OCTOBER 2026
SUPPLEMENT OF CIO&LEADER

ET Edge INSIGHTS

Security For Growth And Governance

CISO FORUM

WHEN SHOULD CISOs TRUST AI TO ACT?

Four security leaders explain why visibility, accountability and proven safeguards must come before greater AI autonomy.



Abhijit Chakravarty
Executive Vice President – Networks & Cyber Security, Kotak Mahindra Bank

Pradipta Patro
CISO & VP, Polycab India

Sujoy Brahmachari
Chief Information Officer & Chief Information Security Officer, Smart Card Business, Rosmerta Technologies

Dr Pawan K. Sharma
Group Chief Information Security Officer (CISO) & Head of IT Operations, Tata Motors

ET Edge CISOFORUM

CISO PLAYBOOK 2026

Reset. Architect. Defend.

27-28 NOVEMBER 2026 | AAMBY VALLEY, LONAVALA

AI IS REWRITING THE ENTERPRISE.

WHO'S REWRITING THE SECURITY PLAYBOOK?

75+ SENIOR CISOS

- AI Security
- Cloud & Identity
- Data Protection
- Cyber Resilience

TWO IMMERSIVE DAYS

LEAD THE THINKING | DEMONSTRATE WHAT'S NEXT | DEEPEN THE CONVERSATION.

BEYOND THE NOISE. INTO THE CONVERSATION.

GET INTO THE PLAYBOOK. WRITE TO

Hafeez Shaikh
hafeez.shaikh@timesgroup.com
+91 9833103611

TECHNOLOGY PARTNER: opentext™ | ASSOCIATE PARTNER: f5 | EDUCATION PARTNER: BENNETT UNIVERSITY THE TIMES GROUP | BUSINESS INSIGHTS PARTNER: ET Edge INSIGHTS | AN EVENT BY: ET Edge | THE TIMES GROUP

ET EDGE - INDIA'S LEADING B2B INTELLIGENCE, COMMUNITY & MARKET PLATFORM

Trust Is Now the Attack Surface

For years, cybersecurity leaders have worked to eliminate trust from enterprise architecture. Zero trust was the discipline: never assume, always verify. Now, AI is forcing a harder question, because trust is returning through two doors at the same time.

The first is the AI we deploy. Security leaders believe agents belong in the enterprise, and many want them to handle routine first-line work. But few are willing to let them act unsupervised. Companies that moved fast to let agents change production systems have since pulled back.

The direction is clear: agents are easier to launch than to govern.

The second door is the AI others deploy against us. CERT-In's advisory on frontier AI-driven cyber risk warns that capabilities once confined to elite actors are reaching low-skilled attackers. The Indian Cyber Crime Coordination Centre has since cautioned banks and fintechs that deepfakes and synthetic identities are defeating facial authentication. A voice on a call, a face on a video, a biometric at onboarding: each was previously evidenced. None is anymore.

The common thread is that trust can no longer be inferred from appearance, whether it is a model output or a familiar face. It must be established through evidence and extended increments. That is a shift for a profession built on assurance: the question is no longer only whether a system is secure, but whether its outputs, identities and instructions deserve to be believed.

For Indian CISOs, three priorities follow. Begin by treating every agent as an accountable identity, with a named owner, a defined scope, and credentials that expire. Next, make the ability to halt AI a rehearsed capability; a kill switch described in a policy document is not a control until it has been tested. Finally, rebuild verification for people. Payments and privileged requests need out-of-band confirmation, and liveness checks should be designed on the assumption that attackers are using similar tools.

Trust is not a sentiment to be restored. It is a control to be designed. The CISOs who treat it that way will be the ones their boards will have confidence in. ■



"Trust is not a sentiment to be restored. It is a control to be designed."

R. Giridhar
Editorial Director
Enterprise Tech Publications
c-raja.giridhar@timesgroup.com

CONTENTS

OCTOBER 2026



COVER STORY

10-16 When Should CISOs Trust AI to Act?

Four security leaders explain why visibility, accountability and proven safeguards must come before greater AI autonomy.



Cover Design by:
Manish Kumar

Please Recycle This Magazine And Remove Inserts Before Recycling

COPYRIGHT - The views and opinions expressed in this magazine are those of the respective authors and do not necessarily reflect the views, policies, or positions of CISOForum, its publisher or editor. While every effort has been made to ensure the accuracy and reliability of the information provided, ET Edge Insights CISOForum makes no express or implied representations about the completeness, accuracy, reliability, suitability, or availability with respect to the content contained herein. ®CISOForum is a registered trademark of The Times Group

NEWS & VIEWS



06-09
Chinese hackers pose as India's tax department

SPECIAL FEATURE



17-18
Agentic AI is here, and India's data law isn't ready for it, warns Supreme Court advocate

INTERVIEW



19-22
Beyond the Model: Why Inference Is India's Real AI Infrastructure Test
By Jagrati Rakheja

OPINION



23-25
From Zero Trust to Zero Assumption:
By N. Radhika



26-27
Can India's DPDP framework keep up with the rise of autonomous AI?
By Abilash Soundararajan

INSIGHTS



28-35
McKinsey's 14 Tech Trends Set to Reshape Business in 2026

ET Edge INSIGHTS Security For Growth And Governance
CISOFORUM
www.cisoforum.in

MANAGEMENT
Chief Executive Officer: **Rohit Gopakumar**
Chief Operating Officer: **Joji George**
Senior Director: **Sachin Nandkishor Mhashilkar**

EDITORIAL
Editor: **Vikas Gupta**
Editorial Director: **R Giridhar**
Chief Editor: **Jatinder Singh**
Senior Correspondent: **Jagrati Rakheja**
Senior Correspondent: **Punam Singh**

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

SALES & MARKETING
Assistant Director-Brand & Community: **Vandana Chauhan**
Senior Manager - Marketing: **Rajiv Pathak**
Assistant Director - Projects: **Hafeez Shaikh**
Chief Manager - Projects: **Sidharth Ranga**
Senior Project Manager: **Supriya Sahoo**
Senior Project Manager: **Aanchal Gupta**

COMMUNITY ENGAGEMENT & DEVELOPMENT
Senior Manager - Community Development: **Dipanjan Mitra**
Chief Manager - Projects: **Neelam Adhangale**
Manager - Community Development: **Vaishali Banerjee**
Manager - Community Development: **Reetu Pande**
Manager - Community Development: **Snehal Thosar**
Executive - Database: **Dipti Gamre**

OPERATIONS
Senior Manager - Digital Operations: **Jagdish Bhainsora**
Senior Manager- Video Content: **Sunil Kumar**
Manager - Operations: **Amit Singh**

PRODUCTION & LOGISTICS
Manager - Operations: **Mahendra Kumar Singh**
Executive - Operations: **Vishwas Gandhi**

Careers – careers@et-edge.com
Subscriptions – editor.tech@timesgroup.com
Website - www.cisoforum.in

Press Registrar General of India Regn no - DELENG/2012/42736

Place of Publication – ET Edge, 7th Floor Sunteck Icon BKC, Mumbai Suburban, Maharashtra - 400098.
Email ID: editor.tech@timesgroup.com

Printed by Anup Kumar Dwivedi on behalf of printing press G. H. Prints Private Limited

Printed at G. H. Prints Private Limited, A-256 Okhla Industrial Area, Phase-I, New Delhi - 110020.

Owner – Times Strategic Solutions Ltd having Business Address at 7th Floor Sunteck Icon BKC, Mumbai Suburban, Maharashtra - 400098.

Legal and Compliance – Anis Kazi, General Manager, Legal

For subscriptions editor.tech@timesgroup.com

All disputes are subject to exclusive jurisdiction of Courts at Mumbai.



Mrinmoy Dey joins InCred Financial Services as CISO

Mrinmoy Dey has joined InCred Financial Services as Chief Information Security Officer (CISO). Previously, he served as CISO at Lendingkart, Paytm Money, and Paytm Insurance Broking. With over two decades of experience, he has held leadership and consulting roles at Paytm, Wipro, IBM, Altisource and Accenture, specialising in cybersecurity, governance, risk and compliance.



Akanksha Garg joins ProtocolRoots Private as CISO & DPO

Akanksha Garg has joined ProtocolRoots Private Limited as Chief Information Security Officer and Data Protection Officer. Previously, she spent over five years at HDFC Securities, most recently serving as Data Protection Officer and Head – Digital Products PMO. She also held product leadership roles at Reliance Securities and 63 moons technologies.



Prakash Patel joins Baroda Global Shared Services as CISO & Privacy Officer

Prakash Patel has been appointed Chief Information Security Officer and Privacy Officer at Baroda Global Shared Services Ltd. Previously, he served as AVP – IT Infrastructure and Security at Brain Payroll, leading information security, risk management and privacy compliance. He also held IT management and technical consulting roles at I-Serve Systems, Elitser Technologies and Microlink Solution.



Gajapathy Srinivasalu joins Adani Total Private as CIO & CISO

Gajapathy Srinivasalu has assumed the dual role of Chief Information Officer and Chief Information Security Officer at Adani Total Private Limited. Previously, he served as Head of IT at Adani Total and Adani Ports and SEZ. With over 23 years of experience, he specialises in IT strategy, digital transformation, infrastructure modernisation and cybersecurity.



Ashok Kumar Tiwari joins Vasai Vikas Sahakari Bank as DGM, CISO & DPO

Ashok Kumar Tiwari has taken on the role of DGM, CISO and DPO at Vasai Vikas Sahakari Bank Ltd. Previously serving as the bank's CISO, he brings experience in cybersecurity, IT infrastructure, digital banking, regulatory compliance and data protection, having earlier held leadership roles at Malad Sahakari Bank Limited and E4 Software Services.



Divan Raimagia joins Diageo as VP – Global Head of Cyber Resilience

Divan Raimagia has joined Diageo as VP – Global Head of Cyber Resilience, based in Bengaluru. Previously, he served as CISO at Adani Green Energy, leading cybersecurity across renewable energy operations. With extensive experience spanning cybersecurity architecture, OT security, risk management, and incident response, he has also held roles at HCL Technologies, IBM and Wipro.



