ET Edge INSIGHTS

OCTOBER 2026
VOLUME 15 • ISSUE 07 • ₹150

CIO&LEADER

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF

SHOULD AI MAKE DECISIONS?

10 CIOs and technology leaders on how enterprises can balance AI autonomy, trust, accountability, and human control. PG.10

A TIMES GROUP PUBLICATION
cioandleader.com
cioandleader
cioandleader

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

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.

But a new name does little to answer the harder question: Are we prepared for systems that can act, rather than simply respond?

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, CIO&Leader 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,
Enterprise Tech Publications
ET Edge-The Times Group
jatinder.singh1@timesgroup.com

CONTENT OCTOBER 2026



COVER STORY

10-22

Should AI Make Decisions?

10 CIOs and technology leaders on how enterprises can balance AI autonomy, trust, accountability, and human control.



Cover Design by:
Shokeen Saifi

Please Recycle this Magazine and Remove Inserts before Recycling

Disclaimer - 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 CIO&Leader, its publisher or editor. While every effort has been made to ensure the accuracy and reliability of the information provided, ET Edge Insights CIO&Leader makes no express or implied representations about the completeness, accuracy, reliability, suitability, or availability with respect to the content contained herein. ®CIO&Leader is a registered trademark of The Times Group.

















NEWS & VIEWS

06
Airtel Business unifies SMS, WhatsApp, RCS and voice on one platform

08-09
Harness acquires Augment Code assets to connect AI coding with software delivery

INSIGHT

23-24
Can India's DPDP framework keep up with the rise of autonomous AI?

25-27
India leads in sovereign procurement intent, but has no plan to back it up

30-31
Tokenization can be India's next digital public revolution

32-33
When India celebrates, payment infrastructure cannot afford to rest

TECH TALK

34-36
The next phase of enterprise IoT is about...

BY PRAVEEN ARORA

37-40
Resilience is one mission, not two scorecards

BY SUBHALAKSHMI GANAPATHY

ET Edge INSIGHTS

CIO&LEADER

TRACK TECHNOLOGY • BUILD BUSINESS • SHAPE SELF

www.cioandleader.com

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
Art Director: **Shokeen Saifi**
Assistant Manager - Design: **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**
Manager Database: **Neelam Adhangale**
Manager - Community Development: **Vaishali Banerjee**
Manager - Community Development: **Reetu Pande**
Manager - Community Development: **Snehal Thosar**
Executive - Database: **Dipti Gamre**

OPERATIONS
Senior Manager - Digital Operations: **Jagdish Bhainsora**
Senior Manager- Video Content: **Sunil Kumar**
Manager - Operations: **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.cioandleader.com

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.

CIO MOVEMENTS



Abhishek Gupta Appointed Chief Information Officer at Akasa Air

Abhishek Gupta moved from DishTV, where he led corporate IT for over seven years, to become Chief Information Officer at Akasa Air, bringing IT leadership.



Dr. Devam R Shah Appointed Chief Agentic Security Officer at Fynd

Dr. Devam R Shah moves from Locus, where he served as CISO and DPO, to become Chief Agentic Security Officer at Fynd, leading security transformation.



Sachin Agarwal Appointed Group Chief Information Officer at Gold Plus Glass Industry Ltd.

Sachin Agarwal moves from G R Infraprojects, where he led IT for 16 years, to become Group Chief Information Officer at Gold Plus Glass Industry.



Sanjay Deshmukh Appointed Vice President, APAC at OpenAI

Sanjay Deshmukh moves from Elastic, where he led Asia Pacific and Japan, to become Vice President, APAC at OpenAI, after earlier senior regional leadership roles.



Gajapathy Srinivasalu Appointed CIO and CISO at Adani Total Private Limited

Gajapathy Srinivasalu expands his role at Adani Total Private Limited from Head of IT to CIO and CISO, overseeing technology strategy, cybersecurity and resilience programmes.



Prakash Patel Appointed CISO and Privacy Officer at Baroda Global Shared Services Ltd.

Prakash Patel moves from Brain Payroll as AVP – IT Infrastructure and Security to become CISO and Privacy Officer at Baroda Global Shared Services Limited.



Dr. Ashish Shrivastava Appointed Chief Information Officer at Radico Khaitan Limited

Dr. Ashish Shrivastava moves from Financial Software & System, where he led technology, to become Chief Information Officer at Radico Khaitan, driving AI-led business transformation.



Kiran Belsekar Appointed Head – Technology Governance at Mahindra Manulife Insurance

Kiran Belsekar moves from Bandhan Life, where he holds senior technology and security leadership roles, to become Head – Technology Governance at Mahindra Manulife Insurance.



Rejin Surendran Appointed Chief Digital & Information Officer at Britannia Industries

Rejin Surendran moved from Wipro Enterprises, where he served as Global CIO, to become Chief Digital & Information Officer at Britannia Industries, leading digital transformation.



Yogesh Kulkarni Appointed Group CISO – Global at STL

Yogesh Kulkarni moves from Crompton Greaves Consumer Electricals, where he served as Group CISO, to become Group Chief Information Security Officer – Global at STL.



Akanksha Garg Appointed CISO and DPO at ProtocolRoots Private Limited

Akanksha Garg moved from HDFC securities, where she was Data Protection Officer and Head of Digital Products PMO, to become CISO and DPO at ProtocolRoots.



Samip Mutha Appointed Global Chief Digital & Information Officer at Tata Chemicals

Samip Mutha moves from TechnoPro Holdings, leading business in Europe and the Middle East, to become Global Chief Digital and Information Officer at Tata Chemicals.



Mrinmoy Dey Appointed Chief Information Security Officer at InCred Financial Services

Mrinmoy Dey moves from Lendingkart, where he served as CISO, to become Chief Information Security Officer at InCred Financial Services, following earlier roles at Paytm.



Sudeep Agrawal Appointed VP – IT & Digital at Olam Agri

Sudeep Agrawal moves from Reliance Group, where he served as CIO, to become Vice President – IT & Digital at Olam Agri, leading its P&VA business.



Ninad Naik Promoted Deputy GM - Cybersecurity, AI and Data Privacy Governance & Assurance at Birla Opus

Ninad Naik is promoted from Assistant General Manager at Birla Opus to Deputy General Manager – Cybersecurity, AI and Data Privacy Governance & Assurance, globally.



Abhishek Gupta Appointed as Managing Director at Kong Inc.

Abhishek Gupta moves from SailPoint, where he served as Managing Director, India, to become Managing Director, India at Kong Inc.

Airtel Business unifies SMS, WhatsApp, RCS and voice on one platform

Airtel IQ Gateway unifies enterprise messaging channels with intelligent routing, delivery visibility and fraud protection.

By **CIO&Leader** | editor.tech@timesgroup.com

Airtel Business has launched Airtel IQ Gateway, a platform that lets enterprises manage customer communications across SMS, WhatsApp, Rich Communication Services (RCS), business voice, and number verification from one place.

Announced on October 5, 2026, the platform is positioned as a Digital Personal Data Protection (DPDP)-compliant control plane for enterprise communications. Airtel says it combines its fraud protection capabilities with network intelligence to give businesses greater visibility into message delivery, vendor performance, and potential fraudulent traffic.

One platform for multiple channels

Enterprises often juggle multiple service providers and tools, which Airtel says makes it hard to track whether messages reach customers and adds operational and regulatory challenges.



IQ Gateway brings these channels together in a single dashboard with real-time analytics and network-level delivery reports. It supports automatic switching to backup providers and intelligent routing across SMS, WhatsApp, RCS, and voice, so enterprises have fallback options when a channel or provider fails to deliver. It also integrates with Airtel's native message template system, Distributed Ledger Technology (DLT), reducing the manual effort of managing pre-approved templates.

Focus on fraud and artificial traffic

IQ Gateway draws on Airtel's network of intelligence. Airtel IQ Verify evaluates network, device, and usage signals to identify potentially fraudulent account creation and protect legitimate customers. The platform aims to detect artificial traffic spikes linked to fraudulent SMS pumping, which can cause unexpected billing.

Silent Authentication can replace the SMS One-Time Password (OTP). It validates whether the number entered during login or SIM-binding matches the active SIM in the device, which Airtel says can reduce risks from SIM swaps and social engineering.

Deployment options for enterprises

IQ Gateway integrates existing communication workflows and vendor ecosystems. Enterprises can deploy it on their own premises or as a Software as a Service (SaaS) offering hosted on Airtel's cloud, depending on security needs.

The platform was launched nationwide following trial deployments with select banks, financial institutions, and other industry collaborators. ■

IBM-Yotta platform brings sovereign control to enterprise agentic AI in India



IBM and Yotta's sovereign AI platform combines India-based infrastructure, agent orchestration and governance to help enterprises scale agentic AI while retaining control over data and inference.

By **CIO&Leader** | editor.tech@timesgroup.com

As Indian enterprises move from AI pilots towards production-scale deployments, the question of where data is stored is increasingly being joined by questions around where AI inference runs, who controls the underlying infrastructure, and how agents are governed. IBM and Yotta Data Services are responding to this shift with the general availability of a Sovereign Agentic AI Platform designed for Indian organisations.

The platform combines IBM's watsonx Orchestrate with Yotta's Shakti Cloud and Shakti Studio. The proposition goes beyond data residency by bringing compute, GPU infrastructure, networking, security, model development, inference, agent orchestration and governance into an India-based environment.

IBM watsonx Orchestrate serves as the agentic control plane, providing capabilities to secure, govern and manage AI agents as organisations scale their deployment. Yotta's Shakti Cloud provides the underlying compute and infrastructure, while Shakti Studio supports the development, fine-tuning, deployment and inference of open-source and proprietary AI models.

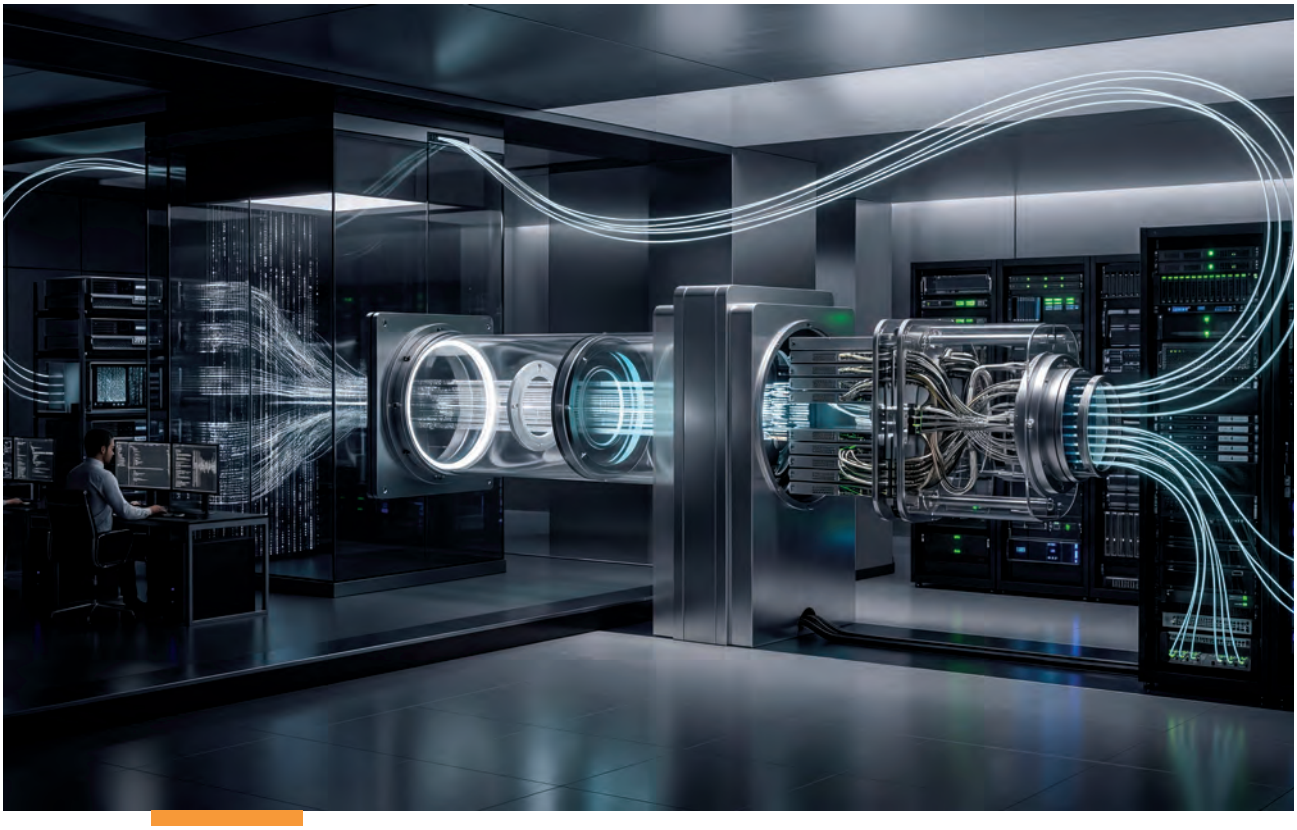
For CIOs, the architecture addresses a growing operational challenge. As enterprises introduce multiple AI agents across functions, controlling individual models or applications is unlikely to be sufficient. Organisations also need visibility into how agents access enterprise data, where inference takes place, what permissions they have and how their actions are governed.

