as the dominant barriers to AI deployment. Both manifest directly in the 2026 pilot failure analysis: data quality issues are the leading cause of pilot failure at 42%, and lack of business ownership — the organizational change management failure that prevents operationalization even when technical results are positive — accounts for 31%.

The central finding — 52% of respondents report sub-25% pilot-to-production conversion rates — points to governance rather than technology. Data quality, unclear business cases, and weak business ownership remain the biggest barriers. Organizations with higher conversion rates succeed through stronger governance, clear success metrics, and business accountability from the outset. Meanwhile, 57% rely on commercial frontier models, creating vendor concentration that will require active management as GenAI becomes mission-critical infrastructure.

Bottom Line: The GenAI phase is at its most operationally demanding point. Function-level deployment is not a destination — it is where the production discipline problems become visible.


The 52% pilot conversion rate is the metric that most honestly describes the state of enterprise GenAI in India: not failing, but not yet systematic.

4.1

Enterprise Use of GenAI

Function-level deployments — GenAI in use within specific business units — is the most common state, cited by 36% of respondents. Enterprise-wide adoption across multiple functions follows at 23%. Employee experimentation without structured deployment accounts for 22%, and pilot deployments for 14%. Only 6% have embedded GenAI into core enterprise workflows. The distribution suggests that the early wave of GenAI adoption has moved beyond purely experimental usage, but has not yet achieved the breadth and depth that the enterprise-wide productivity narrative would imply.

Adoption of GenAI



Generative AI adoption is progressing steadily, with most deployments concentrated within individual business functions rather than enterprise-wide workflows.

CIO Takeaway: If you are at function-level deployment, the priority question is not which new functions to add. It is whether the functions you have deployed are genuinely embedded in daily workflows, or sitting as optional tools that enthusiastic users demonstrate and everyone else ignores. Embedded means the default process uses the AI output; optional means the AI output is available if someone chooses to use it. Count how many of your current deployments are truly embedded. That number is your real GenAI maturity indicator.


4.2

Use Cases Delivering Value

Analytics and reporting use cases lead at 23%, followed closely by customer support automation (20%) and code generation (19%). Knowledge management and enterprise search comes fourth at 17%, and employee productivity copilots fifth at 11%.

The prominence of analytics and reporting reflects GenAI's early strength in text-to-insight tasks. Customer support automation reflects the heavy investment in AI-assisted contact centers across BFSI, telecom, and retail. Code generation's placement at third reflects both the density of IT Services respondents in the Survey 2 sample (21%) and the near-universal adoption of AI coding assistants among development teams.

GenAI Use Case Delivering Value



Analytics, customer support and software development are delivering the strongest measurable value from enterprise GenAI initiatives.

CIO Takeaway: Analytics and reporting, customer support automation, and code generation lead on measurable value — and these are the three use cases with the clearest output metrics by design. If your highest-value use cases are not generating measurable returns, examine the measurement approach before questioning the use case.

The use cases delivering value in this data are not technically superior; they are easier to measure. Apply the same measurement rigor to your own portfolio before concluding that a use case is underperforming.


4.3

The Pilot-to-Production Gap

Fifty-two per cent reported that fewer than one in four pilots makes it to production. Fifteen per cent reported a 25 to 50% conversion rate. Eleven per cent said pilot outcomes are not formally tracked. Only 8% report that more than 75% of pilots reach production.

A sub-25% conversion rate means that for every four AI pilots initiated, three are quietly abandoned — their infrastructure costs written off, and the organizational learning they generated undercaptured.

Pilot-to-Production Conversion Rate



Converting pilots into production remains enterprise AI's biggest execution challenge, with most organizations reporting very low success rates.

CIO Takeaway: When a pilot fails to reach production, ask one question before any post-mortem: did we have a named business owner committed to operating this capability in production before the pilot began? In most cases the answer is no — the pilot was run by IT or the AI team with business stakeholder observation, rather than business stakeholder ownership. The production gap is largely an accountability gap. Close it at the intake stage, not the post-mortem stage.


4.4

Why Pilots Fail to Scale?

Data quality issues, cited by 42%, connect directly to the barrier identified in Section 2 as the primary obstacle to scaling. The persistence of data quality as both the scaling barrier and the pilot failure cause, across independent survey populations, confirms that this is a systemic infrastructure deficit rather than a project management problem.

Unclear ROI or business case is the second cause at 35%, followed by integration complexity (32%), lack of business ownership (31%), regulatory and compliance concerns (24%), and cost overruns (20%). The 31% attributing failure to lack of business ownership is a structural insight: pilots without an engaged business sponsor who will operationalize the output face abandonment risk even when the technical results are positive.

GenAI Use Case Delivering Value



Poor data quality, weak business cases and lack of business ownership are the leading reasons AI pilots fail to reach production.

CIO Takeaway: Data quality (42%) and lack of business ownership (31%) are both solvable before the pilot begins. A two-condition pilot intake gate — a named business owner with operating accountability and a data readiness assessment against the specific use case — would eliminate the majority of failure causes before resources are committed. This is not a technology intervention. It is a governance intervention at the program management level, and it is available to you now.