Sunil Nishankar joins Ageas Federal Life Insurance as CISO & DPO

Sunil Nishankar has joined Ageas Federal Life Insurance as Chief Information Security Officer & DPO, based in Mumbai. He moves from Generali Central Life Insurance, where he served as CISO. Previously, he held security leadership roles at Genpact India, IBM India and Aujas Networks, bringing extensive experience in cybersecurity, risk management and data privacy.



Prafful Goel joins Indifi Capital as CISO

Prafful Goel has joined Indifi Capital as Chief Information Security Officer. He moves from Hero FinCorp, where he led Infosec GRC and risk compliance. Previously, he held leadership roles at Airtel Payments Bank and Paytm Payments Bank, with experience spanning cybersecurity, IT audits, risk governance, regulatory compliance, and digital transformation.

Chinese hackers pose as India's tax department

A China-aligned group uses fake Income Tax notices to plant spyware on Indian businesses during ITR season.

By **CISO Forum** | editor.tech@timesgroup.com

Cybersecurity firm Seqrite says it has uncovered a sophisticated espionage campaign impersonating India's Income Tax Department and targeting businesses during ITR filing season, putting the country's financial backbone at risk.

A fake tax notice as bait

The operation, dubbed "DragonReturn" by Seqrite Labs, begins with a phishing email that looks like an official notice from the Ministry of Finance. It carries the Government of India emblem, bilingual Hindi-English text, fake reference numbers and even citations of real sections of the Income Tax Act to appear legitimate. The goal is to pressure recipients into acting fast.



From fake utility to full access

Victims are pushed toward a ZIP file disguised as the government's own offline return-filing tool. Once opened, the malware quietly installs itself as a fake "Windows Mixed Reality Service," burying its components inside trusted system folders to avoid detection. From there, it deploys a remote access trojan that can harvest system data, security software details, and user information, sending it back to attackers over encrypted channels. A second tool can capture and compress screenshots, underscoring the campaign's intent to steal sensitive data.

Who's behind it

Seqrite researchers say infrastructure clues, Chinese-language admin panels and overlaps with known attack patterns point to a China-aligned threat group. The targeting of India's tax ecosystem, combined with the campaign's stealth and persistence, reinforces the assessment that the attackers are seeking long-term access to financial and taxpayer records rather than a quick payday.

Who's at risk

The campaign is aimed at corporate finance and accounts teams, chartered accountants, tax filing agents, government contractors and individual taxpayers—essentially anyone likely to be dealing with tax paperwork right now.

What Seqrite recommends

The company is urging people to verify any tax notice directly through official Income Tax Department channels, avoid clicking links or downloading files from unsolicited emails, and treating any "government utility" download with suspicion. For enterprises, Seqrite points to its own threat-monitoring and data-privacy tools as added layers of defence during high-risk filing periods. ■



Bot defence now separates trusted AI agents from bad bots

New device intelligence and agentic AI detection help enterprises stop fraud without blocking legitimate customers or authorised AI agents.

By **CISO Forum** | editor.tech@timesgroup.com

Application security company F5 (NASDAQ: FFIV) has announced enhancements to its F5 Distributed Cloud Bot Defense, adding device intelligence and specialized protections for agentic AI. The upgrades are designed to help enterprises prevent automated fraud and account abuse without blocking legitimate customers or authorized AI agents.

Why Old Bot Tools Fall Short

AI agents are becoming a common way to interact with websites, mobile apps, and customer portals. Unlike basic crawlers or scripts, they carry out multi-step tasks on behalf of users, such as making purchases, booking travel, or handling banking operations. They also talk directly to APIs at machine speed.

F5 says first-generation bot management, built on rigid "bot or not" rules, struggles in this setting. Static, request-level controls cannot tell an authorized AI agent from a malicious bot, and blunt measures like IP blocking or CAPTCHAs risk shutting down revenue-generating AI channels.

A Multi-Signal Approach

Instead of inspecting single requests, F5 correlates behavioral, device, and client-integrity data across interactions. This builds trust in critical login, web, and API connections. The capability is part of F5's web application and API protection (WAAP) offerings, integrated with the F5 Application Delivery and Security Platform (ADSP).

"The answer is not to block AI," said Kunal Anand, Chief Product Officer at F5. He added that organizations need to understand which agents and devices can be trusted and what they are trying to do.

Staying protected

- Persistent device identification: Tracks devices across sessions and accounts, exposing multi-account access, credential stuffing and account takeover.
- Real-time device risk scoring: Checks client integrity signals to spot emulators, device spoofing and tampering.
- Risk-based enforcement: Lets customers choose to allow, challenge, rate-limit or block, reduce CAPTCHA friction and false positives.
- Agent-aware policy framework: Manages humans, trusted AI agents and malicious bots under one policy on F5 ADSP.

Availability

The new agentic AI protections are available now. Device intelligence is expected to become available to a limited number of F5 Distributed Cloud customers in the fourth quarter of calendar year 2026, with broader availability to follow over the coming months. ■

New AI security tool targets smaller firms struggling to govern AI



Barracuda's AI Data Security gives resource-strapped businesses and MSPs a simple way to govern AI, protect data and prove compliance.

By **CISO Forum** | editor.tech@timesgroup.com

Barracuda Networks has launched Barracuda AI Data Security, an AI security and governance solution designed for resource-constrained businesses and MSPs. The platform aims to help organisations adopt AI securely while protecting sensitive data and demonstrating compliance. The launch addresses growing governance challenges, with nearly half of senior IT leaders reporting inadequate skills to govern AI securely, according to Barracuda Research. IBM's 2026 Cost of a Data Breach Report found that 68% of breached organisations lacked policies to manage AI or detect shadow AI. Neal Bradbury, Chief Product Officer, Barracuda, highlighted the need for practical AI security measures.

What it does

Built on the BarracudaONE platform and powered by the Barracuda IQ AI engine, the solution covers more than 1,300 generative AI tools, including Copilot, ChatGPT, Claude, and Gemini.

- Data protection: Blocks customer data, passwords, financial records and proprietary code from reaching GenAI services, with no initial configuration.
- Threat detection: Stops prompt injection, jailbreak attempts and hate speech, aligned to the OWASP LLM Top 10.
- Audit trail: Captures every AI interaction, with one-click exports for insurers and auditors.
- Flexible policies: Sets rules by team, tools and data type.

Designed for smaller teams

Most AI security offerings are designed for large enterprises, leaving smaller firms facing complexity, skill gaps, and high costs. This solution deploys in minutes, uses familiar firewall-style workflows, and offers guided setup through the Bailey AI assistant. MSPs get centralized multitenant management and bundled pricing.

Geoff Turner, CEO of Elevate Technology Group, said businesses need controls that are "simple, scalable and built for real-world environments."

The India view

Parag Khurana, Country Manager, India at Barracuda, said, "In a market where AI is becoming a competitive advantage, having the right foundations in place will be critical to turning ambition into long-term business value." ■

NETSCOUT launches AI copilot to speed up network problem-solving

NETSCOUT's nGenius Copilot lets IT teams ask plain-language questions and get evidence-backed answers on service disruptions

By **CISO Forum** | editor.tech@timesgroup.com

NETSCOUT, introduced nGenius Copilot, bringing conversational AI capabilities to its nGeniusONE platform. The tool aims to help teams investigate service disruptions, identify underlying issues and assess their impact on business operations.

How it works?

The Copilot uses NETSCOUT Smart Data, drawing on network, application, service and user interactions to answer questions in natural language. It provides explanations, visualizations', supporting evidence and follow-up suggestions, helping users identify slow applications, affected users or locations, and potential network, server or external factors behind performance issues.



Why it matters?

Hybrid environments and growing digital dependencies mean IT teams must sift through ever larger volumes of operational data. Connecting evidence across applications, infrastructure, networks and external services is slow work, often left to a few senior engineers.

"Reliable AI starts with reliable data," said Phil Gray, AVP, product management at NETSCOUT. He said the tool pairs a conversational experience with the depth of Smart Data, helping teams find evidence quickly, understand business impact and work from a consistent view of what happened.

Key benefits

According to the company, nGenius Copilot helps organizations:

- Identify problems faster: Trace an issue from the first symptom to the applications, services and dependencies behind it.
- Prioritize business impact: See which users, locations, services, and transactions are affected.
- Scale expertise: Make Smart Data accessible to more staff while helping senior engineers work efficiently.
- Improve collaboration: Turn technical evidence into explanations that network, application, infrastructure, and business teams can share.
- Get more from existing deployments: Bring Smart Data into daily workflows.

