A 9.9 GROUP PUBLICATION

DECEMBER 2025 • VOLUME 15 • ISSUE 09 • ₹150

# CIO&LEADER

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF

A 9.9 GROUP PUBLICATION

cioandleader.com cioandleader cioandleader/


## AI GOVERNANCE

### DO YOU KNOW WHAT YOU ARE ACCOUNTABLE FOR?

PG 12


CIO&LEADER studiotalks

# CIO&LEADER STUDIOTALKS—
# WHERE TECHNOLOGY MEETS THE SPOTLIGHT!

CIO&Leader proudly presents StudioTalks—a premium platform where India's most influential CIOs and CTOs take center stage. Captured with high-production aesthetics, sleek visuals, and dynamic backdrops, StudioTalks transforms leadership insights into an engaging cinematic experience, and brings India's most influential CIOs and CTOs into the spotlight.
This exclusive series explores visionary leadership, emerging technologies, and strategic transformation—all presented in a format that blends deep insights with the visual polish of a professional studio production.

## WHY JOIN STUDIOTALKS?

Engage in powerful conversations that shape the future of enterprise IT.

Share your expertise in a high-impact, TV-style format.

Be featured among India's top technology leaders.

Be the voice of transformation. Be part of CIO&Leader StudioTalks.

**SECURE YOUR SPOT NOW!**

For more information
**Jatinder Singh**
Executive Editor - Enterprise Tech
jatinder.singh@9dot9.in, +919718154231

For Business Proposal
**Hafeez Shaikh**
National Sales Head, B2B Tech,
hafeez.shaikh@9dot9.in, +91 9833103611

**Follow us: @CIOandLeader**

# AI Governance comes of age in India

As enterprise AI ecosystems grow across both organisational and government environments, technology leaders are realising that scaling AI is far more complex than simply deploying models. The real challenge lies in achieving compliance, trust, and operational efficiency across the entire AI lifecycle. In practice, three barriers continue to slow enterprise-wide adoption. These include trust and governance concerns such as explainability, ethics, and bias, poor data readiness caused by fragmented and low-quality data, and gaps in skills and operating models that make it difficult to move AI from pilots into production at scale.

Among these, governance stands out as the biggest risk. Without clear guardrails, AI is not ready for enterprise-wide impact. Across industries, most AI-related failures have not occurred because algorithms were flawed, but because governance was missing. Unclear accountability, weak data management, limited explainability, and insufficient human oversight have been the real causes of concern.

Governments across the world, including India, have begun to recognise this risk. The Indian government's new AI Governance Guidelines encouraged wider AI adoption while reducing associated risks. The intent is to protect individuals, safeguard social interests, and uphold democratic values, while also supporting long-term growth and sustainability of India's AI ecosystem.

Recent incidents underline why this shift is necessary. A legal case in San Francisco highlighted how AI-generated content containing inaccuracies was used in judicial writing, leading to citation errors in legal rulings. Although the judge later corrected the record, the episode raised serious questions about the use of AI in high-stakes environments without proper oversight or accountability.

This issue brings together key highlights of the emerging AI governance framework, how CIOs are viewing its growing importance, and why stronger guardrails are essential. More importantly, it explores how the right governance approach can strengthen trust while still giving enterprises the confidence to innovate responsibly. While the guidelines are not yet law, they clearly signal what organisations are expected to do next! ■

**"Across industries, most AI-related failures have not occurred because algorithms were flawed, but because governance was missing"**

**Jatinder Singh**
Editor
jatinder.singh@9dot9.in

# content December 2025

COVER STORY

## 12-25

# AI Governance: Do you know what you are accountable for?

India's new governance guidelines are shaping how organisations are expected to manage AI, and where accountability now sits.


Cover Design by:
**Manish Kumar**

Please Recycle This Magazine And Remove Inserts Before Recycling

NEWS & VIEWS





INSIGHT






EVENT REPORT



TECH TALK








# CIO&LEADER

**www.cioandleader.com**

**MANAGEMENT**
Managing Director: **Dr Pramath Raj Sinha**
Printer & Publisher / CEO & Editorial Director (B2B Tech): **Vikas Gupta**
COO & Associate Publisher (B2B Tech): **Sachin Nandkishor Mhashilkar**

**EDITORIAL**
Group Editor: **R Giridhar**
Editor: **Jatinder Singh**
Principal Correspondent & Editorial Coordinator – CIO&Leader: **Musharrat Shahin**
Senior Correspondent: **Jagrati Rakheja**

**DESIGN**
Creative Director: **Shokeen Saifi**
Assistant Manager - Graphic Designer: **Manish Kumar**

**SALES & MARKETING**
Senior Director - B2B Tech: **Vandana Chauhan**
Head - Brand & Strategy: **Rajiv Pathak**

National Sales Head - B2B Tech: **Hafeez Shaikh**
Reginal Sales Head - North: **Sourabh Dixit**
Senior Sales Manager - South: **Aanchal Gupta**

**COMMUNITY ENGAGEMENT & DEVELOPMENT**
Head - Databases: **Neelam Adhangale**
Senior Community Manager: **Vaishali Banerjee**
Senior Community Manager: **Reetu Pande**
Senior Community Manager: **Snehal Thosar**

**OPERATIONS**
General Manager - Events & Conferences: **Himanshu Kumar**
Senior Manager - Digital Operations: **Jagdish Bhainsora**
Manager - Events & Conferences: **Sampath Kumar**
Senior Producer: **Sunil Kumar**

**PRODUCTION & LOGISTICS**
Senior Manager - Operations: **Mahendra Kumar Singh**

For editorial queries write to:
**editor@cioandleader.com**

For sales/business queries write to:
**responses@cioandleader.com**

**OFFICE ADDRESS**
**9.9 GROUP PVT. LTD.**
(Formerly known as Nine Dot Nine Mediaworx Pvt. Ltd.)
121, Patparganj, Mayur Vihar, Phase - I
Near Mandir Masjid, Delhi-110091
Published, Printed and Owned by 9.9 Group Pvt. Ltd.
(Formerly known as Nine Dot Nine Mediaworx Pvt. Ltd.)
Published and printed on their behalf by Vikas Gupta. Published at 121, Patparganj, Mayur Vihar, Phase - I, Near Mandir Masjid, Delhi-110091, India. Printed at G. H. Prints Private Limited, A-256 Okhla Industrial Area, Phase-I, New Delhi - 110020.

Editor: **Vikas Gupta**

9.9 GROUP

# CIOmovements

### Balaji Narayanamurthy Appointed Chief Data Officer at IndusInd Bank

**Balaji Narayanamurthy** is appointed Chief Data Officer at IndusInd Bank to lead enterprise data and AI strategy

### Dr. Yusuf Hashmi Joins Dell Technologies UAE as Chief Cybersecurity Advisor

**Dr. Yusuf Hashmi** joins Dell Technologies UAE as Chief Cybersecurity Advisor to advance Zero Trust and cyber resilience.

### KRC Murty Joins SBI Life Insurance Co. Ltd. as Executive Vice President – Head IT (BS)

**KRC Murty** takes over as EVP & Head IT (Business Systems) at SBI Life to drive technology and digital operations

### Kanodia Group Appoints Amman Walia as Group Chief Information Officer

**Amman Walia** becomes Group CIO at Kanodia Group, leading digital transformation and enterprise IT modernization

### NEXT100 Winner Saurabh Nigam Appointed Chief Executive Officer at Glimmer Technologies

**Saurabh Nigam** is appointed CEO of Glimmer Technologies to steer innovation, scale, and long-term business vision

### Subram Natarajan joins Larsen & Toubro Vyoma as Chief Technology & Innovation Officer

**Subram Natarajan** joins L&T Vyoma as Chief Technology & Innovation Officer to drive cloud and enterprise innovation.

### Sameer Ratolikar Elevated to Group Head & Chief Information Security Officer at HDFC Bank

**Sameer Ratolikar** is elevated to Group Head & CISO at HDFC Bank, strengthening enterprise-wide cybersecurity leadership

### Khaitan & Co Appoints Dr. Vimal Choudhary as Chief Operating Officer, Strengthening Leadership

**Dr. Vimal Choudhary** becomes COO at Khaitan & Co to drive strategy, operations, and digital innovation

## Pavankumar Shukla Takes Over as Head of Information Security & DPO at IDfy

**Pavankumar Shukla** becomes Head of Information Security & DPO at IDfy, leading cybersecurity and data protection.

## Shivaji Manwadkar Promoted to Chief Security Officer at SBFC Finance Limited

**Shivaji Manwadkar** is promoted to Chief Security Officer at SBFC Finance to strengthen enterprise security frameworks.

## Riddhi Adlakha Takes Charge as Region Head – Marketing, Sales & Ecommerce IT for Asia, Oceania & Africa at Nestlé

**Riddhi Adlakha** joins Nestlé as Region Head – Marketing, Sales & Ecommerce IT (AOA) to drive digital commerce innovation

## Chaitra Shetty Appointed as Head of Marketing – India at Workday

**Chaitra Shetty** becomes Head of Marketing – India at Workday to accelerate brand and enterprise growth.

## Ninad Raje Appointed as Group CIO at The Times Group – Driving Digital Innovation and Transformation Across Media & Technology

**Ninad Raje** is appointed Group CIO at The Times Group to lead digital innovation and technology transformation.

## Eveready Industries Appoints Sreekanth Neriyanuri as AVP & Head – IT to Drive Digital & Enterprise Transformation

**Sreekanth Neriyanuri** joins Eveready Industries as AVP & Head – IT to drive digital and IT modernization.

## Sachinkumar K. Patel joins Thomas Cook India Limited as Vice President – Information Technology

**Sachinkumar K. Patel** joins Thomas Cook India as VP – IT to lead technology modernization and security initiatives.

## NEXT100 Winner Dr. Vijay Choudhary joins JK Jajoo Ventures as the Chief Information Officer (CIO)

**Vijay Choudhary** is appointed CIO at JK Jajoo Ventures to shape tech strategy across real estate and aviation.

# India's new labour law 2025: Your workday, salary & rights on reset

India's new labour law of 2025 marks the biggest overhaul of wages, social security, industrial relations and workplace safety, legally binding on all states.

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

India has entered a decisive new phase in its economic journey. With the enforcement of four consolidated labour codes on 21 November 2025, the country has attempted what policy-makers have debated for nearly three decades: simplifying a maze of 29 antiquated labour laws into a coherent, modern framework. At its core, this reform is an effort to align India's regulatory architecture with the realities of a young, mobile, digitally native, and increasingly diverse workforce that earns, works, and aspires.

To appreciate the scale of this shift, it helps to glance at where India's labour laws come from and what they were never designed to handle.

## A System Built for an Older India

Most of India's labour laws were crafted in the mid-20th century, when industrialisation followed predictable patterns: factories, mills, mines, plantations, and office establishments. Work was stable, full-time, and overwhelmingly male. Over the decades, laws multiplied in response to specific demands, minimum wages here, contract labour there, maternity benefits elsewhere, until employers, workers, and even regulators found themselves navigating a labyrinth of overlapping rules.

By the late 1990s and early 2000s, the system was visibly misaligned. India was opening its economy, global investors were assessing risk, and organised industry was expanding faster than regulation could keep up. At the same time, an entirely new ecosystem was emerging: IT services, BPOs, e-commerce, and eventually gig and platform work. Millions of workers had entered the labour force through channels the law did not even recognise.