4.5

AI Model Strategy

Forty-nine per cent use a hybrid approach combining multiple strategies. Commercial APIs are used by 35% of organizations. Fine-tuned open-source models account for 25%, proprietary models for 24%, and pre-built SaaS AI solutions for 20%. In a separate question on model types, 57% use commercial frontier models (GPT-4, Claude, Gemini) as part of their model portfolio.

Most organizations are not building foundational models; they are combining commercial frontier capabilities with cloud-provider hosted models and, in more sophisticated cases, fine-tuned open-source variants for domain-specific tasks. The 24% pursuing proprietary model development represents primarily large BFSI, IT Services, and manufacturing organizations with the data, infrastructure, and ML engineering capacity to justify the investment.


AI Model Strategy



Enterprises are increasingly adopting hybrid model strategies that combine commercial, open-source and proprietary AI models.


CIO Takeaway: Fifty-seven per cent are using commercial frontier models as part of their portfolio. That is an efficient choice and a concentration risk. A small number of international vendors control the capabilities you are building operational dependencies on. Their pricing, their terms of service, and their model deprecation decisions can materially affect your AI operations with limited notice. Assess your top five AI vendor dependencies against three criteria: switching cost, data exposure, and operational criticality. Do this now, while the dependency is manageable, not after a pricing change makes it urgent.

[illegible]



Agentic AI — The Emerging Frontier





Executive Summary

Agentic AI represents a categorical shift — not merely an incremental advance — in the nature of enterprise AI deployment. Where generative AI tools produce outputs that humans then act upon, agentic systems act themselves: decomposing complex goals into multi-step plans, executing those plans across enterprise tools and data sources, and adapting their approach based on intermediate outcomes. The governance implications, the control requirements, and the operating model changes that agentic AI demands are architecturally different from anything that generative AI alone has required.

Survey 2 of the State of AI in Indian Enterprises 2026 study captures the enterprise agentic AI landscape as of mid-2026 — and the picture it presents is ahead of most external market assessments. Sixty-four per cent of the GenAI survey cohort has moved past pure evaluation into active piloting or limited production deployment of agentic systems. Seventeen per cent have multiple AI agents in production. These are not aspirational figures; they represent real deployment exposure, real failure mode discovery, and real operational learning accumulating at a distinct pace among the more AI-mature organizations in this study.

The 2025 CIO&Leader State of Enterprise Technology survey provides a point of comparison on the infrastructure and deployment maturity underpinning this shift. That study found AI workload orchestration tools as the top AI infrastructure investment priority — ahead of AI-specific MLOps pipelines, cloud-native AI stacks, and on-premise high-performance compute. Sixty-six per cent of 2025 respondents expected AI to optimize cloud operations within the following twelve to eighteen months. The infrastructure readiness required to support AI workload orchestration is also the foundational requirement for running multi-step

agentic workflows across enterprise systems — the two investment agendas are directly interdependent, and the 2025 infrastructure build is what enables the 2026 agentic deployment picture this section documents.

The function-level deployment picture in this section follows a risk-adjusted logic that reflects the operational judgment of the CIOs driving these decisions. IT operations is the dominant target function at 66%, followed by customer service (50%) and cybersecurity operations (48%). The 2025 survey found IT operations leading AI adoption maturity across all business functions, and cybersecurity operations registering AI usage among 78% of respondents. The 2026 agentic deployment pattern extends this established foundation: agents are being directed first at the domains where IT has operational control, where the ROI case is already validated, and where error recovery is available within the same team responsible for the deployment.

The autonomy threshold data is the most forward-looking dimension of this section. The 62% who currently require human approval or active human oversight of all agent actions are adopting the appropriate governance posture for the current state of enterprise agentic maturity—but it is a transitional posture, not a permanent one.


Bottom Line: Sixty-four per cent of the GenAI survey cohort is already in active piloting or production deployment of agentic AI. If your organization is still at the evaluation stage, the gap is not catastrophic — but the learning that comes from early production deployment is accumulating at peers who started earlier. The governance frameworks you need for agents are architecturally different from those for GenAI tools, and most organizations do not have them yet.

5.1

Agentic AI Deployment

Twenty-seven per cent of respondents are evaluating use cases. Another 27% have limited production deployments in place. Twenty per cent are piloting with human oversight. Seventeen per cent have multiple AI agents in production — a figure that is higher than most market assessments of enterprise agentic readiness would predict. Only 8% report no current plans. Taken together, 64% of the Survey 2 population has moved past pure evaluation into active piloting or production deployment of agentic systems.

Agentic AI Deployment Status



Agentic AI has moved beyond evaluation, with a majority of organizations already piloting or deploying autonomous AI agents.