This becomes particularly relevant as India's regulatory environment evolves. The platform is positioned to help organisations maintain data, inference and governance controls within the country, potentially simplifying the requirements for enterprises operating with sensitive or regulated workloads.

The platform is available through Yotta's Panvel and Greater Noida cloud regions and can support use cases including security operations, document processing and HR automation.

However, sovereignty alone does not determine the enterprise value of an agentic AI deployment. CIOs will still need to assess model performance, interoperability, workload economics, GPU availability, integration with existing systems and the controls available for human oversight. The ability to run AI infrastructure within national boundaries addresses one layer of enterprise risk; it does not by itself resolve questions around agent reliability, accountability or business outcomes.

The IBM-Yotta announcement therefore reflects a broader transition in enterprise AI: sovereignty is becoming an architectural consideration alongside performance, security and cost. As agentic systems gain greater access to enterprise workflows and decision processes, the infrastructure and governance layer underneath them is becoming as important as the AI capabilities themselves. ■



Harness acquires Augment Code assets to connect AI coding with software delivery

Harness is integrating AI coding with software delivery, security and production feedback to automate engineering workflows from initial requirements to merge-ready code.

By **CIO&Leader** | editor.tech@timesgroup.com

Harness has acquired select assets from Augment Code to expand automation across the software development life cycle (SDLC). The acquisition includes Cosmos, the Auggie CLI products, the Code Context Engine, and related technology. The teams behind these products will join Harness as the company works to automate more stages of software development.

Cosmos will be introduced as the Harness Cosmos Software Factory Agent, responsible for taking engineering tasks from an initial idea to merge-ready code. Harness plans to connect it with its existing agents for software delivery, security testing, runtime protection, and cost management, linking code creation with the processes required to move software into production.

Harness buys Augment Code assets, including Cosmos, to link AI coding with software delivery, security and production feedback.

The announcement centres on connecting codebase knowledge with information from software delivery and production systems. This approach is intended to help AI agents make informed changes, identify problems, and address issues using feedback from later stages of development.

"AI's impact on software will not be measured by how much code it generates, but by how much valuable software reaches customers," said Jyoti Bansal, founder and CEO of Harness.

He added that the integration would allow downstream agents to send issues identified during testing, security checks, or deployment back to Cosmos, along with the context needed to fix them.

Harness Cosmos aims to automate engineering tasks

The Harness Cosmos Software Factory Agent is intended to handle tasks triggered by written requirements, assigned tickets, and reported bugs. A fleet of agents plans changes, writes code and tests, and opens pull requests for review.

Each agent operates in an isolated virtual machine (VM). Engineers are involved when human judgment is required, such as approving a design or completing the final merge.

The system follows a continuous workflow covering tickets, code, and reviews. If a reviewer raises a concern or an automated check fails, an agent can make corrections on the same pull request.

Teams can also adapt prebuilt agents called Project Builder, PR Author, Deep Reviewer, and PR Fixer to their codebases and deploy them across their organisations.

The Code Context Engine maintains a live map of the codebase to help agents account for existing code and dependencies. Model routing assigns tasks to appropriate artificial intelligence (AI) models, while integrations with GitHub, Jira, and Slack connect the workflow with commonly used engineering tools.

Shared memory allows the system to carry lessons from earlier reviews into subsequent changes. Versioning and budget controls are also part of the offering.

Connecting code context with delivery data

The acquisition brings Augment Code's Code Context Engine together with Harness's Software Delivery Knowledge Graph. The latter connects information about builds, deployments, infrastructure, security, and costs.

Harness says Cosmos will use delivery information before making changes, including relevant test cases, security requirements, previous incidents, and their resolutions. Once code is created, downstream agents can use codebase information to assess dependencies, risks, and required validation.

If testing or deployment identifies a problem, the agents can return the issue to the coding workflow for correction and verification.

Igor Ostrovsky, co-founder and CTO of Augment Code, said joining Harness was a step toward the company's goal of bringing AI across the software development life cycle, beginning with context-aware coding.

He said the combined effort could make more engineering work autonomous while allowing engineers to focus on decisions requiring human judgment.

Extending automation from development to production

Harness already offers agents covering several stages beyond code creation. Its Software Delivery Agent handles the movement of changes into production, while the Security Testing Agent addresses risks before release. The Runtime Protection Agent focuses on applications, application programming interfaces (APIs), and AI in production. The Cost Management Agent manages cloud and AI spending.

The integration plan also covers policies, permissions, and approvals already used across the Harness platform. The stated aim is to apply the same controls to code creation that enterprises use when releasing software.

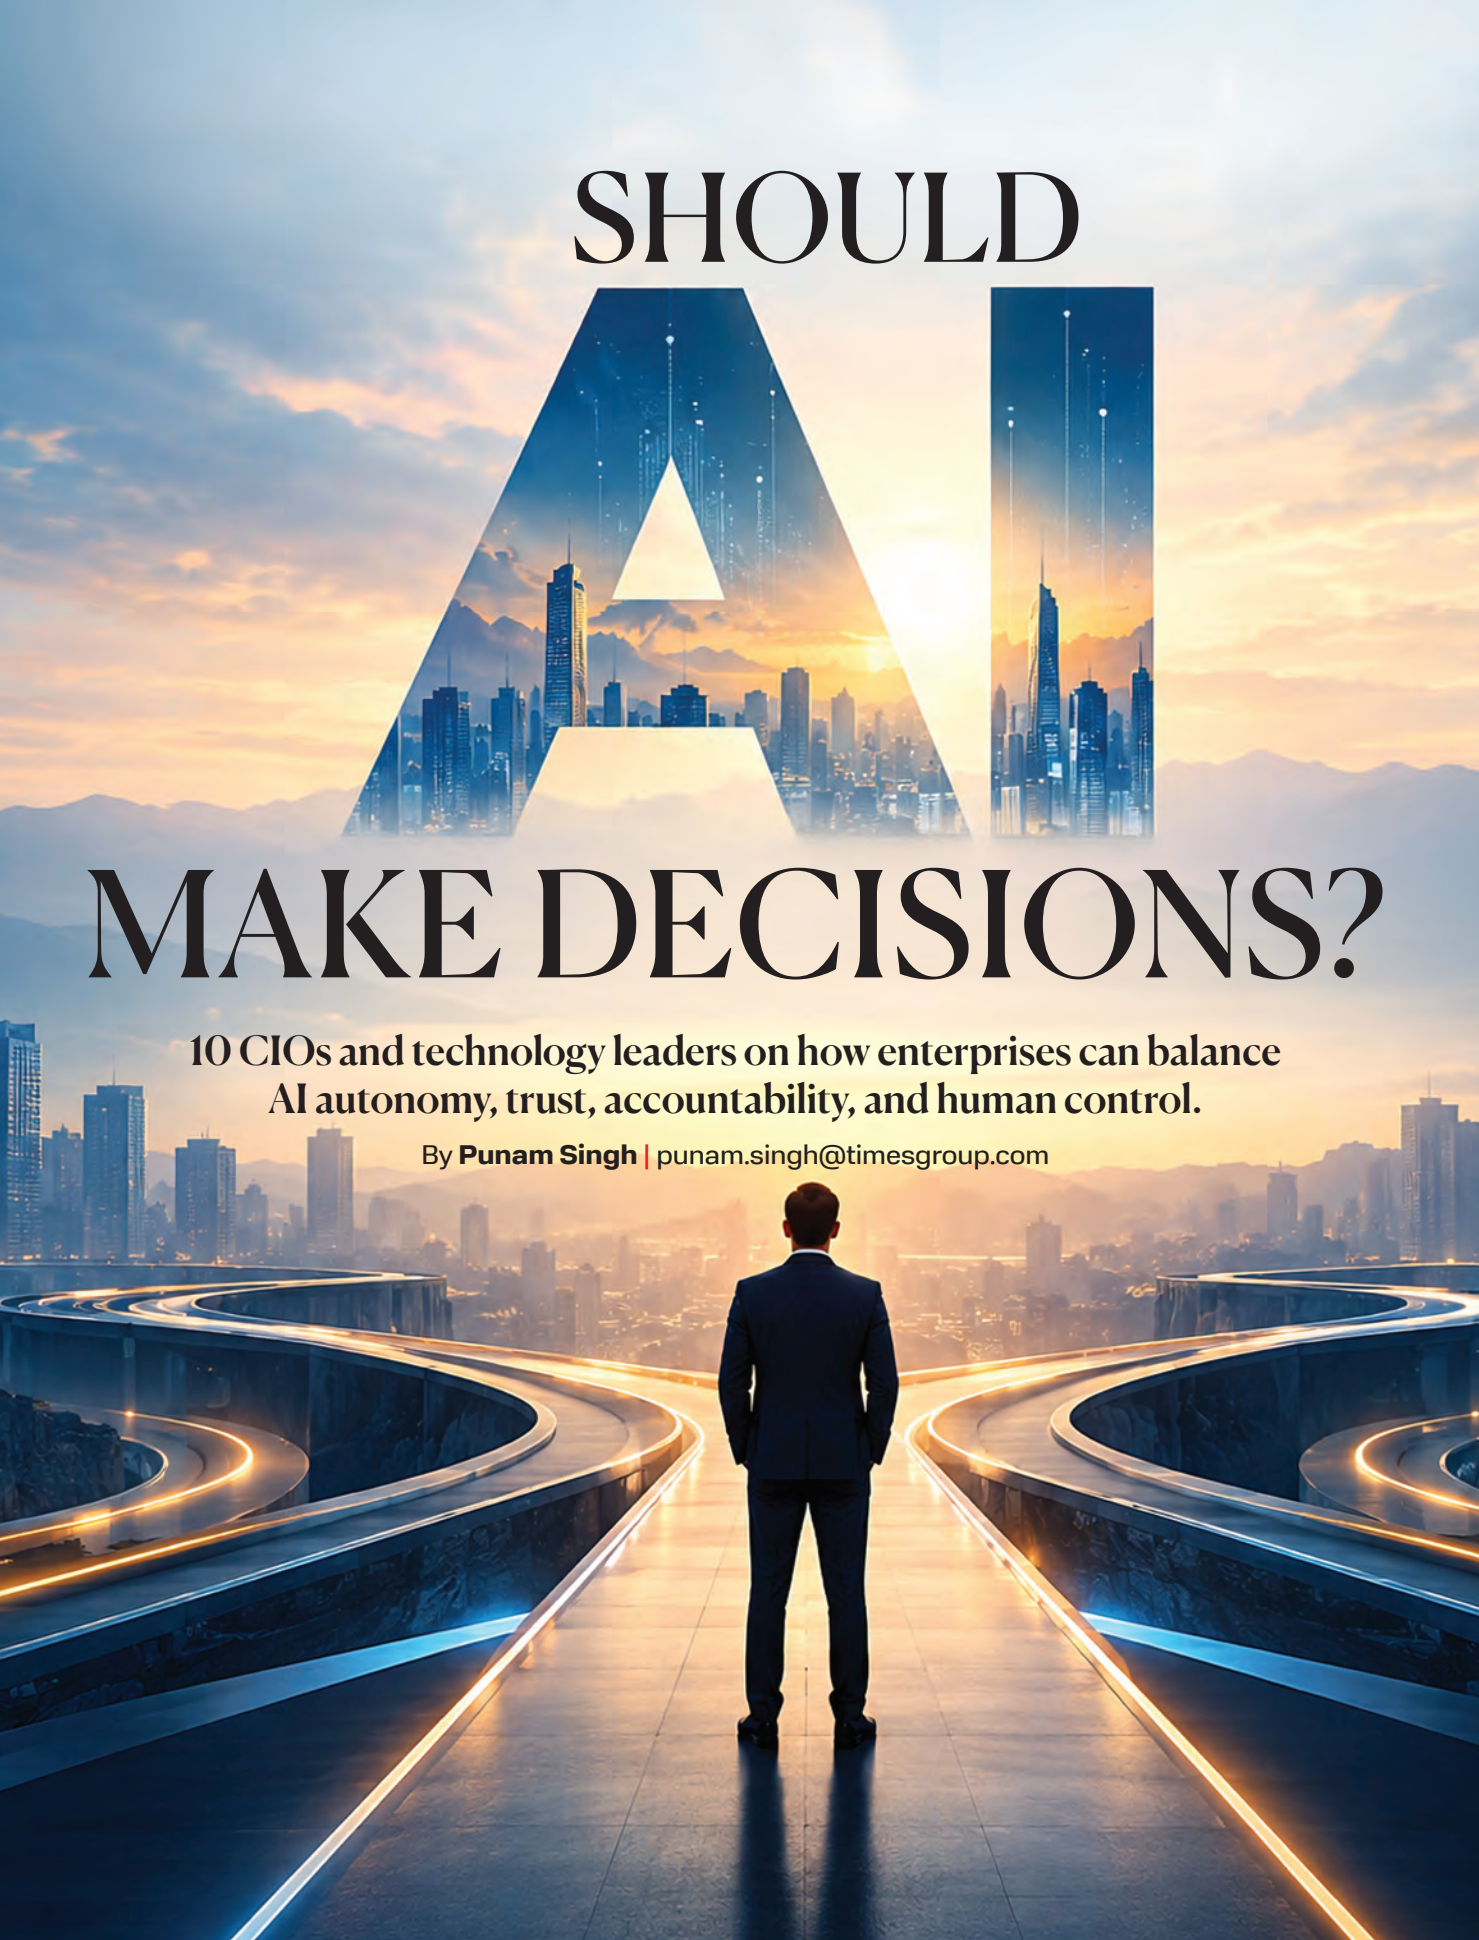
The announcement outlines a broader approach to autonomous software development in which AI agents handle repetitive engineering tasks and exchange information across the development pipeline. Engineers would retain responsibility for decisions requiring judgment, while feedback from testing, security, and production could guide subsequent code changes.