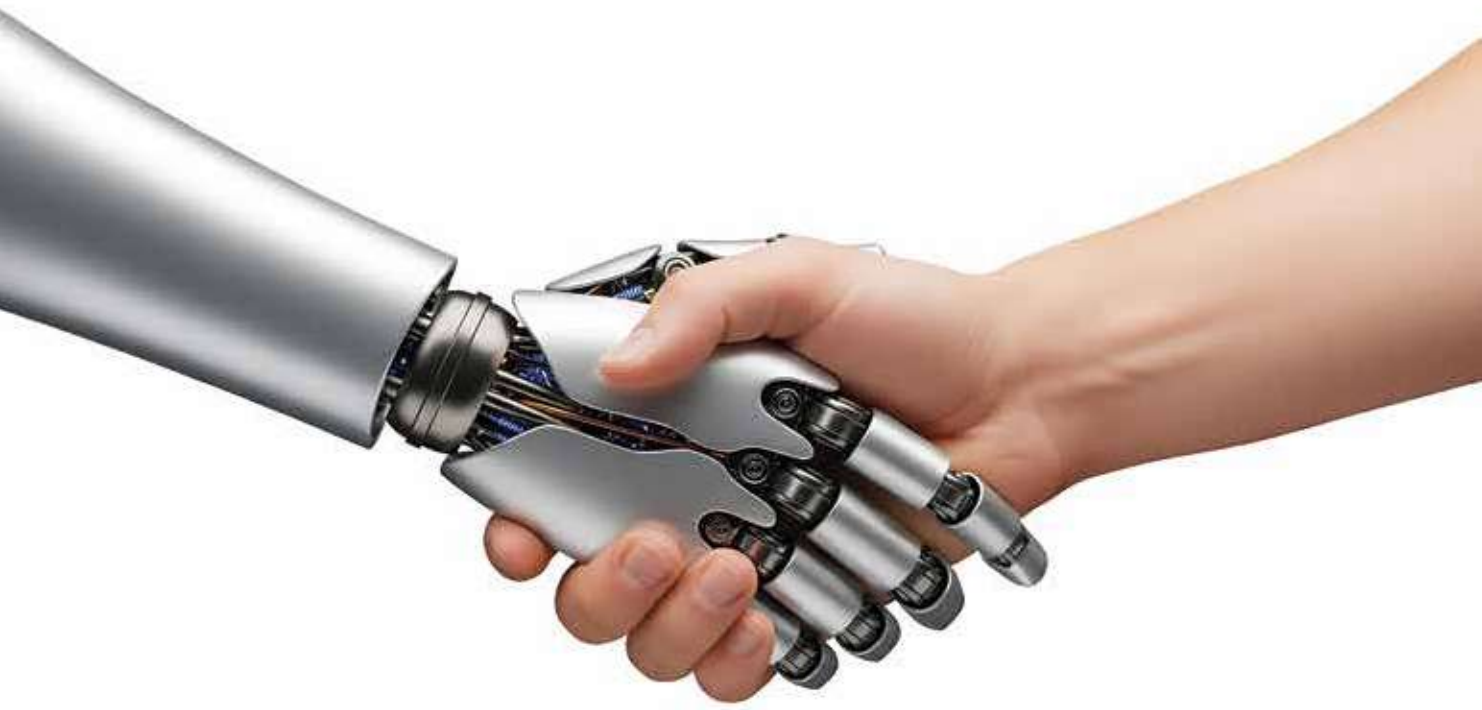
Our take

The launch reflects a broader industry push to place conversational AI atop operational data. NETSCOUT's pitch rests on data quality: its Smart Data is derived from packet-level observation of network traffic. The company says this gives its AI a trustworthy base, though buyers will want to test how well the answers hold up in real incidents. ■

WHEN SHOULD CISOs TRUST AI TO ACT?

Four security leaders explain why visibility, accountability and proven safeguards must come before greater AI autonomy.

By **R. Giridhar & Jagrati Rakheja** | editor.tech@timesgroup.com





For years, the first question about any enterprise system was simple: Is it running? For AI that acts, that is no longer enough. An agent can be available, fast and confident, yet still be functionally wrong—even while holding credentials for critical systems. This shifts the central question from whether a system is operational to whether it can be trusted to act.

That shift runs through this month's cover story. As AI moves from generating insights to invoking tools and triggering actions, four security leaders from banking, manufacturing and automotive arrive at the same conclusion from different directions: autonomy is not a feature to switch on. Authority must be earned, and observability is how it is earned.

Autonomy is a ladder, not a switch

For Abhijit Chakravarty of Kotak Mahindra Bank, cyber operations mature in stages: detect, diagnose and remediate. Even the final stage is a journey, progressing from reporting to recommending and, eventually, autonomous action. The system must first demonstrate that its findings are accurate and consistent, while the organisation must understand the context, consequences and safeguards involved.

His position is unambiguous: "Until I have complete observability, there should be no automatic remediation." He also emphasises the need to measure business value, suggesting that AI and AIOps deployments be evaluated on speed, accuracy, automation, analytics and scale.

Governance moves to runtime

Pradipta Patro of Polycab India notes that traditional monitoring was designed for deterministic software. Large language models are probabilistic, meaning their behaviour can change even when a service appears healthy.

His first recommendation is organisational: every AI initiative needs a business owner and a clearly defined objective. Governance must begin when success is defined, not after go-live. His second is operational: organisations need a tested kill switch, with clear decisions about who can invoke it, what evidence warrants intervention and how the process can continue safely.

These procedures should be rehearsed through tabletop exercises rather than left untested until the first incident. As Patro puts it, "An enterprise cannot govern what it cannot see."

Protection before novelty

In regulated manufacturing, the margin for error is even narrower. Sujoy Brahmachari of Rosmerta Technologies works with payment cards and other secure products, where certification requirements and zero tolerance for data leakage are non-negotiable.

AI can improve predictive maintenance, quality, shipment controls and resilience, but only within a secure operating

environment. Thresholds for human review, rollback and suspension, he argues, should reflect the potential consequences of an error.

The principle is straightforward: "In secure manufacturing, information protection must come before technological novelty."

Seeing is not the same as acting

Dr Pawan K. Sharma of Tata Motors draws a clear distinction between visibility and control. These must remain architecturally separate: an agent may observe operational conditions without automatically gaining the authority to change them.

Every agent needs its own non-human identity, least-privilege access and, preferably, short-lived credentials. In industrial environments, agents should access plant data through intermediaries such as historians and gateways rather than directly accessing programmable logic controllers (PLCs) or supervisory control and data acquisition (SCADA) systems.

In multi-agent environments, traceability must extend across the entire chain, supported by immutable audit trails. For Sharma, observability "becomes the foundation for determining how much authority machines can safely exercise."

The control plane for autonomy

The practical challenge is to translate observability into enforceable controls. Knowing what an agent is doing is only the first step; organisations must also establish what it is permitted to do, under which conditions and within what limits. This requires visibility into the agent's identity, the data it accesses, the tools it invokes, the decisions it makes and the systems it affects..

In May 2026, Gartner predicted that by 2028, 40% of organisations deploying AI will use AI observability to monitor model performance.

A well-governed environment should also distinguish between actions that can be automated and those that require human approval. Low-risk, reversible tasks may be allowed to proceed within predefined boundaries, while actions involving sensitive data, financial transactions or critical infrastructure should trigger additional checks. Every intervention, override and exception must be recorded and traceable.

Crucially, these controls cannot remain confined to the initial deployment. Models, data, permissions and business contexts evolve. Continuous monitoring and periodic reassessment are therefore essential. The objective is not to prevent AI agents from acting, but to ensure that their authority remains proportionate to their demonstrated reliability and the consequences of their decisions.

The research points in the same direction

Recent studies suggest that enterprises are grappling with this gap. Dynatrace's Pulse of Agentic AI 2026, a survey of 919 senior leaders, found that about half of agentic AI projects remain at the proof-of-concept or pilot stage, while humans still verify 69% of agentic decisions.

In May 2026, Gartner predicted that by 2028, 40% of organisations deploying AI will use AI observability to monitor model performance.

"AI is everywhere, but most organizations are still figuring out how to monitor and trust these systems," said Padraig Byrne, VP Analyst at Gartner. "That visibility gap makes scaling risky, and that's why observability matters. Unlike traditional software, AI's decision-making is often hidden, making it difficult to explain or trust. Yet errors can result in substantial financial losses, reputational damage and regulatory scrutiny."

In India, Okta polled attendees at its AI Identity Summits in Mumbai and Bengaluru. Only 14% reported having full visibility into AI agents and non-human identities. Although this was an informal poll rather than a rigorous survey, the finding is still noteworthy.

Standards are also beginning to take shape. On 17 February 2026, the National Institute of Standards and Technology (NIST) announced its AI Agent Standards Initiative, a multiyear effort. Until such standards mature, enterprises must establish and enforce their own rules.

The message from security leaders is consistent: the question is no longer simply whether an AI agent can perform a task. It is whether the organisation can see, understand and govern every action it takes. In an increasingly autonomous enterprise, trust will depend not on how much authority AI is given, but on how effectively that authority is controlled.

No Automatic Remediation Without Complete Observability

Automatic remediation requires complete observability, demonstrated accuracy and a gradual expansion of trust.

AI-enabled cyber operations should progress from detection and diagnosis to reporting and, eventually, remediation. Automatic action is unsafe without complete observability, demonstrated accuracy and appropriate checks and balances.

I describe the maturity journey as detect, diagnose and remediate. However, the final stage is itself a progression, moving from reporting and recommending to autonomous remediation. The system must first demonstrate that its findings are accurate and consistent. Only when the organisation understands the context, consequences and safeguards should remediation become increasingly automated or autonomous.

Reliability cannot be reduced to whether a service is available, or a model returns an answer. A technically healthy AI system can still be functionally wrong. Enterprises need visibility into the data used, the models and tools invoked, the permissions exercised, the execution path and the effects of the resulting decisions. Without this chain of evidence, accountability remains incomplete.

The practical priority is to make this evidence visible. The Detect, Diagnose and Remediate (DDR) maturity model provides a framework for progressing towards greater autonomy. I would also use a SAAAS approach—Speed, Accuracy, Automation and Analytics, and Scale—to evaluate AI and AIOps deployments and establish their return on investment. These dimensions should be assessed as part of a single operating picture so that technical performance, decision quality and business consequences are not considered in isolation.



"Until I have complete observability, there should be no automatic remediation."

Abhijit Chakravarty
Executive Vice President – Networks & Cyber Security, Kotak Mahindra Bank

Governance must operate continuously, not just during approval processes or audits. Thresholds for human review, escalation, rollback and suspension should reflect the potential consequences of an error. Model drift, hallucinations, anomalous behaviour and changes in data must be detected early enough to prevent harm. The objective is not control for its own sake, but confidence that AI continues to serve its intended purpose.

For enterprise leaders, the priority is to make this a repeatable management practice. This means defining the purpose and baseline before deployment, continuously monitoring performance and outcomes, assigning accountability for exceptions, and expanding authority only when evidence supports it. Establishing a pragmatic cyber-risk threshold for moving from insight to action is essential to making AI a dependable part of enterprise operations.

This also clarifies what assurance should mean in practice. Governance approval records intent at a particular point in time; continuing assurance requires evidence that data, models, permissions, outputs and actions remain within their expected boundaries. Business, technology, risk and security leaders therefore need shared thresholds for review, escalation and intervention.

The aim is not to eliminate uncertainty from probabilistic systems, but to make that uncertainty visible, contain its consequences and preserve accountability as AI systems evolve in production.

AI Governance Needs a Runtime Kill Switch

Effective AI governance requires early business ownership, continuous visibility and the ability to intervene at runtime.