CIO Takeaway: The 17% with multiple agents in production are not ahead on vision — they are ahead on failure mode discovery. They know what breaks, what control assumptions do not hold, and what the organizational change requirements are for human-agent workflow integration. If you are at the evaluation stage, the highest-value action is not to accelerate your own deployment but to find structured ways to access the learning from the organizations ahead of you — through peer networks, vendor case studies, or consortium benchmarking. Deploy informed, not first.

5.2

Enterprise Functions Targeted for Agentic AI

IT operations is the dominant target function, cited by 66% of respondents — reflecting both the high volume of repetitive tasks in IT service management and the risk appetite that comes from operating within a domain that IT leaders control. Customer service follows at 50%. Cybersecurity at 48% reflects the growing interest in autonomous threat detection and response. HR and procurement (40%), knowledge management (39%), compliance and legal operations (33%), software engineering (33%), and finance operations (32%) complete the function profile.

The distribution follows a risk-adjusted logic: agents are being directed toward functions where IT has operational ownership, where the ROI case is well-established, and where the downside of an agent error is recoverable.




CIO Takeaway: IT operations at 66% is the rational starting domain — the CIO controls it, the processes are well-defined, and errors are recoverable within the same team that catches them. Start there. The functions where early agentic deployment creates material risk are finance, compliance, and legal — where errors have regulatory or financial consequence and where the audit trail requirements are stringent. If your first agentic deployment is in one of those functions, your control and oversight design needs to be proportionally more rigorous, not proportionally more aspirational.

5.3

Autonomy Thresholds for AI Agents

Thirty-two per cent require human approval for all agent actions. Thirty per cent are comfortable with semi-autonomous execution with human oversight. Twenty-six per cent restrict agents to a recommendation-only role. Ten per cent are comfortable with full autonomy for low-risk tasks. Only two per cent are comfortable with fully autonomous operations. The 62% who require either human approval or active human oversight reflect a governance posture that is appropriate given the current maturity of enterprise agentic systems.

The gap between stated comfort level and actual deployment practice may be narrower than this question suggests, since agents operating within a tight operational scope can function autonomously in practice while technically operating under a human-approval policy.



CIO Takeaway: The 32% requiring human approval for every agent action is the right posture today, but it is a transitional posture, not a permanent one. Design your autonomy escalation path before you need it: identify which action categories have consistently recoverable errors, build the audit trail for those categories, and define the evidence threshold for moving them to semi-autonomous execution. Organizations that do not design this path in advance will either stay at human-approval forever — which eliminates most of the productivity case for agents — or escalate autonomy reactively after an incident forces the conversation.

[illegible]



AI Governance and Risk Management





Executive Summary

AI governance is the discipline that separates organizations building sustainable AI programs from those accumulating unmonitored exposure. It is also, by the evidence of Survey 3 of the State of AI in Indian Enterprises 2026 study, the dimension of enterprise AI that most organizations know requires improvement and have found hardest to operationalize. The 102 enterprise technology leaders who completed Survey 3 reveal a governance landscape where formal frameworks are being constructed, where shadow AI is more prevalent than most organizations are measuring, where risk management capability remains in early development, and where universal concern about data privacy has not produced proportionate action.

The shadow AI data establishes the context. Forty-eight per cent of respondents estimate that more than 10% of their workforce uses unsanctioned AI tools; a further 7% cannot estimate shadow AI usage at all, meaning their governance frameworks are operating on an incomplete picture of what they are attempting to govern. The 2025 CIO&Leader State of Enterprise Technology survey found that 93% of respondents were concerned about AI misuse in cybersecurity specifically — a figure consistent with the near-universal data privacy concern (81%) in the 2026 governance data. The through-line between the two surveys is that concern about AI misuse and unauthorized usage is not new; what remains inadequate is the organizational response to it.

The 2025 survey also recorded that regulatory and ethical risks were cited by 83.6% of respondents as significant concerns in AI deployment — tracking closely with the 81% data privacy concern, 45% lack-of-governance concern, and 40% regulatory uncertainty in the 2026 risk priority map. The 2026 data adds a notable dimension: hallucinations and inaccurate model outputs are cited by 46% of

the governance cohort, signaling that reliability is increasingly being treated as a risk management issue rather than a technical quality problem to be handled by model teams alone.

Governance framework maturity reveals significant variation beneath the aggregate. Twenty-nine per cent of organizations report a mature enterprise-wide AI governance framework — a higher proportion than informal assessments of Indian enterprise governance readiness would suggest, likely reflecting the concentration of BFSI and IT Services respondents in Survey 3. But 43% operate with limited policies or informal guidelines, and 9% have no formal governance structure at all. An AI governance framework that does not operationally cover live model inventory, data classification, vendor due diligence, usage monitoring, and AI-specific incident response is not a governance framework in any meaningful operational sense. The gap between framework existence and framework operational effectiveness is where most Indian enterprises currently sit.

The 19% with no formal AI risk framework are not protected by the absence of detected incidents. They are unmonitored. The distinction matters more as AI deployments scale, because incidents that surface in unmonitored environments tend to be larger by the time they are detected than those that governance-mature organizations identify early and contain systematically.