This mismatch produced a longstanding consensus: India needed fewer, clearer, more adaptable labour laws. But political sensitivity around labour rights repeatedly delayed the transition. The enforcement of the four codes marks the culmination of that long arc.

## The Four Codes: What's Actually Changing?

The new framework is structured around four pillars: wages, industrial relations, social security, and workplace safety. Each code introduces updates designed to correct historical gaps and reflect new forms of work.

### 1. CODE ON WAGES: A UNIFIED FLOOR OF FAIRNESS

The Code on Wages brings together four earlier laws, including the Minimum Wages Act and the Payment of Wages Act, into a single code.

Its most consequential shift is the introduction of a national floor wage, below which no state may set minimum wages. The aim is to reduce wide state-wise disparities and prevent wage suppression in low-income regions.

Another significant change is the standardised definition of "wages", which now determines

how employers structure salaries and calculate PF, gratuity, and bonuses. This is likely to reshape CTC packages across sectors and push more income into formal, benefit-linked components.

### 2. INDUSTRIAL RELATIONS CODE: FLEXIBILITY MEETS REGULATION

The most debated element of the reforms is found here: companies with up to 300 workers (up from 100) can now retrench or close without prior government approval. Employers argue this will make hiring less risky, potentially increasing job creation. Trade unions, meanwhile, fear reduced job security and weakened bargaining power.

Alongside this, the code revises definitions of strikes, lockouts, and dispute-resolution mechanisms, a move intended to reduce sudden workplace disruptions and promote structured negotiation.

### 3. SOCIAL SECURITY CODE: RECOGNISING THE INVISIBLE WORKER

The most forward-looking component is the inclusion of gig and platform workers within social security frameworks. This is a historic expansion of formal recognition to a segment that has grown sharply over the past decade: food-delivery partners, ride-hailing drivers, freelancers, beauty-service professionals, warehouse pickers, and more.

The code enables these workers to access health cover, life and disability insurance, and maternity-related benefits through government contributions and digital platforms. In a country where nearly 90% of the workforce is informal, this marks a shift toward broader safety nets.

### 4. OCCUPATIONAL SAFETY, HEALTH AND WORKING CONDITIONS CODE: A CLEANER, SAFER WORKSPACE

**"The Code on Wages brings together four earlier laws, including the Minimum Wages Act and the Payment of Wages Act, into a single code"**

The OSHWC Code creates uniform safety standards across industries, updating requirements that had long fallen behind modern workplace realities.

Key changes include:

- Mandatory annual health check-ups for specified categories of workers
- More explicit norms for ventilation, sanitation, drinking water, and welfare facilities
- The option for women to work night shifts with strict safety conditions
- A streamlined licensing framework for contractors operating across multiple states

In industries such as manufacturing, logistics, construction, and mining, these changes demand measurable improvements in compliance.

## Why This Moment Matters

The significance of the new labour codes lies not only in what they change, but in when they arrive.

India's workforce is one of the world's youngest. Gen Z now accounts for a significant share of new labour-market entrants, often seeking flexible, gig-based, or hybrid employment rather than traditional full-time roles. Older laws did not reflect this shift.

The rise of e-commerce, logistics networks, and platform-based services has created millions of jobs, many outside conventional employer–employee definitions. Recognising these workers is an essential step toward formalisation.

India's new labour codes represent a significant step toward modernising the country's employment laws. They streamline complexity, expand social protections, and update the rules of engagement between employers and workers. ■

# Ready or not, AI will power every IT job by 2030: Gartner

## According to a Gartner survey, every IT role will involve AI by 2030.

By **CIO&Leader** | editor@cioandleader.com

A new Gartner survey of over 700 CIOs reveals that by 2030, all IT work will involve artificial intelligence in some capacity, 75% of it will be done by humans augmented with AI, and 25% by AI alone. The study underscores a dual challenge for organizations:

**Gartner concludes that AI success will depend not just on adopting advanced tools, but on transforming people and processes.**

achieving both AI readiness (technological preparedness) and human readiness (workforce adaptability) to find and sustain business value from AI.

According to Gartner analysts Arun Chandrasekaran and Galliopi Demetriou, most enterprises still lack a balanced approach to both. While AI technologies are advancing rapidly, human capability to leverage and manage them effectively lags behind. Gartner predicts AI's net impact on jobs will remain neutral through 2026, turning positive by 2027 as AI creates more roles than it displaces.

To harness AI value, Gartner advises CIOs to restrain new hiring for repetitive roles, reskill employees for high-value tasks, and guard against "skills atrophy" from over-reliance on AI. The report also highlights cost, capability, and vendor strategy as key readiness factors urging firms to assess total ownership costs, match AI tools to their maturity, and choose vendors aligned with long-term sovereignty goals.

Ultimately, Gartner concludes that AI success will depend not just on adopting advanced tools, but on transforming people and processes. Organizations that build a workforce capable of partnering with AI will be best positioned to capture and sustain its full potential value. ■

# The AI privacy equation: India shows the world how it's done

## Zoho study finds 71% of Indian organisations strengthened privacy post AI adoption.

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

A new report titled The AI Privacy Equation: India Market Report, conducted by Arion Research LLC and commissioned by Zoho, reveals that India is rapidly advancing in responsible AI adoption. The study shows that 93% of Indian organisations now use AI, and 71% have strengthened their privacy measures after integrating AI into their operations. This reflects a proactive approach toward safeguarding data while scaling AI capabilities.

According to the findings, Indian enterprises demonstrate a mature understanding of AI-related privacy and ethics, with 90% acknowledging the broader privacy implications of AI and 92% having dedicated privacy teams or officers, surpassing global averages. A significant 65% of organisations allocate more than 20% of their IT budgets to privacy protection, underscoring a strong commitment to data governance. Critical high-risk areas identified include cloud storage of customer data, biometric data handling, and training AI models on customer interactions.

Ethical AI governance is also gaining ground, with 61% establishing ethics committees, 56% following data minimisation practices, and 55% conducting regular privacy audits. Experts note that Indian enterprises are disproving the notion that privacy slows down AI deployment, instead showing that strong governance accelerates trust and adoption.

The report further highlights India's progress in AI maturity, with 46% of businesses achieving advanced AI integration across software development, customer service, product development, and decision support. However, challenges persist, including data quality issues (44%), regulatory complexity (39%), and workforce skill gaps.

Overall, the study positions India as a global leader in responsible AI, showcasing how privacy-focused governance can serve as a strategic advantage in scaling AI innovation sustainably. ■

# AI is supercharging ad fraud and enterprises are paying the price in 2025

## New data shows why traditional fraud controls are collapsing as AI reshapes digital advertising at scale.

By **CIO&Leader** | editor@cioandleader.com

As enterprises accelerate AI adoption across marketing, the Ad Fraud in 2025: Beyond the Linear Lens report reveals a stark reality: the same intelligence powering automation and optimization is also enabling fraud at unprecedented scale and sophistication. Fraud is no longer confined to isolated bots or suspicious clicks—it now behaves like a coordinated, AI-driven system that blends seamlessly into genuine user activity.

The data is eye-opening. mFilterIt's 2025 validation analysis shows that fraud sophistication has tripled in just two years, driven largely by AI-based behavioral mimicry. Between 30–45% of programmatic traffic labeled "valid" failed deeper validation checks, while even closed platforms were not immune—9–18% of activity on walled gardens showed signs of invalid or incentivized behavior. In high-cost environments, these small percentages translate into disproportionately high financial losses.

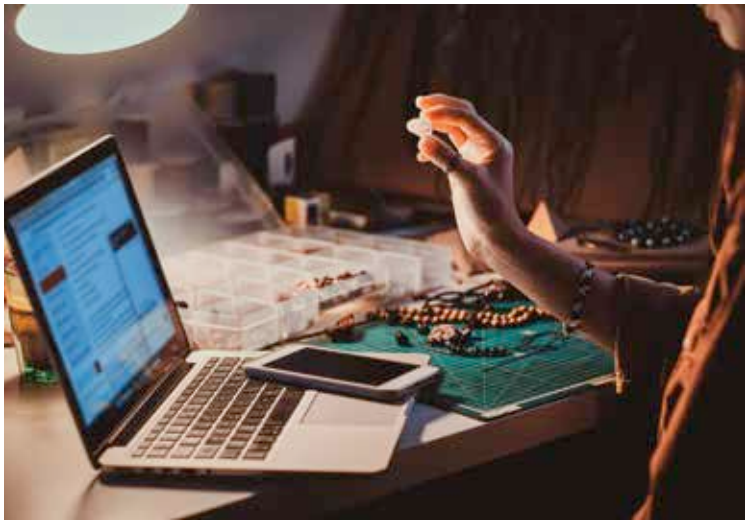

**AI-driven manipulation is also distorting performance metrics that enterprises rely on for decision-making**

AI-driven manipulation is also distorting performance metrics that enterprises rely on for decision-making. The report found that 45–55% of app installs can show anomalies such as abnormal click-to-install times or device duplication, despite appearing "clean" in attribution dashboards. Similarly, up to 35% of affiliate traffic showed signs of bot involvement or misattributed organic actions, inflating conversions while weakening downstream business outcomes.

Even exposure controls are breaking down. 15–18% of display and video impressions breached assigned frequency caps, leading to overexposure, ad fatigue, and wasted budgets—often without detection due to fragmented identity systems.

The report's message to enterprise leaders is clear: AI has transformed fraud from a tactical issue into a strategic risk. In a non-linear, AI-powered ecosystem, siloed verification no longer works. Enterprises must shift toward full-funnel, behavior-led validation to protect both marketing performance and long-term digital trust. ■

# 60% of India's tech graduates unfit for telecom jobs: TSSC and DIDAC join hands

## As the telecom revolution accelerates, this collaboration could define how India builds its next generation of tech talent.

By **CIO&Leader** | editor@cioandleader.com

India's telecom industry is at a critical inflection point. With only 40% of computer science, IT, and math graduates deemed job-ready, the sector faces an acute skill crisis that threatens to slow its 5G and AI ambitions. To tackle this widening gap, the Telecom Sector Skill Council (TSSC) has joined hands with the India Didactics Association (IDA) to drive workforce readiness through DIDAC Skills 2025, Asia's first exhibition and conference dedicated to skill development.

The collaboration aims to align education with industry needs by introducing future-ready learning modules, immersive training, and real-world skill demonstrations. Through this partnership, TSSC seeks to bridge the disconnect between academia and employers by promoting reskilling and upskilling in emerging technologies such as 5G, cloud, IoT, and AI.

DIDAC Skills 2025, organized by Messe Stuttgart India and IDA, will be held from 18–20 November 2025 at Yashobhoomi (India International Convention & Expo Centre), Delhi. Co-located with Didac India 2025, Asia's largest education exhibition, the event will bring together policymakers, educators, industry leaders, and employers to discuss and act on workforce transformation. The platform will feature conferences, live workshops, and industry-led challenges designed to enhance employability through hands-on learning.

Highlighting the collaboration's significance, Gaurav Sood, CEO & MD of Messe Stuttgart India, said, "DIDAC Skills 2025 is a transformative catalyst for India's skilling landscape. Partnering with TSSC strengthens our commitment to fostering industry-ready talent and bridging the telecom skill gap."