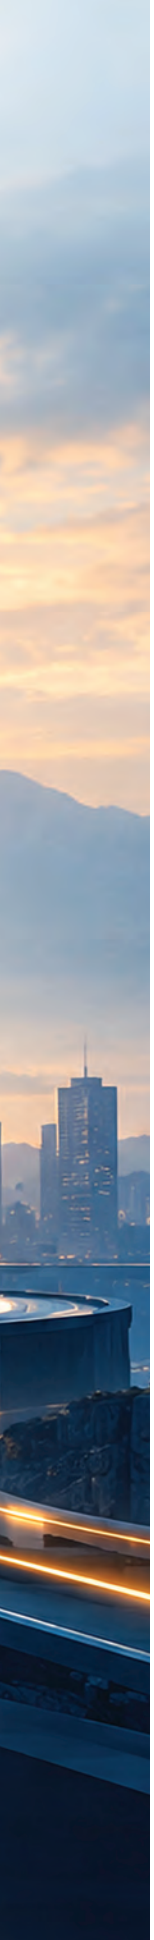
The announcement did not specify the financial terms of the acquisition or provide a timeline for the availability of the integrated offering. ■

SHOULD AI MAKE DECISIONS?

10 CIOs and technology leaders on how enterprises can balance AI autonomy, trust, accountability, and human control.

By **Punam Singh** | punam.singh@timesgroup.com





Enterprises have moved AI from pilot to production faster than they have built the means to supervise it. Assistants handle customer queries, models score credit and claims, and software agents increasingly take actions such as raising tickets, approving requests or changing system settings. Each of these is now a business process, with the usual obligations around accuracy, compliance and cost.

Most organisations are supervising these systems with instruments designed for different software. A dashboard that reports uptime and response times can confirm that an AI application is available. It cannot confirm that its answers are correct, that its decisions are consistent, or that it costs what the business case assumed. The gap between a system that is running and a system that is performing is where risk accumulates.

AI observability is the discipline meant to close that gap.

Why AI defeats conventional monitoring

Conventional monitoring rests on determinism: the same input produces the same output, so any deviation is a defect. Application performance monitoring, which tracks availability, speed and error rates, works well on that premise. AI weakens it in several ways at once.

Outputs are probabilistic. A large language model, which generates text by predicting likely continuations, can respond differently to identical prompts, so quality has to be judged statistically across many interactions rather than verified case by case. Failures are also silent. When a model fabricates a figure or misreads a policy, the application still returns a response promptly and without an error code. The fault is semantic rather than technical, and tools that measure only technical health cannot see it.

The architecture compounds the problem. An enterprise AI workflow typically combines a model, often supplied by a third party, with a retrieval layer that draws on internal documents, a set of connected tools and the instructions, known as prompts, that steer the whole. A change to any one of these alters outcomes, and several sit outside the enterprise's direct control. Providers can update or retire model versions on their own schedule, which means the system validated at launch may not be the system running months later.

Finally, agency raises the stakes. As systems move from advising to acting, errors become executed transactions. A flawed recommendation can be caught by a reviewer. A flawed action taken at scale may have propagated long before anyone looks.

The central supervisory question therefore changes. It is no longer whether the system is up, but whether it is behaving as intended and whether the enterprise can evidence that.

What observability adds

Observability is the capability to understand a system's internal state from the data it produces. Its raw material is telemetry: logs, which are time-stamped records of events, metrics, which are numerical measurements, and traces, which follow a single request through every component it touches. For AI, telemetry has to be extended beyond performance to substance. That means capturing the prompt, the documents retrieved, the tools called, the model and prompt versions in use, and token consumption, since tokens, the units of text a model processes, drive both cost and latency.

Traces are the most useful of these for diagnosis. An end-to-end trace shows what a model was given, what it did with it and what it returned, so that a wrong output can be traced to its cause, whether a stale document, a faulty tool call or a changed model version, rather than argued about. Without traces, root-cause analysis in AI systems tends to be guesswork, and remediation tends to be a change made in hope rather than in evidence.

Traces explain individual incidents, but risk is a property of patterns, which is why model behaviour must be measured in aggregate. Useful indicators include accuracy against reference answers, groundedness, meaning the degree to which an answer is supported by the sources retrieved, refusal and escalation rates, and the frequency with which human reviewers override the system. These are evaluated continuously in production, not only before release, because pre-launch testing samples a world that production will not stay faithful to.

Continuous measurement is what makes drift detectable. Drift is the gradual divergence between the conditions a system was validated against and those it now faces: customer language evolves, source documents age, regulations change and models are revised underneath. Performance erodes without any outage to announce

it. The remedy is to establish behavioural baselines at launch and alert when production behaviour departs from them beyond agreed thresholds.

Explainability completes the picture, and it is frequently overstated. For many models a full account of internal reasoning is not currently achievable. What is achievable, and what regulators, auditors and customers generally need, is an account of context: the inputs, rules and evidence behind a decision and the route it took. That is built directly from the telemetry above, and it is what turns an AI decision from an assertion into a record.

Illustrative example: A drift caught early

Consider a mid-sized lender that uses an AI assistant to summarise loan applications for credit officers. After the model provider releases an update, the assistant begins to give less weight to income-verification documents when summarising files. Infrastructure dashboards remain green, because response times and error rates are unchanged, and the summaries still read fluently, which is precisely why the change goes unremarked by users.

The problem surfaces through behavioural indicators. Groundedness scores for a particular document type fall, and credit officers start overriding summaries more often. Trace review on a sample of affected files isolates the cause to the new model version's handling of retrieved documents.

The resolution is mechanical. The lender pins the previous version while the update is tested, adds the failing cases to its regression suite, and uses the stored traces to identify and re-review the applications summarised since the change. The exposure is bounded to a defined set of files rather than an unknown portion of the book, and the lender can show its regulator which decisions were affected, how they were found and what was done. The same stored traces also let the team confirm, after the fix, that behaviour has returned to its baseline rather than assuming it has.

Without that telemetry, the same fault would most likely have emerged through rising defaults, long after the point at which it was cheap to fix, and with no dependable way of separating affected decisions from sound ones.

Trust, risk, accountability and return

These mechanics matter at board level because they answer four executive concerns directly.

On trust, observability supplies evidence rather than assurances. Customers, regulators and employees extend more autonomy to systems whose behaviour can be demonstrated. On risk, it moves detection from the customer complaint or the audit finding to the dashboard, shortening the interval between a fault and its correction. On accountability, it establishes who owned a system's behaviour and what was known at the time of each decision. Leaders should assign named business ownership for every AI system that affects customers, money or compliance, and set the evidence standard before launch: what must be recorded, for how long, and subject to what privacy controls.

> **"The gap between a system that is running and a system that is performing is where risk accumulates."**

On return, observability exposes the true cost of a task. A pilot that removes effort at one step may add it elsewhere, in model usage, integration work and the human time spent checking outputs. Measuring cost per completed outcome, with review effort included, is the only basis on which a business case can be validated or refuted. Procurement should reflect the same logic, with supplier terms that provide notice of material model changes and access to the records needed to investigate them.

Looking ahead, scrutiny of automated decisions is likely to intensify, and authority granted to AI will increasingly be expected to track the evidence that it behaves. Organisations with the telemetry in place will be able to extend that authority faster.

AI systems cannot be supervised by checking that they are switched on. Their risks lie in behaviour, and behaviour can only be governed if it is observable. Enterprises that build this capability alongside their AI deployments will detect faults earlier, answer regulators faster and know what their AI actually costs and delivers. Those that do not will be running consequential processes on trust alone, and will discover the limits of that approach at the point where it is most expensive to learn them. The sequence is straightforward: instrument first, baseline early, assign ownership and extend autonomy only as the evidence supports it. ■

Disclaimer: This cover story features insights from 10 industry leaders, originally published in the Enterprise AI Strategists book, a thought-leadership initiative by Dynatrace, AWS, and CIO&Leader.

No Automatic Remediation Without Complete Observability

Automatic remediation requires complete observability, proven accuracy and graduated trust.

AI-enabled cyber operations should mature from detect and diagnose to reporting and only then to remediation. Automatic action is unsafe without complete observability, demonstrated accuracy and checks and balances.

I describe the maturity journey as detect, diagnose and remediate, but the final R is a journey from report, recommend, and then move on to autonomous remediation. The system must first demonstrate that its findings are accurate and consistent. Only after the organization 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 model and tools invoked, the permissions exercised, the reasoning or execution path, and the effect of the resulting decision. Without that chain, accountability remains incomplete.

The practical discipline is to make evidence visible. In this context, the relevant signals include DDR (Detect, Diagnose, Remediate) maturity model. Adopting a SAAS (Speed, Accuracy, Automation & Analytics, and Scale) model to deploy AI/AIOps for justifying ROI is critical. These should be evaluated as part of one operating picture, so technical performance, decision quality and business consequences are not reviewed in isolation.

Governance must operate continuously, not only during approvals or audit time. Thresholds for human review, escalation, rollback and suspension should reflect the consequence of error. Drift, hallucination, anomalous behaviour and changing 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 turn this into a repeatable management practice. That means defining the purpose and baseline before deployment, observing performance and outcomes continuously, assigning accountability for exceptions, and expanding authority only when evidence supports it. A pragmatic cyber-risk threshold for moving from insight to action. This is how AI moves from an interesting capability to a dependable part of enterprise operations.

This perspective also clarifies what assurance should mean in practice. A governance approval records intent at a 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 uncertainty visible, contain its consequences and preserve accountability as the AI system changes in production. ■



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

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

When the Agent Becomes Part of the Workload

Dependency-aware observability connects infrastructure failures with downstream business disruption impacts.



> "A technically small dependency can cause a large disruption across downstream applications."

Aseem Joshi
Head - IT Infrastructure,
Aditya Birla Capital

In a multi-cloud financial environment, AI can correlate events, predict degradation and expose dependencies that conventional component monitoring misses. Safe automated remediation requires reliable signals, suppressed false alarms, tested runbooks and visibility into the resources consumed by agents.