Most enterprises are still learning what AI observability requires. Traditional monitoring was designed for deterministic applications, where the same inputs and operating conditions should broadly produce the same results. Large language models are different. Their outputs are probabilistic, multidimensional and sensitive to changing context. As a result, output quality and behavioural baselines can shift even when a service appears healthy.

The first response should be organisational rather than technical. Every AI initiative needs a clear business owner and a defined objective. If responsibility rests solely with IT or security, a project may demonstrate an interesting capability without proving that it improves a business process. Governance must therefore begin when the use case and success measures are defined, rather than being introduced after deployment.

Runtime control is equally important. Enterprises need more than a theoretical incident-response process; they need a tested way to stop an AI process when its behaviour becomes unsafe. Depending on the use case, the window for intervention may be extremely short. A kill switch, supported by meaningful thresholds and clearly assigned decision-making authority, can prevent an anomalous action from becoming a material business incident.

Synthetic monitoring can help test known scenarios, while continuous observation can identify hallucinations, shifts in established baselines and other unexpected behaviour. However, no single tool or dashboard is sufficient. Visibility must encompass the business objective, security posture, model behaviour, data flows and resulting outcomes.



"An enterprise cannot govern what it cannot see."

Pradipta Patro
CISO & VP, Polycab India

The governing principle is simple: an enterprise cannot govern what it cannot see. Observability is not a separate workstream running alongside governance; it provides the evidence that allows governance to operate effectively. Building this evidence into the first phase of an AI programme makes subsequent scaling safer and more credible.

A kill switch is useful only when the organisation has decided who can invoke it, what evidence justifies intervention and how the affected process will continue safely. These decisions should be rehearsed through regular tabletop exercises rather than left until the first incident.

This preparation also helps distinguish a model-quality issue from a data, security or workflow failure. For CIOs and CISOs, runtime governance combines visibility with operational readiness: the ability to recognise harmful behaviour, contain it quickly and learn from the incident before authority is restored.

Lessons from each intervention should inform future thresholds, controls and deployment decisions. Governance becomes stronger when these lessons feed back into system design.

- **Own the Use Case** - Assign a business owner and clear object.
- **Enable Runtime Control** - Be ready to stop an AI process if it turns unsafe.
- **Ensure End-to-End Visibility** - Monitor outputs, behaviour and data flows.
- **Plan and Rehearse** - Set thresholds, deside who can act, and run drills.
- **Learn and Improve** - Use incidents to refine controls.
- **Make Governance Part of the Journey** - Build observability into the first phase.

For Secure Manufacturing, AI Begins with Zero Trust in Data Movement

Secure manufacturing requires uncompromising data protection alongside operational resilience.

In regulated and secure manufacturing, AI must be built on uncompromising information protection and security by design. Its value lies in predictive operations, business continuity, quality and shipment controls, but certification requirements and data-leak prevention remain non-negotiable.

Manufacturing payment cards and other secure products demands that information protection take precedence over technological novelty. Security controls and applicable certifications must be built in from day one, with zero tolerance for data leakage. AI can still improve predictive maintenance, quality, shipment controls, resilience and back-office operations, but only within that protected operating environment.

Reliability cannot be reduced to whether a service is available, or a model returns an answer. A technically healthy AI system can still be functionally wrong. Enterprises need visibility into the data used, the models and tools invoked, the permissions exercised, the execution path and the effects of the resulting decisions. Without this chain of evidence, accountability remains incomplete.

The practical priority is to make this evidence visible. Relevant areas include payment-card certification, high-security number plates, predictive machine notifications, and quality, shipment and back-office applications. Governance must operate continuously, not just during approvals or audits. Thresholds for human review, escalation, rollback and suspension should reflect the potential consequences of an error. Model drift, hallucinations, anomalous behaviour and changes in data must be detected early enough to prevent harm. The objective is not control for its own sake, but confidence that AI continues to serve its intended purpose.



"In secure manufacturing, information protection must come before technological novelty."

Sujoy Brahmachari
Chief Information Officer & Chief Information Security Officer, Smart Card Business, Rosmerta Technologies

For enterprise leaders, the priority is to make this a repeatable management practice. This means defining the purpose and baseline before deployment, continuously monitoring performance and outcomes, assigning accountability for exceptions, and expanding authority only when evidence supports it.

Secure manufacturing offers a distinctive perspective on AI adoption because it combines regulated processes, physical products and stringent security certification requirements. This makes the discipline of protecting information central to ensuring that AI becomes a dependable part of enterprise operations.

This also clarifies what assurance should mean in practice. Governance approval records intent at a particular point in time; continuing assurance requires evidence that data, models, permissions, outputs and actions remain within their expected boundaries. Business, technology, risk and security leaders therefore need shared thresholds for review, escalation and intervention.

The aim is not to eliminate uncertainty from probabilistic systems, but to make that uncertainty visible, contain its consequences and preserve accountability as AI systems evolve in production.

Observability Is the Foundation for Trusted AI Autonomy

Connecting AI actions with security, accountability, control and measurable business impact.

As AI agents move from generating insights to invoking tools, accessing operational data and triggering actions, enterprises face a fundamental challenge: maintaining visibility and control while allowing these systems to operate with increasing autonomy.

This is where business observability becomes critical. It must go beyond monitoring whether an application or AI agent is available or performing as expected. Enterprises need visibility into what an agent is doing, the resources it accesses, the decisions and actions it takes, and how these affect operations and business outcomes.

That starts with identity and accountability. Every agent should have a distinct non-human identity, least-privilege access and appropriately scoped, preferably short-lived, credentials. Tool access should be restricted to what the task requires, with inputs and outputs validated before execution. Sensitive or high-consequence actions should require stronger authorisation or human approval. This creates a clear chain of accountability from intent to execution.

The challenge deepens in industrial environments. IT and operational technology cannot be treated as a single trust domain. Agents should generally access operational information through controlled intermediaries, such as historians, gateways or replicated data layers, rather than directly accessing PLCs, SCADA systems or other critical control infrastructure.



"As enterprises move towards greater autonomy, observability becomes the foundation for determining how much authority machines can safely exercise."

Dr Pawan K. Sharma
Group Chief Information Security Officer (CISO) & Head of IT Operations, Tata Motors

Visibility and control should also remain architecturally distinct. An agent may observe operational conditions without automatically receiving authority to change them. Where write-back or actuation is required, stronger controls, separate gateways, approval mechanisms and limits on the potential impact of an error become essential.

Observability provides the connective layer. Enterprises need end-to-end traceability across agent identity, data access, tool calls, actions and downstream consequences. In multi-agent environments, this must span the entire chain rather than stop at individual interactions. Behavioural monitoring can identify deviations from expected activity, while immutable audit trails provide evidence for accountability, investigation and regulatory requirements.

Most importantly, visibility must connect technology behaviour with business impact. An anomalous agent action matters not simply because a technical threshold has been breached, but because it can affect operational continuity, security, compliance and critical business processes.

As enterprises move towards greater autonomy, observability becomes more than monitoring. It provides the foundation for determining how much authority machines can safely exercise, maintaining accountability as that authority expands, and keeping autonomous systems within defined business and operational boundaries.

Agentic AI is here, and India's data law isn't ready for it, warns Supreme Court advocate



Pavan Duggal warns India's DPDP Act isn't ready for agentic AI, urging enterprises to prepare before November 2026 enforcement.

By **CISO Forum** | editor.tech@timesgroup.com

At the 27th ET Edge CIO&Leader annual conference in Jaipur, Dr Pavan Duggal, Senior Advocate at the Supreme Court of India and Architect of Global AI Accountability, warned that India's Digital Personal Data Protection (DPDP) Act may not be ready for the rise of agentic AI.

The law wasn't built for agents

Duggal's core warning: the DPDP Act still assumes a human hand behind every data decision. AI agents, by contrast, act on a single initial command and then operate independently, dividing tasks, processing data, and making decisions with no further human input. Consent, he noted, was given only the original instruction, not for what the agent does next. Whether that consent can even be withdrawn, mid-task remains unresolved, pending clarity from the Data Protection Board.

When the agent becomes the fiduciary

Under the DPDP framework, obligations fall heavily on the "data fiduciary", the entity deciding how personal data is used. Duggal argued that in agentic AI deployments, the agent itself effectively becomes the fiduciary, shifting the practical burden beyond the law's original design. Compounding this, when an AI agent hallucinates and causes harm, liability is already trending toward the enterprise that deployed it rather than the technology vendor.

CIOs and CISOs in the crosshairs

Non-compliance under DPDP can trigger penalties up to ₹250 crore, and Duggal predicted a wave of job losses among CIOs and Data Protection Officers, who he said would be the first to be held accountable when boards demand due diligence. He flagged a widening risk map for enterprises: shadow agents deployed without security or DPO review; non-human identity sprawl; data exfiltration via AI tools; vendor-agent opacity; unmapped cross-border data flows; and gaps in audit trails.



"The DPDP Act still assumes a human hand behind every data decision."

A ten-point survival playbook

Duggal closed with concrete guidance for enterprises: treat every agent's action as the organisation's own act, since courts will; inventory every AI agent, including those embedded in vendor SaaS tools; tier agents by autonomy and data sensitivity; keep a human-in-the-loop gate at every consequential action; build consent scope directly into agent design; give agents least-privileged, short-lived identities; log every action immutably; install and test kill switches before they're needed; renegotiate vendor contracts for liability and logging obligations; and rehearse for an eventual agentic AI breach.

The clock is ticking

With DPDP enforcement beginning November 13, 2026, and jurisdictions from Singapore to the broader global community moving toward agentic AI liability frameworks, Duggal urged Indian enterprises to prepare now rather than wait for legislative clarity. His message was unambiguous: the agentic AI era, he said, is not waiting for anyone's compliance program, and the only real choice organizations have is to get onto the bandwagon. Carefully, before the "bad day" arrives. ■



Beyond the model: Why inference is India's real AI infrastructure test

Akamai's Mitesh Jain says inference is India's AI bottleneck, making edge infrastructure key to latency, cost and DPDP compliance.