Aditya Gupta, CEO of India Didactics Association, added, "Skilling must move beyond classrooms. Together with TSSC, we aim to create meaningful career pathways and empower India's youth to lead the next wave of digital innovation."

Founded by the Cellular Operators Association of India (COAI), Indian Cellular and Electronics Association (ICEA), and National Skill Development Corporation (NSDC), TSSC serves as the key link between the government, academia, and industry. Its partnership with DIDAC Skills 2025 marks a major stride toward building a skilled, future-ready telecom workforce capable of powering India's digital economy. ■

# AI GOVERNANCE

## DO YOU KNOW WHAT YOU ARE ACCOUNTABLE FOR?

India's new governance guidelines are shaping how organisations are expected to manage AI, and where accountability now sits.

By **Musharrat Shahin** | musharrat.shahin@9dot9.in

At 5:18 AM, while the city was still asleep, the CIO of a large Indian enterprise glanced at an operations dashboard that suddenly didn't make sense.

Overnight, one of the company's core AI systems, responsible for thousands of automated decisions every hour, had shifted its behaviour. Approval rates changed. Routing priorities adjusted. Risk scores moved in unexpected directions.

There was no outage.

No alert. No obvious failure.

Yet the CIO knew something was wrong. In AI-driven operations, the most serious risks rarely arrive with alarms. They surface quietly. Within minutes, data science and security teams were on the call.

"What changed?" the CIO asked.

"The model updated itself," an engineer replied. "It may have discovered a new optimisation path."

"Is this drift? Bias? Or something worse?"

"We don't know yet."

Then came the question that increasingly defines the modern CIO's role:

"Under India's AI governance guidelines, does this qualify as a reportable incident?"

The room went silent. Someone pulled up

*"The guiding principle that defines the spirit of the framework is simple, 'Do No Harm'. The IndiaAI Mission will enable this ecosystem and inspire many nations, especially across the Global South."*

**—Prof. Ajay Kumar Sood**
Principal Scientific Advisor to the Government of India

the guidelines. Another checked the DPDP Act. The CISO revisited CERT-In's six-hour reporting rule. A data analyst flagged the requirement for human oversight.

No one had a clear answer. But everyone understood the stakes.

This is the new reality for Indian enterprises. AI now runs core operations, and AI governance is becoming central to accountability. But the ambiguity is intentional.

The government has deliberately avoided defining a rigid concept of "AI incident reporting". AI risk varies by context: what is critical in banking may be routine in retail. A one-size-fits-all mandate would either overwhelm regulators with noise or discourage innovation through over-compliance.

At the same time, silence carries its own risk. Under-reporting is increasingly seen as a governance failure, not a technical oversight. For CIOs, the message is clear: AI may be automated, but accountability is not. ■

## 7 Sutras of AI Governance Guidelines.

01
**Trust is the Foundation**
Without trust, innovation and adoption will stagnate.

02
**People First**
Human-centric design, human oversight, and human empowerment

03
**Innovation over Restrain**
All other things being equal, responsible innovation should be prioritised over cautionary restraint.

04
**Fairness & Equity**
Promote inclusive development and avoid discrimination.

05
**Accountability**
Clear allocation of responsibility and enforcement of regulations.

06
**Understandable by Design**
Provide disclosures and explanations that can be understood by the intended user and regulators

07
**Safety, Resilience & Sustainability**
Safe, secure, and robust systems that are able to withstand systemic shocks and are environmentally sustainable

Source: India AI Governance Guidelines

## The New Governance Roadmap: Why It Should Be Taken Seriously

When India's Ministry of Electronics and Information Technology (MeitY) recently unveiled the AI Governance Guidelines, it sent a clear signal: we know AI has risks, but we are not going to regulate blindly.

These guidelines function as quasi-regulatory signals. They do not carry penalties on day one, but they:

- Set expectations for responsible behaviour
- Shape how future laws will be interpreted
- Become reference points during audits, disputes, or investigations

If something goes wrong with an AI system, the first question regulators or courts are likely to ask is, 'Did the company follow published government guidance?'

Ignoring these guidelines increases legal and reputational exposure, even in the absence of a dedicated AI law.

For CIOs, CTOs, and CISOs, this is not just another policy PDF. It is a new operating context. The Guidelines assume that AI will be deeply embedded in decision-making, infrastructure, and citizen-facing services and then ask:

- Who is accountable when things go wrong?
- How should risk and liability be graded across the AI value chain?
- What institutional structures and technical controls are expected inside the enterprise?

MeitY's Guidelines prioritize how AI is used over how it's built, encouraging graded oversight based on impact rather than one-size-fits-all rules.

## What the guidelines actually expect from enterprises

Despite the absence of explicit penalties, the guidelines are unambiguous on one point: accountability rests with the organization deploying AI, not the algorithm, the vendor, or the data science team.

This marks a subtle but important shift. AI is no longer treated as an experimental technology or a specialist tool. It is positioned as an enterprise system, subject to the same expectations of oversight, control, and responsibility as core IT, cybersecurity, or financial infrastructure.

For CIOs, this reframes the question from "Is our AI compliant?" to "Can we defend our AI decisions if asked?"

The guidelines expect Indian enterprises to do four things:

- Treat AI as strategic infrastructure, not a side project.
- Classify AI systems by risk, and scale controls (testing, documentation, oversight) with that risk.
- Build formal governance structures: steering committees, ethics boards, AI Ops/MLOps, and integrated risk and grievance processes.
- Embed trust, transparency and inclusion into the design of AI products, especially those affecting livelihoods or public discourse.

## Accountability is explicit, even without penalties

The guidelines make it clear that automated decision-making does not dilute responsibility. If an AI system produces biased outcomes, causes harm, or behaves unpredictably, the organisation that deployed it is expected to explain:

- Why the system was introduced
- What decisions it influences or automates
- What safeguards and oversight mechanisms exist
- Who is responsible for monitoring and intervention

The absence of a formal AI law does not create a liability vacuum. Instead, the guidelines establish a standard of reasonable care: one that regulators, auditors, and courts can reference when assessing enterprise behaviour.

For CIOs, this effectively removes the defence of "the model did it on its own."

## Risk-based thinking replaces one-size-fits-all control

Rather than prescribing uniform controls for all AI systems, the guidelines implicitly push organisations toward risk-based classification.

Not all AI is treated equally.

- A recommendation engine optimising user experience carries limited risk
- An AI system approving loans,

### Existing legal framework

| | |
|---|---|
| 01 | **DPDP Act (2023):** Governs data use, consent, purpose limitation, and accuracy—forming the backbone of privacy and data governance for AI. |
| 02 | **IT Act & IT Rules (2021):** Apply to accountability, cybersecurity, content moderation, and online harms, including risks from AI systems. |
| 03 | **Copyright & Consumer Protection Laws:** Help address intellectual property issues and deceptive or unfair practices involving AI outputs. |

screening candidates, or determining eligibility for benefits carries materially higher risk

The expectation is not that CIOs govern every model identically, but that they know which systems matter most and apply proportionate controls.

This approach reflects India's broader regulatory philosophy: governance should follow impact, not technical complexity.

## Human oversight is no longer optional for high-impact AI

One of the clearest signals in the guidelines is the emphasis on human-in-the-loop mechanisms, particularly for systems that affect rights, access, or outcomes.

In practice, this means:

- AI-driven decisions must be reviewable
- Escalation paths must exist
- Overrides must be possible and documented

For CIOs, this is not merely a design principle—it is an operating model decision. AI systems that run unattended in high-impact environments increasingly represent a governance risk, not an efficiency gain.

## Transparency over perfection

The guidelines do not demand full mathematical explainability of every model. What they expect is reasonable transparency.

Can the organisation explain, at a high level:

- What the system is designed to do
- What data it relies on
- Why a particular decision or output occurred

Black-box systems with no internal understanding are strongly discouraged, especially in regulated or citizen-facing contexts. The emphasis is on defensibility, not academic purity.

These articles were first published in the Futurescape book in collaboration with Dynatrace, and they will continue to offer valuable insights ahead.

## Documentation becomes a control, not paperwork

One of the least visible but most consequential implications of the guidelines is the importance of documentation.

Model change logs, data lineage records, decision rationales, and risk assessments are no longer just internal best practices. They are increasingly viewed as evidence of governance maturity.

In the absence of such records, organisations may struggle to demonstrate that they exercised appropriate oversight, even if no harm was intended.

## Where this leaves the CIO

AI Governance Guidelines do not impose new technology requirements but impose new expectations of leadership.

CIOs are now expected to:

- Classify AI systems by risk and impact
- Ensure clear ownership and accountability
- Embed human oversight into critical workflows
- Prepare the organisation to explain its AI decisions

The guidelines do not ask CIOs to slow down AI adoption. ■

## CIO Action Agenda

| | |
|---|---|
| **Designate AI as Strategic Infra** | Treat AI like cloud or cybersecurity |
| | Long-term investment, not pilots |
| **Determine AI Risk Levels** | Classify systems: low → high risk |
| | Controls scale with risk |
| **Deploy Governance Architecture** | Steering Committee |
| | Ethics & Responsible AI Board |
| | AI Ops + MLOps |
| | Risk & Grievance Pathways |
| **Design for Trust** | Explainability |
| | Transparency |
| | Inclusion & public-impact sensitivity |

# How kotak balances speed, voice, and trust in digital banking

Sankar Sukhla, who heads IT Infrastructure and Technology at Kotak Mahindra Bank, describes a landscape where AI has become both the face and the spine of modern digital banking. With close to 3,000 branches and a deeply embedded digital footprint, Kotak increasingly relies on voice bots and AI agents to manage routine interactions, balance enquiries, EMI timelines, service requests, and self-service workflows, while human agents step in only at the final decision or resolution stage. The result is a hybrid operating model that closely mirrors the Guidelines' People First and Understandable by Design principles: machines absorb the repetitive load, but customers retain access to clarity, context, and human judgment when it truly matters.

The same logic applies to financial inclusion at scale. For customers in tier-3 and tier-4 towns particularly senior citizens, Sukhla sees voice, not chat, as the most intuitive gateway to digital finance. This insight aligns with India's broader push toward multilingual, accessible AI through platforms such as Bhashini, where familiarity of language and speech lowers adoption barriers. When executed thoughtfully, AI becomes an enabler that widens access to banking services rather than a filter that excludes those who are not app-native or digitally fluent.

Yet governance is never far from the conversation. Sukhla is acutely aware that BFSI operates in a high-liability environment where personally identifiable information is pervasive, fraud risks are persistent, and customer trust is sometimes fragile. His approach leans heavily on the Guidelines' Trust as the Foundation sutra: rigorously verifying caller identity, surfacing trusted entity names, announcing known identifiers such as CRN numbers, and ensuring every AI workflow is vetted by GRC (Governance, Risk, and Compliance) and risk teams, including compliance with RBI outsourcing norms. Unsurprisingly, he characterises Kotak's posture as an "80–20" bet on AI 80% focused on speed and scale, with the remaining 20% firmly reserved for security, compliance, and governance. ■

**"In banking, AI can drive speed and inclusion, but governance is what makes it deployable at scale."**

**—Shankar Shukla,**
VP–Head, IT Infrastructure & Technology, Kotak Mahindra Bank

# How newsrooms are defending trust against deepfakes

AI is no longer limited to summarising content or recommending articles; it is increasingly becoming part of the integrity layer that helps distinguish what is real from what is manipulated or fake.