Bottom Line: Governance is the dimension that separates organizations that will scale AI sustainably from those that will face an incident that forces reactive remediation. The 19% with no formal AI risk framework are not protected by the absence of detected incidents. They are unmonitored. The distinction matters.


6.1

Shadow AI: The Unsanctioned Reality

Forty-five per cent estimate that 0-10% of their workforce uses unsanctioned AI tools. Twenty-four per cent estimate 11-25%. Thirteen per cent estimate 26-50%. Seven per cent estimate more than 75% of their workforce is using shadow AI. A further 7% say they cannot estimate. The 7% who cannot estimate is as informative as the distribution itself: organizations that lack visibility into shadow AI usage are unable to govern it.

The broader point is that the 0-10% estimates are likely underestimates the problem — in the absence of active monitoring, shadow AI estimates will systematically under-count actual usage.

Employees Using Shadow AI




Shadow AI has become widespread, highlighting the growing gap between employee AI adoption and enterprise governance.

CIO Takeaway: If your estimate of shadow AI usage is below 10%, it is almost certainly a product of measurement gaps rather than workforce restraint. Ask your CISO one specific question: what percentage of browser-based generative AI usage — ChatGPT, Gemini, Copilot, Claude — are our endpoint monitoring tools currently capturing? The answer determines whether your shadow AI estimate is a measurement or a guess. For most organizations, the answer will prompt an immediate review of monitoring scope.

6.2

AI Governance Framework

Twenty-nine per cent of organizations have a mature enterprise-wide AI governance framework — a higher figure than many assessments of Indian enterprise governance readiness would predict, likely reflecting the concentration of BFSI and IT Services respondents in the sample. A further 20% have governance under development. But 31% have only limited governance policies, 12% rely on informal guidelines, and 9% have no formal governance structure at all.




CIO Takeaway: A mature AI governance framework is not an ethics statement or an AI policy document. It is an operational control system that covers five things: a live model and tool inventory, data classification standards applied to AI inputs and outputs, a vendor due diligence process for AI suppliers, a usage monitoring capability, and an incident response procedure specific to AI failures. If your governance framework does not address all five in operational rather than aspirational terms, it will not hold under an audit or an incident.

6.3

Concerns About Enterprise AI

Data privacy and compliance is the near-universal concern, cited by 81% of respondents. Security risks follow at 73%. IP and data leakage is cited by 55%, hallucinations and inaccurate outputs by 46%, lack of governance by 45%, and regulatory uncertainty by 40%. The high prominence of hallucinations at 46% is worth noting — this is a reliability concern, not a governance one. The gap between the 81% who are concerned about data privacy and the 19% who are highly prepared for DPDP compliance (Section 7.2) is the most pointed expression of the governance lag in this report.




CIO Takeaway: The 81% who cite data privacy as their primary concern are identifying the right risk — but the highest-exposure channel is frequently not the one they are monitoring. Most organizations have stronger controls on structured personal data in their CRM or core banking systems than on the unstructured personal data flowing into GenAI tools through employee prompts. Map your actual data exposure through AI tool usage, not your theoretical exposure through data classification policy.

6.4

Technical Controls in Place

Twenty-seven per cent report full AI monitoring combined with data loss prevention and policy enforcement. Almost one-quarter report partial monitoring and controls. Twenty-three per cent rely on employee awareness policies only. Nineteen per cent have limited visibility into AI usage. Only six per cent have a dedicated AI incident response plan, and four per cent report no formal controls. The 4% with no formal controls is directly correlated with the shadow AI challenge: without detection capability, shadow AI management is impossible.

Technical Controls for Enterprise AI




Organizations are strengthening AI guardrails through monitoring, access controls and governance mechanisms, though maturity varies widely.

CIO Takeaway: The 23% with partial monitoring represent a more dangerous position than the 4% with no controls, because partial monitoring creates institutional confidence that is not fully warranted. Before your next governance review, map specifically: which AI tools are covered by your monitoring, which user groups are visible, and which channels have no coverage at all. Present the gaps as a risk map, not an aspiration list. Gaps that are named and owned are manageable; gaps that are assumed not to exist are not.

6.5

AI Risk Management Maturity

Forty-two per cent describe their AI risk management as early-stage development. Eighteen percent have defined and operational risk frameworks. A further 19% have advanced frameworks with continuous monitoring. But 19% report no formal AI risk framework at all, and 3% describe their approach as ad hoc. The 19% with no formal AI risk framework represents a material exposure gap — organizations whose AI deployments are not subject to systematic risk assessment, and whose boards may lack the information needed to form an accurate view of AI-related risk.



CIO Takeaway: If you are among the 42% at the early-stage development level, the fastest path to a functional AI risk management capability is extension rather than construction. Your enterprise risk framework already has assessment categories, escalation paths, review cadences, and board reporting structures. Map AI system risk onto that existing framework as a new risk category rather than building a parallel AI risk system from scratch. The architecture is there. The vocabulary and the specific AI failure modes need to be added.