In a multi-layered, multi-cloud environment, a failing third-party API can disrupt several downstream applications even when the core system appears healthy. Correlating events and predicting service degradation can shorten diagnosis. Runbooks may then automate basic troubleshooting, but false alarms must be suppressed and agent workloads themselves must be observed and resourced.

The practical test is whether operational signals can be connected to a business journey. Latency, model calls, token use and infrastructure consumption matter, but they become meaningful only when related to customer experience, revenue, productivity, risk, cost-to-serve or recovery. This demands baselines before deployment and continued measurement after the model enters production.

The practical discipline is to make evidence visible. In this context, the relevant signals include third-party API failure affecting downstream applications; predictive service-health correlation; runbooks; agent resource telemetry. These should be evaluated as part of one operating picture, so technical performance, decision quality and business consequences are not reviewed in isolation.

Leaders should also account for the cost of failure, not just the cost of running AI. An inaccurate recommendation can create rework, delay, customer friction, regulatory exposure or safety risk. Business observability provides the feedback needed to decide whether a use case should be scaled, redesigned or stopped, and whether apparent technical success is producing durable enterprise value.

For enterprise leaders, the priority is to turn this into a repeatable management practice. That means defining the purpose and baseline before deployment, observing performance and outcomes continuously, assigning accountability for exceptions, and expanding authority only when evidence supports it. Dependency-aware observability for both AI-run operations and the agents creating new workloads. This is how AI moves from an interesting capability to a dependable part of enterprise operations.

The leadership implication is that measurement must follow the full business journey rather than stop at the technical event. A useful baseline should show what the process cost, how reliably it performed and what customers or employees experienced before AI was introduced. Post-deployment signals can then reveal whether an apparent improvement is durable or simply shifts cost and risk elsewhere. This connection between operating evidence and business consequences gives leaders a defensible basis for scaling, redesigning or discontinuing the initiative. ■

Solve the Business Problem Before Scaling AI

Successful enterprise AI begins with a measurable business need and grows through dependable data, trust, governance and ownership.

A successful proof of concept shows an AI application works under selected conditions, not that it can improve a business process across an enterprise. Organizations must recognize this early. Scaling AI requires strategy, trust, and control to move together. Strategy ensures AI solves the right business problem; trust ensures people use its recommendations; control ensures AI scales securely and responsibly without unacceptable risk.

The first question is the business problem. Leaders should define what they expect to improve, agree with a measurable baseline, and decide who judges the result. For me, AI strategy starts with outcomes, not use cases: are we improving revenue, productivity, quality, cost, customer experience or risk? Otherwise, an impressive pilot can leave its practical value uncertain.

Data often exposes the gap between pilot and production. A pilot may use prepared data; scaling can require records from several systems with inconsistent quality, definitions, or labels. Cleaning and connecting it is less visible than model development but determines whether results can be trusted. Trust also means people know when to rely on AI recommendations, when to challenge them and when human judgement, or Human-in-Loop, must prevail.

Manufacturing and supply chain decisions show why context matters. Sales forecasts inform production plans, material requirements, and procurement. AI may recommend reducing safety stock based on demand patterns, but if it overlooks current stock, supplier reliability, lead times or operational constraints, an apparently efficient recommendation can increase production risk. Accuracy alone is not enough; AI must operate within business context and defined guardrails.

I see governance as part of process design, not a final pre-launch check. Teams must decide what the system can recommend or execute, and when a business owner should approve. Control should be proportionate to risk, with security and privacy designed from the start. As AI agents access enterprise data and applications, their permissions must match assigned tasks, making agent based access critical.

Functional business participation is essential throughout. IT can build and support a capability, but the business must define the requirement, validate the output and own adoption. That shared responsibility helps distinguish genuine value from a technically successful demonstration.

Sales forecasting, predictive maintenance, quality control, energy optimization and finance automation help build evidence through measurable results. For us, the route to scale is a sequence of sound business decisions: solve a meaningful problem, strengthen data foundations, prove the outcome, and expand with appropriate control. STRATEGY tells us where to apply AI, TRUST whether people adopt it, and CONTROL whether we scale it responsibly. When these come together with measurable outcomes, an organization starts becoming truly AI-mature. ■



"I am investing in business problems, not technology for technology's sake."

Atul Bhatia
Chief Digital & Information Officer, Greenlam Industries

Build Guardrails Before the Pilot Becomes Production

Guardrails designed during pilots enable confident scaling across complex factories.



"Guardrails are not a brake on scale; they are what make scale possible."

Madhavi Kanumoory
Chief Digital & Information Officer,
Birla Carbon

Industrial AI should be selected and governed through a value-cost-risk lens. Enterprises need to map process interdependencies, define expected outcomes and build guardrails during pilots so models can be trusted, rolled back and scaled confidently across factories.

In complex manufacturing, the process and its interdependencies must be mapped before AI is introduced. We evaluate opportunities through value, cost and risk. Guardrails and rollback cannot be postponed until after a promising pilot; they are what allow a model to earn trust and move confidently into production across factories.

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 model and tools invoked, the permissions exercised, the reasoning or execution path, and the effect of the resulting decision. Without that chain, accountability remains incomplete.

The practical discipline is to make evidence visible. In this context, the relevant signals include factory-process interdependencies; custom AI for high-impact areas; user demand for preferred LLMs; rollback; pilot-to-scale challenge. These should be evaluated as part of one operating picture, so technical performance, decision quality and business consequences are not reviewed in isolation.

Governance must operate continuously, not only at approval or audit time. Thresholds for human review, escalation, rollback and suspension should reflect the consequence of error. Drift, hallucination, anomalous behaviour and changing 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 turn this into a repeatable management practice. That means defining the purpose and baseline before deployment, observing performance and outcomes continuously, assigning accountability for exceptions, and expanding authority only when evidence supports it. Guardrails as a scale enabler in complex continuous manufacturing. This is how AI moves from an interesting capability to a dependable part of enterprise operations.

This perspective also clarifies what assurance should mean in practice. A governance approval records intent at a 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 uncertainty visible, contain its consequences and preserve accountability as the AI system changes in production. Intended business outcomes realisation must remain part of this continuing assurance, alongside value, cost and risk. ■

In Banking Outcome Integrity Comes Before Autonomy

Banking AI must demonstrate outcome integrity, human control and economics.



"For banks, trust and value have to be demonstrated together."

Pankaj Khare
Chief Information Officer,
IndusInd Bank

Banks have used AIOps for some time to correlate events, identify security incidents and generate actionable alerts. AI introduces a more demanding question: is the system merely functioning, or is it producing an outcome that the business can trust? Technical monitoring, security monitoring and business-outcome monitoring serve different purposes, and enterprises need all three.

In banking, data security and governance sit at the centre of every AI decision. When an agent produces an answer, the discussion cannot stop at whether the result looks useful. We must ask whether the underlying data is trustworthy, whether the output is consistent with the intended use case, and whether the institution is prepared to stand behind the resulting decision.

The stakes become clear in underwriting or credit scoring. A fluent response is not sufficient evidence for a lending decision. Hallucination, incomplete context or compromised data can translate directly into financial, regulatory and customer consequences. Human approval will therefore remain necessary wherever the consequence of error is high, even as agents assume more of the analytical workload.

Monitoring must cover the full chain of information feeding the decision, not only the health of the pipeline. It should help teams understand what changed, which source influenced the result and when intervention is required. This is where outcome assurance becomes as important as system availability.

Economics must be governed with the same discipline. An organisation can save money in one activity and spend two or three times more elsewhere on models, infrastructure, integration and control. The right measure is not the apparent efficiency of one step, but the total cost of producing a safe, accurate and usable business outcome. For banks, trust and value have to be demonstrated together.

This discipline should begin with explicit baselines for decision quality, turnaround time, customer impact and total operating cost. It should continue through periodic testing and live monitoring, with thresholds that reflect the materiality of the decision. A low-risk recommendation and a credit decision should not receive the same degree of autonomy or scrutiny. By aligning monitoring with consequence, banks can increase the useful role of AI without treating efficiency as a substitute for institutional accountability. It also enables management to compare the promised benefit with the complete cost of controls, exceptions and remediation. That comparison should continue as volumes, models and customer behaviour change. ■

Why Enterprise AI Is Still Waiting for Its ERP Moment

Enterprise AI needs infrastructure visibility, data control and production confidence.



"Success should be measured through production outcomes, not the number of pilots launched."

Rajeev Singhania
Group Chief Information Officer,
Jindal Poly Films

Enterprise AI in India currently resembles the early phase of ERP adoption. Interest is widespread and the underlying technology is advancing quickly, but production success remains uneven. The constraint is no longer access to AI tools. It is whether enterprises have the visibility, data protection and operating confidence required to place those tools inside core workflows.

AI agents change the risk profile because they do more than generate responses. They may access databases, use enterprise context and take action across systems. That requires a level of transparency that many organisations are still building. Moving a pilot into production is difficult when leaders cannot clearly see how specialised resources are being consumed or how enterprise data travels through the workflow.

Infrastructure visibility is a practical example. Conventional cloud dashboards have made CPU and memory consumption relatively easy to understand. AI workloads depend heavily on GPUs, yet the dashboards and environments available to many end-user organisations do not offer the same clarity. Limited visibility makes it harder to manage utilisation, performance and cost at scale.

Data security creates a second barrier. Enterprises need confidence that proprietary information will remain within approved environments and will not be exposed through public services or shared contexts. This is why tenant-controlled deployments are often viewed more favourably than open-ended use of public tools. The decision is less about a particular product and more about control over data, identity and access.

India's ERP journey eventually matured as implementation practices, controls and leadership understanding caught up with the software. AI is likely to follow a similar path. Enterprises will move faster when infrastructure consumption becomes visible, data boundaries are clear and success is measured through production outcomes rather than the number of pilots launched.

The next stage of maturity will require AI infrastructure to become as operationally legible as other enterprise platforms. Leaders should be able to connect workload demand, specialised compute, data movement and model behaviour with the business process being served. That visibility supports better capacity decisions, clearer accountability and more credible economics. It also gives boards and business leaders a practical basis for distinguishing a promising demonstration from a production capability that can be governed, funded and trusted over time. Operational transparency is therefore part of the adoption model, not a technical concern to be resolved after deployment. ■

Production Monitoring Starts in the Pilot

Production assurance begins by demonstrating monitoring during the pilot itself.

AI now sits across a wide range of enterprise activities, from factory and operations data to underwriting, loan documentation and employee productivity. Much of this capability arrives through partners and technology providers, and teams often understand it at different levels. A common enterprise playbook is therefore essential: it creates a shared language for what should be measured, approved and monitored.

My central requirement during a proof of concept is straightforward: show me the monitoring. I do not want a system to enter production before the organisation has worked out how it will observe its behaviour. Monitoring should be demonstrated during the pilot, alongside the capability itself, because the pilot is where assumptions about data, workflows and control can still be tested safely.

Financial services already provides a useful discipline. Risk, compliance and audit sign-offs are established parts of deployment. AI should sit inside that operating model rather than be treated as an exception because it is new. What changes is the evidence required: teams must be able to see how prompts, applications, workflows and outputs behave over time.

It is also important to be precise about what enterprises build. Very few organisations develop foundation models themselves. Even when proprietary data is used, the core model is generally supplied by a small group of providers. Enterprise co-creation usually happens in the application and workflow layers: how data is retrieved, how prompts are designed, what tools are called and how an output enters a business process.

That distinction helps focus observability where the enterprise has responsibility and control. Leaders need visibility into their own data flows, applications, decisions and outcomes, while also understanding dependency on the underlying model provider. Production assurance starts by making this division of responsibility explicit - and by proving the monitoring before the use case is allowed to scale.

The pilot should therefore be treated as a rehearsal for the complete operating model. Alongside accuracy and usability, teams should test ownership, escalation, access control, cost visibility and the evidence available when an answer is challenged. This also helps expose gaps between a provider's responsibility and the enterprise's responsibility before they become production disputes. A use case is ready to scale only when both the capability and the mechanisms for understanding its behaviour have been demonstrated under realistic conditions. Monitoring readiness should be a formal production gate. ■



"Show me the monitoring during the pilot, before the system reaches production."

Ramesh Narayanaswamy
Group President - Digital & Data Intelligence, Hinduja Group

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

Secure manufacturing requires zero-tolerance data protection and predictive operational resilience.



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

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

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

Manufacturing payment cards and other secure products place information protection ahead of novelty. Security 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 envelope.