For Ninad Raje, Group CIO at the Times Group, the Guidelines act as a necessary check on what he describes as AI "going out of the roof." In newsrooms, AI has already delivered clear benefits: analytics are faster, insights are deeper, and teams can finally use data at scale to understand audiences, content performance, and operational efficiency—capabilities that were difficult to achieve even a few years ago. At the same time, these very advances have also accelerated the creation of deepfakes and synthetic media, bringing governance concerns directly into day-to-day editorial operations.

Before any video is aired or a breaking story is published, Raje's teams now use AI-driven tools to detect manipulation, regardless of whether the content is a short clip or a long-form video. This reflects the kind of practical safeguards outlined in the Guidelines, including content provenance, watermarking, automated verification, and strong human editorial oversight. Together, these measures help reduce the risk of harmful or misleading content reaching the public.

Raje also highlights a challenge that boards are increasingly facing. As governance frameworks become stricter, experimentation can slow down. When early-stage AI projects carry potential regulatory or reputational risk, organisations may hesitate to test bold ideas. While the Guidelines support innovation over excessive restraint, adapting organisational culture to balance speed with responsibility remains a real challenge.

**"In the media, AI is no longer just a productivity tool, it's part of the trust infrastructure that decides what is real and what is not."**

**—Ninad Raje,**
Group CIO, Times Group

On employment, Raje offers a balanced view. AI is more likely to augment human roles than replace them, but only for those who are willing to learn and adapt. The bigger risk, he argues, is not AI eliminating jobs, but organisations failing to adopt and govern AI responsibly and losing relevance as a result. ■

# Why GPU-scale infrastructure demands responsible design

If AI is often compared to electricity, then hyperscale data centres form the grid that powers it. Rajesh Garg, President and Group CIO at Yotta, operates at the core of this grid, overseeing what he describes as the largest AI infrastructure in Asia today, supported by nearly 10,000 GPUs. From his perspective, the Guidelines largely formalise what serious AI infrastructure providers already understand: at this scale, strong governance is not optional but essential.

Yotta has therefore built a multi-layered governance structure. This includes an AI steering committee led by the CEO and supported by the CIO and CISO, a dedicated risk and compliance function embedded within GRC, an AI Ops capability within core operations, and a responsible AI and ethics board reporting to the CISO. Together, these bodies reflect the governance framework recommended in the Guidelines, with clear accountability for decisions, defined ownership of risk, and structured ethical oversight.

Garg is equally focused on the physical demands of AI infrastructure. While a traditional CPU rack may consume around 6 kW of power, a GPU rack can draw anywhere between 30 and 60 kW, requiring advanced liquid cooling systems and ultra-low-latency networking. This makes the Guidelines' focus on safety, resilience, and sustainability highly practical rather than theoretical. Without careful planning for power density, cooling, and network design, both the environmental impact and operating costs of AI infrastructure can rise quickly.

From a C-suite perspective, Garg's experience highlights a critical point: AI governance is not only about data, algorithms, and models. It also depends on long-term capacity planning, sustainable operations, and the resilience of the physical infrastructure that supports AI at a national scale, especially as demand, regulation, and public expectations continue to rise.

Leaders must therefore treat infrastructure decisions as strategic choices, not back-end concerns. Without this alignment, even well-governed AI initiatives can face operational bottlenecks or fail to scale reliably. ■

**"At GPU scale, AI governance isn't theoretical, it's embedded in how you design, power, and secure the infrastructure itself."**

**—Rajesh Garg,**
President & Group CIO, Yotta Data Services

# AI in the background: When agents become colleagues

In manufacturing and consumer businesses, AI is increasingly becoming part of everyday operations in a quieter but no less significant way. At Asian Paints, Associate Vice President Deepak Bhosle describes the change almost in anthropological terms. AI systems today can "think like humans, reason like humans, translate voice to text, analyse images, and converse like humans." In effect, this creates a new kind of workforce—autonomous software agents that operate alongside employees.

When these agents start interacting with employees, customers, and other systems across the business ecosystem, the risk landscape changes. Bhosle emphasises the importance of governance: defining clear protocols for interactions, setting strict boundaries for what agents can trigger, and ensuring they do not have unchecked authority to make decisions that could harm the business, employees, or customers. He also notes that regular audits, monitoring, and performance reviews of these AI agents are essential to prevent unintended consequences and ensure alignment with organisational goals.

He also points out that AI-related grievances should be managed through established complaint mechanisms, but with upgraded expertise. Someone needs to assess whether an issue arises from a model error, poor data quality, misuse of the tool, or unrealistic expectations of automation. Each incident then serves a dual purpose: managing risk in real time and creating a learning loop that strengthens policies, improves AI design, and tightens guardrails over time.

This approach reflects the Guidelines' emphasis on internal grievance redressal mechanisms—not as bureaucratic red tape, but as an adaptive safety layer. In a world where machines are increasingly "acting" rather than merely "advising," this ensures accountability, builds trust, and allows organisations to safely scale AI across operations, while continuously improving both human and machine collaboration. ■

**"The moment AI agents start acting, not just advising, governance becomes an organisational necessity, not a technical choice."**

**—Deepak Bhosale,**
Associate Vice President – IT, Asian Paints

# Why AI accountability starts with enterprises, not regulators

The AI Governance Guidelines rest on two key pillars: people and responsibility. Chhetry emphasises that organisations must invest significantly in training and upskilling teams to understand AI, while also embedding clear governance structures to ensure that models do not inadvertently encode or amplify bias. Developing internal expertise is essential, as human oversight remains the cornerstone of responsible AI adoption. It also helps build a culture where employees understand both the potential and the limitations of AI, enabling better decision-making and accountability across the organisation.

For Chhetry, regulators are only one part of the accountability picture. The primary responsibility lies with enterprises themselves. They manage customer data, trust, and outcomes, which makes it critical to implement robust processes that protect sensitive information while continuously monitoring algorithms for bias. This is particularly important in high-impact areas such as pricing, creditworthiness, eligibility, or access to services, where algorithmic errors or unfair decisions can have significant consequences. Embedding regular audits, scenario testing, and cross-functional review mechanisms further ensures that AI systems behave consistently and fairly.

He also offers a realistic view of the near-term AI landscape: the pace of innovation is only accelerating, with new use cases emerging daily. This means that governance cannot be static; it must grow, adapt, and strengthen over time. For C-suite leaders, the message is clear: AI governance is not a one-time board agenda item but an ongoing strategic capability. It requires continuous oversight, iterative policy updates, and a proactive approach to ensure that AI systems remain ethical, compliant, and aligned with both business goals and societal expectations, while also fostering trust among customers, employees, and regulators. Leaders must treat governance as a living function, integrating lessons learned into policies, workflows, and training to keep pace with emerging technologies. ■

**"AI governance isn't just about compliance—it's about protecting customers from invisible bias as systems scale."**

**—Aashish Kshetry,**
Vice President-Information Technology, Asian Paints

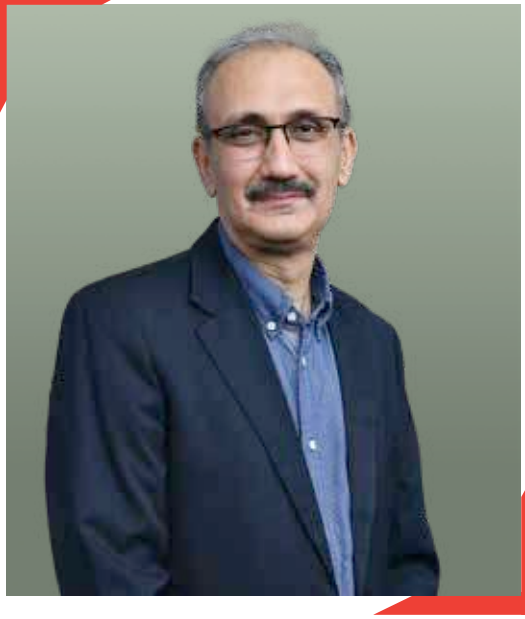

# The threat is already here: security, hallucinations and AI vs. AI

For Pradipta Patro Global CISO and Head of IT at RPG Group, governance begins with a straightforward acknowledgment: AI is already deeply embedded in the cybersecurity threat landscape. Cyber risks today are no longer theoretical; AI is both a tool for attackers and an essential instrument for defenders, making governance and oversight critical.

Patro views AI adoption as a journey rather than a single solution. The first step is ensuring high-quality, reliable data—reducing noise and minimizing the risk of hallucinated or inaccurate outputs. This is followed by continuous monitoring of models and outcomes, with each use case rigorously tested in its specific domain before deployment. This stepwise, iterative approach mirrors the Guidelines' emphasis on continuous monitoring, reassessment, and incident reporting for AI systems, particularly in sensitive or high-risk areas.

He stresses that trust in AI is still evolving. For the time being, human intelligence and AI must operate together to build confidence, case by case and KPI by KPI. Some applications will mature faster than others, and none should be relied upon as fully definitive without careful validation, ongoing testing, and clear accountability measures.

Patro also highlights a pressing reality for CISOs: in cybersecurity, it is increasingly AI versus AI. Attackers use AI to probe systems, launch phishing campaigns, automate attacks, and scale threats, while defenders deploy AI to detect anomalies, correlate events, and respond in real time. Opting out is no longer an option; the strategic choice lies in how responsibly, effectively, and ethically organisations leverage AI to protect themselves, maintain resilience, and stay ahead in this rapidly evolving digital arms race. This also requires ongoing investment in skilled personnel, adaptive threat intelligence, and robust incident response plans. Organisations must foster collaboration across teams and with external partners to anticipate emerging attack vectors, ensuring that AI-powered defence systems remain agile, transparent, and accountable at every step. ■

**"AI can be misutilised or utilised appropriately and that can have catastrophic effects."**

**—Pradipta Patro,**
Head of Cyber Security & IT Platform,
RPG Group.

# Building the AI value chain and taking responsibility

If the Guidelines provide macro principles, Sujoy Brahmachari offers a practical operating model for how enterprises can align.

He breaks the AI value chain into three roles:

- Developers who design and train models, ensuring accuracy, fairness, and compliance.
- Deployers who integrate these models into production manage infrastructure and monitor performance.
- Users who apply AI outputs, provide feedback, and follow governance guidelines.

On Risk, Brahmachari is clear that systemic risk—where a failure can have large-scale, interconnected consequences—is the most difficult to assess and therefore requires the greatest caution. He interprets the Guidelines' graded liability framework in a straightforward way: the intensity of controls, documentation, and oversight should match the level of risk posed by the system.

For high-risk AI systems that affect livelihoods, critical infrastructure, or essential services, this means implementing multiple layers of safeguards:

- Rigorous data quality checks to ensure accuracy and completeness.
- Regular model evaluations, performance testing, and explainability assessments.
- Structured bias and fairness testing to prevent unintended discrimination.
- Strong human-in-the-loop oversight to catch errors or anomalies before decisions are final.
- Detailed post-deployment monitoring with full audit trails to track performance and enable accountability.