This image shows a full page of white paper with horizontal dotted lines. The lines are evenly spaced and run across the width of the page, providing a guide for handwriting practice. There are no margins, text, or other markings on the page.





Executive Summary

The regulatory environment for AI in India is no longer in formation — it is in force. The Digital Personal Data Protection Act creates present obligations for organizations that process personal data through AI systems, and sector-specific guidance from RBI, SEBI, and CERT-In layers additional compliance requirements on BFSI, capital markets, and critical infrastructure operators. Survey 3 of the State of AI in Indian Enterprises 2026 study documents how prepared Indian enterprises are for this environment. The answer, for the majority, is inadequately.

The DPDP readiness finding is the most pointed single statistic in this report. Eighty-one per cent of Survey 3 respondents identify data privacy and compliance as their primary AI concern. Nineteen per cent describe themselves as highly prepared for DPDP-related AI compliance obligations. The 62-percentage-point gap between concern and preparation is the largest such divergence in any dimension of this study, and it reflects a structural challenge: organizations recognize the compliance obligation in principle but have not yet translated that recognition into operational controls.

The 2025 CIO&Leader State of Enterprise Technology survey is relevant here on two dimensions. First, data security and privacy (91.7%) was the top-ranked barrier to AI deployment in 2025 — the only barrier to cross the 90% threshold and the only one rated as a major concern by a clear majority of respondents. Twelve months later, that concern has not produced a proportionate step-change in compliance readiness. Second, the 2025 survey identified GenAI-enabled phishing and deepfake-powered social engineering as emerging AI-specific threat vectors actively discussed by enterprise CISOs — consistent with the 2026 finding that 10% of organizations have confirmed AI-related security incidents in production,

including prompt injection attacks and model manipulation.

The security incident data in this section requires careful reading. Fifty-seven per cent report no incidents with active monitoring in place — a confident and reassuring position. But 26% report no incidents with limited monitoring, meaning that the absence of detected incidents reflects limited visibility rather than confirmed security. The 22% with limited monitoring and the 9% with confirmed incidents together mean that approximately two in five organizations in this survey cohort has a known or likely AI security exposure — against a backdrop where only 9% have a dedicated AI incident response plan in place.

The governance challenge data identifies third-party AI vendor risk at 49% — the second most prominent challenge after consent and DPDP compliance (60%). This finding reflects a structural feature of enterprise AI that is systematically underweighted in risk frameworks: most enterprise AI deployments route personal data through commercial APIs, cloud-provider models, and vendor-embedded SaaS features. The data handling practices of those third parties are governed by the policies of those vendors, which may not align with DPDP requirements. Third-party AI risk is, in operational terms, supply chain risk for the data layer, and it warrants the same structured management discipline applied to other forms of third-party operational dependency.


Bottom Line: The compliance posture of most Indian enterprises on AI sits below what DPDP and sector regulators now require. The 19% highly prepared figure is the most actionable number in this report for most CIOs — not because 81% are non-compliant, but because most of the journey from moderate preparation to high preparation is a matter of documented process, not technology investment.

7.1

AI-Related Security Incidents

Forty-nine per cent reported no incidents with active AI security monitoring. Twenty-two per cent reported no incidents with limited monitoring, meaning that the absence of observed incidents reflects limited visibility rather than confirmed security. Thirteen per cent reported suspected but unconfirmed incidents. Nine per cent reported confirmed incidents in production — including data breaches, model manipulation, or prompt injection. Eight per cent indicated that AI security is not yet a priority. The 22% with no incidents and limited monitoring are the most analytically ambiguous segment — unobserved, not necessarily secure.

AI-related Security Incidents



Confirmed AI-related security incidents remain relatively limited, but suspected incidents indicate growing enterprise exposure.


CIO Takeaway: Nine per cent confirmed incidents and 13% suspected unconfirmed means that 22% organizations in this cohort have a known or likely AI security event in the last period — and the 22% with limited monitoring have an unknown status that should not be reported to the board as clean. The specific threat vectors most commonly missed in current enterprise AI monitoring are prompt injection in customer-facing AI applications, data exfiltration through AI-assisted coding tools, and model manipulation through adversarial inputs. If your security incident reporting does not include these categories, your incident picture is incomplete.

7.2

DPDP Act Readiness

Forty per cent describe themselves as moderately prepared. Twenty-four per cent are minimally prepared. Nineteen per cent are highly prepared. Thirteen per cent are still assessing requirements. Four per cent are not prepared. The finding that only 19% describe themselves as highly prepared — against a backdrop where 81% of the same survey population cited data privacy and compliance as their primary AI concern — is the sharpest expression of the concern-action gap in this report. The DPDP Act is not a future risk; it is a present obligation.

DPDP Compliance Readiness



Despite high awareness, only a small proportion of organizations consider themselves fully prepared for DPDP-related AI compliance.