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 model and tools invoked, the permissions exercised, the reasoning or execution path, and the effect of the resulting decision. Without that chain, accountability remains incomplete.

The practical discipline is to make evidence visible. In this context, the relevant signals include payment-card certification ecosystem; high-security number plates; predictive machine notifications; quality, shipment and back-office applications. These should be evaluated as part of one operating picture, so technical performance, decision quality and business consequences are not reviewed in isolation.

Governance must operate continuously, not only at approval or audit time. Thresholds for human review, escalation, rollback and suspension should reflect the consequence of error. Drift, hallucination, anomalous behaviour and changing 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 turn this into a repeatable management practice. That means defining the purpose and baseline before deployment, observing performance and outcomes continuously, assigning accountability for exceptions, and expanding authority only when evidence supports it. A rare perspective combining regulated manufacturing, physical products and security certification. This is how AI moves from an interesting capability to a dependable part of enterprise operations.

This perspective also clarifies what assurance should mean in practice. A governance approval records intent at a 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 uncertainty visible, contain its consequences and preserve accountability as the AI system changes in production. ■

Why We Moved from Buying AI to Building Context

Trusted enterprise context creates defensible value as model intelligence becomes commoditised.

Enterprise AI is evolving rapidly. The first phase was about access to intelligence; the next is about context. As foundation models become more capable and accessible, differentiation will come less from the model than from proprietary data, domain knowledge, workflow integration and governance.

Enterprises are moving beyond SaaS experiments to build proprietary context on governed data foundations. The real opportunity lies in owning not the model, but the context in which it operates.

Our own journey reinforced this. Buying AI solutions accelerates learning, but long-term value requires an internal AI foundation. Our early AI-powered hiring experiments were insightful, but we realised sustainable value needs a central data foundation, selective use of internal and external models, and disciplined focus on one scalable workflow rather than disconnected pilots.

Most proofs of concept succeed in controlled environments. Real AI agent deployment must address process reimagination, data quality, governance guardrails, change management, security and cost discipline. Moving from pilot to production is an operating-model challenge, not just a technology one.

Measurement must change too. Many AI programmes still track activity: users, prompts or tools deployed. These indicate adoption, not value. Value lies in business outcomes: faster decisions, better customer experience, stronger compliance, reduced risk and measurable productivity gains. AI return on investment must be judged by business value, not technology usage.

As powerful models become universally available, advantage shifts to enterprise context built through trusted data, institutional knowledge, historical decisions, domain expertise and workflow integration. The path to AI success is not from models to outcomes, but from enterprise systems to trusted data, intelligence and, ultimately, trusted autonomy.

Indian enterprises have a significant opportunity. Businesses operate across diverse environments, complex supply chains, multilingual workforces and fast-evolving markets. AI creates value when grounded in these realities and connected to trusted data.

The fundamental question for leaders is: How do we operate AI as a trusted enterprise capability?

Trust cannot be assumed; it must be engineered. Trusted data, clear guardrails, decision traceability and measurable outcomes form the foundation of operational trust and the confidence to scale AI enterprise wide.

Enterprise context becomes the true differentiator. Organisations that combine AI with proprietary knowledge, trusted data and operational discipline will build hard-to-replicate advantages.

The future belongs to organisations that combine intelligence with trust. ■



"Model intelligence may commoditise, but trusted enterprise context creates defensible value."

Tarun Anand
Group CIO,
Jakson Group

From Data to Intelligence to Autonomous Action

Autonomous enterprises turn intelligence into action. The real challenge is defining AI's authority while keeping humans accountable.



"Autonomy is the final stage of a longer journey from data to intelligence."

Viraj Deshpande
Chief Information Officer & Senior Vice President, Reliance Industries

Autonomous enterprises do not emerge simply by deploying AI agents. They evolve from data and automation, to intelligence, and ultimately autonomous actions. Dependable autonomy needs data discipline, integration across ecosystems, effective governance, and a workforce equipped to work alongside intelligent systems.

In a large industrial ecosystem, API-based automation across trade partners, shipping lines, logistics networks, and government portals is the foundation. The next level requires disciplined data capture, contextual intelligence, and AI skills.

The strategic question is not "Can we automate this task?" but "How much authority should we delegate to the machine?" Delegation requires defined decision rights, boundaries, accountability and a way to return control to human judgement when risks or consequences change.

Autonomous decision-making depends on bringing the right signals together into one operating picture: sales and service operations, trade and logistics partners, government portals, data lakes, customer data, protocols such as MCP, and workforce capabilities. Technical performance, decision quality and business outcomes must be viewed together; an agent that performs well technically but delivers poor business outcomes is merely automated, not autonomous.

The goal is not to place an agent on an existing process but to redesign the workflow with business, operations, risk, and security. Wrapping an inefficient process in an agent automates delay, ambiguity, and poor decisions. Sustainable autonomy comes from simplifying the process, clarifying ownership, establishing trusted data, defining decision rights, embedding controls, and enabling intelligent execution. Human intervention is not a failure of autonomy; it should be designed for exceptions, ambiguity, and high-consequence decisions.

Autonomy should be a repeatable management discipline, not a one-time initiative. Before deployment, leaders should define purpose and baseline, decision boundaries, evidence requirements, guardrails, observability, exception of ownership and the evidence-based conditions for expanding authority. Transitions between maturity stages, from API automation to autonomous operations, should rest on evidence, not enthusiasm.

As enterprises adopt agentic AI, focus should shift from model costs to business value per unit of intelligence consumed, linking token consumption to revenue, margin, productivity and risk. The objective is not to minimize tokens, but to maximize business value per token.

CIOs and business leaders must decide what AI may support, recommend or execute, what evidence must precede expanded authority, and where human judgement remains central. The real test comes at the boundaries: incomplete data, conflicting objectives, or consequences beyond the workflow. That is how AI becomes dependable, accountable, and scalable. ■



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

India's DPDP framework demands more than static consent and periodic audits. Continuous monitoring, runtime controls and traceable accountability are key to responsible AI governance.

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

Enterprises are moving from generative AI assistants that answer prompts to 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. An agentic DPIA, revisited as agent scope and tool access change, keeps the assessment a living record rather than a one-time document.

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) imply it. Agents with broad, standing access to data work against both. Consent-based access control, where consent and purpose 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. Synthetic data is not private by default, so empirical leakage checks remain necessary.

AI Assessment: Test Before Deployment

Section 8(3) requires completeness, accuracy and consistency where personal data is used to make a decision 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.



Abilash Soundararajan
Founder & CEO, PrivaSapien

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. A security gate can score prompts for jailbreak and injection patterns, and a safety layer can enforce refusal or redaction on outputs. Logging of these events supports breach detection and intimation under Section 8(6). It also supports the log-retention and monitoring expectations in the Rules' security safeguards provisions.

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. Autonomy makes purpose and data flow dynamic, so evidence must be dynamic too. Rights under Sections 11 to 14 illustrate this. A Data Principal's request for access, correction or erasure is only answerable if the organisation can trace which agents touched which data, for what purpose, and where inferences were stored.

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 assert it. ■

India leads in sovereign procurement intent, but has no plan to back it up



Everpure's Global Data Sovereignty Report 2026, based on 2,100 enterprise organisations across eight markets, reveals a defining disconnect between strategic priority and operational readiness.

By **Punam Singh** | punam.singh@timesgroup.com

"We talk about data sovereignty a lot internally now, but nothing has really changed, to be honest." A phrase attributed to a senior banking executive in Germany, offered to researchers conducting Everpure's Global Data Sovereignty Report 2026 captures the current state of enterprise data sovereignty more honestly than any boardroom presentation.

90 percent of respondents to the report consider data sovereignty a business concern. 64 percent have no formal data sovereignty strategy in place. The gap between those two figures is where most enterprises currently live — aware of the risk, unable to translate that awareness into organised action.

Why does the gap exist

Data sovereignty — where data is stored, who can access it, and which legal jurisdictions govern it — has shifted from a compliance consideration to a board-level strategic question. Escalating geopolitical instability, new AI workloads that require knowing precisely where data lives before it can be used and expanding regulatory frameworks like India's DPDPA have combined to make this question urgent in a way it was not three years ago.

Yet the report reveals that most organisational action is focused on immediate tactical measures rather than addressing long-term operational risk. The top three drivers of data sovereignty action are cybersecurity risk at 47 percent, new regulatory and compliance requirements at 39 percent, and AI adoption at 34 percent. Board-level strategic direction, listed by only 19 percent, ranks last. Organisations are responding to fires rather than designing fireproof buildings.

The tactical nature of current responses shows clearly in what enterprises have deployed. Encryption and access controls — implemented by 60 percent of respondents are the most widely used sovereignty measures. Sovereign or private cloud environments, which require deeper architectural commitment, are in place at only 53 percent. More tellingly, 56 percent maintain no mitigation strategy at all for data exfiltration or service disruption of geopolitical origin — despite identifying geopolitical instability as a concern.

What AI has changed about the stakes

If there is one force accelerating the urgency of this conversation, it is AI. The report makes the connection explicit: 92 percent of respondents view data sovereignty as a critical part of their overall AI strategy. 87 percent will prioritise local or sovereign environments for AI initiatives or workloads considered high risk or sensitive in the next one to three years.



Matthew Oostveen
CTO & VP Asia Pacific and Japan, Everpure

"We're in this interesting chasm between awareness and action. It's one thing for a thought leader inside an organisation to be aware of this, but you need to educate and cascade that information across quite a depth of people to come on board. It's not an initiative executed by an individual — it's something done in collaboration."

— Matthew Oostveen,
CTO & VP Asia Pacific and Japan, Everpure

The reason is structural. AI systems, and agentic AI systems in particular, require knowing what data exists, where it lives, who has access to it, and whether it can legally be used in a given pipeline — before it is fed into a model. The data discovery and classification work that data sovereignty demands are not separate from AI readiness. It is the same work.

46 percent of respondents regard sovereignty to ensure confidence and trust in AI adoption. 45 percent agree that data sovereignty directly influences where AI models and data are hosted. And 39 percent acknowledge that sovereign AI introduces additional operational complexity — particularly the challenge of migrating legacy systems to sovereign clouds without downtime, closing the gap between compliance expertise and cloud architecture capability, and finding sovereign providers that match the feature depth of hyperscalers.

India's contradictory position

India presents what is perhaps the most instructive case study in the report. Among all eight markets surveyed, Indian organisations lead in treating sovereignty as a significant consideration in vendor selection and RFP processes at 44 percent, compared to the global figure of 32 percent. India's enterprises are the most procurement-aware in the study.

The report also reveals that 95 percent of Indian respondents believe sovereignty strategy is very critical — yet 57 percent of those same respondents have no idea how to embed sovereignty into their own organisation or lack a strategy to do so. 40 percent are actively limiting the use of SaaS platforms that rely on non-domestic hosting infrastructure. 9 percent in direct tension with the stated commitment — cite reducing dependency on global hyperscalers as a top priority.

He pushed back gently on the framing of hyperscaler dependency as a binary choice: "I don't necessarily think people are walking away from the hyperscalers. I think people are putting contingency plans in place because we're continuing to see wonderful growth on the hyperscale side. It is possible to have an A and a B option."

The picture this draws of India's sovereignty posture is one of high intent, active early-stage action, and significant execution gaps — particularly in formalising strategy and building the internal skills to implement it.

The execution problem is a skills problem

Across all eight markets, the implementation picture shares a common anatomy. 57 percent of respondents struggle to balance sovereignty with wider strategic priorities. 56 percent lack the internal skills and capacity to deliver data sovereignty initiatives. 55 percent cannot keep pace with changing regulatory requirements. Only 19 percent report that their data sovereignty strategy is being driven by the board.

The investments, however, are rising sharply. 86 percent of organisations increased their sovereignty investments in the past six to twelve months. The data repatriation trend is accelerating just 1 percent had moved the majority of their enterprise data to local or sovereign infrastructure twelve months ago. Today that figure is 28 percent. In the one-to-three-year projection window, 41 percent plan to host the majority or all their data in local or sovereign environments.

The pilot-to-production problem

On the question of what is stopping AI pilots from reaching enterprise scale, Oostveen identified data quality as the single largest constraint.

"There's a really high failure rate taking an AI pilot to production. In my experience, it's about 75 percent. It comes down to tokenomics and economics — we got something to work in a lab, and then we went to run it inside an environment, and the cost of doing so was prohibitive. The second point is trust issues. While we might get something working in a pilot, hallucination rates could be too high because we haven't solved the data problem appropriately. We've got issues of concurrency — different versions and copies of data across an environment. One might be plus one day, minus two days. That adds complexity for AI to find out what the real data set is, so we're starting to see poor outputs."