For lower-risk applications, processes can be lighter but should never be completely absent. In every case, maintaining clear evidence of due diligence becomes the organisation's strongest defence, whether facing regulators, courts, or public scrutiny. ■

**"High-risk systems must face stricter validation, audits, oversight and documentation; low-risk ones can be streamlined, but everyone must maintain audit trails."**

**— Sujoy Brahmachari,**
CIO & CISO, Rosmerta Technologies

# Moving beyond pilots to enterprise accountability

Harnath Sahu emphasizes that India's AI Governance Guidelines represent a fundamental shift: enterprises can no longer treat AI as isolated projects but must manage them as full end-to-end lifecycles with continuous accountability and oversight. He argues that governance maturity begins with clearly defined roles, developers must embed ethics, explainability, and fairness into models; deployers are responsible for enforcing operational, security, and compliance controls; and users must exercise judgment rather than relying blindly on AI outputs.

Sahu stresses that explainability and bias mitigation are essential, especially in high-impact domains such as hiring, credit decisions, and pricing models. "If you cannot defend an AI decision, you should not deploy it," he notes, reinforcing the Guidelines' emphasis on transparency, fairness, and accountability. For him, the greatest challenge lies in assessing systemic risks, which do not arise from individual models alone but from complex, interconnected digital ecosystems, supply chains, and cascading operational dependencies.

He also warns that compute scalability could become a significant bottleneck unless organisations modernize data pipelines, adopt distributed training frameworks, and ensure secure access to India's growing pool of GPU resources and national datasets. As AISI standards evolve, Sahu believes governance teams must develop engineering-grade rigour, tracking data lineage, documenting risk-related decisions, and demonstrating measurable safety outcomes rather than merely stating intent.

Finally, Sahu highlights grievance mechanisms as the most human layer of trust. The ability to hear, diagnose, and resolve AI-related harm, he explains, will distinguish organisations that implement responsible AI from those simply pursuing automation. Proper grievance management, he adds, also strengthens accountability, informs governance improvements, and builds confidence among employees, customers, and regulators alike. ■

**"If an organisation cannot explain and defend an AI decision, it has no business deploying that system."**

**—Harnath Babu,**
Partner & CIO, KPMG India

## The CIO's new reality: challenges and gaps

The Critical Analysis is candid: The Guidelines are ambitious but high-level, leaving CIOs and CISOs with significant interpretation work.

### Key pain points include

| | | |
|---|---|---|
| 01 | **Operational ambiguity** | Broad definitions of "AI systems" and "risk frameworks" without detailed thresholds make it hard to know what exactly is in scope. |
| 02 | **Operational Uncertain liability** | Graded liability sounds good, but "due diligence" is undefined; boards can't easily estimate legal exposure. |
| 03 | **Operational Voluntary vs. mandatory** | Many controls are "voluntary" today but may become mandatory later, making it hard to justify investments or design long-term architectures |
| 04 | **Operational Talent gaps** | AI governance talent is scarce; training can cost ₹5–10 lakh per employee annually, and programmes currently cover only a tiny fraction of India's IT workforce. |
| 05 | **Operational Multi-regulator complexity** | CIOs might need to align with MeitY, CERT-In, sector regulators, and upcoming bodies like AIGG and AISI, each with evolving expectations. |

## CIO Playbook: A Practical Framework for AI Governance Readiness

Translating the Guidelines into action, a pragmatic seven-step CIO playbook emerges – closely aligned to the report's own practical guidance and the Critical Analysis.

### 1. DISCOVER – BUILD AN AI INVENTORY

- Catalog all AI/ML systems – from simple scoring models to generative assistants and agentic workflows.
- Map data sources, vendors, DPI integrations, and business owners.
- Tag systems that affect livelihoods (hiring, lending, pricing) or safety-critical domains.

### 2. ASSESS – RISK, FAIRNESS, BIAS, AND SECURITY

- Classify systems using a risk lens that includes malicious use, transparency, systemic risk, and loss of control.
- Run bias and fairness assessments on high-impact models, especially in HR, credit, and healthcare.
- Conduct security reviews for data poisoning, model theft, and adversarial input risk.

### 3. DESIGN – EMBED EXPLAINABILITY, OVERSIGHT, AUDITABILITY

- For high-risk use cases, demand explainable models or strong post-hoc interpretability.
- Design human-in-the-loop checkpoints, especially where law or ethics demand human accountability.
- Architect logging and audit trails for model decisions, inputs, and overrides.

### 4. OPERATIONALISE – BUILD PROCESSES AND COMMITTEES

- Set up AI steering and governance committees on the Yotta model.
- Integrate AI risk into enterprise GRC frameworks – from risk registers to board-level reporting.
- Clarify RACI (who is Responsible, Accountable, Consulted, Informed) across CIO, CISO, legal, business.

### 5. MONITOR – WATCH FOR DRIFT, ANOMALIES, AND MISUSE

- Implement continuous monitoring for model drift, data quality degradation, and abnormal behaviour.
- Use red-teaming and adversarial testing, especially for generative and public-facing models.

### 6. REPORT – ALIGN WITH CERT-IN, DPDP, SECTORAL RULES

- Prepare incident playbooks that integrate AI harms into existing breach and outage processes.
- Design reporting workflows that can meet aggressive windows

(e.g., six hours) where required.
- Ensure grievance channels exist for employees, customers, and partners – with clear SLAs and escalation.

### 7. IMPROVE – CLOSE THE LOOP

- Use grievances, incident learnings, and audit findings to refine models, datasets, and processes.
- Track KPIs like reduction in bias metrics, false positives, and grievance volumes.
- Feed learnings back into training, policy updates, and board conversations.

## What "Ready" Really Means for the C-Suite

Taken together, India's AI Governance Guidelines and these CIO perspectives converge on a simple but demanding message for leadership:

Being "ready" is not about having a few pilots or a policy PDF on a shared drive. It means:

- You have a clear inventory of AI systems, mapped to business owners and risk tiers.
- You have formal governance forums—steering, ethics, risk, and AI Ops—where AI decisions are discussed, documented, and owned.
- Your organisation can explain and justify AI-assisted decisions in high-impact areas like lending, hiring, pricing, healthcare, or content.
- You treat grievances and incidents as inputs to improve models and guardrails, not as one-off crises.
- You see trust, explainability, and fairness not as compliance burdens but as differentiators in a market where customers and regulators are visibly losing patience with opaque AI.

The Guidelines set the direction. The CIOs are already moving. The question is whether your organization is willing to do the slow, sometimes uncomfortable work of turning AI from a clever tool into a governed, trusted part of how you run the business.

Because the age of "unchecked AI" in India is over. And what comes next will be defined not just by how much intelligence you deploy but by how well you govern it. ■

# Cybersecurity workforce faces growing pains as industry ages

Cybersecurity faces a talent crunch, rising stress, aging workforce, and growing demand for adaptable professionals with strong soft skills.

By **CIO&Leader** | editor@cioandleader.com

ISACA's State of Cybersecurity 2025 report reveals a profession at a crossroads, with mounting stress, an aging workforce, and shifting priorities creating new challenges for the industry.

## The Aging Workforce Crisis

One of the most concerning findings is the graying of cybersecurity professionals. The largest group of survey respondents (35%) is between 45 and 54 years old, while the number of younger workers under 35 has declined slightly. With many experienced professionals nearing retirement and fewer young people entering the field, organizations may soon face a critical talent shortage. Only half of the respondents manage staff with less than three years of experience, raising questions about who will replace retiring managers.

## Stress Levels Remain High

Despite being in high demand, cybersecurity professionals are experiencing burnout. Sixty-six percent report that their roles are more stressful now than they were five years ago. The main culprit? An increasingly complex threat landscape, though fewer professionals cited this as a problem compared to last year (63% versus 81%). High stress levels are pushing people to leave their jobs, yet surprisingly, one-quarter of organizations aren't taking any steps to address burnout.

## Adaptability Tops the Wishlist

When hiring, employers now value adaptability above all else—61% say it's essential. This marks a shift from previous years when hands-on experience was king. The importance of prior cybersecurity experience has decreased significantly, from 73% to 60%, indicating that employers are looking for professionals who can adapt quickly in a rapidly changing environment.

**47% of cybersecurity professionals are now involved in developing AI policies, up from 35% last year, indicating a more secure and responsible approach to AI.**

## Soft Skills Are the Biggest Gap

The report identifies soft skills as the most significant deficiency among cybersecurity professionals, with a notable increase from 51% to 59% in just one year. Critical thinking, communication, and problem-solving are among the most essential skills. This gap may explain why boards sometimes fail to prioritize cybersecurity adequately—professionals struggle to communicate their value to non-technical leadership effectively.

## Budget Pessimism Grows

Only 41% of respondents believe their cybersecurity budgets will increase in the next year, down from 47% in the previous year. Meanwhile, 18% expect cuts—a significant jump from 13% last year. This budget uncertainty, combined with declining employer benefits like certification fee reimbursement (down from 65% to 54%), paints a challenging picture.

## AI Adoption Increases

On a positive note, organizations are increasingly using AI for security operations, particularly for automating threat detection and endpoint security. More importantly, 47% of cybersecurity professionals are now involved in developing AI policies, up from 35% last year, indicating a more secure and responsible approach to AI implementation ahead. ■