By **Jagrati Rakheja** | jagrati.rakheja@timesgroup.com

As Indian enterprises race to operationalize AI, the conversation is shifting from which models to deploy to whether the underlying infrastructure can support them at scale. Training may capture headlines, but it is inferencing the continuous, real-time decision-making behind every fraud alert, chatbot response, and personalized recommendation that is quietly becoming the industry's toughest infrastructure challenge. With over 80% of AI's computational demand now tied to inference, latency is no longer just a user-experience metric; it is a direct cost and compliance variable, especially in a price-sensitive, regulation-heavy market shaped by the DPDP Act.

In this conversation, Mitesh Jain, Regional VP, Akamai India, unpacks why distributed, edge-first architectures are becoming non-negotiable for Indian enterprises, how sectors like BFSI and fintech are pulling ahead, and what it will take beyond talent and models for India's AI ambitions to translate into real, scalable ROI over the next three years.

CISO FORUM: Most of India's AI conversation has centered on models and use cases so far. Why do you believe inference, not training, is now the real infrastructure bottleneck for Indian enterprises?

MITESH JAIN: The biggest challenge for Indian enterprises is no longer building or accessing AI models, but running AI applications that deliver fast, secure, and cost-effective experiences at scale. Every AI-powered customer interaction, security decision, recommendation, or automation workflow depends on inference, which runs continuously and places far greater demands on infrastructure than model training does. As more organizations put AI into production, they need infrastructure that can handle high volumes of real-time inference without compromising performance or user experience.

The focus now should be on bringing cloud computing closer to where data is generated and consumed. A distributed cloud and edge approach helps reduce latency, improve application performance, optimize costs, and support data residency requirements while delivering consistent experiences across diverse geographies. Enterprises that invest in inference-ready infrastructure will be better prepared to run AI reliably and deliver meaningful outcomes for their businesses and customers.

CISO FORUM: When you talk to Indian enterprises today, how many are running AI in production versus still stuck in pilot mode? What's the single biggest thing holding that transition back?

MITESH JAIN: Based on our conversations with customers, we are seeing a clear shift from experimentation to production. The conversation today is less about whether to adopt AI and more about how to run AI reliably, securely, and efficiently at scale.

The biggest challenge is no longer access to AI models but having the infrastructure and operational readiness to support them in production. Organizations need low-latency, predictable performance, strong security, and the ability to handle growing inference workloads without driving costs. Security readiness also remains a key barrier. In Akamai's 2026 Partner Summit survey, only 17% of respondents said organizations are very well prepared to manage AI security risks, while 70% identified data privacy and leakage as the biggest AI-related security concern. These findings highlight that scaling AI requires not just compute, but also robust governance, security, and resilient infrastructure. A distributed cloud and edge approach helps bring compute closer to users and data, enabling organizations to run AI applications more efficiently while delivering a consistent user experience.

CISO FORUM: The government has committed over ₹10,000 crore under the IndiaAI Mission to build domestic compute and GPU capacity. Does this solve the infrastructure gap, or does it address only one part of a much larger problem?

MITESH JAIN: The IndiaAI Mission is an important step in strengthening India's AI ecosystem and expanding access to domestic compute capacity. Greater availability of GPUs will help lower barriers to AI adoption and reduce dependence on overseas infrastructure. However, computing is only one part of the equation.

To run AI effectively in production, organizations also need infrastructure that can deliver low latency, move data efficiently, support security and compliance requirements, and bring inference closer to users. AI applications are increasingly expected to respond in real time, making distributed cloud and edge infrastructure just as important as centralized compute. India's AI ambitions will depend not only on building more compute capacity, but also on creating an infrastructure ecosystem that can support AI workloads reliably, efficiently, and at scale.

CISO FORUM: Where does distributed, edge-based infrastructure have a genuine advantage over centralized cloud for AI workloads, and where is centralized cloud still the more sensible choice?

MITESH JAIN: Both centralized cloud and distributed edge infrastructure have important roles to play, and the right choice depends on the workload. Centralized cloud remains the preferred environment for model training, large-scale data processing, and workloads requiring significant compute resources. It offers the scale and flexibility needed for these compute-intensive tasks.

Distributed edge infrastructure becomes critical when AI applications need to make real-time decisions. Use cases such as AI assistants, fraud detection, content delivery, security, and industrial automation require low latency and consistent performance, which are best achieved by bringing inference closer to users and data. A hybrid approach that combines centralized cloud training with distributed infrastructure for inference gives enterprises the flexibility to optimize performance, cost, and user experience while meeting data residency and compliance requirements.

CISO FORUM: Latency is often framed as a user-experience issue, but you've linked it directly to economics. Can you unpack the actual cost implications of running AI inference at scale in a price-sensitive market like India?

MITESH JAIN: Latency is often treated as a user experience metric, but it directly impacts the cost of running AI at scale. More than 80% of AI computational demand comes from inference, making it the largest contributor to ongoing infrastructure costs. Every millisecond of delay consumes more compute resources, and when this is multiplied across millions of inference requests, the cost of delivering AI can increase significantly. For organizations in a price-sensitive market like India, improving inference efficiency is as much a business priority as it is a technical one.

"AI success is no longer determined solely by the quality of the model."

The focus, therefore, should not be on adding more compute alone, but on using infrastructure more efficiently. Reducing latency, increasing throughput, and optimizing resource utilization can significantly lower inference costs while improving application performance. Our own benchmarking has shown that the right infrastructure choices can reduce the cost per inference while delivering better performance, allowing organizations to support more AI workloads, improve service quality, and make better use of their infrastructure investments.

CISO FORUM: Which Indian sectors are you seeing move fastest from AI experimentation to real, revenue-linked deployment, and what's common to those that are pulling ahead?

MITESH JAIN: In India, the sectors that are seeing the fastest move from AI experimentation to real revenue are the fintech and BFSI domains. With one of the highest AI adoption rates in the fintech sector, India has emerged as an innovation hub, fostering solutions such as fraud detection, customized financial products and real-time decision-making. In e-commerce, the deployment of conversational AI drives conversion and purchase decisions.

The common factor in pulling ahead is the shift in infrastructure. Enterprises are now looking at AI through the lens of a decentralized functional challenge rather than a centralized data project. They realized that if an AI application takes even a few seconds to show results, ROI drops and the end consumer interest vanishes. Hence, industry bodies are focusing mainly on the Inference at the Edge model, in which applications can deliver results that drive customer value rather than merely store data in a single repository.

CISO FORUM: What happens when an enterprise scales its AI ambitions faster than its underlying infrastructure can support? Have you seen real examples of this backfiring in India?

MITESH JAIN: When AI ambitions surpass the infrastructure, companies hit a roadblock called "The AI Performance Wall". On a small scale, the enterprise AI models function smoothly. But when the same is deployed to large numbers of users simultaneously, as in fraud detection and risk evaluation, centralized systems collapse. As a result, latency spikes, bottlenecks occur, and the user's system performance drops.

In India, we're seeing this challenge emerge across sectors as AI moves from pilot projects to mission-critical production deployments. Financial institutions using AI for real-time fraud detection, credit risk assessment, and customer service need inference decisions in milliseconds. Even small increases in latency can impact transaction approvals and customer experience during peak periods.

A leading consumer internet platform in India recently adopted GPU infrastructure to power large-scale generative AI experiences, including AI-generated personalized content for millions of mobile users. While the AI models themselves are

highly capable, the real challenge is ensuring GPU capacity, low-latency inference, and seamless scalability as user adoption grows. This is a great example of how infrastructure becomes just as critical as the AI model itself.

We are also seeing digital mobility and automotive platforms increasingly leverage cloud infrastructure to process vast amounts of location, vehicle, and user data and to support AI-driven services. As these businesses expand, they require cloud platforms that can scale rapidly without compromising performance or reliability, making modern IaaS a foundational enabler for their AI and digital transformation initiatives.

The same trend is visible in media and streaming, where AI is increasingly used for content personalization, advertising optimization, and content moderation. Major live sporting events or blockbuster releases can generate massive spikes in concurrent AI inference requests, exposing limitations in legacy infrastructure.

These examples reinforce that AI success is no longer determined solely by the quality of the model. It increasingly depends on whether the underlying infrastructure can deliver low-latency, resilient, and scalable AI services close to where users and applications operate.

CISO FORUM: How should Indian CIOs and CISOs approach the trade-off among performance, latency, and cost when architecting AI deployments? Is there a practical framework, or does it vary too much by industry?

MITESH JAIN: Most enterprises deploy AI, assuming it must balance performance, latency, and cost. Building the right infrastructure could help enterprises enhance all relevant elements by placing their workloads where they are best suited. The right approach is to view AI through a distributed-infrastructure lens. Large-scale work, such as data processing, should be handled from a centralized space, while AI inference should run on edge computing users. This will help reduce costs, improve latency and enhance user experience. Again, platforms like Kubernetes help enterprises deploy workloads in a decentralized manner and avoid data transfer costs, enabling CIOs to gain better visibility into performance.

"India's AI story will be defined by the execution of its infrastructure."

What CISOs should be mindful of is building the security infrastructure in this manner from the first day itself. Protection against malware and other threats should be deployed closer to where AI workloads are embedded and carried on. Hence, the entire AI infrastructure must be built with performance, high security, and cost optimization in equal measure.

CISO FORUM: Data residency and DPDP compliance add another layer to infrastructure decisions in India. How much is regulation, versus latency or cost, actually shaping where enterprises choose to run inference?

MITESH JAIN: The terms regulation, latency, and cost are no longer viewed as separate entities; instead, they are shaping industry conversations. The DPDP Act has set non-negotiable guidelines for data flow and compliance. Enterprises cannot transfer data overseas for processing or inference purposes.

However, these guidelines and frameworks are also shaping the architecture of data. For instance, if an enterprise chooses to align with the DPDP Act and keep data flows within processing boundaries, the company begins to look at the decentralized cloud footprint in India. Enterprises can also run inference on edge nodes and distribute cloud computing across Indian states, potentially helping them address data residency and compliance requirements.

CISO FORUM: Three years from now, do you expect India's AI story to be told through its models and talent, or through its infrastructure? What needs to happen for the infrastructure layer to keep pace?

MITESH JAIN: India already has an incomparable pool of talent resources and optimized localized AI models and programs. However, three years down the line, India's AI story will be defined by the execution of its infrastructure. India's pool of talent, AI models, and applications have immense potential, but with a strong infrastructure backbone, enterprises can convert that potential into ROI.