— Matthew Oostveen, CTO & VP Asia Pacific and Japan, Everpure



"A lot of deployments are scaling, and people are taking the learnings from their pilot phase to the scaling phase right now. The difference you must understand is that the type of data which is getting used into these AI models is way too different from what was possibly existing five years ago. So, the sort of governance which you need to embed into these data formats is much different from what people have done in the past."

— Matthew Oostveen, CTO & VP Asia Pacific and Japan, Everpure

The urgency is real

The Everpure report's conclusion frames the next decade as a test of whether enterprises can close the execution gap while managing three simultaneous pressures: international regulatory demands, cloud scalability, and secure AI deployment.

Those three pressures are not in tension with each other in theory. In practice, they are in tension every day — in budget conversations where sovereignty competes with other priorities, in vendor selections where sovereign providers lag on features, and in skills markets where the combination of regulatory expertise and cloud architecture capability is genuinely scarce.

The organisations that resolve this tension first, move from the holding pen of awareness into the disciplined execution of a formal sovereignty strategy, will not just be more compliant. They will be more AI-ready, more trusted by their customers, and more resilient against the geopolitical volatility that shows no sign of diminishing. ■



AI agents touched sensitive data in 87% of firms, while only 19% caught it in real time

Delinea's 2026 report exposes a critical AI security gap: 87% of firms saw agents access excess sensitive data, but only 19% detected violations in real time.

By **Punam Singh** | punam.singh@timesgroup.com

Enterprise AI adoption has moved faster than the governance designed to contain it. That much is understood. What Delinea's 2026 Identity Security Report makes uncomfortably concrete is how wide that gap has become — and how consistently organisations on both sides of it are underestimating the exposure.

The report surveyed over 4,500 respondents across two distinct groups: IT and security leaders, and the non-IT employees who actually spin up enterprise AI tools day to day. The contrast between what each group believes is happening is the report's most important finding.

Two organisations, one security problem

On the IT side, 99.7% of leaders say a formal AI data-access policy exists at their organisation. On the employee side, 30% aren't sure a policy exists or cannot describe what it covers. IT says 19% of employees find workarounds without formal approval. Employees report the actual figure is 41%. IT says 96% could demonstrate to a regulator that AI stays within enforced boundaries. Meanwhile, 44% of employees say they've used AI tools that accessed more information than expected — and one in three of those never reported it.

These are not marginal discrepancies. They describe two organisations operating with fundamentally different pictures of the same risk environment.

The enforcement gap is structural

76% of employees say they have bypassed IT approval to use AI tools. 48% say they do so always or regularly. The reasons are not recklessness — they are systemic. When business deadlines require AI tools to move faster than governance processes allow, fewer than two in five employees say they'll stick to the approved path. An aggregate of 36% of respondents cite friction in the governed path itself — IT or legal is too slow, approved tools aren't effective enough, or governance rules are simply unclear.

When the compliant route is harder than the workaround, the workaround wins. That is not a culture problem. It is a design problem.

The report also surfaces a finding that should concern CIOs directly: the biggest shadow AI risk sits in the C-suite. 81% of C-level executives bypass AI access approval always or regularly, compared to 33% of intermediate staff. Despite having the highest awareness of any group that approval is required, executives are nearly three times more likely than intermediates to use unapproved AI because IT or legal responds too slowly.

AI agents add a layer of governance

Beyond human behaviour, the report identifies a deeper structural issue: AI agents are creating a new category of standing privilege that most identity programmes are not equipped to handle.

87% of IT leaders confirm that an AI tool or agent accessed sensitive data beyond what the task required in the past year. Yet only 19% were able to detect a scope violation in real time. For 55% of organisations, when an AI tool last accessed out-of-scope data, it took a full day or longer to detect. For 30%, that detection window stretched to four days or more.

Organisations are checking access at the door, then losing visibility entirely once an agent is inside the system.

The data AI agents touch is not trivial. Employee data, customer data, financial records, security logs, legal and compliance information — all appear among the most commonly accessed categories. And 42% of IT organisations have no automatic mechanism to revoke agent access when a session ends. Credentials, in many cases, remain live until manually disconnected or until an audit catches them.

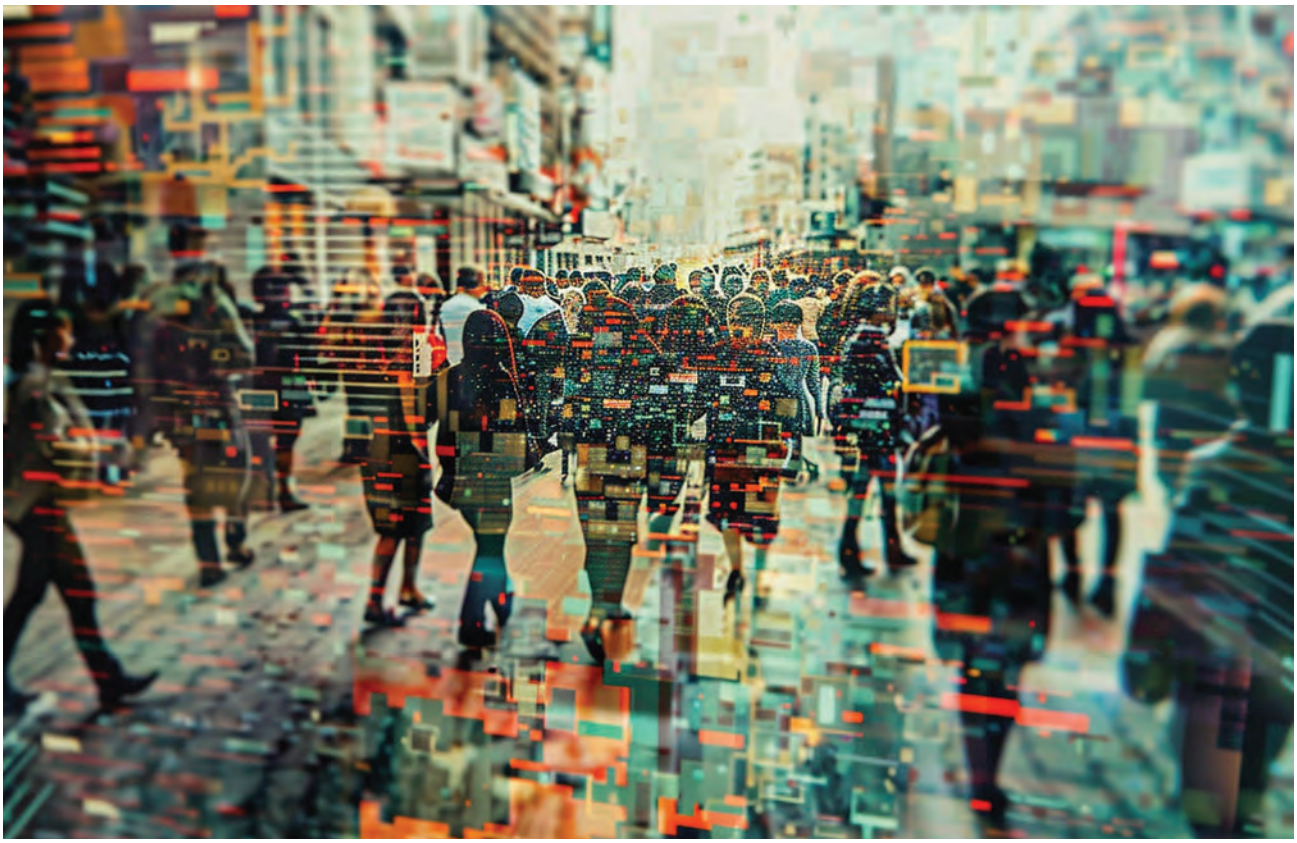
This is the enforcement gap in its most consequential form. Organisations are checking access at the door, then losing visibility entirely once an agent is inside the system.

What CIOs need to act on

The report's recommendations centre on a principle that identity security programmes have long accepted for human accounts but have not yet extended to agentic ones: scope authority to the task, and extend authorisation beyond the point of access to the moment of action.

For CIOs, the practical translation is clear. Policies alone do not constitute governance. Real-time enforcement, just-in-time access, and runtime authorisation — the ability to evaluate and act on what an agent is doing, not just what it was permitted to do at login — are what the agentic environment actually demands. The EU AI Act's enforcement deadline in December 2027 moves audit capability for AI-initiated activity from best practice to legal requirement. The window to close the enforcement gap before that deadline is narrower than most organisations currently appreciate. ■

Tokenization can be India's next digital public revolution



Tokenization could expand investment access, unlock capital for infrastructure and SMEs, and modernise financial markets, positioning India to shape the next phase of digital finance.

By **Rohit Dalmia** | editor.tech@timesgroup.com

India's digital transformation has demonstrated how technology, combined with visionary public policy, can reshape an economy. From Aadhaar establishing digital identity to UPI revolutionising payments, India has built digital infrastructure that has become a global benchmark. Today, tokenization offers another opportunity to redefine investment, ownership, and capital formation.

While blockchain discussions have often been overshadowed by cryptocurrency speculation, tokenization deserves a different lens. It is not about creating new forms of money, but modernising how assets are owned, financed, and traded.

Simply put, tokenization converts ownership rights in real-world assets into digital tokens recorded on distributed ledger technology. These assets can include commercial real estate, infrastructure, renewable energy, corporate bonds, commodities, intellectual property, and government securities. Tokens represent verifiable ownership interests, enabling more efficient and cost-effective transfers.

For India, this is an opportunity to democratise wealth creation.

Access to premium investment opportunities has traditionally been limited to institutions and affluent investors. High entry costs have kept ordinary citizens out of wealth-generating assets. Tokenization enables fractional ownership, allowing investors to participate in commercial property or infrastructure projects with smaller investments rather than committing crores.

This is particularly relevant as India's financial inclusion expands. Millions have bank accounts, digital payment capabilities, and growing financial literacy. The next step is broadening access to quality investment opportunities.

The benefits extend beyond individual investors' benefits. India's ambition to become a developed economy requires substantial long-term capital for infrastructure, manufacturing, logistics, clean energy and urban development. Tokenization can unlock new funding sources by making these projects accessible to a wider pool of domestic and international investors.

It could also help small and medium enterprises (SMEs) overcome financing constraints. Tokenized financial instruments can create more efficient capital-raising mechanisms, reducing dependence on traditional lending channels, and connecting businesses with investors seeking diversified opportunities.

Transparency is another advantage. Transactions recorded on distributed ledgers create auditable records that are difficult to alter unilaterally. This can improve trust, reduce reconciliation costs, and strengthen investor confidence. In sectors burdened by documentation and verification requirements, tokenization can simplify administration and enhance regulatory oversight.



Rohit Dalmia
Chairman & Managing Director,
CineNow

Efficiency gains are equally compelling. Financial markets often depend on multiple intermediaries, extensive paperwork and settlement cycles lasting days. Tokenized assets could enable near-real-time, round-the-clock settlement, reduce operational costs and counterparty risks while improving liquidity. Faster settlement also allows capital to be deployed more efficiently.

Real estate illustrates this potential. Despite being a preferred Indian investment class, property carries high transaction costs, limited liquidity and substantial entry barriers. Fractional ownership through tokenization could broaden access to commercial real estate and expand its investor base while retaining regulatory oversight.

Renewable energy presents another opportunity. Meeting India's climate commitments requires significant private capital alongside public investment. Tokenized financing could enable citizens and institutions to participate in funding solar parks, wind farms and green infrastructure, aligning investment opportunities with national sustainability objectives.

However, technology alone is insufficient. India's digital successes have relied on thoughtful regulation, interoperability and public trust. Tokenization requires a robust framework covering investor protection, legal certainty, cybersecurity, and anti-money laundering compliance. Regulatory clarity will encourage responsible innovation and build confidence among institutional and retail investors.

India already has important advantages: established digital public infrastructure, a mature financial ecosystem and globally respected technology talent. These provide a foundation for developing a trusted tokenization ecosystem.

The question is not whether tokenization will reshape global finance, but whether India will lead this transformation or adapt to innovations developed elsewhere.

India did not merely digitise payments; it reimagined them through UPI, creating an interoperable platform admired worldwide. Tokenization presents a similar opportunity. By establishing clear regulations, encouraging responsible innovation and fostering collaboration among government, regulators, financial institutions and technology companies, India can help shape global standards.