CIO Takeaway: The practical DPDP-AI compliance gap for most organizations at the moderate preparation level closes to three specific workstreams: documenting consent mechanisms for personal data processed by AI systems, building data principal rights workflows for AI-generated outputs that involve personal data, and creating a register of third-party AI data processors with contractual DPDP alignment. None of these requires new technology. All of them require legal, privacy, and technology teams to work from a shared definition of what constitutes AI-processed personal data in your specific operational context. That definitional conversation is frequently the bottleneck.

7.3

AI Governance Challenges

Consent and DPDP compliance tops the list at 60%, followed by third-party AI vendor risks (49%), AI explainability (48%), monitoring and auditability (45%), shadow AI usage (41%), cross-border data flows (38%), and regulatory uncertainty (35%). The prominence of third-party AI vendor risks at 49% reflects a structural feature of enterprise AI that is often underweighted: most enterprise AI deployments involve commercial APIs, cloud-provider models, vendor-embedded features, and SaaS solutions. Data processed by these external systems is subject to the governance and security policies of those vendors, which may not align with DPDP requirements. Third-party AI risk is, in effect, supply chain risk for the data layer.


AI Governance Challenges



Third-party AI risk, data governance and regulatory compliance have emerged as the most pressing AI governance challenges.

CIO Takeaway: Third-party AI vendor risk at 49% is the governance challenge that most organizations are managing least systematically. Your vendor risk management process almost certainly covers third-party software and services. It probably does not adequately cover the personal data your AI vendors process on your behalf, the sub processors those vendors use, or your contractual rights in the event of an AI-related incident involving your data. Extend your existing vendor risk framework to explicitly cover AI suppliers, starting with your top five by data exposure, not by contract value.

[illegible]



Future AI Outlook





Executive Summary

The most reliable forward signal in any research program is convergence. When independent respondent populations — answering different questions from different operational vantage points — arrive at the same conclusion about the future, the finding carries more weight than any single survey can generate. The cross-survey agentic AI consensus in this report is precisely that kind of convergent signal: 38% of the GenAI deployment cohort and 46% of the governance and risk cohort independently identify agentic AI as the emerging capability most likely to reshape enterprise operations over the next eighteen months. The practitioners focused on deployment and the leaders focused on risk are looking at the same future.

Survey 2 and Survey 3 of the State of AI in Indian Enterprises 2026 study both posed this forward-looking question, and this section synthesizes both responses alongside the Survey 3 data on how enterprise technology leaders envision the future enterprise operating model. The synthesis reveals a future that the majority of respondents describe as hybrid and agentic — not fully automated, not human-supervised throughout, but a deliberate combination of human judgment and machine execution with explicitly designed boundaries between the two.

The 2025 CIO&Leader State of Enterprise Technology survey provides a useful baseline for reading these forward projections. Seventy-five per cent of 2025 respondents expected GenAI for workload optimization to be enterprise-ready within two years — a timeline that places the expected maturity point squarely in the 2026–2027 window this report's respondents are looking toward. Ninety-three per cent planned to increase AI and analytics spending. And the 2025 survey described AI as having crossed from "niche capability" to the new enterprise

operating layer — running in production across IT monitoring, financial forecasting, cybersecurity response, and content generation. The 2026 agentic outlook is a continuation of that trajectory, not a departure from it. What is changing is the nature of the capability: from generative outputs that humans act on, to autonomous execution that must be governed before it is deployed.

The operating model projection data in Section 8.2 is instructive in its nuance. Thirty-four per cent expect AI agents to automate many enterprise workflows — but the majority do not expect a binary shift from human to machine execution. Twenty-six per cent envision hybrid human-AI operating models, and 22% expect AI copilots to become standard workplace tools. These are not competing predictions; they are different points on the same autonomy spectrum, and together they describe an enterprise where the default question is no longer whether AI is involved in a given process, but how much autonomy it holds, on which decisions, and under what conditions.

The implications that close this section are structurally grounded in the data gaps documented across the preceding seven sections. Data quality, ROI measurement, DPDP compliance, and agentic governance design are the four workstreams that must run in parallel. The forward trajectory is visible. The discipline required to navigate it is what the data consistently reveals as the most consequential variable separating Indian enterprises that will scale AI sustainably from those that will not.

Bottom Line: The convergence of two independent respondent cohorts on agentic AI as the highest-impact upcoming capability is the most reliable forward signal in this report. When the people deploying AI and the people governing AI arrive at the same answer independently, the signal is strong enough to plan from, not just note.


8.1

Agentic AI is the Capability Consensus

Agentic AI leads the emerging capability ranking in both surveys. In the GenAI cohort (Survey 2), 38% of respondents identify agentic AI as the capability with the biggest enterprise impact over the next eighteen months. In the Risk and Governance cohort (Survey 3), the figure is 46%.

Multimodal AI is second in both surveys (17% and 15% respectively). This cross-survey agreement is more significant than either figure in isolation. Respondents to Survey 2 were thinking primarily about deployment — their answer reflects what they expect to be deploying next. Respondents to Survey 3 were thinking primarily about risk — their answer reflects what they are preparing to govern next. When practitioners with different vantage points arrive at the same conclusion, the signal is strong.