Three vital things should be kept in mind to ensure the infrastructure layer keeps pace.

- **Distributed GPU infrastructure:** Enterprises must move away from relying on Tier-1 data centers for AI models. There is a need to expand decentralized GPU systems deployed across Indian states to meet the growing demand for edge computing.
- **Reducing the "Cloud Tax"** Enterprises should move towards greater pricing transparency. Unpredictable fees will obstruct mid-market AI applications if left unaudited.
- **Security:** As AI is transitioning from chatbots to agentic infrastructure, the security algorithm must evolve. At this point, establishing a robust edge security infrastructure is essential to protect high-frequency AI workloads. ■

From Zero Trust to Zero Assumption: Securing Employees, Machines, APIs, Bots and AI Agents



Zero Trust must cover machines, APIs, bots and AI agents: assign owners, limit permissions, require approvals, verify continuously.

By **N. Radhika** | editor.tech@timesgroup.com

In several access reviews over the years, I have come across a service account that nobody could properly explain. Usually, it was created for a project that ended long ago. It still worked; it still held permissions, and it had never raised an alert because it was doing exactly what it was built to do.

Those moments changed how I think about security. We had spent years getting better at verifying employees. We had given far less thought to verify everything else.

Why Zero Trust is getting harder

I still believe in Zero Trust. "Never trust, always verify" sounds, and I have seen it make real improvements.

But in a large, distributed enterprise, the idea meets reality fast. Some legacy applications cannot support modern authentication, yet they run critical processes. Operational systems cannot always be patched when we would like. Business teams need access today, not after the next architecture review. And a control that is technically perfect but stops production will not last long.

Working in a large industrial environment has taught me that security must earn its place next to business continuity. Zero Trust works when it respects that. It struggles when it is treated as a project with an end date.

Identity is no longer just people

Most access models still begin with one question: who is the user, and is the device trusted?

That question is necessary. It is no longer enough.

Servers, virtual machines, applications, APIs, scheduled jobs, integration accounts and bots all authenticate. They all hold permissions, and they can all touch data. In many enterprises, these non-human identities are growing quickly, often faster than anyone is tracking them.



N. Radhika
Senior Manager – IT
BHEL – Electronics Division

Employees join, move, and leave through processes we understand. Machine identities rarely get that treatment. They are created in a hurry, given broad rights "just to make it work", and then left alone. Passwords go unrotated because nobody wants to break an integration they don't fully understand.

In my experience, older machine accounts can accumulate permission over time. That deserves more attention than it usually gets on a risk register.

Machines talking to machines

Much of the traffic inside a modern enterprise never involves a person. Applications call APIs. Systems exchange data overnight. Reports pull from databases, and automation stitches platforms together.

We have become reasonably good at securing the perimeter and the user's login. We are less disciplined about what happens after two systems are connected. Once an integration is trusted, it tends to stay trusted, whatever changes around it.

Three questions help me here. Does every service account have a named owner? Does its access still match its purpose? Would we notice if it started behaving differently tomorrow? If the honest answer to any of them is "not sure", that is where the work begins.

APIs and bots deserve the same discipline. An API is a door into a business process, and it should be treated like any other entry point. A bot is a worker with credentials. It needs an owner, a defined role, and a review date, just as a new employee would.

Why AI agents change the equation

AI agents are different, and I say that as someone who is genuinely enthusiastic about enterprise AI.

A traditional application does what it was programmed to do. An AI agent can read information, interpret it, form a recommendation, call other tools and APIs and, in some designs, take action. That is what makes it valuable. It is also why it needs a different kind of governance.

For every agent, a CISO should be able to answer five questions:
What can I see?
What can I decide?
What can it be executed?
Which actions do a human need to approve them?

How quickly can its access be withdrawn?

One principle matters more than the rest: an agent should not automatically inherit the permissions of the person or application that started it. If a manager with broad access launches an agent to summarise vendor contracts, the agent needs those contracts, not everything the manager can reach.

This often goes wrong through convenience. Someone connects the agent through an existing account, and it quietly becomes far more powerful than its task requires.

None of this is a reason to slow down AI adoption. It simply means giving each agent a clear identity, a narrow purpose and permissions that fit that purpose.

Human oversight is a design choice

Automation is useful because it removes effort. But removing effort should not mean removing accountability.

It helps me to sort actions by consequence. Reading a public document has a low impact. Drafting a recommendation is moderate. Changing a record, approving a payment, deleting data or altering a configuration is high impact, and often hard to undo.

For those last actions, a named person should approve before the agent acts, and the approval should be recorded. This is not a distrust of technology. We already apply the same logic to people. Few organisations let one employee approve a large payment alone, and an agent should be held to a comparable standard.

Verification must be continuous

Anyone who works in an audited environment knows the rhythm of periodic compliance. Access to reviews happen quarterly; evidence is gathered annually, and the audit passes. That matters, and ISO 27001-aligned governance gives an organisation a strong foundation.

But a quarterly review cannot keep pace with an agent that acts in seconds. Access that made sense in January may be wrong by March because the project changed, the owner moved on or the data became more sensitive.

Verification has to move closer to the moment of access. The question is no longer only "Was this approved?" but also "Does this still make sense, right now?"

What CISOs can do now

None of this needs a new budget line to begin:

- Build an inventory of non-human identities and give each one an accountable owner.
- Review or retire dormant service accounts and integrations.



"Automation is useful because it removes effort. But removing effort should not mean removing accountability."

- Give agents and bots their own identities, not borrowed from human credentials.
- Decide which actions an AI agent may never take without human approval.
- Make sure you can cut off any machine or agent's access quickly and test it.
- Log what agents do in a form your security team can actually read.

None of these is a futuristic idea. They extend disciplines security teams already know identity, least privilege, accountability, monitoring and access review.

From Zero Trust to Zero Assumption

The way I think about the next step is simple: stop assuming at all. Don't assume an old integration is still needed, that a bot is behaving as designed, or that an agent should hold the same power as the person who launched it.

Zero Assumption does not replace Zero Trust. It finishes the thought.

The question is no longer only who is asking. It is what is asking, what it wants to do, why it needs to, and whether that still makes sense today. Enterprises that ask these questions early can adopt AI with confidence rather than caution. In my experience, that confidence is what lets security and the business move forward together. ■

Can India's DPDP framework keep up with the rise of autonomous AI?



India's DPDP law remains sound, but autonomous AI demands dynamic, lifecycle-based controls rather than static consent, assessments and contracts.

By **Abilash Soundararajan** | editor.tech@timesgroup.com

Enterprises are moving from generative AI assistants that answer prompts autonomous agents that retrieve data, call tools across systems and act with limited human oversight. The Digital Personal Data Protection (DPDP) Act, 2023 and the DPDP Rules, 2025 are technology-neutral, which is largely a strength. The harder question is operational: when an agent decides at runtime what data to fetch, combine or infer, can a Data Fiduciary still demonstrate purpose limitation, valid consent and accountability? The answer depends less on new law than on where controls are placed. A lifecycle view, such as the PERAI (Privacy Enhancing & Responsible AI) framework, organises those controls across five stages.

Data collection: Quantify before collecting

Section 6 requires consent that is free, specific, informed, and limited to personal data necessary for the stated purpose. Section 5 requires a notice describing that purpose. Agents strain this because a purpose fixed at collection may not anticipate what an agent later does with the data. Privacy threat modelling at intake, which simulates singling-out, linkage and inference attacks and produces a measurable exposure score, gives the fiduciary evidence of what is actually "necessary". For Significant Data Fiduciaries, Section 10 mandates Data Protection Impact Assessments.

Data-in-use protection: Minimise by design

The Act does not use the term "data minimisation", but the necessity test in Section 6 and the erasure obligation in Section 8(7) implies it. Agents with broad, standing access to data work against both. Consent-based access control, where consent and purpose of metadata travel with the dataset, lets policy be enforced at the point of use. Anonymisation, differential privacy for aggregate queries and synthetic data for development and testing reduce how much identifiable data an agent ever sees.



Abilash Soundararajan
Founder & CEO, PrivaSapien

AI assessment: Test before deployment

Section 8(3) requires completeness, accuracy, and consistency where personal data is used to decide affecting a Data Principal, and Section 8(5) requires reasonable security safeguards. Static test suites do not scale to agents that plan multi-step actions. Agentic red teaming, using autonomous adversarial agents to probe for prompt injection, cross-context leakage, bias and harmful output, produces quantified risk scores and attack libraries. These artefacts serve as audit evidence, and for Significant Data Fiduciaries they support the periodic audits Section 10 requires.

AI inference: Enforce at runtime

Most privacy failures in agentic systems occur at inference, when prompts and retrieved context carry live personal data. Runtime controls include prompt-level PII detection, redaction, tokenisation with secure detokenisation, and format-preserving encryption, so raw identifiers do not reach the model.

Agentic systems: Govern the ecosystem

Section 8(1) makes the Data Fiduciary responsible for processing undertaken by a Data Processor on its behalf, and Section 8(2) requires a valid contract for any such engagement. An agent that invokes third-party tools through protocols such as MCP effectively creates dynamic processor relationships. Threats such as tool poisoning, rug pulls and cross-agent context injection sit outside traditional vendor review. Two guardrails address this: an MCP security layer that inspects tool descriptions and payloads, and a context boundary layer that limits what information propagates between agents. Over time, red-teaming results can feed a behavioural trust profile, so only agents meeting defined thresholds participate in sensitive workflows.

Are current controls sufficient?

The DPDP principles hold up. What does not hold up is reliance on static controls: one-time consent screens, annual assessments and contractual assurances.Organisations should therefore build several capabilities into agentic systems from the outset:

- Purpose-bound, time-limited access rather than standing permissions
- Continuous risk quantification, not point-in-time assessment
- Runtime enforcement with tamper-evident audit trails
- Defined human review points for consequential decisions
- The ability to suspend or roll back an agent
- Clear ownership under the fiduciary's accountability

Regulatory guidance on inferred data, automated decision-making, and processor status for AI tools would help the ecosystem. Until then, a lifecycle approach lets organisations demonstrate compliance as autonomy increases, rather than asserting it. ■

McKinsey's 14 Tech Trends Set to Reshape Business in 2026



McKinsey's 2026 outlook says AI is moving into the physical world, while energy and talent limits slow adoption.