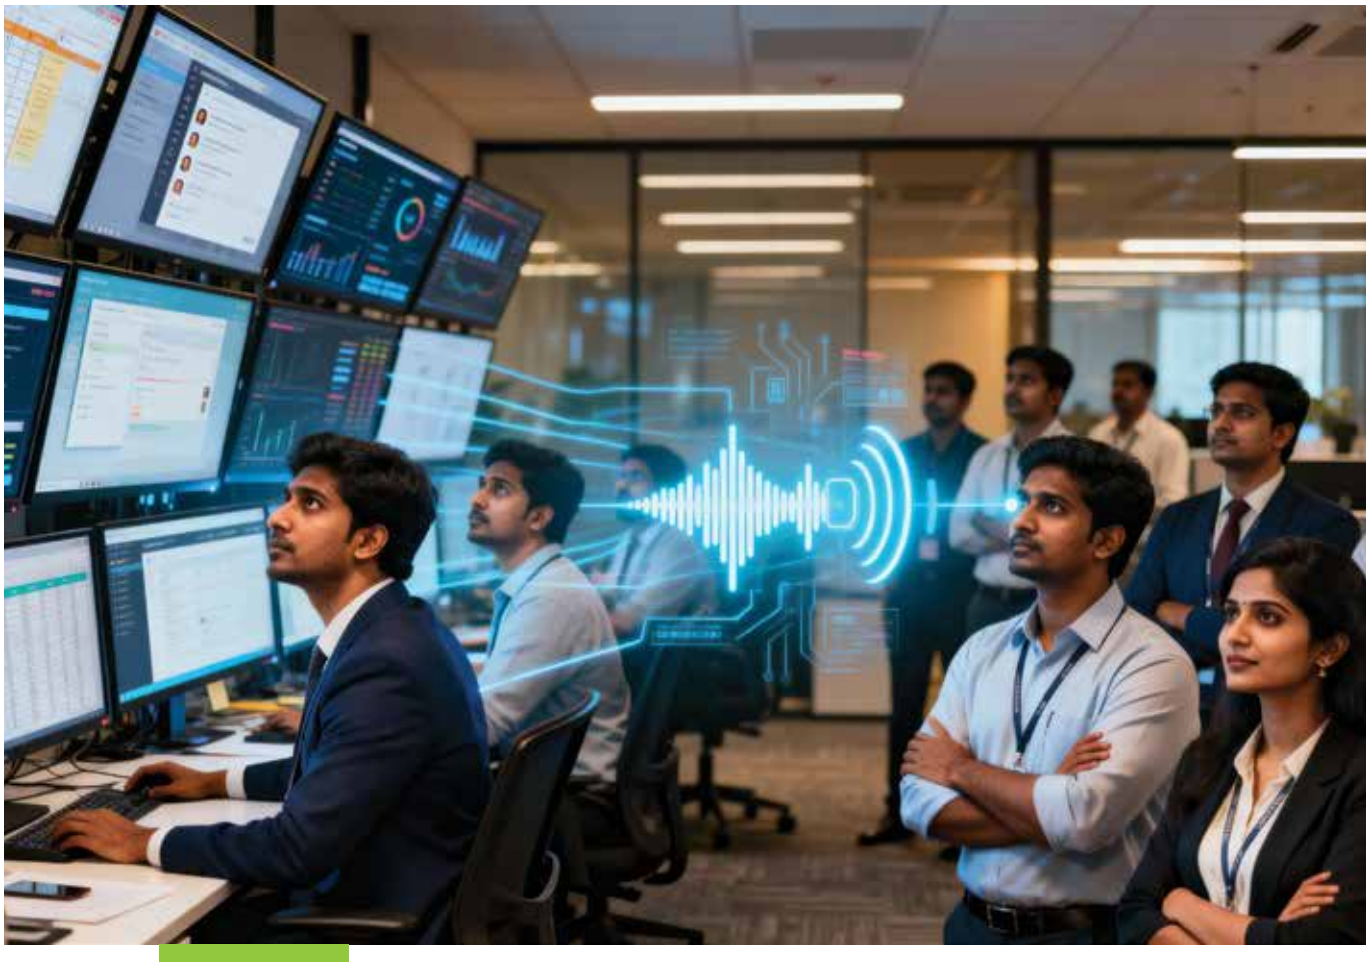

# Why Indian businesses can't scale without voice

Kaushal Bansal, Co-Founder and CEO of Callerdesk, on the API-driven communication stack and why Indian businesses can't scale without voice in the workflow.

By **Kaushal Bansal** | editor@cioandleader.com

## From click to conversation in seconds

In high-velocity markets, speed wins. A lead that lands from Meta Ads, a website form, or WhatsApp is most valuable in the first few minutes. Callerdesk's lead-to-call API closes that gap by triggering an instant, two-way connection between the agent and the prospect the moment a lead appears. Response time drops from minutes to seconds, lead leakage falls, and conversion rates rise. For lenders, insurers, education providers, and local services, that single change often determines who books the business.

## IVR that adapts to your operations

Queues swell at lunch, festivals, and flash sales.

Traditional IVRs are slow to change and depend on manual switches. With Callerdesk, IVR menus, routing rules, and failovers are programmable through API, so teams can set policies that adjust themselves. During peak windows, calls can overflow to backup agents; premium segments can jump the queue; noncritical lines can be paused during outages. The payoff is fewer abandoned calls and a smoother experience without babysitting dashboards.

## Voice and CRM as one motion

With NASSCOM forecasting India's voice-AI market to hit US$ 1.82 billion by 2030, Indian firms must embed voice APIs directly into their communication workflows, voice is no longer an add-on but a core, scalable layer of engagement.

Sales and support live inside the CRM. If calls sit outside that system, context goes missing and reports become guesswork. Callerdesk ships bi-directional integrations with platforms like Zoho, Freshsales, and HubSpot so profiles are fetched before the call, outcomes are logged the second the call ends, and follow-ups fire automatically. Missed calls can open tickets, recordings can attach to the contact, and tasks can be scheduled without human entry. When voice and CRM move as one, managers finally see the full journey and can coach to real numbers.

## Privacy built for marketplaces

Marketplaces and platforms handle sensitive phone numbers for buyers, sellers, couriers, patients, and owners. Callerdesk's workflow-based number masking connects parties without exposing personal details, while keeping calls auditable. Rules can be tailored to the use case. A delivery can allow a return-call window that expires after completion. A rental inquiry can stay anonymous until a booking is confirmed. The result is compliant, secure conversations that do not slow down operations.

**Kaushal Bansal**
Co-Founder and CEO of Callerdesk

## Missed calls that still move work forward

India's missed-call culture is a useful signal. With Callerdesk's missed-call APIs, that signal becomes automation. A single ring can trigger an acknowledgement SMS, create a CRM record, queue a callback, or launch an onboarding flow. Customers feel heard without waiting on hold, and agents spend time talking rather than typing.

APIs are the new operations team. They don't sleep, they don't miss leads, and they never wait for instructions. Callerdesk brings that reliability to voice" ~ Kaushal Bansal, CEO & Co-founder, Callerdesk

## Turning voice into a data asset

Phone conversations carry cues that forms and emails miss. Callerdesk Smart Analytics exposes real-time call logs, first-response and handle times, connect rates by campaign, and agent performance without manual exports. As teams mature, keyword and sentiment layers can flag risk and training needs. The path is pragmatic. Start by making logs reliable and outcomes tagged; then layer insights. Over time, voice shifts from an opaque cost to a measurable driver of revenue and satisfaction.

Indian businesses don't scale by adding more agents. They scale by automating the moments between a click and a conversation. At Callerdesk, our API-driven voice stack turns every workflow into a real-time, intelligent trigger so teams respond in seconds, not minutes." — Rajesh Dimania, CTO & Co-founder, Callerdesk

## One customer, many touchpoints, one stack

Most apps still rely on SMS for OTPs and updates, even when delivery is patchy. Callerdesk's voice APIs let apps and web flows trigger verification calls, proactive reminders, and service notifications when the moment calls for something faster and more personal. A loan app can place a verification call in low-signal zones. A healthcare app can deliver pre-op instructions and capture confirmation. A service app can trigger an automated reminder before a scheduled visit. When voice, web, and app channels are unified behind APIs, brands meet customers on the medium that works right now, not the one that is easiest to blast.

## The scale advantage

Indian businesses are competing on speed, trust, and unit economics. An API-driven voice stack anchored by Callerdesk improves all three. It connects intent to conversation in seconds, keeps data consistent across systems, protects privacy in complex workflows, and turns every call into searchable context. ■

# Meet L&T Vyoma: L&T's biggest digital leap for India's tech landscape


Welcome Keynote by **Ms. Seema Ambastha, CEO – L&T Datacenter and Cloud Services**

Larsen & Toubro Vyoma marked the start of a defining new chapter.

By **CIO&Leader** | editor@cioandleader.com

The Vyoma identity was officially unveiled by **Mr. S.N. Subrahmanyan, Mr. Puneet Chandok,** and **Mr. Prashant Jain.** Derived from the Sanskrit word meaning limitless sky, Vyoma reflects the vision of building digital foundations for scale, sovereignty, intelligence, and purpose.

**The Keynote :** The launch day featured insight-rich discussions with leaders from India's technology ecosystem. In the keynote "AI: The Future Already in Motion," **Puneet Chandok** spoke on the evolution of intelligence, the role of human creativity, and shared five predictions shaping the future of work and enterprise transformation.

Lighting of the Lamp by Leaders


Brief update on Business and invite CMD on stage by **Prashant Jain**

Panel discussion hosted by CIO&Leader


Address by S N Subrahmanyan – Chairman & MD (represented by Puneet Chandok.)

**Amit Luthra**
Managing Director, Lenovo ISG, India

# Building for what's next: AI, infrastructure, and the road to 2026

In conversation with CIO & leader **Amit Luthra, Managing Director, Lenovo ISG, India,** on how Indian enterprises scaled AI in 2025, why hybrid infrastructure became the default, and what CIOs must prioritise next.

By **Jatinder Singh** | jatinder.singh@9dot9.in

**CIO&Leader: How has the Indian enterprise technology landscape evolved in 2025, and what key trends stood out among your customers?**

**AMIT LUTHRA:** In 2025, Indian enterprises moved decisively from experimentation to execution. It was no longer enough to pilot AI or test emerging technologies; organizations asked deeper questions about how AI could tangibly improve productivity, support faster and more informed decision-making, and scale reliably across large and complex operations.

We have seen hybrid environments combining edge, on-premise, and cloud becoming the standard approach. These setups allowed companies to keep AI workloads close to the data while balancing performance, governance, and security requirements, making operations more efficient and resilient.

Another notable trend was that Indian enterprises thought strategically about infrastructure, not just technology. According to Lenovo's CIO Playbook 2025, nearly 63 % of organizations preferred hybrid or on-prem setups for AI workloads. This reflected careful planning around scalability, operational control, and energy efficiency; decisions that were critical when moving AI from pilot projects into production-scale operations.

Overall, 2025 was a year where strategy met execution. Enterprises considered how AI and hybrid infrastructure could solve real business problems, rather than adopting technology rather than deploying technology without a clear business outcome in mind.

**CIO&Leader: Which technologies or solutions saw the highest adoption this year, and what challenges did they help your clients overcome?**

**AMIT LUTHRA:** This year, generative AI and agentic AI tools saw some of the highest adoption across Indian enterprises. These technologies proved particularly useful in automating routine tasks, enhancing decision-making, and creating new opportunities for employees to focus on higher-value work.

Enterprises increasingly prioritized solutions that were integrated across devices, infrastructure, and enterprise systems. Integration is no longer optional; it is essential for ensuring consistent performance and reliability, and for deriving measurable business outcomes.

Hybrid AI platforms played a key role in helping organizations

scale AI from pilot programs to full production deployments. Modern infrastructure solutions not only provided performance and security but also optimized energy usage, cooling, and compute utilization, addressing operational challenges that can otherwise limit AI adoption.

The Lenovo CIO Playbook 2025 highlights that integration of complexity and demonstrating ROI remain significant barriers to adoption. Enterprises that focused on unified platforms, ones that brought together AI tools, workflows, and data management, have been able to overcome these challenges effectively. These platforms helped businesses move beyond experimentation to measurable business impact, which is now expected by the senior leadership

**CIO&Leader: Looking ahead, which emerging technologies or sectors do you expect to see the most growth in 2026?**

**AMIT LUTHRA:** Looking into 2026, agentic AI and hybrid AI platforms are likely to see continued growth, particularly where they enable better productivity, decision-making, and operational efficiency. These technologies are moving from pilot stages into real-world business applications, where they can impact everything from workforce efficiency to customer experience.

Infrastructure modernization, including optimized storage, efficient compute, and hybrid deployments, will be critical to sustaining AI at scale. Organizations are recognizing that scaling AI is not just about software; it requires the right underlying architecture and infrastructure to support growing workloads while managing costs and energy efficiency.

The Lenovo CIO Playbook 2025 suggests that planned investments in hybrid workloads are particularly strong, which indicates that

**Investing in flexible hybrid infrastructure is critical for supporting growth while maintaining operational stability.**

scalability, operational efficiency, and reliability will continue to drive growth across sectors such as manufacturing, financial services, telecom, and healthcare. In essence, the focus for 2026 will be on practical, measurable adoption of AI underpinned by infrastructure that can scale effectively.