The country's digital journey has always been about expanding opportunities, not merely adopting technology. Tokenization aligns with that vision by democratising investment, improving capital formation, strengthening transparency, and making markets more efficient.

When policy, technology and public purpose converge, transformational change follows. Tokenization could become the next chapter in India's digital journey, positioning the country not only as a fast-growing major economy but also as an architect of the future global financial system. ■



When India celebrates, payment infrastructure cannot afford to rest

Festive transaction spikes test banks' payment resilience. Dynamic scaling, intelligent routing, and real-time monitoring are critical to sustaining transaction success and customer trust.

By **Chirag Rajda** | editor.tech@timesgroup.com

India's festive season is one of the most concentrated stress tests for payment infrastructure. During Navratri, Diwali and Christmas, transaction volumes across ATMs, card networks, point-of-sale terminals, UPI channels and digital banking platforms rise sharply. Purchases are concentrated within compressed timeframes; month-end payroll credits coincide with advance shopping, and cash demand increases. Every channel a bank operates is tested simultaneously, often when transaction density is highest and tolerance for failure is lowest.

The scale of festive transaction growth illustrates this pressure. During the Dhanteras-to-Diwali period in 2025, average UPI transaction volumes rose to 737 million transactions per day from 568 million during the corresponding period a year earlier, a 30 percent increase. While UPI

represents one part of the wider payments ecosystem, this growth demonstrates how sharply transaction loads can rise within a concentrated festive window.

For banks, the challenge extends across the payments ecosystem. Cards, ATMs, point-of-sale networks and digital payment channels can all come under pressure, making infrastructure scalability and transaction performance critical.

Uptime is where this conversation should begin, because conventional measures of infrastructure resilience can give a misleading sense of readiness. A payment switch running at 99.9 percent availability sounds acceptable. In practice, that translates into more than 40 minutes of monthly downtime. During most months, this may pass without significant consequences. If it coincides with a peak festive evening or the first day of a major sale, however, it becomes a service failure experienced simultaneously by customers, merchants and banking partners.

For banks, headline uptime cannot be the sole measure of resilience. The real test is how infrastructure performs under peak transaction volumes and simultaneous pressure across channels. How quickly can it detect degradation and reroute transactions? Can it absorb sudden spikes without response times exceeding what customers and merchants will tolerate?

The cost of failure extends beyond the original transaction. A failed card payment can trigger complaints, reversals, and exception management. An unsuccessful ATM transaction creates repeat attempts and reconciliation requirements, while a timeout during a high-value digital payment may require dispute handling across two banks and a payment network.

When daily transaction volumes reach hundreds of millions, even a small decline in transaction success rates can generate substantial numbers of failures, increasing operational workloads, and customer complaints. Transaction success rates, therefore, deserve greater importance alongside uptime as measures of payment infrastructure resilience.

Scalability is equally important. Infrastructure designed around average transaction volumes may struggle when festive demand spikes unexpectedly or remains elevated for extended periods. The result can be slower response times, higher decline rates, and cascading failures as processing queues exceed system capacity.

Banks must therefore design peak demand, not just average volumes. Dynamic scaling, which redistributes workloads across available capacity in real time, helps infrastructure withstand festive surges without compromising performance.



Chirag Rajda
Business Head, Writer Payments

Intelligent routing at the switch layer is central to this approach. When multiple card networks, payment rails and settlement pathways are available, detecting a degraded network and redirecting transactions to a healthier route can help maintain success rates. During festive surges, adaptive rather than static routing can directly influence transaction outcomes, and the volume of operational escalations banks must manage.

Real-time monitoring enables these capabilities. Systems that report problems only after they occur offer limited scope for prevention. Continuous visibility into transaction success rates, switch response times, network health and channel availability allows operations teams to identify and address degradation before it affects customers.

During peak periods, early detection can be as important as recovery after failure.

India's festive season is predictable. Banks and their infrastructure partners know when demand will rise and what it requires. Festive readiness should therefore be treated as an infrastructure-design priority, not merely an operational exercise undertaken when transaction volumes begin climbing.

Banks that build for peak resilience will be better positioned to protect transaction success, operational efficiency and customer trust when their infrastructure faces the greatest pressure.

At the scale of India's digital payments ecosystem, the difference between infrastructure designed for average conditions and infrastructure built for peak demand is no longer a technical detail. It is a strategic choice about how reliably a bank serves its customers when demand is highest. ■

The next phase of enterprise IoT is about intelligence, not just connectivity

Praveen Arora, Vice President-IoT, Tata Communications, explains how enterprises can turn connected data into business value through edge intelligence, AI, secure connectivity and increasingly autonomous industrial operations.

By **Punam Singh** | punam.singh@timesgroup.com

Enterprise IoT is evolving from connecting devices to enabling intelligent, predictive and increasingly autonomous operations. Across manufacturing, utilities, mobility and infrastructure, enterprises are integrating connectivity, edge computing, AI and security into unified operational frameworks.

In this conversation, **Praveen Arora, Vice President-IoT at Tata Communications,** discusses how enterprises can turn connected data into business value, balance edge and cloud workloads, strengthen security and prepare for AI-led industrial operations while retaining human oversight of critical decisions.

CIO&Leader: Enterprise IoT in India has shifted from isolated pilots to mission-critical deployments. What are the primary macro trends driving large-scale adoption across Indian and global enterprises today?
PRAVEEN ARORA: Enterprise IoT has evolved from isolated pilots into a core pillar of digital transformation. Organisations are replacing siloed deployments with integrated platforms connecting people, assets, machines and operations. By combining sensors, connectivity, cloud, edge computing and analytics, enterprises can turn operational data into real-time insights and measurable outcomes. Technologies such as eSIM and Remote SIM Provisioning (RSP) are also making it easier to deploy and manage connected devices across geographies and networks. Adoption is accelerating in proven use cases, including connected vehicles, smart manufacturing, asset tracking, worker safety and AI-powered video analytics, which improve visibility, reduce costs and enhance safety and efficiency. Enterprises are also looking for connectivity that can be managed consistently across locations, networks and device types, rather than treating each deployment as a separate project.

Another key trend is the convergence of IoT with AI and edge intelligence. Processing data closer to its source enables real-time decisions, predictive maintenance and operational optimisation while reducing bandwidth and cloud-processing costs.

CIO&Leader: Tata Communications provides cross-border fabric across connectivity, cloud, and edge. How does your global network infrastructure differentiate Tata Communications' IoT proposition from standalone system integrators or telco providers?
PRAVEEN ARORA: Our differentiation lies in bringing together devices, networks, applications, edge infrastructure and operational intelligence in a unified ecosystem. Rather than a one-size-fits-all approach, we combine LPWAN and Private 5G to address varied industrial and utility connectivity needs.

For large deployments such as smart metering, utilities and cities, LPWAN provides long-range, low-power, cost-effective connectivity across wide geographies. Our infrastructure supports some of the world's largest smart utility deployments and is designed to connect millions of



"The conversation has shifted from 'How do I connect assets?' to 'What business value does connectivity create?"

devices without placing unnecessary load on cellular networks.

For industrial environments, Private 5G extends this fabric to the edge, delivering ultra-low latency, high bandwidth, deterministic performance and secure on-premises connectivity. It supports mission-critical applications such as machine vision, video analytics, robotics, autonomous operations and connected-worker solutions.

Unlike standalone integrators or traditional connectivity providers, we go beyond connectivity by bringing together the Interaction fabric, network fabric, edge intelligence and operational platforms in a managed environment. Customers can deploy, secure, monitor and scale connected operations through one partner rather than coordinate multiple vendors. Our objective is to connect data, intelligence and operational workflows to deliver measurable business outcomes at scale. This integrated approach also helps enterprises simplify deployment and management as their connected environments grow, while maintaining a consistent view of performance across the technology stack.

CIO&Leader: How has the conversation around IoT monetisation shifted in the C-suite—are enterprise leaders primarily prioritising operational cost reduction (preventive maintenance, energy management) or net-new revenue streams enabled by connected products?

PRAVEEN ARORA: The conversation has shifted from "How do I connect assets?" to "What business value does connectivity create?"

Operational outcomes remain the primary driver of IoT investments. Predictive maintenance, energy optimisation, asset utilisation and worker safety continue to deliver the fastest and most measurable returns.

As connected ecosystems mature, enterprises are also exploring new revenue opportunities. In connected vehicles, for example, OEMs are evaluating data monetisation models involving infotainment, usage-based services and ecosystem partnerships. The common thread is that enterprises establish operational value first, then use connected data, services and ecosystems to develop new revenue streams and business models.

CIO&Leader: As ultra-low-latency applications grow, how is Tata Communications architecting the handoff between edge compute (processing data near the asset) and central cloud systems?

PRAVEEN ARORA: It is not edge versus cloud, but a balanced architecture that optimises network load and ensures rapid response times.

Real-time, low-latency workloads, such as computer vision for safety and productivity, can be processed locally at the edge, close to the camera or asset. AI model training, which requires learning at scale, is supported by the cloud.

Private 5G connects cameras, sensors, autonomous machines and connected-worker devices to on-premises edge compute with low latency, high bandwidth and reliable performance. Real-time decisions can therefore be made locally without depending on the public internet. Dedicated, secure connectivity also keeps high-volume video streams within the plant, while only relevant events and frames travel to the cloud over Tata Communications' network backbone. This reduces wide-area network load while supplying data for model training and enterprise integration.

Private 5G helps route each workload appropriately, with edge devices capturing incidents locally and sending relevant data to the cloud for longer-term storage, analysis and integration. Managing this through a single operational view balances workloads, optimises data flows and enables faster production-line responses. The architecture also avoids sending every data point to central systems when an immediate local response is required, while preserving the cloud's role in broader analysis and model improvement.

CIO&Leader: IoT deployments often generate massive volumes of unstructured or semi-structured telemetry data. How do you help enterprises ingest,

clean, and pipe that stream to data lakes or AI models without bandwidth bottlenecks?

PRAVEEN ARORA: IoT deployments generate data from thousands of sensors, connected-worker devices and operational systems. We create an integrated environment where edge intelligence identifies deviations and relevant events locally, so only required data is transmitted to central platforms.

Depending on the use case, we use Private 5G for high-volume workloads such as video and machine vision, and LPWAN for low-power sensors. We also collect connectivity insights, including network QoS data, and correlate them with operational and OEM data to improve business and operational efficiency. This edge-cloud approach optimises data flows, supports AI/ML-driven insights and ensures the right data reaches the right system at the right time, rather than moving everything across the network.

CIO&Leader: Connecting thousands of unmanaged endpoints significantly broadens an enterprise's attack surface. How do you address device-level vulnerability, mutual authentication, and zero-trust security from the edge device through the network core?

PRAVEEN ARORA: Security can no longer be treated as a perimeter challenge. We embed it across the IoT ecosystem, from edge devices and connectivity to platforms, applications and data. Secure device identities, certificate-based authentication, encryption, continuous monitoring and anomaly detection help protect connected environments end to end.

For global deployments, Tata Communications MOVE provides a regulatory-compliant connectivity framework supporting data privacy and local data-storage requirements across markets. IoT Safe further secures the data path by enabling protected communication between devices, networks and applications.

"It is not edge versus cloud, but a balanced architecture designed to optimise network load and ensure rapid response times on the line."

Our goal is an end-to-end zero-trust framework in which devices, networks, applications and data are continuously authenticated, encrypted, monitored and protected, allowing enterprises to scale IoT securely and confidently. This is particularly important as deployments expand across regions with different privacy obligations and data-localisation requirements.

CIO&Leader: How is the convergence of IoT and Generative AI/Agentic systems transforming industrial automation—are we moving from simple threshold alerting to autonomous machines that diagnose and resolve telemetry anomalies independently?

PRAVEEN ARORA: The convergence of IoT and generative AI is accelerating the shift from threshold-based alerts to intelligent, predictive operations. Instead of reacting after an issue occurs, organisations can analyse machine, sensor, video and connected-worker data to identify patterns, detect risks early and take preventive action.

AI's value today lies in strengthening human decision-making. Machine and vibration data can support predictive maintenance, computer vision can improve quality inspection and safety compliance, and combining data sources can reveal risks that might otherwise go undetected. These capabilities enable faster decisions, better efficiency and improved worker safety.