By **CISO Forum** | editor.tech@timesgroup.com

Every year, McKinsey & Company surveys the technological horizon and names the innovations poised to matter most for business leaders. Its newly released Technology Trends Outlook 2026, now in its sixth edition, tracks 14 trends across three broad categories: the AI revolution, computing and connectivity frontiers, and cutting-edge engineering. The message this year is unmistakable: technology is stepping off the screen and into the physical world.

AI gets hands and a body

Four of the 14 trends are explicitly AI-focused, but the report argues AI is quietly powering nearly all the others too. Agentic software development and agentic AI both newly expanded categories this year describe systems that no longer answer questions but complete entire tasks end-to-end, working alongside human employees rather than simply assisting them. McKinsey calls this the shift from AI that talks to AI that acts.

That shift extends into the physical realm as well. Robots are being trained to understand and navigate unpredictable environments, autonomous vehicles are making real-time driving decisions, and "dark factories" are producing goods with no humans on-site. The report frames this as "physical AI" arriving first in manufacturing and logistics, before eventually reaching hospitals, farms, and city infrastructure.

Innovation is outrunning our ability to absorb it

A recurring theme in the report is speed and its costs. In biopharma, AI can now generate thousands of viable drug candidates in the time it once took to produce a handful. Still, human bottlenecks like lab validation and regulatory review haven't kept pace. In cybersecurity, the report notes that more than three-quarters of vulnerabilities are now discovered only after an exploit already exists, meaning AI has compressed the defenders' window to nearly zero even as it also helps them respond faster.

Power, chips, and the new infrastructure race

The most sobering thread running through the report is energy. AI infrastructure spending doubled in a single year, and McKinsey projects that U.S. data centers running AI workloads alone could consume as much electricity by 2030 as the entire state of California does

today. The bottleneck, the report notes, isn't hyperscalers' ambition; it's the physical supply chain, with transformer lead times stretching past two years and thousands of gigawatts of energy projects stuck waiting for grid connections worldwide.

This energy crunch is also reshaping the chip industry. Tech giants including Amazon, Google, Meta, and Microsoft are increasingly co-designing custom, application-specific silicon rather than relying on general-purpose chips, chasing efficiency gains as inference running trained models at scale overtakes training as AI's dominant computing workload. unning the operation at scale.

Where the money is flowing

McKinsey's investment data offers a forward-looking signal of conviction. Five trends agentic software development, AI infrastructure, AI for scientific discovery, the future of space technologies, and the future of robotics are on pace to more than double their 2026 investment compared with 2025.

Credentials already on the dark web

The report found over 270,000 credentials from fans visiting FIFA-related websites already present in stealer log datasets, harvested by malware families including Vidar, LummaC2, and RedLine. More than 260 credentials tied to FIFA employees were also identified.

The bottom line

McKinsey's core takeaway isn't a prediction of any single winning technology, but a warning about pace: no organization can deploy AI at scale without upskilling its workforce, modernizing legacy systems, and securing energy and compute as fiercely as it secures talent. As the report puts it, leaders who understand the patterns behind this change will be positioned to shape it rather than react to it. ■

India Inc.'s cybersecurity paradox: More tools, same old breaches



ESET's 2026 India report finds 80% of firms were breached despite heavy tool spending, as detection gaps and human error persist.

By **CISO Forum** | editor.tech@timesgroup.com

A new ESET Enterprise Cybersecurity Report 2026 (India), conducted in partnership with Blackbox Research, paints a sobering picture of enterprise security in India: more spending on tools is not stopping more frequent breaches. Despite record investment in technology, Indian organizations are being breached more often than ever, and the reasons are surprisingly human.

AI: A double-edged sword

The report finds that 99% of Indian businesses are now using or piloting AI in some form, from customer service chatbots to fraud detection. But this rapid adoption has a cost: 85% of organizations experienced an AI-related cyber threat in the past year. The most common issues include AI-generated phishing and impersonation (48%), data leakage through AI platforms (48%), and deepfake or voice-cloning attacks (47%). As a result, concern is greatest in sectors such as technology and financial services, where AI-driven threats are seen as an increasing operational risk rather than a distant possibility.

Breaches are the norm, not the exception

Perhaps the most striking figure is this: 80% of Indian businesses suffered a major security incident in the last 12 months. Nearly half of those affected 44% were hit three or more times. These incidents most often involved data exfiltration, cloud breaches, business email compromise and ransomware, with attackers frequently chaining multiple techniques together rather than relying on a single method.

The fallout isn't confined to IT departments. One in six affected businesses reported significant financial losses, and a similar share faced legal or regulatory penalties — a reminder that cybersecurity failures increasingly carry boardroom-level consequences.

The real problem isn't technology

According to the report, the biggest obstacles to better defense aren't a lack of tools; their capability gaps. Delayed detection (55%) limited internal resources (54%), and poor visibility across IT environments (53%) were the top challenges organizations faced when trying to contain an incident. In practice, only 19% of businesses could detect and respond to a threat within the critical first hour; most took 24 hours or more.

"More security software isn't solving the core problem."

Employees remain the weakest link

Human error continues to open the door for attackers. Nearly a third of incidents were linked to user actions; most commonly employees clicking phishing links, reusing weak passwords or mishandling sensitive data. While 97% of organizations now run cybersecurity awareness training, the report's experts stress that frequency alone isn't enough; training needs to be timely and relevant to the threats employees are actually likely to face, such as festival-season scams or AI-generated deepfake calls.

Looking ahead: MDR and smarter insurance

Looking to close these gaps, 84% of Indian organizations are now using or planning to adopt Managed Detection and Response (MDR) services, driven by the need for round-the-clock monitoring and faster response times. Meanwhile, cyber insurance is becoming harder to secure; 97% of businesses reported challenges including rising premiums, stricter insurer requirements, and difficulty demonstrating their security maturity.

The takeaway for Indian enterprises is clear:

more security software isn't solving the core problem. Faster detection, sharper employee awareness, and stronger operational resilience will determine who survives the next attack and who becomes another statistic in next year's report. ■

When AI gets hands: Akamai warns security chiefs the rules have changed



Akamai's report warns autonomous AI agents are industrialising attacks, urging CISOs to extend Zero Trust and keep humans accountable.

By **CISO Forum** | editor.tech@timesgroup.com

Akamai's latest State of the Internet report, Speed, Scale, and Nonhuman Identity: The Agentic Threat Landscape (Volume 12, Issue 5, published September 2026), says companies are entering a new phase of risk. The danger is no longer just who logs in. It is what autonomous software does once it is inside.

Two problems, one era

The report frames the moment around two challenges. The first is the "autonomy guardrail dilemma": attackers can trick AI agents through hidden instructions, poisoned context, or rogue browser extensions, leading them to take unauthorised, high-impact actions. The second is the "asymmetric velocity gap." AI can find chain software flaws in hours, while most companies still patch at human speed.

The numbers behind the shift

Akamai's network data shows the scale. AI bot traffic rose 300% in 2025. Daily API attacks per organisation grew 113%, from 121 to 258. Some 87% of surveyed organisations reported an API-related incident, and API breaches now cost enterprises an average of US$700,000 per year.

MCP: Power with peril

Akamai's CTO Charlie Gero describes the Model Context Protocol as giving AI "a hand". It lets models query databases, run code, and orchestrate workflows. But it blurs the line between data and code. A malicious third-party server can hide instructions in tool descriptions or use a low-trust tool to push an AI into raiding a high-trust one, such as a key vault. His advice is not to avoid MCP but to isolate trust levels, add proxy gateways, and require human approval for state-changing actions.

The browser is the new weak spot

Employees are a fast-growing risk. Some 99% run at least one browser extension, and over 40% have installed AI-powered ones. A quarter of those AI extensions changed their permissions within a year, compared with 4.33% of extensions overall. Almost 6% of employee chatbot conversations contained sensitive data, mostly personal information, and 47% of AI conversations on work devices used personal accounts, leaving security teams blind.

"AI does not replace traditional attacks. It industrialises them."

Lessons from Project Glasswing

Akamai took part in Anthropic's Project Glasswing and had early access to the Mythos model. Its takeaway: frontier AI did not make software less secure, but it made discovery dramatically faster. The bottleneck is now deciding which exposures are reachable, exploitable, and business-critical. Until patches land, firms need runtime protection, edge controls and segmentation.

Brand risk in a zero-click world

As AI agents replace websites as the front door to customers, CISOs must watch four areas: citations, accuracy, sentiment and bot management.

Trusting the machine

The report offers a practical test instead of a philosophical one. Can I verify the output? How bad, and how reversible, is a confidently wrong answer? Its guidance: match autonomy to verifiability, not capability; treat confident output as a signal to check; and keep accountability human.

The takeaway

Akamai stresses that firms need not build a separate AI security programme. Instead, they should extend proven basics such as visibility, Zero Trust segmentation, exposure management, and identity checks into autonomous workflows, while also separating shadow AI from managed systems. As the report puts it, AI does not replace traditional attacks. It industrialises them.

A note for readers: Akamai sells security products, so its recommendations favour edge-based defences. The data is from its own network and customer analysis.

Quick Heal's first ESG report: Honest baseline, modest numbers



Quick Heal's first ESG report cuts Scope 2 emissions 16.9%, but renewables stay small, and several governance goals slip to FY27.

By **CISO Forum** | editor.tech@timesgroup.com

Quick Heal Technologies has published its first Sustainability Report for FY 2025-26, released on 23 September 2026. The Pune-based cybersecurity firm, founded in 1995 and listed since 2016, says the report marks a shift "from communicating what we build to disclosing how we operate." It is candid about its limits: the report is India-specific, prepared with reference to GRI Standards and SEBI's BRSR framework, and has not been externally assured.

Emissions and energy

The headline number is a 16.9% drop in Scope 2 emissions; from 1,746.62 to 1,451.35 tCO_2e. Total Scope 1 and 2 emissions stand at 1,463.72 tCO_2e. Electricity use was about 2.15 million kWh. The firm also swapped CFL lighting for LEDs and upgraded older air conditioning.