Emerging AI Capabilities



Agentic AI has emerged as the technology leaders' strongest consensus on the next major enterprise AI capability.


CIO Takeaway: Thirty-eight and 46 per cent across two independent cohorts identifying agentic AI is not a prediction — it is a current assessment of where the highest expected value lies.

8.2

AI Deployment in the Enterprise

Survey 3 asked respondents which statement best reflects their view of the future enterprise. Thirty-four per cent believe that AI agents will automate many workflows. Twenty-six per cent envision enterprises evolving into hybrid human plus AI operating models. Twenty-two per cent expect AI copilots to become standard workplace tools. Ten per cent believe AI will primarily augment employees without fundamental operating model change. Eight per cent expect that humans will increasingly supervise AI systems. The distribution reveals pragmatic realism about the pace and form of AI-driven transformation.

Vision for Future



Organizations expect AI adoption to accelerate from pilots toward broader enterprise deployment over the next 12–18 months.

CIO Takeaway: The 34% who expect agents to automate many workflows and the 26% who expect hybrid human-AI operating models are describing the same future at different points on the autonomy spectrum. Before your next agentic deployment, define three things explicitly: which decisions agents will make autonomously, which decisions will require a human to approve an agent recommendation, and which decisions will remain entirely human regardless of what an agent suggests. That classification is not a technology architecture decision — it is an operating model decision that your business leaders, not your AI team, need to own.

8.3

Implications for Enterprise Tech Leaders

The data quality imperative is not new, but the survey confirms it has not been resolved. Data quality is simultaneously the top barrier to AI scaling and the leading cause of pilot-to-production failure. Organizations that have not treated data readiness as a foundational AI investment will continue to encounter this barrier regardless of how much they invest in models, platforms, and tooling. This reinforces a simple reality: AI outcomes are ultimately constrained by enterprise data maturity. Investments in governance, metadata, lineage, and data stewardship will increasingly determine which organizations scale AI successfully.

The ROI measurement gap is becoming a governance issue. The 57% of respondents who either have no measurable ROI or cannot determine whether returns exist cannot defer this question indefinitely. Building an AI return measurement framework—even an imperfect one—is a near-term priority for CIOs who want to sustain investment authorization. Boards and business leaders will increasingly expect AI initiatives to demonstrate measurable business outcomes, making standardized KPIs, financial benchmarks, and value-tracking mechanisms essential to future funding decisions.

DPDP compliance is a present obligation, not a future preparation. With only 19% of Survey 3 respondents describing themselves as highly prepared, the majority of organizations deploying AI on personal data are operating below the compliance threshold the Act implies. The most immediate practical requirements—consent mechanisms, data principal rights workflows, and third-party vendor due diligence—are buildable with existing governance infrastructure. Organizations that embed privacy-by-design into AI programs today will also strengthen customer trust and reduce long-term regulatory risk.

Agentic AI requires a different governance model than GenAI. The autonomy thresholds that Survey 2 respondents are drawing—human approval, human oversight, recommendation-only—are appropriate for the current state of enterprise agentic systems. But as agent capabilities improve and deployment scope expands, the governance frameworks designed for GenAI tools will be insufficient. The 64% already in active Agentic AI territory should be building agentic-specific controls now. These controls should include agent identity, decision traceability, execution limits, escalation protocols, and continuous monitoring to ensure autonomous systems remain aligned with business objectives and risk policies.

CIO Takeaway: The four priorities the data identifies—data quality investment, ROI measurement, DPDP compliance, and agentic governance design—are not a sequential roadmap. They must run in parallel because they address different failure modes that will emerge simultaneously. Organizations that integrate these priorities into a single enterprise AI transformation program, backed by executive sponsorship and cross-functional ownership, will be better positioned to scale AI responsibly and consistently. The one dependency that is genuinely sequential: without trusted data, every subsequent AI investment becomes significantly harder to scale and sustain.

Conclusion

Across three independent surveys, the State of AI in Indian Enterprises 2026 resolves into a coherent set of findings that are specific enough to act on and broad enough to apply across the enterprise technology community.

The AI maturity curve remains steep. Two-thirds of respondents describe themselves as Explorers or Adopters — not because AI investment is absent, but because the path from experiment to enterprise-wide scale is longer and more operationally demanding than early AI strategies anticipated. This is not a pessimistic finding. It is a calibration. Organizations that understand the difficulty of the journey are better positioned to navigate it than those operating on the assumption that maturity should have arrived by now.

The pilot-to-production problem is the most actionable finding in this study. A 52% sub-25% conversion rate is not a technology problem. It is a discipline problem — specifically, the discipline of selecting pilots that can be productions, engaging business sponsors who will operate the output, and measuring the right outcomes during the pilot phase. The organizations with greater than 75% conversion rates (8% of the survey population) are not building better AI; they are building better AI programs.