As generative AI and agentic systems mature, they will increasingly coordinate actions across systems, recommend corrective measures and automate routine workflows. However, in industrial settings where safety and business continuity are critical, human oversight remains essential. The goal is not simply fully autonomous machines, but intelligent systems that understand context, predict issues, recommend actions and sometimes automate responses while keeping people in control of critical decisions. This can deliver safer, more efficient and resilient industrial operations.

CIO&Leader: Looking across the next 24 to 36 months, what are your major investment priorities for Tata Communications' IoT portfolio regarding platform capabilities, partner ecosystems, and new market expansion?

PRAVEEN ARORA: Over the next 24–36 months, our priority is to evolve the IoT portfolio from individual solutions into a unified, intelligence-led industrial operating platform connecting people, assets, facilities, utilities and operations through a common digital fabric. The focus is to help enterprises progress from connectivity and visibility to intelligence, prediction and, ultimately, autonomous operations.

We are investing in AI, generative AI, machine learning, computer vision, digital twins, unified data fabrics and edge intelligence to create a common operational layer. Across worker safety, asset visibility, factory operations, utility management and smart infrastructure, the objective is to bring data into a single intelligence layer that delivers predictive insights, anomaly detection, operational recommendations and automated decision-making.

Ultimately, our vision is to move customers from connected operations to intelligent and eventually autonomous operations, while keeping people, safety, resilience and sustainability at the heart of every transformation. ■

Resilience is one mission, not two scorecards

Subhalakshmi Ganapathy of ManageEngine explains why unified data, AI-assisted response and shared accountability are essential to building trust and strengthening resilience.

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

Attackers do not respect org charts, and neither do outages. Yet many enterprises still keep IT operations and security in separate tools, data, and narratives, leaving early warning signs unconnected until damage is done. In this conversation with CIO&Leader, **Subhalakshmi Ganapathy, Chief IT Security Evangelist at ManageEngine, Zoho Corp.,** argues that convergence begins with shared incidents, not restructured teams. She explains why a common data foundation matters, where automation must hand over to human judgment, and why mindset, not tooling, is the hardest barrier. Her larger point: resilience must become a shared metric, built on trust in teams, data, and AI.

For insurers, however, greater autonomy comes with an equally important responsibility: ensuring that AI-driven decisions remain transparent, secure and accountable. The challenge is no longer simply how much can be automated, but where autonomy creates value, and where human judgment must remain in the loop.

CIO&Leader: When IT operations and security teams have worked in silos, what gaps or blind spots does that create for enterprise resilience?

SUBHALAKSHMI GANAPATHY: Adversaries don't respect silos, and neither do incidents. Attackers often target operations to disrupt systems, and early indicators of a breach may appear as operational symptoms such as unexplained latency, system downtime, failed jobs, or performance degradation. When IT and security work separately, IT may see a performance problem and try to restart or scale a system. In contrast, security sees a faint anomaly without the operational context to prioritize it. Neither team has the complete picture, and the underlying cause can go undiagnosed until the impact becomes significant. The blind spots extend beyond incident response. Duplicate tools can tell different stories, post-mortems can have conflicting timelines, and unclear ownership can create gaps between teams. This fragmentation slows detection, weakens response, and can ultimately affect business trust. Enterprise resilience requires a shared view of what is happening across the environment, rather than separate operational and security narratives.

CIO&Leader: What does "unified operational visibility" concretely look like in practice, and how does it build trust in shared data? Is it a shared dashboard, a shared team structure, or something deeper?

SUBHALAKSHMI GANAPATHY: It's deeper than a shared dashboard, although that may be the most visible layer. True unified visibility means IT and security teams work from the same telemetry, including logs, network flows, endpoint signals, and application performance data, normalized into a common data fabric. Both teams see the same event with the same context, whether it initially appears as a performance issue or a security anomaly.

Trust in that data comes from three things: a single source of truth (no duplicate tools telling different stories), agreed standards for data ownership and quality, and joint review processes where both teams validate what they are seeing.

Structurally, this can take the form of a fusion center or a unified SecOps/ITOps function. But the deeper shift is cultural, where operational and security signals are treated as two lenses on the same reality rather than as competing narratives.

CIO&Leader: For an enterprise still running IT and security as separate functions, where would you advise them to start the convergence journey, and how should they build trust along the way?

SUBHALAKSHMI GANAPATHY: Start with a shared incident, not a shared org chart. Pick a high-impact scenario such as ransomware or a major outage and run a joint tabletop exercise where both teams work from the same data, timeline, and playbook. This quickly surfaces the real points of friction, including visibility gaps, conflicting priorities, and unclear hand-offs.

Compliance can also provide a practical starting point. A shared audit-logging initiative or a unified breach notification playbook aligned with DPDP, GDPR, or CERT-In requirements gives both teams a concrete reason to work together, often with executive sponsorship already in place.

From there, integrate the tools teams already use so they can access the same underlying context. For example, integrating a SIEM with an ITOps observability platform can ensure that a security alert carries relevant performance context. This does not necessarily require a single platform; well-integrated systems can work equally effectively. The next step is to establish a shared incident-command model with clear roles.

Trust is built through small, visible wins like a jointly resolved incident, a shared post-mortem, or a KPI that both teams own. Structure should follow shared work, not the other way around.

CIO&Leader: How is AI changing the speed and shape of incident response when IT and security teams operate on the same data and tooling, and how does it affect trust in the response?

SUBHALAKSHMI GANAPATHY: AI becomes more effective when it can learn from richer and more connected data. When IT and security operate on the same telemetry, operational and security signals can be analyzed together rather than as separate events. An anomalous login, a latency spike, and an unusual data transfer, for example, can move from being three separate alerts in different queues to becoming one coherent narrative that AI can reason about. This can significantly compress incident response times by reducing the cross-team coordination that would previously have been required. Trust in AI-driven response, however, depends on explainability. IT and security teams need to understand why something was flagged, what data informed the recommendation, and how confident the system is in its recommendation. The risk is over-reliance. AI should accelerate human judgment, not replace it, especially for actions affecting production systems or customer trust.



"Start with a shared incident, not a shared org chart."

CIO&Leader: What role does automation play in reducing mean time to detect and mean time to respond once teams are unified, and where should automation stop and human judgment take over to preserve trust?

SUBHALAKSHMI GANAPATHY: Automation is particularly valuable in detection and the initial stages of response. It can handle enrichment and correlation, isolate an endpoint, revoke a session, or open a ticket with the relevant context. These capa-

bilities can help bring MTTD and MTTR down from hours to minutes. When teams are unified, automation can extend across the incident lifecycle because handoff friction is reduced. However, automation should stop where consequences become irreversible or require significant business context, for example, taking down production systems, communicating with regulators or customers, attributing intent, or determining acceptable business risk. Those need human judgment, ideally from IT and security together. The principle should be simple: automate the known, escalate the novel. Trust in automation is built through transparent playbooks, clear rollback paths, and regular reviews of the system's autonomous actions. Hence, humans stay accountable even when they are not in every loop.

CIO&Leader: Culturally, what's the hardest part of breaking down the IT-security silo: is it tooling, reporting lines, incentives, trust, or something else?

SUBHALAKSHMI GANAPATHY: It is more about mindset and identity than tooling or reporting lines. IT teams typically see themselves as builders and enablers, measured on delivery and uptime, while security teams see themselves as guardians, measured on risk reduction and control. Those perspectives can come into conflict, particularly when a security action could affect a business-critical system. Until leadership creates a shared definition of what "good" looks like for both functions, teams can default to their own priorities. Trust is the second major hurdle. Security has historically been viewed as the "no" function, while IT has sometimes found ways to work around it. Rebuilding that trust takes time, joint wins, and leadership that visibly supports collaboration. Tooling is the easiest part; it can be procured. Reporting lines are a lever, but they are not a solution on their own; even organizations with unified CIO-CISO structures can operate in silos. The real work is redefining a shared purpose, i.e., resilience.

CIO&Leader: How should CIOs and CISOs redesign accountability and KPIs so resilience becomes a shared metric, reinforcing trust rather than creating two separate scorecards?

SUBHALAKSHMI GANAPATHY: Resilience is not a security-only concern; it belongs equally to ITOps. A Full Stack Observability (FSO) view makes this obvious. The same telemetry that shows a service is degraded can help determine whether the cause is a configuration drift, a capacity issue, or a security compromise. KPIs should therefore reflect shared ownership. This could include time to restore critical services (regardless of cause), the percentage of incidents resolved without customer impact, mean time to detect and recover across both cyber and operational events, and a composite resilience score covering redundancy, recovery, and response maturity. Compliance metrics can also be included in this shared scorecard, such as audit readiness, control effectiveness, and breach notification timelines under DPDP, GDPR, or sector-specific mandates. These are inherently cross-functional and give both leaders a board-level reason to co-own the outcome. Individual KPIs like patch compliance and uptime will continue to matter, but they should ultimately contribute to shared outcomes. Two scorecards can reinforce two separate priorities; a shared scorecard helps establish one resilience mission.

CIO&Leader: As enterprise environments grow more distributed, multi-cloud, hybrid work, and IoT/OT convergence, how do unified IT-security operations need to evolve to keep pace?

SUBHALAKSHMI GANAPATHY: The traditional perimeter is disappearing, and so is the notion that any single team can own a technology domain end-to-end. With workloads distributed across multi-cloud, hybrid, and IoT/OT environments, the volume and variety of telemetry data have increased significantly. The only way to keep pace is to consolidate that telemetry into a common data lake that both observability and security teams can draw from. When performance metrics, logs, traces, network flows, identity events, and threat signals all land in the same repository, IT and security spend less time debating whose data is correct and more time responding to what data shows.

This shared foundation is what makes distributed unified operations viable at scale. Instead of stitching together fragmented views from each cloud, region, or edge location, teams get a single correlated picture regardless of where the workload runs. For example, a latency issue in a regional cloud environment and a suspicious login from an OT device on a factory floor are investigated together when the relevant data is connected. Observability platforms provide the operational view and security platforms provide the threat view, but both can work from the same underlying facts. It also addresses the cost and complexity of duplicating telemetry across separate technology stacks. When combined with cloud-native tooling, identity as a control plane, and OT-aware sensors, a shared data foundation can become the back-



"The principle should be simple: automate the known, escalate the novel."

bone of unified operations across a highly distributed enterprise.

CIO&Leader: What's a real-world example (anonymize if needed) where unified IT-security operations meaningfully changed the outcome of an incident and strengthened trust, versus how a siloed response might have played out?

SUBHALAKSHMI GANAPATHY: Consider a generalized scenario involving a mid-sized financial services organization. Its monitoring team detects intermittent latency in a customer-facing application, while the security team, working from the same platform, simultaneously identifies unusual outbound connections from the same server cluster. Because both signals are visible with shared context, the teams can connect them within minutes and identify that the latency is linked to a compromised process attempting to exfiltrate data. The affected systems can then be isolated, the exfiltration blocked, and forensic investigation initiated immediately. In a siloed model, the IT team might initially treat the latency as a performance issue and restart services or scale capacity, potentially giving the attacker more time and affecting the availability of evidence. Security, meanwhile, may lack the operational context needed to link the anomaly to the application issue. The key point is that unified visibility does more than shorten response times. It can change how an incident is understood and enable IT and security teams to respond to the underlying issue rather than treating its symptoms separately.

CIO&Leader: Looking ahead, what will "resilient enterprise operations" look like in 3-5 years, and what should organizations be building toward today to strengthen trust and resilience?

SUBHALAKSHMI GANAPATHY: In the next three to five years, resilience will be defined not only by prevention but by adaptive recovery, meaning the ability to absorb disruption, whether caused by an attack, an outage, or a supplier failure, while continuing to serve customers. We can expect AI-driven autonomous operations to handle more routine detection, response, and remediation; digital twins of critical business services to help organizations simulate and rehearse disruptions; and identity-centric, zero-trust architectures to become increasingly standard. Regulatory expectations around AI governance, critical infrastructure, and sector-specific requirements are also likely to make unified operations important for compliance. IT and security may not become a single function in every organization, but the direction is toward a single resilience discipline with specialized expertise. Organizations can start building toward this today by investing in a unified data platform, establishing shared incidents, resilience, and compliance KPIs, upskilling teams in AI and cloud-native operations, and conducting joint resilience exercises across cyber, operational, and business continuity scenarios. Above all, organizations need to build trust in their teams, their data, and their automation. That is the foundation on which resilient operations will be built. ■



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

RNI No. DELENG/2012/42736

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