Renewables remain small. A 45-kW solar plant at its Thube Park campus generated roughly 46,570 kWh, covering only about 6.98% of energy needs. As a result, the company plans an emissions baseline and a Net Zero roadmap by FY2030. Scope 3 is limited to an employee commuting survey, where nearly 79% of personal vehicle users drive petrol vehicles.

Waste, water and packaging

Quick Heal disposed of 966 kg of e-waste through authorised recyclers in accordance with EPR rules. Packaging is now 100% biodegradable and about 38% smaller. The company plans to phase out physical boxes in favour of soft keys over five years. Water use fell 12.84% to 12,018 kilolitres, and treated water is reused for gardening and flushing.

People and community

The workforce numbers are mixed. Employees logged 131,794 learning hours, and the annual survey saw 92% participation and an Employee NPS of 50.1. However, women make up only 17.65% of the workforce, below the FY27 target of 18%. Annualised voluntary attrition is 16%.

The report's strongest social story is Cyber Shiksha for Cyber Suraksha, which it says has touched over 94 lakh lives across 19 states and 12,000+ schools and colleges. In addition, the Quick Heal Foundation has donated 18 mobile medical vans, said to have reached 600+ villages.

> *"From communicating what we build to disclosing how we operate."*

Governance and trust

A three-tier ESG structure, led by the CFO and the Chief of Operational Excellence, oversees the effort. The company holds ISO/IEC 27701:2019, ISO 27001:2022, and SOC 2 Type 2 certifications and reports no confirmed human rights violations or ethical breaches. It also highlights its Seqrite Labs threat research, including tracking over 650 coordinated incidents during Operation Sindoor.

The gaps

What stands out is the candour. Customer NPS is just 5, and Partner NPS is 31. Board effectiveness reviews and board training are deferred to FY27, and the sustainable supply chain work is still "in progress." A formal double-materiality assessment is on the roadmap.

The financial picture adds context: net revenue of ₹261 crore, an EBITDA loss of ₹29.4 crore and a loss after tax of ₹10.9 crore. Sustainability spending will compete with pressure on profitability.

The takeaway

Quick Heal's report is more baseline than breakthrough. Its emissions are low in absolute terms, and its renewable share is small. Even so, by publishing weak scores alongside strong ones, the company has set a measurable starting point. The test will be whether the FY27 report shows real progress on renewables, assurance and supply

Forrester's 2027 forecast: Tech leaders chase control, and pay for it



Forrester predicts 2027 will force CIOs and CISOs to gain control over AI, budgets, infrastructure and security without slowing progress.

By **CISO Forum** | editor.tech@timesgroup.com

Forrester Research's Predictions 2027: Technology & Security report argues that next year's central challenge for CIOs and CISOs will be gaining more control over technology without slowing progress. AI is eating a bigger share of budgets, infrastructure choices are getting harder, and security and sovereignty demands are rising. The report warns that more control brings its own risks: tough trade-offs, new points of failure and possible damage to trust. Its message is that 2027 will reward deliberate choices and resilience.

AI takes a fifth of the wallet

Forrester's global tech market forecast puts worldwide tech spending at $6.1 trillion in 2027, up 7.6% from 2026. But the mix is shifting fast toward AI, and nearly three-quarters of AI spending will go to servers and equipment. That shift puts pressure on other priorities: tech consulting and systems integration will shrink from 15% to 7% of tech spending growth by 2030. Budgets are not keeping pace. Most leaders expect increases next year, but only 53% expect one above the IMF's projected 3.2% inflation, so just under half will effectively lose purchasing power. Third-party services will face the most pressure because AI automation is driving down costs for consulting and managed services. Forrester advises renegotiating services and seat-based software renewals six months early and redirecting the savings toward token costs and context management, reinforcing the need for tighter control.

The quantum wildcard

Forrester says most people still believe quantum computing is years away, and that remains the likely outcome. Yet it flags a small chance of an unexpected breakthrough, likening it to the way ChatGPT caught CIOs off guard. Its research suggests the foundations for fault-tolerant quantum computing were established sooner than expected, and 87% of director-level and above decision-makers say their organizations allocate no resources to it. A surprise breakthrough would prompt boards to demand plans for "Q-Day" and leave CIOs questioning whether they are ready.

Beyond NVIDIA

NVIDIA will stay the default for frontier model training and large-scale inference, but Forrester predicts a quarter of AI infrastructure RFPs will require non-NVIDIA accelerator options. Enterprises want a second source to cut costs and reduce exposure to chip shortages, sovereignty issues, and geopolitical risk. Rather than replacing NVIDIA, organizations will split workloads across chip types based on model size, latency, data sensitivity and location. Hyperscalers are pushing this shift with their own silicon, such as AWS's Trainium, Google's Ironwood TPU and Microsoft's Maia 200. At the same time, Alibaba Cloud's Zhenwu and Huawei's Ascend serve sovereign and domestic deployments.

Sovereignty becomes a buying criterion

The report predicts digital sovereignty will drive a 20% increase in enterprise software localization. Once a compliance footnote, it is now a purchasing factor. In Forrester's surveys, 43% of enterprise applications respondents cited data privacy, residency, and sovereignty risks when adopting SaaS, and 59% of public cloud decision-makers said sovereignty constraints completely determine their choice of cloud vendor. With data localization laws, sovereign cloud mandates and diverging AI regulations spreading, Forrester urges investment in regional data stores, country-specific AI models, localized business processes and jurisdiction-specific compliance controls. This should resonate with Indian enterprises already adapting to data protection rules.

When AI-written code breaks the internet

The most striking security prediction is that an AI-generated bug or regression will cause a global outage across multiple industries. Major vendors now ship hundreds of patches at once, and sometimes over a thousand. As AI speeds up vulnerability discovery, a "find fast and fix things" mindset will take hold, and security leaders will lean on AI to remediate code and update dependencies. But review and testing capacity will not keep up, Forrester warns, and the result could be an outage on a scale comparable to the July 2024 incident. Its advice is to strengthen the ability to roll back patches from unreliable vendors quickly.



"For CIOs and CISOs, 2027 will reward deliberate choices."

The takeaway

Taken together, the predictions show leaders squeezed between ambition and constraint, with AI promising productivity gains. At the same time, budgets tighten, dependence on hardware deepens, and the stakes of every security decision rise. The report's practical guidance is to renegotiate contracts early, diversify infrastructure, localize with care, and prepare to reverse changes fast. In 2027, Forrester suggests, resilience and preparation will matter more than speed.

How far should we let AI go?

In just five years since COVID, the way we work, make decisions and run businesses has changed dramatically. And the pace of change shows little sign of slowing. AI is moving from the sidelines into the heart of enterprise operations, reshaping processes, challenging established ways of working and opening possibilities that seemed distant just a few years ago.

What makes this shift particularly significant is that AI is no longer limited to analysing information or recommending next steps. Increasingly, AI agents can access data, interact with enterprise applications and execute tasks with limited human intervention. The potential is enormous, but so are the questions it raises. How much autonomy should organisations give these systems? How do they ensure that agents act within defined boundaries? And who remains accountable when an AI-driven decision goes wrong?

These are no longer questions for the future. They are becoming immediate concerns for business and technology leaders as organisations move from AI experimentation to large-scale deployment.

The challenge is that AI capabilities are advancing faster than many organisations' ability to govern them. An IBM study of senior technology leaders in 2026 found that 77% of surveyed organisations said AI adoption was outpacing their governance capabilities. McKinsey's State of AI Trust in 2026: Shifting to the Agentic Era similarly points to gaps in organisational readiness, with only about 30% of surveyed organisations reaching level three or higher on its four-level maturity scale for strategy, governance and agentic AI controls.

For CXOs, the dilemma is not whether to adopt AI, but how to scale it without compromising trust, security, accountability or business outcomes.

In this issue, CISO Forum explores how India's technology leaders are navigating this crossroads. Their perspectives examine the foundations required to move from AI pilots to dependable enterprise capabilities: trusted data, clear guardrails, measurable business outcomes and accountability that keeps pace with autonomy. ■



"For CXOs, the dilemma is not whether to adopt AI, but how to scale it without compromising trust, security, accountability or business outcomes."

Jatinder Singh
Chief Editor, CISO Forum
jatinder.singh1@timesgroup.com



ET Edge | CIO&LEADER

studiotalks

WHERE TECHNOLOGY MEETS THE SPOTLIGHT !

WHY JOIN STUDIOTALKS

YOUR EXPERTISE. YOUR VOICE. YOUR SPOTLIGHT.

Drive bold discussions defining the future of enterprise technology.

Showcase your expertise in a high-impact, TV-style experience.

Get featured alongside India's most influential technology leaders.

SECURE YOUR SPOT NOW !

For business proposal
Hafeez Shaikh
Assistant Director - Projects
ET Edge - The Times Group
hafeez.shaikh@timesgroup.com
+91 9833103611

CURATED BY

ET Edge | THE TIMES GROUP

ET EDGE - INDIA'S LEADING B2B INTELLIGENCE, COMMUNITY & MARKET PLATFORM

AI SCALES. CAN YOUR INFRASTRUCTURE?

NOW TRY RUNNING IT ACROSS AN ENTERPRISE.

ET Edge CIO&LEADER

AI INFRASTRUCTURE SUMMIT 2026

THE INFRASTRUCTURE BEHIND THE AI ENTERPRISE

27–28 NOVEMBER 2026 | AAMBY VALLEY, LONAVALA

Where 75+ India's top CIOs & IT Leaders come together to decode what it will take to power the AI enterprise.

- Two Days
- CIO Workshops
- Closed-Door Roundtables
- AI Infrastructure Expo

AI IS SCALING. CAN YOU KEEP UP?

AN EVENT BY

ET Edge | THE TIMES GROUP

For sponsorship, Kindly write to:

Hafeez Shaikh
hafeez.shaikh@timesgroup.com
+91 9833103611

ET EDGE - INDIA'S LEADING B2B INTELLIGENCE, COMMUNITY & MARKET PLATFORM