**CIO&Leader: What advice would you give CIOs and enterprise leaders as they plan their technology investments and governance strategies for next year?**

**AMIT LUTHRA:** My advice to CIOs and enterprise leaders is to remain grounded in outcomes. Choose technologies that deliver measurable productivity gains and streamline workflows, rather than being driven by hype or the latest trend.

Investing in flexible hybrid infrastructure is critical for supporting growth while maintaining operational stability. This includes planning for performance, energy efficiency, and lifecycle costs. Govern responsibly.

With nearly two-thirds of Indian enterprises now preferring hybrid deployments, planning infrastructure and governance together is essential to enable growth while mitigating risk. Finally, initiatives should empower people, streamline workflows, and deliver tangible business value, ensuring technology investments translate into real outcomes. ■

# The future platform has to be unified, intelligent, and AI-first

**Praful Poddar, Chief Product Officer,** and **Sunil Kumar, Chief Technology Officer, Shiprocket,** on how Shiprocket unifies data, AI, and logistics to empower India's SMBs

By **Jatinder Singh** | jatinder.singh@9dot9.in

Founded in 2017 by Sahil Goel, Gautam Kapoor, and Mohit Gupta, Shiprocket has emerged as one of India's most influential logistics and enablement platforms for small and mid-sized businesses. What began as a logistics aggregation play has steadily evolved into a full-stack commerce and fulfillment platform, helping sellers manage everything from shipping and warehousing to checkout, payments, marketing, and returns.

At its core, Shiprocket solves a simple yet underserved problem: enabling India's SMBs and digital-first entrepreneurs to compete with larger enterprises on speed, cost, and customer experience. By combining scale-driven logistics partnerships with a deeply integrated technology layer, the company has positioned itself as a critical infrastructure player in India's booming e-commerce ecosystem.

Technology has been central to this journey. Shiprocket operates a cloud-native, API-first platform that increasingly leverages data science, machine learning, and generative AI to optimize supply chains, reduce friction, and improve seller decision-making. With more than a third of its business flowing through APIs and a rapidly expanding AI footprint, the company is pushing toward a unified, intelligent platform designed for scale.

In this conversation with Jatinder Singh, Editor, CIO&Leader, Praful Poddar, Chief Product Officer, and Sunil Kumar, Chief Technology Officer, share how Shiprocket is using technology to solve fundamental logistics challenges, how AI is reshaping both internal operations and merchant experiences, and what lies ahead as the company prepares for its next phase of growth.

**CIO&Leader: You are leveraging technology to provide end-to-end logistics services. How are you using these technologies, and what key challenges are you trying to solve in the Indian market?**

**PRAFUL & SUNIL:** Some principles have been core to Shiprocket right from the beginning, when Sahil and Gautam started the company. We have always focused on solving problems for SMBs. This segment is significantly underserved by cost-effective, scalable technology solutions that actually work.

Enterprise-grade solutions exist, but they are expensive and out of reach for most SMBs. That creates an uneven playing field, where smaller sellers cannot compete effectively with large enterprises. Our mission has been to identify the correct problems and solve them in a scalable, technology-driven way.

In logistics aggregation, the original problem statement had two

parts. The first was access. SMBs needed access to multiple courier partners in one place. Once a small seller tries to work directly with tier-1 couriers [such as Blue Dart or Delhivery], it becomes difficult to get attention, negotiate contracts, or secure competitive pricing. Aggregation gives them choice and flexibility.

The second was pricing. Because we operate at scale, we can negotiate better rates, which are then passed on to sellers. Combined with the technology layer, this creates a largely hands-off logistics experience. Traditionally, an SMB would need one or two people to manage couriers, disputes, delays, and daily coordination. We absorb that complexity through technology, operations, and managed services, so logistics feels outsourced by default.

As we continued solving logistics challenges, we kept hearing from sellers about issues beyond shipping. After logistics, there were challenges around returns, cancellations, customer communication, and remarketing. Before logistics, there were gaps in payments, checkout, ads, and commerce workflows.

We chose to expand selectively, focusing only on areas where we had confidence and complementary capabilities, and where we could take a proper zero-to-one approach.

One insight we gained early was that SMBs do not like paying fixed SaaS fees. They are uncomfortable committing to an INR 5,000 monthly subscription. What they are comfortable with is paying per transaction. That led us to build a pay-per-use model in which costs scale with usage and margins are embedded in the transaction flow.

On warehousing, this again came from a seller problem. SMBs selling on marketplaces like Amazon and

**Sunil Kumar,** Chief Technology Officer and **Praful Poddar,** Chief Product Officer, Shiprocket, on how Shiprocket unifies data, AI, and logistics to empower India's SMBs

**"Our largest customer contributes less than 3 percent of revenue. That forces us to build systems that work for extreme fragmentation."**

Flipkart benefit from fast delivery expectations, often one or two days. Sellers operating through their own websites, WhatsApp, or social media struggled to match that experience.

Warehousing enables us to distribute inventory nationwide rather than tying sellers to a single location. If inventory is closer to the buyer, delivery becomes faster. For instance, if a seller places inventory in our Bengaluru warehouse, orders from South India can be fulfilled locally, enabling same-day or next-day delivery. This is why we now operate around 30-35 warehouses across India.

On exports, we currently do not operate our own international warehouses. However, we support both B2C parcel shipping and cargo shipping for Indian sellers shipping to markets like the US, Europe, the Middle East, Africa, Asia, and Australia. Some sellers ship single parcels through dropshipping models, while others move inventory in bulk to international fulfilment centres. We support both use cases through our international shipping offerings.

**CIO&Leader: Technology plays a crucial role in identifying efficient shipping methods and better discounts. With AI and ML becoming more prominent, what innovative models are you working on? Are these built internally or with partners?**

**PRAFUL & SUNIL:** From the outside, Shiprocket appears to be a single end-to-end platform that handles fulfillment, warehousing, shipping, checkout, and value-added services. Under the hood, however, these are distinct products operating in parallel.

### Logistics as a digital platform

Shiprocket's journey highlights a broader CIO lesson: physical industries are increasingly being run by software platforms. Success depends not just on automation, but on orchestration, data unification, and resilience across digital and physical layers.

### Architecture Snapshot

**Inside Shiprocket's Tech Stack**

- Cloud-native, multi-cloud setup on AWS and GCP
- Fully containerized using Kubernetes and EKS
- Event-driven, federated product architecture
- Internal data lake with Snowflake for analytics and AI
- API-first design, with 40 to 50 percent of traffic flowing through APIs

**CIO Insight:** This architecture enables independent scaling while preserving end-to-end visibility.

### AI in Practice

**From ML to GenAI with Guardrails**

**Before GenAI:**

- Courier recommendation across 17+ partners
- COD RTO risk prediction
- Lead scoring across hundreds of thousands of sellers

**With GenAI:**

- Conversational workflows via AI Copilot
- Trends AI powered by large-scale transaction data
- LLM-agnostic MCP layer enabling conversational commerce

**Key Principle:** AI augments decisions; humans retain accountability.

### Strategic outlook

Shiprocket's next phase focuses on:

- Unifying independently built products into a cohesive platform
- Making AI a native layer across experiences
- Reducing friction across the seller lifecycle
- Supporting SMBs without forcing enterprise-level complexity

The challenge is that they must work together seamlessly because it is the same customer and the same order moving through multiple stages, from placement to delivery, returns, or exchanges. This requires a federated data architecture in which products are self-sufficient but communicate via events.

Every order status change triggers events, including warehouse processing, courier assignment, and delivery updates. This orchestration layer is critical.

We have built a strong data foundation using our internal data lake and Snowflake. This allows teams to access real-time, reliable data and enables advanced analytics and AI use cases. Clean and structured data is a prerequisite for effective AI, especially LLMs.

We approach AI across three tracks:

- AI for Shiprocket, focused on internal efficiency
- AI for merchants, focused on simplifying seller workflows
- AI for Bharat, contributing to broader ecosystem needs

We use a mix of internally built models and fine-tuned external models. Use cases include address correction, text-to-speech, speech-to-text, and experimentation with open-source and commercial LLMs such as OpenAI, Claude, and others. We are also working on India-first voice models to address linguistic diversity.

Even before GenAI, we ran machine learning models for courier recommendation across 17-plus partners, RTO risk prediction for COD orders, and lead scoring across hundreds of thousands of monthly signups.

With GenAI, platforms like ours are becoming conversational. Instead of navigating screens, sellers can issue instructions. We built an AI Copilot that automates actions and answers queries within the platform.

We also launched Trends AI, powered by insights from billions of transactions across millions of Indian buyers. It helps sellers understand what sells where, payment preferences, buyer personas, and demand patterns, enabling better marketing and product decisions.

**CIO&Leader: How does your technology architecture help you handle demand spikes during festive seasons or sudden surges?**

**PRAFUL & SUNIL:** We are cloud-native and operate across AWS and GCP, which gives us the elasticity we need. While we are not hybrid yet, our multi-cloud setup allows us to scale workloads dynamically. We are fully containerized, using Kubernetes and EKS. This makes scaling faster and more predictable, as everything runs in small, modular pods.

Monitoring is another critical pillar. We use an ELK stack based on Elasticsearch to process terabytes of logs, with real-time dashboards and alerts. On top of that, Prometheus and Grafana help correlate infrastructure metrics with business performance.

## "Even programmatically, I cannot decrypt sensitive data myself. Access is strictly controlled and role-based."

If a large client expects a three- to four-times spike over a short period, we can scale efficiently while maintaining buffer capacity. At our current scale, we can absorb a 30 to 40 percent upside without issues, purely through elasticity rather than idle infrastructure.

Every new product or feature is stress-tested at five to ten times the projected load. Automation and QA testing run continuously to identify vulnerabilities early.

Warehousing introduces a unique challenge. Unlike software, physical capacity cannot scale instantly. Shelf space, workforce, and processing stations must be forecasted accurately. Data-led planning is critical because both underutilized and overloaded warehouses create problems for sellers.

**CIO&Leader: With an IPO on the horizon, how do you see Shiprocket's next phase of growth? What are your technology priorities over the next two to three years?**

**PRAFUL & SUNIL:** Our long-term vision has always been to become a one-stop shop for SMBs. While enterprises contribute 20 to 25 percent of revenue, our base is highly fragmented. No single customer contributes more than 3 percent of revenue.

Our core focus is on sellers shipping between 500 and 10,000 orders a month, though we also serve micro-entrepreneurs and large enterprises. Across this spectrum, sellers need a range of tools spanning cataloging, payments, marketing, logistics, warehousing, and returns.

Our effort has been to bring more of this under one platform, so data flows better and products work together more effectively. We do not want to build everything ourselves. In areas like website creation, Shopify already does an excellent job, and we integrate deeply with them.

Over the next phase, unifying our products into a single, cohesive platform is a significant priority. Many solutions were built independently to find product-market fit. Now the focus is on integration and orchestration.

AI is the second central pillar. We are moving toward an AI-first platform, rethinking workflows and experiences through an AI lens.

**CIO&Leader: How do you balance AI-driven efficiency with the need for human-centric customer experience?**

**PRAFUL & SUNIL:** We are very conscious of this balance. Not every interaction should be AI-first. For example, a first call about a delayed pickup can be handled by AI. But if the customer calls again within a short window, we route them directly to a human agent.

AI should reduce friction, not increase frustration. The key is knowing when to step back and let humans take over.