The governance gap is narrowing but remains wide. The DPDP Act, sector-specific regulatory guidance, and the increasing sophistication of board-level AI questions are all forcing governance capability to develop faster than it might have done organically. Governance frameworks built in response to regulatory necessity are imperfect, but they are frameworks — and they create the documented accountability structures that enable organizations to scale AI responsibly.

The agentic threshold is approaching faster than governance frameworks are developing. The cross-survey convergence on Agentic AI as the next consequential capability shift — at 38% and 46% across two independent respondent populations — is the most forward-looking finding in this report. Enterprise AI teams that are not yet thinking through the governance, control, and operating model implications of AI agents acting autonomously on behalf of the organization are behind the curve.

The five data points that best summarize the 2026 state of AI in Indian enterprises:

- 66% at Explorer or Adopter stage (across the combined survey population, by maturity self classification)
- 52% of pilots failing to reach production;
- 12% with significant measurable ROI;
- 19% highly prepared for DPDP compliance; and
- 46% identifying Agentic AI as the capability most likely to reshape enterprise operations.

About This Study

Study Objectives

The State of AI in Indian Enterprises 2026 was designed to benchmark how Indian organizations are adopting, operationalizing, governing, and scaling artificial intelligence across business functions. The study set out to address four core questions: Where are organizations on the AI maturity curve? What is the production reality of generative AI deployments? How mature are AI governance and risk management frameworks? And what emerging capabilities do enterprise leaders expect to have the greatest operational impact?

Survey Methodology

CIO&Leader administered three independent surveys, each addressing a distinct domain of enterprise AI. All three surveys shared a common demographics module covering role, functional responsibilities, geographic scope, industry sector, organization size, annual revenue, AI maturity stage, and operating model. The domain-specific modules followed.

Survey 1, covering AI Strategy, Investment, and Enterprise Readiness, was completed by 117 respondents. Survey 2, on Generative AI, Agentic AI, and Production Reality, was completed by 98 respondents. Survey 3, addressing AI Risk, Governance, Security, and Future Outlook, was completed by 102 respondents. The combined survey population was 317 enterprise technology leaders.

Surveys were administered digitally across CIO&Leader's enterprise readership network between May and June 2026. Participation was voluntary. Respondents self-selected into the survey most relevant to their domain. The surveys were not designed to produce a unified cross-tabulated sample; findings from each survey are reported independently, with cross-survey synthesis on overlapping questions and themes where relevant.

Respondent Profile

Respondents are senior technology and digital leaders at India-operating enterprises — primarily CIOs, CTOs, GMs heading IT or digital functions, and CISOs. The respondent population skews toward India-headquartered enterprises (55-56% across surveys) and large organizations, with 30% of respondents from organizations with annual revenues below Rs 1,000 crore and 16-17% from organizations with revenues above Rs 50,000 crore.

Industry representation is led by Banking, Financial Services, and Insurance (BFSI); Manufacturing, Engineering, and Industrial; and IT Services, Software, and Technology Products — together accounting for approximately 60% of responses across the three surveys. This distribution broadly reflects the readership composition of CIO&Leader and should be taken into account when applying these findings to sectors underrepresented in the sample.

A full cross-survey comparative analysis of respondent demographics — including role, functional responsibilities, geographic scope, industry sector, revenue, employee size, AI maturity, and operating model — is provided in Section 9 of this report.

Appendix

A. Survey Instruments

The full survey questionnaires for all three surveys — AI Strategy, Investment and Enterprise Readiness; GenAI, Agentic AI and Production Reality; and AI Risk, Governance, Security and Future Outlook — are available from the CIO&Leader research team upon request. Each survey comprised a shared demographics module (Questions 1 to 8) and a domain-specific module (Questions 9 to 18).

B. Notes on Data Interpretation

All percentage figures in this report have been rounded to the nearest whole number. For multiple-selection questions, percentages sum to more than 100%. Survey 1 had 117 respondents for the demographic module and 101 for the strategy module. Survey 2 had 98 respondents for demographics and 88 for the GenAI and Agentic AI module. Survey 3 had 102 respondents for demographics and 91 for the Risk and Governance module. The variation between demographics and substantive module response counts reflects question-level skip patterns.

C. About ET Edge

ET Edge is India's leading publication and community for enterprise technology leadership. Published by ET Edge, CIO&Leader covers enterprise AI, cloud, cybersecurity, digital transformation, and the evolving role of the CIO. The annual CIO&Leader Conference is the flagship event for India's enterprise technology leadership community.

This image shows a full page of white paper with horizontal dotted lines. The lines are evenly spaced and run across the entire width of the page, providing a guide for handwriting practice. There are no margins, text, or other markings on the page.



© 2026 All rights reserved: Reproduction in whole or in part
without written permission from The Times Group.

7th Floor, The Times Group, Sunteck Icon, CTS 6956 VLC,
Kolekalyan Vimantal, CST Link Road, Kalina, Santacruz (East),
Vidyanagari, Mumbai, Maharashtra India, 400098