**CIO&Leader: Any final thoughts on industry challenges ahead?**

**PRAFUL & SUNIL:** Many logistics challenges remain fundamental. COD efficiency, weight discrepancies, and delivery speed are persistent issues. Technology can reduce information asymmetry, improve transparency, and help sellers make better decisions.

AI, predictive planning, and more intelligent inventory placement can help SMBs deliver faster without excessive duplication. Ultimately, solving simple, fundamental problems well still matters the most. ■

**Bikramdeep Singh**
Country Manager, Proofpoint India

# AI changes the question from 'what' to 'who'

In his conversation with 9.9 Group, **Bikramdeep Singh, Country Manager, Proofpoint India** explained that AI is transforming cybersecurity by shifting the core question from what needs protection to who needs protection.

By **Jatinder Singh** | jatinder.singh@9dot9.in

With cyberattacks becoming increasingly targeted and behaviour-driven, organisations are being forced to rethink long-held assumptions around risk, identity, and resilience. The US-headquartered Proofpoint has played a key role in shaping this global conversation, driving a people-first security model grounded in deep threat intelligence and real-world attacker behaviour.

In a recent conversation with Jatinder Singh, Editor, 9.9 Group, Bikramdeep Singh, Country Manager at Proofpoint India, shared insights on managing human-centric risk in an AI-driven landscape. He discussed the growing influence of agentic AI, the imperative to distinguish human versus AI-driven activity, and the capabilities security leaders must strengthen to stay ahead. He also outlined his mandate in India, driving strategy, expanding customer engagement, and advancing Proofpoint's vision for modern, human-centric cybersecurity.

"Agentic AI signals a massive shift; the relevant question changes from "What is the content?" to "Who created it: human or agent?" he stated.

Excerpts from the interview follow.

**CIO&Leader: With unprecedented change shaking up cybersecurity, how is Proofpoint reshaping its India strategy, and what major milestones are in sight?**

**BIKRAMDEEP SINGH:** Globally, we have a significant lead in email security, addressing what remains the single largest threat vector. We've expanded beyond email to data protection, stitching together a unified narrative that secures data throughout organizations across channels.

In India, we have served customers for years, but the sharpening of data sovereignty demands, especially from the BFSI, government, and IT/ITES sectors, has driven us to accelerate localization. We are establishing data centers in Mumbai, Chennai, and Bengaluru. By Q3, all Proofpoint solutions will fully comply with India's data residency regulations, empowering customers with both security and compliance peace of mind.

We have also launched a development center in Pune focused on product quality and support, while expanding sales and technical capabilities in Mumbai, Delhi, and Bengaluru. We project surpassing 200 people in India this year, reflecting robust demand fueled by rapid technology adoption.

**CIO&Leader: Your research reveals that nearly 90% of cybercrime losses stem from fraud, especially phishing. In an AI era,**

**why does human vulnerability endure, and what new defense models are shifting the balance?**

**BIKRAMDEEP SINGH:** Despite growth in firewalls, EDR, and SIEM, attackers leverage human weaknesses because they remain easier to exploit than hardened tech infrastructure. Email remains the prime attack vector by sheer volume.

Hybrid work complicates this further. Employees work beyond corporate perimeters using diverse communication tools, amplifying the attack surface. AI compounds the challenge: phishing has evolved from mostly content-based attacks to highly personalized, context-aware, multilingual AI-crafted messages, making social engineering vastly more convincing.

While security has evolved in many domains, EDR extending to XDR, analytics emerging from SIEM, email security, and DLP have lagged in transformation. Proofpoint's AI/ML-powered context analysis moves beyond legacy pattern or content filters to confront today's socially engineered threats with better efficacy.

**CIO&Leader: Many vendors claim AI superiority. How do you differentiate yourself through AI-powered prediction, automation, and user defense to stand out in this crowded market?**

**BIKRAMDEEP SINGH:** AI's success depends on data scale and quality. We analyze roughly one-third of global email traffic, granting deep training data unavailable to many competitors.

Running large-scale AI models is costly, and this limits many. Because of our focus on email and data security, we deploy advanced models efficiently and with precision, balancing scale and operational costs.

More importantly, our layered human-centric approach spans email, cloud, and endpoint channels, providing cross-channel visibility and protection. This wide lens, combined with deep AI models, sets us apart significantly.

**CIO&Leader: Which Indian sectors are adopting your tech fastest, and how is agentic AI changing their risk landscape?**

**BIKRAMDEEP SINGH:** Government remains a top priority at both the central and state levels, followed by BFSI, with IT/ITES supporting global clients also heavily engaged. These sectors are early adopters of generative and agentic AI, raising new security risks.

Agentic AI signals a massive shift; the relevant question changes from "What is the content?" to "Who created it: human or agent?" Proofpoint plays an integral role in identifying and curbing AI-driven threats by distinguishing human activity from agent activity.

"Data residency laws have increased the demand for autonomous data classification and protection. Our classification engine not only tracks initial tagging but also reclassifies data at exit points, mitigating leaks even if users manually alter classifications, a critical control for regulated sectors".

**CIO&Leader: Agentic AI promises automation gains, but many pilots falter. How do you see enterprise adoption evolving, especially for human-intensive tasks?**

**BIKRAMDEEP SINGH:** Agentic AI transcends prior automation approaches like RPA and generative AI by enabling autonomous decision-making, real-time adaptation, and complex workflow execution. It can act swiftly, generate benign or malicious outputs, and continuously improve using context.

This fundamentally transforms the threat landscape. Organizations must now distinguish human-driven behavior from agent-driven actions. Our AI models analyze intent, language patterns, and behavioral trends over time to detect subtle anomalies that traditional tools miss.

Though many pilots currently fail, the relentless pursuit highlights agentic AI's compelling value in terms of speed, cost reduction, and enhanced customer experience. Adoption will inevitably accelerate despite challenges.

**CIO&Leader: How are you building the partner ecosystem in India to facilitate broad technology adoption and support?**

**BIKRAMDEEP SINGH:** We are fully partner-driven in India. Cybersecurity success requires not just deployment but deep adoption, which demands skilled partners capable of unlocking the platform's full power.

Over the past quarter alone, we've trained over 100 partner engineers across India and SAARC, ensuring customers are supported locally for deployment, adoption, and scaling. Partners are critical as we expand reach and service quality across all regions.

**CIO&Leader: India faces diverse threats globally and locally. What unique attack patterns and security concerns do you observe?**

**BIKRAMDEEP SINGH:** India faces a full range of global threats, including phishing, BEC, QR-code phishing, telephone-oriented attacks (TOAD), and others. Given our vast population and enterprise base, we see all types of threats at scale. A distinctive demand emerging from Indian CISOs is risk-based security. Not all employees pose equal risk; developers, CXOs, and finance personnel face different threat profiles.

Proofpoint identifies three key personas: Very Attacked People (VAPs), Very Important People (VIPs), and Very Vulnerable People (VVPs). Focused resources on those overlapping profiles, e.g., frequently attacked and prone to clicking malicious links, can enhance defense efficacy significantly.

**CIO&Leader: In many of my conversations with CISOs, a recurring challenge is demonstrating the business value of cybersecurity investments. With evolving threats and complex risks, how do you help boards and business leaders translate human risk intelligence into clear, actionable metrics that justify cybersecurity spending and directly link to business outcomes?**

**"Our threat intelligence, based on the largest dataset, coupled with our AI/ML models, empowers customers to apply insights across their entire organization"**

**BIKRAMDEEP SINGH:** That's a fascinating question and a real problem organizations face. They want examples of attacks that could have threatened them, because the unique challenge with security is that nothing is truly breached until it is identified. Otherwise, you live in blissful ignorance. You don't grasp the loss until the data has already gone.

What we provide to organizations is a "health check" capability that identifies threats landing inside the organization—without disrupting ongoing business operations—that are not detected by existing tools. These threats aren't limited to phishing; they include BEC, TOAD attacks, and internal propagation, such as an employee sending a malicious email to a colleague. We identify these threats and present customers with specific details: if this got through and was clicked, here is the type of data at risk.

With our DLP capability, we not only identify threats but also assess which data may have leaked. Data valuation is critical. For example, a USB drive is only as valuable as the data it holds: a CEO's data has far greater cost impact than a movie or picture. Knowing how to value your data helps quantify threats and their associated costs.

"We provide metrics that help organizations understand the threats and their business impact, enabling them to develop holistic security strategies that address both insider and external risks"

Although we have expanded into DLP, cloud security, and insider risk management, our business remains heavily anchored in email protection.

**CIO&Leader: Many enterprises are moving towards zero trust and consolidated security platforms. How are you preparing to diversify without diluting your core strength?**

**BIKRAMDEEP SINGH:** I would answer that in two ways. First, Proofpoint is the most significant player globally in email security, a position we've held for a long time. On DLP, too, we lead globally by revenue according to Gartner, a fact not widely known in India due to our historic presence gaps.

The key is how we stitch together data from phishing and email security with DLP—email DLP, endpoint DLP, cloud DLP, and DSPM (data security posture management). This covers data categorization and monitoring across channels to prevent data exfiltration via USB, Bluetooth, email, and other vectors.

Our platformized and human-centric approach applies consistently across email, DLP, and all channels, making us unique for organizations seeking integrated, data-centric protection. While providers like Microsoft, Google, and other OEMs offer diverse products, we provide a laser-focused data-centric solution for the largest risk vector—people and their data handling. ■

# Join 28,000+ CIOs Who Power the Most Engaged Tech Leadership Community on LinkedIn.

Join the exclusive **CIO&Leader LinkedIn Group**-a vibrant community where IT leaders like **YOU** come together to connect, collaborate, and share insights. With engagement levels higher than all our leading competitors combined, this is the ultimate platform to keep you informed, inspired, and ahead of the curve.

Discover curated content, leadership Insights, and thought leadership tailored for today's CIOs. Be part of the conversations that matter, learn from industry pioneers and network with the best minds in the industry.

The CIO&Leader community is your gateway to thought-provoking dialogue, cutting-edge tech & trends and actionable strategies.

## Join the CIO&Leader LinkedIn Group today and elevate your leadership journey.

Follow us on LinkedIn

**@CIOandLeader**

Scan the QR Code to follow

You can also visit us at:
**www.cioandleader.com**

For more information, write to:
**editor@cioandleader.com**

RNI No. DELENG/2012/42736

17th Annual Conference

# CISO FORUM

India's Leading Cybersecurity Summit

# Thank You

for making 17th edition of The CISO Forum a Grand Success.

## Our Partners

GOLD PARTNERS

kaspersky · Microsoft · securiti

SILVER PARTNERS

CLAROTY · proofpoint · VERSA

ASSOCIATE PARTNERS

netskope · SentinelOne Secure Tomorrow

EXHIBIT PARTNERS

ANA Cyber Forensic · ARMIS · Barracuda Your business, secured. · Cotelligent A TechDemocracy Company · Cyber Vigilens SECURE • MONITOR • DEFEND · MORPHISEC · NeoSOFT · ProTechmanize · AQUILA I · SOPHOS · MAGNAMIOUS SYSTEMS CONNECTING AIMS TO SOLUTIONS · Trellix

MEDIA PARTNER

CNBC TV18

CONCEPT BY

CISOFORUM Security For Growth And Governance

#TheCISOForum